

The History of Computing at Cornell University

By John W. Rudan

Director Emeritus, Cornell Information Technologies
Cornell University

The History of Computing at Cornell University

By John W. Rudan

Director Emeritus, Cornell Information Technologies
Cornell University

Edited and designed by Cornell Communication and Marketing Services

Published by
The Internet-First University Press

This manuscript is among the initial offerings being published as part of a new approach to scholarly publishing.

The manuscript is freely available from the Internet-First University Press repository within DSpace
at Cornell University at

<http://dspace.library.cornell.edu/handle/1813/62>

The online version of this work is available on an open access basis, without fees or restrictions on personal use.

A professionally printed and bound version may be purchased through Cornell Business Services by contacting:
digital@cornell.edu

All mass reproduction, even for educational or not-for-profit use, requires permission and license.
For more information, please contact dcaps@cornell.edu. We will provide a downloadable version
of this document from the Internet-First University Press.

Ithaca, N.Y.
June 2005

The History of Computing at Cornell University

By John W. Rudan

Director Emeritus, Cornell Information Technologies
Cornell University

Introduction

I was encouraged to compile the story of how computing started, evolved, and grew at Cornell. My first association with computing at Cornell was in the spring of 1959, just six years after the first digital computer was installed on the campus. This association continued until 1996, except for about six years when I held different positions or undertook further studies. Over the decades I accumulated over 50 boxes of records, pictures, and other memorabilia related to computing at Cornell. As I approached retirement, documenting this history seemed like a great project, letting me preserve information that might otherwise be lost while also staying connected with computing technologists and former associates.

This history draws on not only the Computing at Cornell archive that I have assembled but also articles in the *Cornell Chronicle*, published reports including annual reports from different organizations, and reports and other documents. Personal communications with individuals who participated in many of the events documented were often obtained using e-mail, but the recorded personal stories were invaluable in providing facts and anecdotes worth mentioning.

It was easy to write about computing at Cornell up to 1969, when there were no more than a half dozen digital computers on campus. By the end of the 1970s, a decade later, there were several hundred computers on campus, and in another decade perhaps several thousand local area network servers and personal workstations. With the explosion in the number of computers on campus came an explosion in the number of organizational entities that were formed to support them. Since it was impossible to document all these entities thoroughly, this history focuses on the “computing center,” the central computing agency at

Cornell, which has focused on campuswide computing and supporting technologies. However, the organizations that were closely allied with the central organization in staff or in computers or other technology are considered part of the story.

This history is organized by decades, that is the 1950s, 1960s, etc., except for the pre-1950 period, which is given special treatment as it documents as much as could be found about what took place at Cornell for the years up to 1949. The first section and succeeding decades begin with an overview of the most important developments in computing/information technology (IT) during that decade. This overview is intended to provide some insights into the evolution and status of computing/IT in general, as well as what was happening (or not happening) at Cornell. Each overview will also define some relevant acronyms or buzz words from that decade.

The collection of an oral computing history started partway through writing this story when sources were sketchy and incomplete and I knew of the key person or persons who were involved with the particular issue. Other individuals who were aware of the project wrote their own stories and sent them to me. I have now accumulated a large number of these personal stories and recollections about each person’s experience during their career at Cornell. The names of all those individuals who provided information this way are listed in the Acknowledgments. I thank them all not only for the factual information and publications they provided but also for the very interesting personal comments about incidents, some of which are included in the endnotes of the relevant decade. Readers are encouraged to go to the Cornell Information Technologies (CIT) web site, www.cit.cornell.edu/computer/history/, for access to most of these documents.

Giving witness to the fact that counting, enumerating, and calculating (computing) have been important to Cornell for a very long time, Walter T. Federer, Liberty Hyde Bailey Professor Emeritus and early supporter and user of computing at Cornell, makes the following statement:¹

Professor Harry H. Love kept pace with computing in the twenties and thirties. He purchased a computer called a Millionaire, supposedly hot stuff for the time. I believe that it is stored somewhere in Plant Breeding. A display of all the old machines would be illuminating to the students of today. When I arrived in 1948, Dr. Love had made certain that we had many Monroe, Friden, and Marchant calculators around. There were rooms of these machines with one being in Warren Hall under the dictatorship of Onnie Zaharis.

The means to accomplish all this storage, processing, and analysis of information continually improved. The punch card technology came to Cornell in the 1920s and improved the ability of researchers to count and calculate and summarize the results of experiments and other observed data. Until the 1960s laboratories of tabletop calculators were all over campus where numerical calculations formed the basis for problem exercises and examinations in class work. Engineering and the disciplines which relied on statistical analyses were key in this regard. I shall try to pursue some of the important early developments, but this history will focus more on the “digital” computer technology and how it grew and developed.

To put in perspective the profound changes that have occurred in the almost 50 years since computing got its start at Cornell, it is only necessary to give a few examples. In 1953 there was one computer on campus; 50 years later there are possibly as many as 30,000 computers during the school year when the students are on campus. Not only that, but each of these current computers has several orders of magnitude more “power” than that first computer and those of the early years. For example, the IBM 650 on campus in 1956 was rated at 0.001 MIPS (million of instructions per second), while the first Apple Macintoshes in 1984 using a Motorola 68000 chip were rated at 2 MIPS. In the early 1990s, when the IBM or IBM compatible microcomputers were introduced on campus using Intel Pentium chips, the systems were rated at over 100 MIPS!

Richard C. Lesser, first director of the Cornell Computing Center, provides an appropriate closing comment, which expresses my views as well:²

I hope that these remembrances will be of interest to Cornellians, past and present (and others as well), and that my memory has not admitted of too many errors or omissions. These early years were ones of trial, error, discovery, challenge, and fulfillment, and the problems we faced and solved (with limited tools) would be incomprehensible to those who practice computing today.

Acknowledgments

I would like to thank both H. David Lambert, vice president of information technology at the university (1995 to 1997), for offering his strong support and encouragement to start this project, and Polley A. McClure, vice president of information technology (1999 on), for continuing this support. I also thank Professor J. Robert Cooke for his support to publish this history in the Internet First University Press (D-Space at Cornell University).

I also wish to acknowledge the continuing and valuable help of Shayle R. Searle, whose generous comments and encouragement kept me forging ahead. Special thanks go to Richard C. Lesser, Dominic Bordonaro, and Lydell Wadell. Lesser wrote about the formation and early years of the Cornell Computing Center along with important pictures and documents. In addition to contributing to the Oral History, Bordonaro provided written commentary and documentation about Machine Records, the first business systems unit at the university. Wadell provided the history of the Northeast Dairy Records Processing Center, which supplemented other historical information about one of the first computer installations on campus. All these documents, along with other relevant documentation, can be found on the CIT web site.

The following individuals contributed oral histories or their personal stories and other materials that were of importance to compiling this manuscript. I thank them for all their contributions: Dan Argetsinger; Mark H. Anbinder, Ralph Barnard; Dominic Bordonaro; Scott W. Brim; G. V. Chester; Richard C. Cogger; Richard W. Conway; J. Robert Cooke; Dale R. Corson; Cecilia Cowles; Cathy Dove; James F. Doolittle; Tom A. Dimock; Agelia Dumas; Charles V.

¹ Federer, Walter T., “Walter T. Federer’s Recollections, 1948 to present,” History of Computing at Cornell, www.cit.cornell.edu/computer/history/, 2002.

² Richard C. Lesser, Richard C. “Richard C. Lesser’s Recollections, 1953 to 1964, History of Computing at Cornell, www.cit.cornell.edu/computer/history/, 2002.

Evans; Walter T. Federer; Larry Fresinski; Doug Gale; Peter Hoyt; Errol W. Jones; Keith W. Kennedy; Ken M. King; David W. Koehler; Richard C. Lesser; Bruce Lloyd; M. Stuart Lynn; Bob Mack; James F. Manning; Mark Mara; John McKeown; Eric Mintz; Pat Nelson (Searles); H. Michael Newman; Mike Oltz; John S. Ostrom; Ron Parks; Alan Personius; Lynne Personius; Carrie Regenstien; Ruth Sabean; Sid Saltzman; Louise Schulden; Clint Sidle; Pete Siegel; Donna Taber and Kevin Leonard; Douglas E. Van Houweling; Irene Van Zile; Henry G. Vaughan; M. Scott Walters; Steve L. Worona; Joan Winters; Paul Zarnowski; and Gene Zeigler. The help of Simeon Moss and Bill Steele in retrieving information from the Cornell Chronicle archives is appreciated.

I wish to acknowledge the help of Mark Anbinder, Dan Bartholomew, Bob Blackmun, Mark Bodenstein, Dick Cogger, Bob Cowles, Tom Dimock, and Rick Jones, who fielded many questions on technology and events throughout the preparation of this manuscript. I also thank Marcia Poulsen (CIT) for her excellent suggestions for improving the manuscript in both form and substance, and the staff from Communication and Marketing Services, notably Judy Stewart, Susan Pohl, and Barbara Drogo, who provided excellent editing and graphic design in preparing the final book (Dspace) document. I'm also grateful to the many other friends and associates who have encouraged me in this effort and who also took the time to answer questions and make comments about their own experiences. I apologize in advance if I have failed to acknowledge the assistance of anyone and for any errors in omission or fact. That is my fault entirely.

Dedication

This history is dedicated to all the staff who worked in the central computing or information technology organizations over the years—from the Cornell Computing Center and Machine Records in the 1940s and 1950s until today's Cornell Information Technologies. This includes clerks, operators, programmers, analysts, supervisors, managers, directors, and vice presidents, whose hard work and dedication resulted in Cornell's achieving widespread and intelligent use of computers on campus for instruction, research, and administration and also for achieving the high regard of the computer industry and higher education for developing innovative solutions to help students, faculty, and staff use this technology to advantage.

I personally dedicate this work to my mother, Helen (Anhel) Rudan, for her strong belief that education opened doors and her persistence that I should go on to university studies.

I also dedicate this work to my wife, Doreen J.

Rudan, for her encouragement and support not only during my career at Cornell University but also during the preparation of this manuscript.

Lastly, I dedicate this work to my son John, daughter Tessa, and son Michael and his wife, Brenda, and grandchildren Megan, James-Paul, and William Rudan.

The Years up to 1949

Contents

The Years up to 1949—Industry Overview	4
Hardware.	4
<i>Punch Cards and Equipment</i>	4
<i>The First Digital Computers</i>	5
The Years up to 1949 at Cornell	6
1940 and Earlier Years	6
<i>Record Systems on Campus</i>	6
Punch Card Installations on Campus	6
<i>Department of Agricultural Economics</i>	6
<i>Machine Records (Business Data Processing)</i>	7
<i>Dairy Records Processing Laboratory</i>	7
<i>Summary Comments</i>	7

Figures

Figure 1. IBM Punch Card	4
------------------------------------	---

The Years up to 1949—Industry Overview

The following is a brief description of the major punch card technology and commercial developments up to 1949.

Hardware

Punch Cards and Equipment

Punch cards were the basic medium for recording information and for subsequent processing by a variety of equipment that “automated” certain operations and facilitated faster and more accurate processing of the information. Herman Hollerith invented the punch card and first processing equipment in 1888 to improve the speed of processing the 1890 U.S. Census.¹ In addition to improving speed, the use of this technology also improved the accuracy of the census figures. From that time forward, the Hollerith Company produced the cards and equipment for commercial and other use until 1928, when IBM, the International Business Machine Company, acquired it.

Typically, 80-column punch cards from IBM (also referred to as unit records) were laid out into “fields”

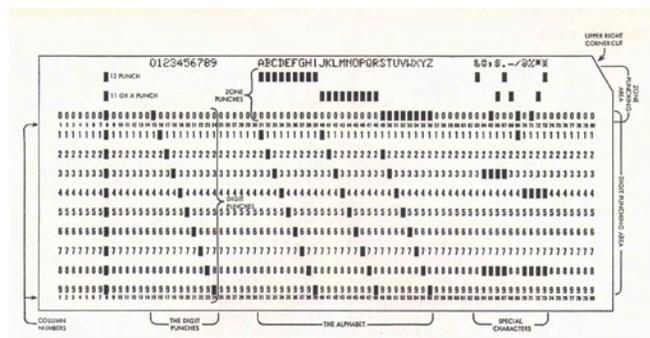


Figure 1. IBM Punch Card

of a fixed number of columns and defined format, e.g., an implied decimal point.

Each column had the capacity to record 12 punches—the digits 0 to 9 and the “11” and “12” zone punches. As shown above, alphabetic and special characters, often referred to as alphanumeric information, were coded by combining the numeric digits with the zone punches and additional digits in a manner that preserved the collating sequence so that on sorting, A came before B came before C, etc. This 48-character coding scheme became known as Hollerith code.

Verifying information that was transcribed onto cards was a feature of those early days, especially for business or otherwise critical information. It was often the case that data being punched onto cards was not

¹ Throughout this document, technical terms are based on definitions from the Encyclopedia of Computer Sciences, 4th edition, or from Wikipedia at www.wikipedia.com.

totally legible because the original document was handwritten, possibly out in some agricultural field or on the factory floor, or because it was being read from smudgy carbon copies of printed material. Prior to 1970, carbon paper inserted between preprinted forms or between sheets of paper was the common method of making copies of business or other information. Often the paper rollers that kept the paper in place left their own impression on the copies, especially if the movement was not perfectly synchronized, or the person writing on the forms was careless. In most cases the transcription of hand-prepared data to cards was accepted as correct on the assumption that the keypunch operators had an acceptably low error rate. When there was concern that the error rate of transcription could make the results suspect or cause significant errors or major repercussions, an additional verification operation was performed. Most often, the verification was visual, where possibly another person compared the print across the top of the card with the original source document. Alternatively, two separate decks were prepared by two different operators and compared for similarity using a collator.

Later on IBM developed special equipment called a *verifier*, which looked like a keypunch but which only sensed the holes instead of creating them. As a result, when a clerk “verified” a deck of cards, the cards were put into the machine and processed in the same order as the initial deck had been punched. When a column mismatch was detected, the verifier put a v-shaped notch on the top of the card to indicate a mismatch in that column. When a card was verified without error, the machine cut a half circle notch on the side of the card. Consequently, after a deck was verified, examination of the side of the deck revealed the cards with errors, which then were examined for the column with the mismatch and corrected by creating a new substitute card.

Using a *sorter*, which sensed the punches in a column set by the operator, cards were sorted into bins by multiple passes over any field, and the sorter equipment automatically counted the number going into each separate bin. Accordingly, for a pure numeric sort, there was a count in all the 10 bins. By moving across a field, one could then sort the cards by the entire field.

Card decks could be merged together using a *collator*, which fed cards from two different feed slots as a given field continued to match or not. The collator was programmed by use of a special board that, when wired, activated the built-in instructions. Such decks were then processed by a *tabulator/printer* to accumulate various sums and counts and to obtain printouts of the results. When connected to a *card-punching* unit, the tabulator also could be made to punch a new set

of master summary cards of accumulated totals for the next cycle of updating the records.

Programming of the tabulators and other equipment having built-in logic was done by wiring boards that performed different functions on selected fields and created the counts or sums and other calculations. Thus one tabulator, for example, could be used for different work by inserting the board programmed for that application. The front of the board was fixed with a particular pin configuration that made contact with the corresponding pins in the machine when the board was inserted in its slot and pushed to engage the pins. The back of the board was typically a maze of wiring—the more complex the counts and calculations, the greater the maze of wires. Programming as such followed much the same steps of programming and debugging that were used later in software-based programming. Covers were supposed to be screwed onto the back of the board once the board wiring was completed, to act as a security measure as well as to discourage on-the-fly programming changes. However, more often than not the wiring was left exposed and changes were made, creating many of the same problems that appeared later when production programs were under the control of the application programmer.

The collection of sorters, collators, the tabulator, and associated equipment became known as the *tab shop*, because the tabulator was the key machine, producing the end result, a tabulation and reports. Due to the way holes in the cards were sensed in the different types of equipment, for each type, cards had to be placed in a particular way in the input hopper. At the hopper there was usually a sign saying “12 edge, face down” or “9 edge, face up,” and failing to follow this practice would stop the machine or produce incorrect results. (The joke at the time about the burial of Tom Watson, founder of IBM, was that he was buried “12 edge, face up.”) The keypunches and verifiers were typically located in a keypunch room/section and referred to as such. Later, this operation became known as data entry as different modes and media were used for recording information, for example, entering data directly onto magnetic tape.

The First Digital Computers

There were several notable achievements during this decade to build a digital computer. Extensive histories were written about those first machines, and what follows are short summaries.²

Charles Babbage is generally credited with first proposing a digital computing device with his Difference

² Scott McCartney, *ENIAC: The Triumphs and Tragedies of the World's First Computer* (New York: Penguin Group, 2001).

Engine, built on ideas from the Jacquard Loom. This took place between 1820 and 1830.

Howard Aiken from Harvard, working with IBM during the 1940s, attempted to implement Babbage's machine using electromagnetic relays. His implementation was named the Mark I, and while it operated successfully, it was so slow that it was deemed a technical dead end and was never developed further.

ENIAC (electronic numerical integrator and computer) was built by John Mauchly and Presper Eckert at the University of Pennsylvania, and it was the first to use electronic tubes. The first unit was finished in the late 1940s and formed the basis for the later commercial version introduced as the Univac (universal automatic computer) in the 1950s.

John Atanasoff at the University of Iowa, working in the 1940s, also developed electronic circuits for calculating machines and built his own computer. There is still controversy today as to who was the first person to build the first digital computer.

It is interesting to note that the work of Aiken, Mauchly and Eckert, and Atanasoff was undertaken as part of the World War II effort to improve ballistics calculations and was funded by the U.S. military.

The Years up to 1949 at Cornell

The use of punch cards and the related processing equipment first started at Cornell in the early 1920s and was increasingly used for record keeping and data analysis up to 1949.

1940 and Earlier Years

Record Systems on Campus

In addition to the introduction and use of punch cards on campus, Dominic Bordonaro, one of the early employees of Cornell involved with punch card and computing equipment, recalls the existence of a variety of card-based equipment in many departments across the campus.³ There were McBee Keysort cards, which were punched and could be selected by sticking a long ice-pick-like device through a column. Dennison Tag Systems made another similar product. Then there was the Remington Rand (later Univac) Cardex System. All these systems could search and select records meeting certain criteria. Bordonaro recalls, but is not sure, that one or more departments on campus had Remington Rand card-punching and sorting, but not tabulating, equipment during this time. While all this sounds rather mundane in the 21st century, at that time punch card technology was so much more powerful than the older methods that expansion of its use was considered inevitable. How this took place at Cornell will be discussed in turn.

Punch Card Installations on Campus

Department of Agricultural Economics

The first known organized effort on campus to automate punch card counting and calculation technology started in the 1920s when Hollerith Company equipment was installed by the Department of Agricultural Economics in the old temporary building in which they were located at the time. (In 2001 Agricultural Economics became the Department of Applied Economics and Management after several intermediate name changes.) The 1957 report on data processing and computing at Cornell describes this founding as follows:⁴

The first punch card tabulating equipment was brought to Cornell by the Department of Agricultural Economics in the middle 1920s. Since that time, this department has continuously rented and operated a variety of tabulating machines that have kept pace with its growing needs and ever-improving state of the art. Records of the Department of Economics of the Household show that it contributed to the support of the machines at least as far back as 1926; the records of the Department of Rural Sociology show the same back to at least 1929....

From the beginning there was paid help in the form of a graduate student, though in about 1939 this was supplemented with the assistance of a full-time statistical clerk. When the last graduate student supervisor left in 1942, the statistical clerk assumed supervision. She remained until 1945, when she resigned

It is most likely that the equipment was IBM equipment, although that is not certain. According to Bernard F. Stanton, retired professor of agricultural economics (personal communication, November 1997), this installation ran many tabulations for surveys and studies of prices and economic index numbers during the 1920s. The basic work of computing the federal government wholesale price index was done in this facility. When Warren Hall was built in 1933, a large part of the basement was used for this punch card processing tab shop. By the late 1940s this shop was continuing to do work for the

³ Dominic Bordonaro, "Dominic Bordonaro, Oral History, 1947 to 1968," History of Computing at Cornell, www.cit.cornell.edu/computer/history/, 2002.

⁴ W. S. Gere, B. McK. Johnson, and R. C. Morris, "An Analysis of Data Processing and Computing at Cornell University" (Cornell University, November 14, 1957, copy number 10)

School of Industrial and Labor Relations (ILR) and the Departments of Agricultural Economics, Rural Sociology, and Economics of the Household, with funding from those units. At about this time, Alma Coles was one of the supervisors of the facility, and her husband, Theron Coles, was one of the IBM customer engineers who maintained the equipment.

Machine Records (Business Data Processing).

In 1946, the Machine Records facility was organized as a tab shop to use punch cards and the associated equipment for administrative record purposes. The 1957 report on data processing and computing at Cornell provides this brief introduction:

In late 1946 at the request of the Registrar, the Machine Records Unit was opened in Day Hall. The Registrar was responsible for its operation. Its first assignment was to process student grade and registration information after first transcribing it from paper to punch cards. As time went on, its operations grew and a variety of work was successfully absorbed.

The beginning of the first administrative information (business data processing) systems unit had a typical Cornell flavor—it was organized as a fee-for-service operation. A number of forward-thinking administrative departments with university-wide responsibilities were faced with increasing numbers of business transactions and record keeping and were looking for ways to improve their operations. They were aware of the advantages being touted for punch cards and the associated equipment that used them for faster processing and accumulations, and they decided to pursue using that technology.

Bordonaro, who joined the university in 1948, confirms the formation of Machine Records in October 1946 to process student records.⁵ However, he believes one of the first applications was also for basic payroll preparation. He recalls that Robert Burghardt was the first manager of the facility, who reported to the registrar.

The location of the first offices for Machine Records remains a mystery. However, after Day Hall was built in 1947, Machine Records was located in the basement and remained there for about 20 years. The first equipment for Machine Records was all from IBM and in 1946 consisted of two keypunches, a sorter, a collator, and a 405 tabulator/printer. This equipment was rented from IBM at a total cost of \$500 per month. By 1951–52, the equipment had grown to three keypunches, two verifiers, two sorters, one collator, one interpreter, one reproducer, one 602A calculator, and two 405 tabulators.

As noted previously, the initial customers of Machine Records were the Registrar's Office and Endowed Payroll. In 1946–47, a pilot program was established to process the statutory colleges' payroll. The success of this program led to the processing of the main payroll for the statutory colleges in 1948. In 1948 there was also a pilot program in the Admissions Office. History is scant after that.

Dairy Records Processing Laboratory

The third documented installation of punch card equipment on the campus was in the College of Agriculture. The Department of Animal Husbandry (later Animal Science) started the Dairy Records Processing Laboratory (later changed to Cornell Dairy Records) in 1947. The laboratory evolved from the Dairy Herd Improvement Association, which took an interest in improving the record keeping of, and in turn the ability to improve, milk production from dairy herds. The laboratory processed herd-test records for individual farms. For each cow in a herd in the record-keeping program, a sample day's milk production was converted to a monthly basis and was accumulated to give a lactation-to-date yield for the month. The percentage of fat in the sample day's milk was used to provide monthly and lactation-to-date yields of butterfat (now referred to as milk fat).

The following historical description of the beginning of the laboratory is taken from Kenneth L. Turk's history of animal husbandry at Cornell.⁶

At that time the first IBM equipment was brought in by the Department of Animal Husbandry to establish the dairy record processing laboratory (DRPL). The initial pieces of equipment were a card punch, sorter, tabulator, and printer, which were lifted through a window to the attic of Wing Hall. Soon it was apparent the space was inadequate, and the processing equipment was moved to the NYABC (New York Artificial Breeders Cooperative) headquarters on Judd Falls Road. The amount of equipment and personnel involved continued to expand, and after a few years the laboratory was moved into one of the old workshop buildings formerly used by the Department of Agricultural Engineering.

⁵ Dominic Bordonaro, "Dominic Bordonaro, Oral History, 1947 to 1968," History of Computing at Cornell, www.cit.cornell.edu/computer/history/, 2002.

⁶ Kenneth L. Turk, *Animal Husbandry at Cornell University—A History and Record of Development from 1868 to 1963*, (Ithaca, N.Y.: College of Agriculture and Life Sciences at Cornell University, 1988).

Computations that were required as part of this effort were done by using the equipment in the Machine Records operation in Day Hall.

Summary Comments

In 1949 there were three known tab shops on campus, two in the College of Agriculture, primarily concerned with processing records but whose value was in the research that could be derived from the collected information, and one concerned with processing business information records.

There was no digital computer on campus.

1950 to 1959

The First Decade of Digital Computing

Contents

1950 to 1959—Industry Overview.	10
Hardware.	10
Software	10
Services	11
Vendors	11
Technology Terms, Acronyms, and Buzz Words	12
1950 to 1959 at Cornell	12
The Cornell Computing Center	12
<i>The Founding of the Computing Center in 1953</i>	
—Richard C. Lesser, Director	12
<i>Research Computing</i>	16
<i>Instructional Computing.</i>	16
<i>The Computing Center in Rand Hall</i>	16
<i>User Groups</i>	17
Business Systems.	18
<i>Machine Records—Dominic Bordonaro, Director</i>	18
<i>Statutory Finance and Business Office.</i>	19
Other Computing Installations	20
<i>Dairy Records Processing Laboratory.</i>	20
<i>Agricultural Economics</i>	21
Status of Computing on Campus	21
Summary	23

Figures

Figure 1. Basic elements of a digital computer	10
Figure 2. Professor R. J. Walker and Dick Lesser at the CPC in Rand Hall	13
Figure 3. Dick Lesser seated at the IBM 650 console with Ann Walbran and Prof. R. J. Walker in Phillips Hall	14
Figure 4. Computer demonstration for Engineering faculty	16
Figure 5. Burroughs 220 computing system in Rand Hall	17
Figure 6. John Rudan and Helen Church at the Burroughs 220 console	17

1950 to 1959—Industry Overview

The following is a brief description of the major hardware, software, and service innovations in the 1950–59 decade when the first commercial computers became available, as well as how different vendors entered and left the market in these areas.¹

Hardware

In the decade from 1950 to 1959, hardware was primarily noted for being big and bulky and requiring lots of air conditioning to dissipate the heat given off by vacuum tube–based components. For example, the Burroughs 220 Datatron unit—which was only the interface between a 240-card-per-minute reader and a 150-line-per-minute printer, was contained in a cabinet that was 12 feet long, 5 feet high, and 3 feet wide. Inside were hundreds of replaceable cards, each 8 inches deep, 6 inches high, and 2 inches wide.

One spoke of instruction times in milliseconds; one talked of mean time to failure in hours. It was not uncommon to have several outages per day, which could vary from minutes to hours and could be caused by hardware failures most of the time, but sometimes by software failures. Neither was easy to detect because the diagnostic tools were not very advanced.

The basic design of the stored program digital computer model, which is the same today, is shown in Figure 1.

This formulation of the digital computer is attributed to John von Neumann, a mathematician from Princeton University who proposed storing the computer instructions with the data in memory and providing an instruction decoder for executing the stored instructions. In later years, papers from the development of the ENIAC suggested that Mauchly and Eckert first proposed this idea to overcome the hardware programming needed for their machine. However, von Neumann was the first to publish a paper.

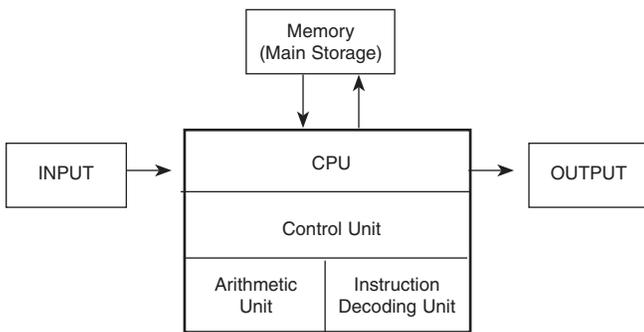


Figure 1. Basic elements of a digital computer

Determining who was first with the proposal is even more controversial in that an argument can be made for John Atanasoff at Iowa State.²

In the 1950s, computer memory was notably small. Different strategies were advanced to improve the speed and capacity of memory. The cathode ray relay gun was used to circulate bits through a long tube, where they were read and written as they passed through the read/write head at one end of the tube. Magnetic drums were another memory medium, with the surface of the drum partitioned into a matrix of rows and columns, each containing a fixed-size word and having as the address the row and column number. A feature of these early devices was a destructive “read” of memory that had to be followed by a “restore,” thus increasing the memory access time. One typically thought of the number of distinct addresses as being in the hundreds, and in the thousands toward the end of the decade. Internal number representations varied from decimal to biquinary to binary.

Input and output devices were crude and relied on the available punch card technology. There was even competition for different styles of cards, with IBM, for example, offering the 80-column card with rectangular holes and Univac offering the 90-column card with round holes. At the time, however, there was a growing market in card equipment since it was heavily used in commercial and census-type activities to record, count, and accumulate different characteristics, which were then coded and punched onto the card^a.

Magnetic tape technology was slowly emerging and evolving, and different vendors offered alternatives. One of its main uses was to expand main memory by acting as a temporary buffer or storage for intermediate results. The fact that the technology was slow and data was stored sequentially mattered little compared with the alternatives. Input console keyboards were of the Teletype variety, which was a proven but not very user-friendly technology. Paper tape, which came as a free add-on of the teletype technology, was initially looked upon as a good alternative to cards but gradually evolved into being used primarily as a backup device to store completed programs.

Software

In contrast to hardware, software referred to all the nonhardware components, that is, the reloadable programs that made the hardware execute its built-in capabilities. In this period, programming the machines

¹ Throughout this document, technical terms are based on definitions from the Encyclopedia of Computer Sciences, 4th edition, or from Webopedia at www.webopedia.com.

² Scott McCartney, *ENIAC: The Triumphs and Tragedies of the World's First Computer* (New York: Penguin Group, 2001).

was done at the basic machine level using the defined operator, operand, and command succession structure built into the control unit, arithmetic unit, and instruction decoding unit that the central processing unit comprised. Typically, there was some bootstrap command that initiated a read of the instructions from a card reader into a defined area of memory, and after reading, a specially coded last card would begin executing instructions from the defined memory location. This was clearly a craft industry, and the greatest concern of programmers was to fit the program and data into memory. The introduction of assembler programs that allowed mnemonic operation codes and symbolic addresses greatly advanced the productivity of programmers. The assembler read the almost-English-language program punched onto the cards, substituted machine level operations, and assigned variable addresses to specific storage locations. When the assembly was complete, the assembler produced a list of programming errors, such as spelling errors and other inconsistencies, so that the program could be easily debugged by removing them.

Most of the programming was one-on-one—you needed a job done, you wrote a program, and you thought later of making it production quality if it was to be used repeatedly. There were almost no applications packages as we think of them today. Further, with punch card technology, a source program—that is, the original program—consisted of a deck of cards, with each card representing a line of code. Provision was made to include a sequence number, but this often was ignored—and sorrowfully regretted if the deck was accidentally dropped and the program had to be reconstructed! Programs were verified by eye when keypunches starting printing the characters at the top of the card, or by running them through a tabulator to print a program listing for examination. Source code decks in machine language were read by the system, and the program was executed directly. Source code decks in assembly language were first converted to object code, that is, machine language, by an assembly program and then executed.

Fortran (formula translation), the first high-level programming language, was defined by John Backus and a group from IBM in 1953³. Fortran allowed programs to be written in an algebraic notation, for example $A = B + C$, where again there was more freedom than was available in assembly languages to define operations on variables without worry about memory allocation, etc. A Fortran program was then run through another program called a compiler, which translated the Fortran language into basic machine language. Depending on the capabilities of the com-

puter being used, this translation could be direct or involve an intermediate step of creating assembler code and then running that through the assembler. Although attempts were made to create tight code, that is, object code equivalent to that produced by a top-flight programmer, it was understood that Fortran code generally executed more slowly than direct programmer-produced code. Nonetheless, the use of Fortran greatly advanced the productivity and use of computers by scientists in all fields and, in effect, started the tradeoff between using computers for producing results directly and using them to assist in producing the programs and the results too.

Services

One person running one job took up the entire machine, and so one could realistically think of billing by the hour. Simple batch might be the short phrase that described this style of service. Debugging a program involved repeated attempts to run the program to completion and then read the console or selected memory locations and commands to diagnose the error. It was not uncommon for programmers to spend many hours or all night at the console reading registers, re-initializing the program, and executing another run. When that failed, the last resort was to step through the program by depressing the single command switch.

All activity took place at the computing site except possibly for the preparation and processing of cards at other remote locations. When cards were prepared off-site, they were then trucked to the computing center for subsequent processing. It was not uncommon to see full, heavy trays containing thousands of cards, kept under pressure to preserve their condition, transported to the computing site and then back to the location of origin for storing.

Although it is hard to be precise, it can be estimated that the number of users of any computer was in the hundreds. At Cornell, professors and their graduate students were the largest user group. Undergraduate students by themselves or in formal course work were in the minority.

Vendors

A variety of vendors, notably IBM and Univac (later Sperry Rand), started to manufacture commercial equipment. The environment was much the same as existed in the late 1990s when vendors were trying to outdo each other and to make their offerings increasingly appealing to users. Being recognized as the standard clearly enhanced a vendor's chances to improve its share of the market. Vendors then were generally referred to as "Snow White and the seven dwarfs," where IBM was Snow White and the seven dwarfs

³ Charles W. Johnson, "Fortran History and Features," www.eskimo.com/~cwj2/kcolleges/cs515/fortran

were Sperry Rand, the next largest, and Control Data, Honeywell, Burroughs, GE, RCA, and NCR.

During this time there was a very important legal case involving the U.S. Justice Department and IBM. The outcome was important to Cornell and higher education in general. In 1954 the Justice Department accused IBM of monopolistic business practices in the marketing of their punch card equipment. In May 1956, IBM signed a consent decree to change its business practices, which required IBM to sell as well as lease its machines. This led indirectly to other business changes, the primary one for Cornell being the elimination of the 60 percent educational allowance on equipment acquisitions. Further restrictions included limitations on the nonuniversity users who could run programs on computers acquired with the educational allowance as well as on the selling of such equipment. Another change, important to all customers, was the ability to purchase equipment and make independent arrangements for maintenance and repair instead of the single option of leasing the equipment with all other costs bundled into the cost of the lease.

Technology Terms, Acronyms, and Buzz Words

Core—A substitute word for main memory, which resulted from the technology that greatly increased the amount of storage. This led to the use of terms such as “in core” instead of “in memory” in references to different methods of doing work.

Data processing—As the transition from punch card technology to computer technology was taking place in the business world, this term evolved to distinguish work that involved the processing, organizing, and accumulating of large volumes of data elements from number crunching, which involved the extensive and repetitive calculations most often performed by researchers.

1950 to 1959 at Cornell

The first digital computer was installed on campus in this decade, and the Cornell Computing Center was organized to provide support for the use of computers by researchers. The first course about computing and the programming of computers was organized, and computers were used in support of other educational programs. The use of punch cards and related processing equipment for business systems continued to increase.

The Cornell Computing Center

**The Founding of the Computing Center in 1953—
Richard C. Lesser, Director**

Dick Lesser provides an excellent commentary on the installation of the first digital computer on the Cornell campus and the formation of the Cornell Computing Center.⁴

Computing at Cornell had its beginning in the spring and early summer of 1953. Professors Robert J. Walker and J. Barkley Rosser of the Mathematics Department and a small group of faculty formed an advisory committee to discuss the establishment of a computing center that would offer services to academic and research staff. The decision to go ahead was made by Dr. T. P. Wright, the vice president for research, and Provost Hill. The Cornell Computing Center was to be set up as a part of the Department of Mathematics with an advisory committee representing interested groups on campus. At this time I was on the staff of the Center for Statistical Services and Scientific Computation at MIT, which provided such services to MIT faculty. The scientific computation area was my responsibility. There was no MIT computation center as yet on the horizon. That summer I became acquainted with a Cornell mathematician, Professor Mark Kac, who was a visiting professor for the summer at MIT, and he recommended me to Professor Walker as a candidate for director when the Cornell center was formed. I visited Cornell in August 1953 and was shortly thereafter offered the position of director of the Cornell Computing Center.

⁴ Richard C. Lesser, “Richard C. Lesser’s Recollections: The Cornell Computing Center, the Early Years, 1953 to 1964,” Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

The CPC Years 1953–1956

The new center had two missions—first, to make computing power available to faculty and graduate students, and second, to assist the campus in gaining computer knowledge and literacy. It is hard to believe today that there was no use of computers in the classroom. Before I moved to Ithaca in September 1953, we placed an order for an IBM card-programmed electronic calculator (CPC) for delivery in the fall. It should be noted that the center was to support itself by selling time to university departments, and, to assuage administration fears of a deficit, the Cornell Aeronautical Laboratory⁵ in Buffalo guaranteed a portion of the budget. The estimated budget for the period December 1, 1953, to June 30, 1954, was a staggering \$19,275.

The time before the delivery of the CPC was spent meeting with faculty and giving talks to any group that was interested. In preparation for the equipment, I spent untold hours wiring and testing the board that controlled everything. The CPC was a composite of three separate machines that were cabled to function as a single unit. Card input and printed output was through an IBM 418 accounting machine (tabulator, or tab for short), card output was via an IBM 527 reproducing punch, while calculations were performed by an IBM 605 electronic calculator (which was a successor to the widely used IBM 602A calculating punch). Decks of cards with instructions were fed into the tab (at the rate of 100 to 150 cards per minute), calculations were performed, intermediate results punched out onto cards, and instructions re-entered in new instruction decks for further work. One instruction card could perform an operation such as $[Ax+B+C=D]$. Sometimes the intermediate results were collated (appropriately, on an IBM 077 collator) and iterated until the desired result was achieved. To invert a 20×20 matrix, for example, took 20 iterations and several hours, longer if a deck

⁵ Cornell Aeronautical Laboratory (CAL and later CALSPAN) was a wholly owned subsidiary of Cornell, established in 1946 as an independent, nonprofit research organization and deriving its financial support solely from the performance of contracts from sponsors. It was located in Buffalo, N.Y. It went through a number of reorganizations before it was sold as an aftermath of student complaints in the late 1960s.

got dropped or out of order in the process. Despite its clumsiness, with ingenuity and perseverance we could do quite complicated computations with the CPC. For example, a square root was calculated by iterating a formula until acceptable convergence was achieved. The internal data storage was limited to registers in the tab and the 605, a problem that was somewhat eased by the arrival of a new storage unit, the IBM 941A. As I sit here with a gigabyte of hard disk, I remember how exciting it was to use the 941 with its capacity to store 16 ten-digit numbers (with sign!). To those who think extension conflicts are tough to resolve, I finally tracked down why an otherwise reliable computation dropped a “1” about every three months. (It turned out that a pilot selector [relay] in the tab dropped out for 10 degrees of its 360 degree cycle and any “1” passing through at that time was lost.)

The center was housed in the east end of the third floor of Rand Hall in spacious, if ancient, quarters. Our nearest neighbors were Buckminster Fuller and his students building geodesic domes. The initial staff was myself, Ann Waymeyer, secretary and keypunch operator, and Dorothy Hartman, programmer, key punch, and computer operator, and in the first summer of 1954, a high school student, Margaret Hunter as jack-of-all-trades. In addition to the CPC, we had keypunches, a verifier, a sorter, and the collator.

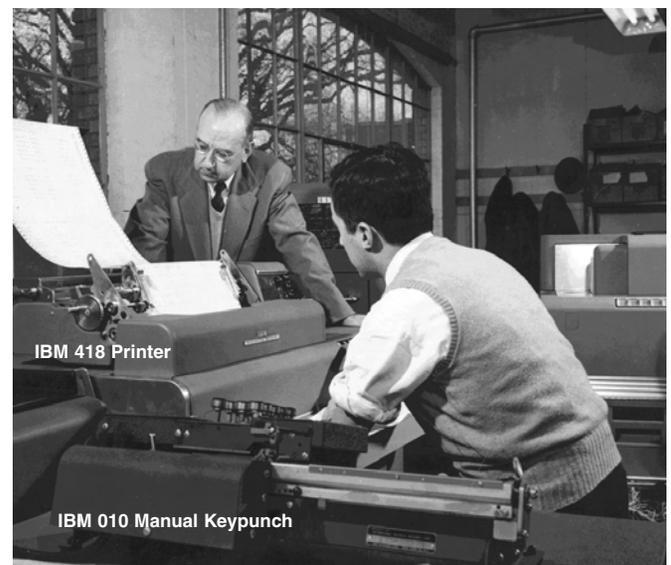


Figure 2. Professor R. J. Walker and Dick Lesser at the CPC in Rand Hall

Finally, on December 15, 1953, at 8:00 a.m., the CPC was powered up for the first time. Much of the following weeks were spent in testing the boards and the first programs. It should be noted that the main board (the tab board) controlled most of what went on in the CPC. It was about 3 foot square and was piled high with over 3 inches of wires, color coded as to length. Making changes or additions was a daunting procedure. Making a board diagram was a must in order to be able to follow what was going on. I have had no liking for documentation of programs to this day as a result.

Our first customer was Professor Lyman G. Parratt of the Physics Department. In the first year many of the jobs were for faculty in the state Colleges of Agriculture and Home Economics, mostly statistical in nature, such as correlations and analyses of variance. Early customers were Professors Urie Bronfenbrenner, of Child Development and Family Behavior, Walter Federer, of the Biometrics Unit of Plant Breeding, M. T. Vittum, of the Geneva Agricultural Experiment Station, Richard Bersohn, of Chemistry, and Arthur McNair, of Civil Engineering. We charged customers at the rate of \$25 per hour.

The IBM 650 Years 1956–1959

During 1954 and 1955, demands for the center's services increased steadily. The staff continued to grow with additional programmers and support staff. In late 1954, we placed an order for an IBM 650 magnetic drum computer for delivery in early 1956. The acquisition of the new machine was furthered by a \$50,000 grant from the National Science Foundation and a 60 percent rental contribution from IBM. This new machine would require more space and air conditioning (but the staff would not), so a new home was called for. With the help of Dean S. C. Hollister of the College of Engineering and Prof. William Erickson, director of the Department of Electrical Engineering, space was found on the first floor of the new Electrical Engineering building, Phillips Hall. In addition to the computer room, we had an auxiliary machine room and an office for the director. Our first research (graduate) assistant,

Virginia A. (Ann) Walbran, joined the staff.

The 650's drum (storage unit) was capable of holding 2,000 words (numbers) of 10 decimal digits and sign. It was our first experience with a stored-program machine, and the days of running program decks over and over on the CPC were soon forgotten. Information was read into and punched out of the 650 via the IBM 533 read-punch unit, and the output was printed on an IBM tab. Instructions and data were interspersed on tracks on the drum, and execution speed was dependent on placing data so that it would be available as the drum rotated. Initially, programming was in machine language, and the process of making sure the data were available was called hand optimization. This was a time-consuming programming task, requiring knowledge of the rules for how soon the drum would rotate to a given location. As a program got larger, one had to keep track of available locations and be careful not to slow the run-time by requiring a full rotation of the drum before data could be accessed.^b

Fortunately, help was on the way in the form of two new ideas, the interpretive and the assembly programs. The first of these was described in *Technical Newsletter #11*, published by IBM in March 1956. (It should be noted that *Technical Newsletter #1* was the first [non-operating manual] literature published by IBM). This interpreter was titled "A Complete Floating-



Figure 3. Dick Lesser seated at IBM 650 console with Ann Walbran and Prof. R. J. Walker in Phillips Hall

Decimal Interpretive System for the IBM 650 Magnetic Drum Calculator” and was developed by Dr. V. M. Wolontis at the Bell Telephone Laboratories. The 650 being a fixed-decimal machine, it was transformed into a “three-address, floating-decimal, general purpose computer, primarily suited for scientific and engineering calculations.” The assembler was the Symbolic Optimal Assembly Program, the famous S.O.A.P., described in IBM’s 650 programming Bulletin 1, dated May 1956. In 650 programming, the operation (OP) code 65 stood for Reset and Add to the Lower Accumulator, abbreviated RAL for the purposes of the assembler. (RAL are also my oldest son’s initials, giving me a lingering link to the 650). SOAP gave us the ability to use mnemonics for operation codes and for operands as well. Programming was performed according to SOAP’s rules, and a program with the resulting OP codes and optimized addresses was assembled at the rate of 50 to 75 cards per minute. This program, when punched out, could be listed on the tab for debugging, or run back through the 650 for execution. Programming had advanced from card programming through machine language to a new level of ease. Our rate for sponsored (by research agencies of the federal government) computing was now \$75 per hour.

The 1957 report “Data Processing and Computing at Cornell University” adds the following perspective to Lesser’s story.⁶

In 1948 the IBM Corporation introduced their model 605 as the first commercially available electronic computer. This led the Mathematics Department to ask that the university provide its use as a research tool. As a result, the Cornell Computing Center and the Advisory Committee of the Cornell Computing Center were both created in 1953. The center opened for business in Rand Hall in September of that year. The Advisory Committee, a cross-campus group of senior faculty members interested in furthering computing, reported administratively to the Cornell vice president for research, and the director

of the center reported to the chairman of the committee. In July 1957 the status of the center was changed from an auxiliary enterprise to an academic division. Because of this, the director of the center and the chairman of the advisory committee both reported directly to the provost. At the beginning, the budget of the center was approved by the Mathematics Department and forwarded directly (i.e., not via the College of Arts and Sciences) to the vice president for research. Now it goes directly from Mathematics to the provost. This fact and the fact that the chairman of the Mathematics Department has had a great deal of interest and experience in computing has led to a special relationship between the director of the center and the chairman of the Mathematics Department. In the strict administrative sense, this relationship is completely unofficial, but in practice, it is quasi-official, as it leads to decisions of the sort that might normally be referred to one’s senior in the chain of command.

To pay its own way, as required by the university, the center accepted both research and purely commercial computing for a fee. To help this young organization, the Cornell Aeronautical Laboratory guaranteed an income of \$8,000 during its first year. In 1956, when the Dairy Herd Improvement Program anticipated the need for a more powerful computer, a grant of \$50,000 from the National Science Foundation made it possible to acquire an IBM model 650 electronic computer without sustaining a heavy loss during its first year of operation. This new machine arrived in June 1956, as the center was moving to its present quarters in Phillips Hall. The past year of shakedown operations has been very successful and the center is now on a firm footing with regard to its service and financial stability. However, the director of the center has been advised orally that now that the center is no longer an auxiliary enterprise, the university will be more willing to underwrite limited deficits.

⁶ W. S. Gere, B. McK. Johnson, R. C. Morris, “An Analysis of Data Processing and Computing at Cornell University,” (Cornell University, November 14, 1957, copy number 10).

Research Computing

This same 1957 report provides a narrow glimpse into the research computing taking place on campus in 1957. It provides a list of departments, “autonomous organizations now operating data processing machines.” The Department of Poultry Husbandry, the Veterinary College, and the Physics Department each had their own card punch equipment, and there were two more keypunches on order whose placement had not been determined.

Instructional Computing

The IBM 650 brought with it the first course in instructional computing on campus. Professor Richard W. Conway from the Department of Industrial Engineering⁷ developed and taught the first course on digital computers at Cornell—Industrial Engineering 3281, Computers and Data Processing Systems, in the fall of 1956. According to Conway⁸, a requirement of the contract for the IBM 650 was the formal teaching of a three- to four-credit computing course. The course was not a course in programming but had to follow an outline supplied by IBM. In looking for ways to meet this requirement, Conway recalls that he and Professor Andrew S. Schultz Jr., then head of the Department of Industrial Engineering, worked together to include this course on the Industrial Engineering course roster and for Conway to take on this assignment. At the time, Schultz was also on the Advisory Committee of the Computing Center. The first class of IE 3281 consisted of 45 graduate and undergraduate students. Conway continued to teach this course for a number of years, and it remained for some time the only computing course at Cornell. As will be evident from this history, Conway played a major role in the development of computing on the campus for the next 20 years or so in a variety of capacities



Figure 4. Computer demonstration for Engineering faculty

The Computing Center in Rand Hall

We continue with Lesser’s narrative.⁹

The Burroughs 220 Years 1959–1962

It seemed almost axiomatic in the early days of computing at Cornell that the capacity of the current equipment was exceeded every three years. So it was that in 1958 we began looking for a successor to the 650. Once again, a search for a new home for the Computing Center was necessary, as we had overflowed the existing space available in Phillips Hall. After a long evaluation of available equipment, it was decided to order an Electrodata Datatron 220 (renamed the Burroughs 220 when Burroughs acquired the Electrodata Corporation). The Burroughs 220, valued at \$601,000, was purchased with a grant of \$250,000 from the National Science Foundation and a purchase grant of \$300,000 from the Burroughs Corporation. All previous equipment having been rented, it was deemed more economical to maintain this machine ourselves, and two trainee engineers were hired and sent to school. The space problem was solved by the availability, once again, of Rand Hall. This time, however, the center was to take up the entire first floor, eventually expanding into the second floor for offices.

The \$100,000 renovation of Rand Hall for the 220 was extensive, as a large air-conditioning system was required for the last of the great tube machines (magnetic core was just on the horizon). In addition, a bright yellow raised flooring was installed so that the multitude of cables connecting the components would be hidden. Installed in 1959, the 220 was a decimal computer with 5,000 ten-digit words of storage. New features to Cornell were magnetic tape and paper tape systems as well as a supervisory

⁷ The Department of Industrial Engineering has gone through a number of name changes since the 1950s. Although it is now known as the Department of Operations Research and Industrial Engineering, for convenience we will stay with the name Industrial Engineering throughout.

⁸ Richard W. Conway, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

⁹ Richard C. Lesser, “Richard C. Lesser’s Recollections: The Cornell Computing Center, the Early Years, 1953 to 1964,” Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.



Figure 5. Burroughs 220 computing system in Rand Hall

printer linked to the console. I claim the “honor” of being the last individual allowed to debug programs at the console. Input was by cards on an IBM collator and, very occasionally, paper tape. Output was printed on an IBM tabulator, or punched by an IBM reproducer, or on paper tape. We charged \$150 per hour for computer time used by sponsored research projects.

User Groups

Lesser makes the following comments about the formation of the first user’s group involving Cornell:¹⁰

Cornell was one of the founders of the original Burroughs user group, the Cooperating Users’ Exchange (CUE), of which Robert Gordon of Stanford was the first president and I was the second. This group had large representation among universities and military research facilities, and through the urging of this group, the next breakthrough in software was achieved. Burroughs Corporation in early 1961 published and distributed the Burroughs algebraic compiler, “a representation of ALGOL for the



Figure 6. John Rudan and Helen Church at the Burroughs 220 console

Burroughs 220 data-processing system.” As had board wiring and machine language programming, assemblers now dwindled in importance. We had been using the assembler CAP (Cornell assembly program), written in 1959–60 for the Computing Center by an undergraduate student, David J. Waks.^c With the availability of the CAP assembler and the BAC (Burroughs algebraic compiler), programming the 220 was greatly facilitated.

¹⁰ Richard C. Lesser, “Richard C. Lesser’s Recollections: The Cornell Computing Center, the Early Years, 1953 to 1964,” Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

As Lesser has noted, in 1959 the Computing Center replaced the IBM 650 with the Burroughs 220. The reasons for this were complicated. In part, according to Lesser, the main reason for going to other vendors was that IBM did not have a competitive successor to the 650. Their 1400 series of equipment was more suited for business data processing, the 1130 was too small a scientific machine, and the 704 was too expensive for the research load at the time. Further, other conditions in the market favored a switch. When IBM reduced the educational allowance as an outcome of the consent decree, the then-large cost of the computer and its associated maintenance more than doubled, rising from 40 percent to 85 percent of list price—a significant cost increase. At the same time, other vendors, notably Burroughs, were aggressively trying to get into the market and unburdened and—perhaps encouraged by the consent decree—were offering substantial discounts to the higher education market. Their sales pitches were easy to accept because the installed user base was small and the conversion problems were reasonably easy compared with the advantages of reduced costs and increased capabilities. Despite this situation, IBM took some exception to losing the Cornell account, and there was some loose talk about how this decision was made.

An added feature of the Burroughs 220 was the beginning of an applications program library for statistical analysis of data. A number of different statistical calculations were packaged into CUSTAT—Cornell University Statistics^d—and made available to users. The library was stored on magnetic tape and invoked with a program card. CUSTAT read the name of the program, retrieved it from the tape, and loaded it into memory. From there the statistical program took control and read both its parameters and data from the card reader and, after calculation, printed out the results on the printer. Data also could be stored and read from magnetic tape if they were large in quantity and required additional processing before being analyzed by the statistical program. The first general purpose linear regression program had a limit of one dependent and up to six independent variables, the limitations coming from the way the punch card was laid out to accept both data and identifying information and sequence number.

The installation of the 220 also brought with it another first for the Computing Center; the hiring of the first computer operator, David W. Pulleyn. Pulleyn had joined Cornell earlier as a technician in Electrical Engineering but transferred to the Computing Center when the staff was expanded to accommodate the need to use the 220 more intensively for customers and eliminate some of the self-service operations. The seven-days-a-week/24-hour operations also required

increased contact with the computer maintenance staff and coordination of their activities. Pulleyn went on to play increasingly more important roles in the future evolution of computing on campus.

Business Systems

Machine Records—Dominic Bordonaro, Director

The 1957 report on data processing and computing at Cornell University¹¹ has this to say about the Machine Records Office:

In 1952 the Purchasing Department took over responsibility for Machine Records. Though operations were quite successful, there were no written procedures, and the entire organization was sustained by the dynamic personality of the supervisor. As a direct result of this casual management, when he left in the summer of 1952, output rapidly fell off. The situation was aggravated as one supervisor after another was hired and released. From this time on, the Machine Records Office has had an unfortunate history. Machine operator turnover was very high. The office was down to half complement at one time, and untrained temporary help was even used, though with poor results. By 1953 there was a two-year backlog of work. At that time charges were made by a flat rate contracted agreement before a job started rather than by an hourly rate, with the hours added up at the end of a job. On a billed basis the deficit in fiscal 1955 was \$13,000. The following year a systems manager was hired but was released for non-performance in June 1956, when the annual deficit was \$12,000.

Since early 1956 Machine Records has made a steady improvement. In that year it was transferred from Purchasing to the Office of the Controller. The effect of higher salaries began to show: personnel turnover dropped, and experience grew. Better equipment was introduced, and more wiring boards were bought so that they could be kept wired for special jobs and filed away between times, thus saving considerable set-up time for high-frequency repeat jobs such as payroll. Even so, the record of the past still dogs the unit. Many

¹¹ W. S. Gere, B. McK. Johnson, R. C. Morris, "An Analysis of Data Processing and Computing at Cornell University" (Cornell University, November 14, 1957, copy number 10).

one-time customers have withdrawn their support, and many potential customers still shy away. The organization still has no written official purpose. Its status in relation to the other machine units on the campus is not yet clear, and there are no requirements that an administrative unit use Machine Records services instead of setting up its own organization. Historically, its workload has been primarily of administrative origin, and so it remains today. Nevertheless, some 5 percent to 10 percent of its work is for research.

Information obtained from Dominic Bordonaro¹² fills in many of the gaps in this summary. It should be noted that it was Bordonaro who took over as director of Machine Records in 1955–56 and whose leadership led to service improvements that lasted many years.

By 1951–52, as Bordonaro documents, Machine Records had grown from two to three keypunches, from one to two sorters, and from one to two 405 tabulators and continued with the single collator. By the end of the decade, it had expanded to three tabulators (two 407s and one 408), a 513 reproducing punch, an interpreter, a 519 collator, a 602A calculator, and six (or more) keypunches. Stated this way, this growth does not reflect the real increases in processing capability. For almost every type of equipment, there was a three- to four-fold increase in processing speed over the course of the decade; for example, the 407 could print up to 150 lines a minute, whereas the 405 could print at best 50 lines per minute. As a result, Machine Records was processing a considerable amount of the university's business information records using punch card technology and planning for the acquisition of a 1401 computer. Applications ranged from billing student tuition and fees for the treasurer's office, producing course rosters, processing registration and student grades, and posting student transcripts for the registrar's office. Also it provided payroll processing and other support functions for the endowed and statutory payroll offices, reports for the General Accounting and Budget Office, and accounting and billing services for the Buildings and Grounds Division.

Another way to look at the growth of Machine Records in the years up to 1959 is to consider the increases in the number of staff. It is important to remember that while the capacity of each piece of equipment was fixed at any point in time, more operators would permit more intensive use of each piece. By taking a job-shop approach to utilizing equipment by overlapping operations and by adding more shifts of operations, the throughput of the whole shop could be increased dramatically. As noted earlier, Machine Records started in 1946 with two staff members, a

supervisor, and an operator. By the early 1950s, the staff numbered seven—one supervisor and six operators—and reached a total of 12 in 1959. Key staff in the latter part of the decade were Bordonaro, director; Les Phillips, supervisor of keypunch and the tab room; and Ted Yeager, manager of the same operations. Irene Van Zile recalls that when she started her employment as a keypunch operator in 1957, she was employee number 11 in Machine Records.¹³

Over the same period the budget rose from about \$12,000 in 1946 to \$145,000 in 1959–60. There are no detailed breakdowns between the components of equipment, staff, and general expenses in the historical materials for this period.

Bordonaro also provides some more factual information about the organization and the leadership. When the 1950s started, Machine Records reported to the registrar (Ernie Whitworth), later to Personnel (Deidrich K. Willers), and next to Purchasing (Wallace B. Rogers). In the period between Bob Burghardt's leaving in about 1951 or 1952, to 1955 when Bordonaro was appointed director, there were six interim directors, each of whom stayed about six months and left. Bordonaro recalls two names from that period—Roger Gettings, who stayed for three weeks, and Stan Gill, who stayed the longest. In that same year, 1955, Machine Records started to report to the newly appointed university controller, Arthur H. Peterson, and continued to do so for about 12 more years.

Statutory Finance and Business Office

When Machine Records processed the main payroll and related payroll functions for the statutory colleges, it dealt with the Statutory Finance and Business Office (SF&BO). At this time SF&BO had wide-ranging responsibilities for operations of the statutory colleges at Cornell. In about 1955, under the direction of Lloyd Slater, SF&BO undertook some of the preparation of input data for their payroll processes. The inventory of sites and equipment from the 1957 report on data processing and computing at Cornell University¹⁴ states that in 1956 this office, then in Roberts Hall, had a duplicating card punch (031 model) and a typewriter punch (824 model). The office continued using Burroughs bookkeeping machines for its principal financial and accounting work until the end of the

¹² Dominic Bordonaro, "Dominic Bordonaro, Oral History, 1947 to 1968," Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

¹³ Irene Van Zile, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

¹⁴ W. S. Gere, B. McK. Johnson, R. C. Morris, "An Analysis of Data Processing and Computing at Cornell University" (Cornell University, November 14, 1957, copy number 10).

decade. This office would become a much larger player in university-wide business systems in the next decade, after 1960.

Other Computing Installations

Dairy Records Processing Laboratory

For completeness of the documented history for the 1950 period, the activity in Animal Husbandry must be included, where the use of punch cards and related equipment started in the 1940s. Lyle Wadell wrote a detailed history of the Dairy Records Laboratory and the information below is from that document. Wadell's story, which starts with the year 1950,¹⁵ expands on Turk's description of the beginnings of the Dairy Records Processing Laboratory (DRPL), discussed in the previous chapter.

In 1950, an IBM 601 calculator was added and the then-existing equipment was operated by two employees. Any printing that needed doing was done using a tabulator at an Agriculture College facility in Warren Hall. In 1951 this changed when the existing equipment was moved to the New York Artificial Breeders Cooperative, with their tabulator being used for the printing.

An IBM 101 unit was obtained in 1952 at the time H. W. Carter took over management of what had now grown to three employees. In 1953 an IBM 602A calculator was added (being the first unit on campus that could multiply and divide) and an IBM 402 tabulator. At this time in history the unit was doing sire and cow research using data submitted on handwritten cards by DHI supervisors. Age-correction factors were developed for the New York area and the original research was completed on herdmate sire proofs. In 1954 a staff of five employees started making parallel runs of daughter-dam proofs and herdmate sire proofs.

In 1955 the group moved into an old Agricultural Engineering building and started publishing AI sire herdmate proofs twice per year. At the same time, H. W. Carter and J. D. Burke started working on a central processing program for the computer processing of DHI records. October 1956 saw the first computer-generated DHI reports go to a small group of volunteer dairymen in Tioga County, N.Y. The computations were done by the university's IBM 650 computer. An Animal Science staff of

six employees completed the other work necessary to get these first reports processed.

Shortly thereafter, in 1957, New York started processing DHI records for other northeastern states (4,000 cows on test on January 1). Extensive research continued in sire evaluation techniques. An IBM 407 tabulator was added, and the staff had grown to eight employees.

During 1958 the Animal Science Department obtained its own IBM 650 computer and started publishing three SI daughter-level reports a year. Type appraisal summaries were also started at this time. The January 1 figures indicated 42,000 cows on test. The DHI records-processing and research effort now involved a staff of 13 workers.

In 1959 George O'Blennen was hired as the administrative supervisor of a staff of 16 workers. It was approximately at this time that the unnamed unit received the name Dairy Record Processing Laboratory. Lyle Wadell was hired to work with the research group coming from Michigan via Iowa. Central processing had reached a level of 108,000 cows, and the sire evaluation techniques had become more refined by making adjustments for the number of daughters.

As noted above, Lyle Wadell joined DRPL in 1959 to work with the faculty and students using the records for research purposes. The research leader was Charles R. Henderson, professor of animal breeding, who, with numerous Ph.D. students, over a 40-year period steadily developed improved methods for using farm and animal production records (especially those of dairy cows) to improve animal production through genetics. A prime feature of this research was that of using dairy bulls' daughter records for optimally selecting the few bulls to be used in artificial insemination programs.

The 1957 report on data processing and computing at Cornell University¹⁶ complements and expands on the information above.

¹⁵ Lyle Wadell, "History of the Northeast Dairy Records Processing Laboratory, 1948–1985, with Additional Comments by J. D. Burke and H. Wilmot Carter," Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

¹⁶ W. S. Gere, B. McK. Johnson, R. C. Morris, "An Analysis of Data Processing and Computing at Cornell University" (Cornell University, November 14, 1957, copy number 10).

Since November of 1956 the Animal Husbandry Department has assumed the responsibility of processing, on a commercial basis, the milking, feeding, and etc. records of the New York Dairy Herd Improvement Cooperative. The intent of the department in undertaking this project was to compile otherwise unobtainable data relevant to present and contemplated research. The maintenance of these records requires the use of a wide range of electronic processing equipment. The card punching, sorting, collating, and tabulating necessary have been accomplished at the existing Animal Husbandry installation. Computation has been performed at the Cornell Computing Center by members of the department's staff. At the present level of operations, approximately 34,000 cows, computing time at the center averages about 20 hours a month. The Animal Husbandry Department has found this arrangement unsatisfactory for at least three reasons. The problem of transporting some hundred thousand cards monthly between the two installations has not been a small one. There has been some problem in obtaining regular and reasonable times to do this work. It is believed that in view of the present and possible near future volume, the Computing Center's charge for those hours has been and is still excessive. It is also pointed out that if the load increases as anticipated, the available capacity of the center will be exceeded. Consequently, a request was made and an order placed for a second 650, which will tentatively arrive during May of 1958.

As a historical footnote, it may be of interest that in 1998 the section of Judd Falls Road where the New York Artificial Breeders' building was located was renamed Pine Tree Road, and in 2003 the site became the home of CISER, the Cornell Institute for Social and Economic Research.

Agricultural Economics

The 1957 report has very little to say about the installation in Agricultural Economics. In the summary of equipment installations on campus, Warren Hall is noted as having one duplicating punch (016), one printing card punch (026), one verifier (056), and one card-counting sorter (075).

Status of Computing on Campus

Although the 1957 report on data processing and computing at Cornell University originated to consider specifically the installation of a second IBM 650 computer on campus, the authors did quite a comprehensive and thorough investigation of the state of computing on campus. They also presented some thoughtful recommendations and summary status information. They contacted 16 deans and academic department heads and 21 heads of administrative departments either through an interview or a questionnaire, depending on that unit's likely immediate use of computing technology. They also sent letters to 20 companies soliciting information about their equipment offerings and received 12 replies. They had the guidance of a steering committee made up of Andrew Schultz Jr., head of the Department of Industrial Engineering and the committee chair, T. P. Wright, vice president for research, and A. H. Peterson, controller.

In addition to the information already quoted from the report in earlier sections, the report summarizes the computing situation on campus as follows:

Present Picture

A complete list of Cornell's computing and data processing machinery is given in Table 2, Appendix F [of the original report]. There are eight autonomous organizations now operating data processing machines, and there is an order for the machinery to create a ninth. It now costs about \$99,000 a year for equipment rental, and there are orders outstanding that could raise this to about \$190,000. If supervisory and operator salaries are included, the present annual cost comes to \$178,000 **and the anticipated annual cost would be about \$359,000.** No charges have been included for space and utilities. Nor have second and third shift rentals been included. If IBM should start to charge for rental for second and third shift use, as they now contemplate doing, these figures would be even higher. (List rental for the second shift is half of the list rental for the first shift; so is list rental for the third shift. The usual discounts are given from list.) For the sake of comparison, **the total Ithaca campus budget for both state and endowed schools is about \$37.5 million,** including expenses for sponsored research.

The bold type in the above quote emphasizes both the magnitude of the dollars involved and the concern

expressed later in the report that the total costs for computing would require the university to allocate 1 percent of its total budget for this activity.

After considering the information gathered and the report writers' evaluations of the computing situation on campus, the report summarized the findings and concerns in 14 different conclusions. These conclusions dealt with general policies and issues such as centralization and decentralization, the priority of different types of work, financial support for academic uses, and the request for another 650 computer. Considerable attention was given to the Machine Records unit and the potential for applications of data processing technology in other administrative operations. It is evident that a new system for the Alumni and Development Office was an obvious high priority, for much of the discussion of potential new administrative systems focused on this office, and Appendix B to the report was a detailed discussion about this office. The appendix (42 pages) dealt with an evaluation of the present system, a proposed new system, and three implementation alternatives and their respective costs. Appendix C was a detailed analysis of further applications in administrative operations.

The report made 12 recommendations:

Regarding centralization of machine facilities:

1. It is recommended that the university policy be (a) to have a centralized data processing facility for business purposes, this facility being the Machine Records Unit, and (b) to provide a computing facility for education and research that will meet the needs of all academic divisions of the university that do not have their own resources. To the maximum extent consistent with satisfactory academic performance, this computing center shall serve as a centralized computing and data processing resource for teaching and research.

Regarding centralization of control:

2. It is recommended that the board of the Cornell Computing Center be supplemented by personnel technically competent in data processing; be given responsibility for the periodic review of the effectiveness, utilization, and need for all academic installations; and be given responsibility for recommendations concerning equipment additions in an academic (nonbusiness) installation.

Regarding deficiencies in business data processing:

3. It is recommended that there be added to the controller's staff a competent methods and procedures specialist.

Regarding Machine Records in particular:

4. It is recommended that the policy with regard to responsibility of the Machine Records Unit in data processing be as follows: The primary responsibility of the Machine Records Unit is the provision of business data processing for the non-academic agencies of the university. To the extent that greater economy can be achieved by increasing machine utilization in slack periods, there may be undertaken routine, repetitive sponsored research data processing jobs that do not interfere with business data processing deadlines or require a great deal of technical consultation. Such jobs should always be cleared with the Cornell Computing Center.

5. The staff of the present Machine Records Unit should be supplemented by two machine room foremen.

6. It is recommended that the Machine Records Unit be allocated considerably more area in which to carry out its operations. Assuming that certain other recommendations concerning extension of machine applications are carried out, the need will be in excess of 3,000 square feet.

7. It is recommended that the rate structure be so revised so that the expense of low utilization is not borne by the departments serviced by the installation, but rather by the university.

Regarding the Cornell Computing Center:

8. It is recommended that the Cornell Computing Center be charged with the responsibility of providing computing and data processing resources for all academic divisions of the university; with the responsibility of maintaining information on the loads and capacity of any decentralized facilities for the purpose of suggesting actions to balance loads; and with responsibility for providing advice, consultation, and information on the subject to interested academic departments. In addition, the Computing Center should be directed

to maintain sufficiently close relationships with the Machine Records Unit.

Regarding the Dairy Herd Improvement data processing:

9. It is recommended that the acquisition of a new IBM 650 by the Animal Husbandry Department be postponed as long as possible, that a sliding scale of rates for the use of the Computing Center be applied to this project, and that a serious effort be made to ease the card handling problem.

Regarding recommendations for potential new machine applications:

10. It is recommended that a policy be adopted to provide adequate priority for academic use of the equipment in the Computing Center for purposes of instruction and unsponsored research even to the extent of excluding potential commercial income or sponsored work and including budget appropriations to underwrite the resultant deficit. It is further recommended that discussions be held with IBM to investigate the possibility of educational discounts for every academic machine facility.

11. It is recommended that the Development Office actively pursue development of a punch card data processing procedure for record keeping, address maintenance, and addressing, and this development be closely coordinated with the development of the ability of the Machine Records Unit to meet deadlines and handle an increasing load. It is further recommended that as soon as this ability is demonstrated, the Development Office convert to this new system.

12. It is recommended that further detailed systems and procedures analyses be made with the intent of improving data processing methods and discovering possible economic machine applications for the following areas:

Admissions	Payroll (endowed and state)
Personnel	Purchasing and Stores
Registrar	Residential Halls
State Business Office	Auxiliary Enterprises

While not all the details of the report are provided here, the recommendations were quite sweeping in scope. The report made clear that these were the recommendations of the authors and subject to review and acceptance by the university administration. There is no record of how these recommendations were accepted and acted upon, but it will become obvious that some of these recommendations came to pass—although not as soon as customers and practitioners desired or needed—some were delayed until the 1960 decade and later, and some were ignored.

Summary

Looking back now, it is clear that using computing technology in innovative ways was very much evident in these early days and started initiatives that have prevailed until this day. We note in particular that Cornell was building its own software (e.g., the development of CAP) and was a founding member of users groups (e.g., the Burroughs users group), and the central facility was an itinerant camper in different buildings every few years (from Rand Hall, to Phillips Hall, and then back to Rand).

Although Dick Lesser comments about various financial items, such as costs of the computers and renovations and rates for computer time, very little mention is made about the tenuous financial underpinnings of the center. From the start, because the university treated the Computing Center as an enterprise operation—if not directly, as the 1957 report states, then indirectly—the expense budget had to match the level of recoverable earnings from billings to users. The financial model was constructed so that the estimated costs were divided by the saleable/billable hours, yielding a charge rate. This was in keeping with Circular A-21 from the federal Office of Management and Budget, which established the policies governing cost recovery from sponsored research contracts. The model basically resulted in all the nonbillable hours being considered as spare hours, whose marginal cost was zero. Consequently, use of the computers for research and instruction was encouraged because it seemed wasteful to have an expensive and very useful resource sit idle. While such a rational argument to use an idle and important resource was difficult to argue against, this thinking led to some interesting future policy and financial questions when the university was asked to cover the costs of increasing instructional and nonsponsored research use out of general purpose funds.

Endnotes

^a Although recycling of office paper is a fairly recent phenomenon, recycling of computer cards was practiced in those very early days. The cards were made of special and expensive paper, and so it was advantageous for scrap recyclers to collect used cards and pay the department directly. This money went into a special fund that was used for staff picnics and other such staff social activities for which university funds could not be used in those early days. The fund accumulated money as long as there were cards to recycle and was finally depleted in the late 1980s.

^b I first encountered computing at Cornell in the spring of 1959 when I took Professor Richard Conway's course Industrial Engineering 3281, Introduction to Digital Computers. I still recall the first night the class went to the Computing Center in Phillips Hall. Working as pairs, we punched program cards for our demo program and waited our turn to load our program deck into the 650 reader and then into the 650 itself. We spent a few minutes reading the console output, which was lit using the biquinary numbering system, a challenge in itself. When our time ran out (we had 10 minutes or so), we punched out the answers onto cards. We then took those to the 402 tabulator/printer to get them printed. If you've never seen a 402 printer, you have missed the experience of a lifetime! It had 120 long vertical print bars, each bar being about 18 inches long and containing the full character set. On receiving a print line, all the bars rose simultaneously to their respective letter or number or character position, and then with a mighty whack, all 120 hammers struck the bars to create an imprint of the line on paper. The floor and building literally shook each time a line was printed.

I enjoyed programming the 650. One mapped the rotating drum memory into this big matrix of memory locations, and with a card that gave the instruction times, one then scattered instructions and data in memory. When one instruction ended, the next instruction was coming up under one of the read/write heads in another column and it could be read at that time. Otherwise, one had to wait one full drum rotation for that memory location to come to a read head again. This was quite a challenge. Timing the speed of a bubble sort program made up a significant portion of the course grade. The completed program was submitted on a deck of cards, and the instructors ran this program against their own data streams and timed the execution. I had the second fastest sort in the class, and I claim, right or wrong, that this led to my selection as a graduate assistant at the Computing Center in the fall of 1959!

^c The development and production use of CAP was quite an experience. The programmers were initially avid users but then grew more cautious because the program had numerous errors. The current process of alpha and beta testing and field testing was not a routine practice. You wrote a program, you ran it. So with CAP. It was not uncommon for Mr. Waks to come to the Computing Center after his classes, find a number of error complaints, and go to the console and create fixes and a new version, which became the production version at the completion of his time. And so it went! Eventually CAP became more reliable, and an increasing amount of programming was done using it.

^d In response to a request from Professor T. C. Liu in the economics department, and with his help and that of a U.S. government publication, I wrote a general-purpose regression program for the Burroughs 220. When program names came to be used, it was called MUREG, for multiple regression, and went through many updates and improvements in response to users who requested more incidental statistics. MUREG was incorporated into CUSTAT as the library of such programs increased.

1960 to 1969

The Batch Decade

Contents

1960 to 1969—Industry Overview	27
Hardware	27
Software	28
Services	28
Vendors	29
Technology Terms, Acronyms, and Buzz Words	29
1960 to 1969 at Cornell	31
1960 to 1965	31
<i>The Development of CORC, the Cornell Compiler</i>	31
<i>Control Data 1604 Installed; Burroughs 220 Sold</i>	31
<i>Cornell Computing Center—John W. Rudan, Director</i>	32
<i>Changes in Rates and Operating Practices</i>	32
Instructional Computing	33
CORTAN	33
<i>More about the Use of CORC</i>	33
<i>Instructional Statistics</i>	34
Research Computing	35
Business Systems Data Processing	35
<i>Machine Records Installs IBM 1401 Computer</i>	35
<i>Machine Records Staff—Dominic Bordonaro, Director</i>	36
<i>Statutory Finance and Business Office; IBM RAMAC Installed</i>	36
<i>1965 Business Systems Studies</i>	36
<i>User Groups—CUMREC and CAUSE</i>	38
1965 to 1967	38
<i>The Office of Computer Services (OCS)—R. W. Conway, Director</i>	38
<i>The Founding of the Department of Computer Science in 1965</i>	39
<i>Central Computing Moves Off-Campus to Langmuir Laboratory</i>	40
<i>Coping with the Failure of the IBM 360/67; 360/65 Installed</i>	40
<i>Operating System Changes and Modifications</i>	41
<i>Campus Terminal Facilities for Accessing the 360/65</i>	42
<i>Academic Support Staff Move to Campus</i>	43
<i>Formation of Computer Activities Group—CAG</i>	43
<i>Finger Lakes Colleges Computing Center</i>	44
<i>Machine Records Moves to Langmuir</i>	44
<i>Plan to Remove 1401 Computer</i>	45
<i>Accounting and Billing for Computer Resources</i>	45
1968 to 1969	46
<i>OCS Transitions 1968 to 1969—Erik D. McWilliams, Director</i>	46
<i>CLASP Deployed; CLAPTRAP Report</i>	47

Instructional Computing	47
<i>The Development of PL/C</i>	48
Research Computing	48
Business Systems Data Processing	48
<i>New Public Affairs and Library Systems</i>	48
<i>Administrative Systems Planning and Control Board</i>	48
<i>New Financial Accounting and Payroll Systems</i>	48
<i>Dealing with Campus Disruptions</i>	49
Closing Out the 1960s	49
Other Computing Organizations on Campus	50
<i>Dairy Records Processing Laboratory</i>	50

Figures

Figure 1. Control Data 1604 console and magnetic tape units	32
Figure 2. A brief pictorial of the CORC process for student programs	34
Figure 3. The IBM 360/65 computer at Langmuir Laboratory	41

Tables

Table 1. Cornell Computing Center: Number of CORC Programs Run, 1962 to 1966	34
Table 2. Cornell Computing Center: Billed Hours to Research Computing, 1960 to 1965	35
Table 3. Department of Computer Science Courses, Enrollments, and Credit Hours; 1965 to 1969	48

1960 to 1969—Industry Overview

The following is a brief description of the major hardware, software, and service innovations in the 1960–69 decade as well as how different vendors came into and left the market in these areas. Large computers requiring extensive air-conditioned space were still the norm, but smaller minicomputers started to come out on the market and could tolerate a wider range of conditions. IBM introduced their 360 line of computers, which were the first computers that used a byte as the addressable memory element and had compatibility in hardware and software to provide nondisruptive upgrades over a range of different processor sizes and speeds. Operating system software became an important way to get the maximum efficiency out of any hardware configuration. Telecommunications became important as more and more of computing work was done at remote sites and not at the computing center itself.

Hardware

The computing engine began to shrink in size and increase in speed with the introduction of transistor technology. Air conditioning was still necessary, but not at the same high level. Replaceable component cards were now 4 inches deep, 4 inches high, and 1/2 an inch wide on a plastic substrate with visible interconnecting wiring to the small button-sized transistors and other components. One spoke of instruction execution times in microseconds during the decade.

Computer memory got much larger and much faster with the introduction of core memory. Memory was fabricated by creating a frame of horizontal and vertical wires and placing a round magnetic doughnut—a core—at each intersection of such wires. This formed the write logic, since a core could be switched from 0 to 1 or 1 to 0 by appropriate current at the intersection of a horizontal/vertical wire pair. Another diagonal wire, threaded through each doughnut, acted as a read unit and was not destructive, so memory did not have to be restored after each read. Frames were stacked together to form a memory unit, the number of frames depending on the word length being represented. One could think of memory in units of thousands.¹

Improvements in hardware technology were summarized by Grosch's Law, which came into vogue in the 1960s. In 1953 Herbert Grosch from IBM proposed as an observed law that if one spends twice as much for a replacement computer, one can expect performance to be four times greater. Another way of stating this is that overall computer performance is proportional to the square root of the price. This law, in effect, was no more than a reflection of economies of scale. But look-

ing back, it is not clear that the law may have been a result of IBM marketing practices—every three years IBM increased the performance of computing by a factor of two and reduced the price by a factor of two—in effect, achieving a cost/benefit ratio of four!

IBM changed the whole paradigm of computing hardware when it introduced the System 360 family of computers in 1964. First, this introduced the concept of scalability, where one could buy the size of computer needed for the job and then seamlessly upgrade to larger systems as needed, because each model was built using the same instruction set and operated the same way. Second, the design of the 360 memory introduced the revolutionary concept of the byte as the basic unit of memory, which then allowed both instructions and data to take different numbers of bytes. One no longer thought of fixed-length words.²

Input and output devices also continued to increase in speed and reliability. Card readers were uniformly able to read more than 500 cards per minute (cpm) and some were able to read more than 1,000 cpm by the middle of the decade. Printers using raised print characters on print wheels (rotated on a horizontal axis) or print trains (on a vertical axis) approached speeds of 1,000 lines per minute. However, at this speed only the first top copy was really legible. Multiple copies using 14 7/8-inch by 11-inch standard or preprinted forms were obtained by using three-, five-, or seven-part carbon paper.

Paper tape faded away. One-half-inch magnetic tape became the standard for temporary or long-term data storage and data transfer between systems. IBM, which introduced the first magnetic tapes in 1952, brought out their family of model 2401 drives that recorded at a density of 1,600 bits per inch (bpi), but could also read the earlier models that recorded at 200 or 800 bpi.³ It is interesting to note that the first magnetic tapes in the 1950s, a 10.5-inch reel of tape recorded at a density of 100 bpi, could hold the equivalent of 35,000 punched cards! Console or direct data entry-type terminals moved away from the teletype unit to the IBM 3270, still in use by the end of the millennium.

Magnetic disks became available for commercial use for recording data and making random retrieval of information possible. Retrieval from a magnetic disk was much faster than from magnetic tape, but one needed to be concerned about the placement of data

¹ Throughout this document, technical terms are based on definitions from the *Encyclopedia of Computer Sciences*, 4th edition, or from Webopedia at www.webopedia.com.

² "History of IBM," www-1.ibm.com/ibm/history/.

³ R. Bradshaw and C. Schroeder, "Fifty Years of IBM Innovation with Information Storage on Magnetic Tape," *IBM Journal of Research and Development*, vol. 47, no. 4, July 2003.

on disks, because delays due to waiting for disk rotation would significantly increase the processing time for a job. Disks also started to be used as a temporary extension of main storage to support virtual memory system operations. The first disks at Cornell from IBM were the 2311, units capable of storing 7.2 megabytes on a removable disk assembly with 10 platters and 18 read heads.

Software

High-level programming languages such as Fortran, COBOL, ALGOL, and APL and their various offspring like BASIC became the common way to prepare programs. In 1963 Fortran IV was introduced and remained a standard definition and implementation of the language for many years. COBOL was defined in 1959 when the Committee on Data Systems Languages (CODASYL) accepted the specification for this COmmon Business-Oriented Language. Implementations however, only took place in the 1960s.

ALGOL (ALGORITHM-Oriented Language) was first proposed in 1958 by a joint European and American group, but it was later revised to ALGOL60, which was actually implemented. ALGOL was defined to be suitable for algorithms and programs on computers of different types and capabilities. To this end, it had three representations—reference, publication, and hardware. In 1968 ALGOL68 was defined, and it has remained the standard. Shortly after the introduction of the 360 hardware, IBM introduced the PL/I (Programming Language One) as the replacement for COBOL and Fortran and all other such languages. APL (A Programming Language) was first defined as a notation by Professor Howard Aiken at Harvard and was first implemented in 1965 at IBM. Ken Iverson directed this effort. This implementation, which turned into a general purpose programming language and was widely used on time-sharing systems in IBM, even had special operators, which took advantage of the character set on IBM Selectric typewriters used as input terminals.

BASIC (Beginners All-purpose Symbolic Instruction Code) was developed by Kemeny and Kurtz at Dartmouth to allow students to write simple programs and learn about computing. It ran on the Dartmouth time-sharing system and used teletype-like terminals and input and output. In 1967 the University of Waterloo in Waterloo, Ontario, Canada, developed the WATFOR compiler, which provided fast in-core compilation and execution of Fortran programs on the IBM 360/370 series of equipment. WATFOR became very popular at educational institutions for student use in course work. Later in the decade, this version was succeeded by a version named WATFIV (but still pro-

nounced watfor) that continued to be used for instruction on a variety of different hardware platforms.

The availability of these high-level languages greatly improved the ability to write programs and, as a result, increased the number of programmers. Assembly language was only used by those specialists who coded the often-used components such as compilers themselves or subroutines where the tightness of the code led to substantial improvements in speed and less use of still-precious memory. Operating systems began to play a more prominent role in order to automate the various subfunctions such as input/output, program succession, and error detection. Writing and maintaining operating systems created a breed of programmers called system programmers, who became the high priests of the trade. Given the dominance of IBM, OS/360, the operating system for the 360 family of computers, became a focus for these programmers.

At the end of the decade, in June 1969, IBM made another major policy change and unbundled all software products with the exception of the operating system that came with the computer. This created not only a new income stream for IBM but also a whole new industry, because others could now develop and sell their products.

Services

Accessing the computer with batch (as in the 1950s, with one program using the whole computer at any time) was the dominant mode at the beginning of the decade and continued to be popular for production work, that is, long-running research or administrative reports, throughout the decade. With an increase in the number of users, day hours were typically dedicated to program development, testing, and debugging, and very short time limits were imposed on the length of each run so that the largest number of users could access the computer. This was automated so that program decks or decks read onto magnetic tapes could be cycled through the system under operator control. Unless trained operators were on hand, selected users were trained and allowed to use the system on nights and weekends.

Time sharing, also multi-user computing—the concurrent use of the computer by more than a single user by having the operating system cycle in some fashion through the active users—was discussed early in the decade. One of the first such systems, CTSS (compatible time-sharing system), was developed in 1963 by Project MAC (multiple access computer)⁴ at MIT. An improved version called Multics came later in the

⁴ MIT, Laboratory for Computer Science (formerly Project Mac) from Webopedia.com.

decade. Toward the end of the decade, time-shared systems became the rage and were a dominant force. The lack of quality terminals for input, remote communications facilities, and the inability of operating systems to do everything demanded of them hampered the spread of time sharing.

The modem (modulator/demodulator) continued to be the most often used means to link a terminal to the central computer, although the acoustic coupler was popular. With the coupler, which included a built-in modem, a regular telephone handset could be used to dial the computer telecommunications controller; on receiving acknowledgment of the connection, the handset was inserted into foam cushions on the coupler so that two-way communications could take place. Typically, the coupler operated at 110 baud, or bits/second.

The introduction of time sharing, the improvements in batch services, and the development of new application packages and new communication services gave rise to the concept of the computer utility.^a This concept was essentially patterned after the electrical utility concept, whereby a large and expensive central generator of electricity distributed electrical cycles to any number of different customers using different devices. At this time, the prevailing situation in the computer industry was that the maximum “bang for the buck,” as stated by Grosch’s Law, was achieved by obtaining the largest possible computer within the constraints of the budget. As a result, given its large size, the analogous large central computer of the computer utility concept could distribute its computing cycles for whatever services a particular customer needed—time sharing, batch, data analysis, business data processing, etc. Further, it was argued that the large expense of the computing complex called for intense use of the computing resource to generate the funds to recover or justify the costs. There was much discussion about the validity of this argument, but since “bang for the buck” often took precedence over the needs of most users, which often were thought to be easily met with software, the concept became prevalent during this period. In a way, Digital Equipment Corporation took advantage of this thinking when it produced computers aimed at a particular audience such as the research community, whose members had moved beyond computing and were interested in data recording and interacting with experiments, something that batch computing could not do.

Vendors

The major vendors were now IBM plus the BUNCH—Burroughs, Univac, NCR, Control Data, and Honeywell. However, the upstart Digital Equipment Corporation, or DEC, was formed and became a significant player in the market.

Technology Terms, Acronyms, and Buzz Words

Core—a substitute word for main memory coming from the technology that greatly increased the amount of storage. The use of core led to terms such as in-core instead of in memory in references to different methods of doing work.

Baud—the number of bits per second transmitted over a communication line. The number of bits per second transmitted, for example, between a typewriter terminal and a computer depended on the coding scheme of the number of bits per character. So, for example, if 10 bits are used to represent a character, then a 300-baud line would transmit 30 characters per second. At higher data transfer rates, the term bits per second is used.

Data base management system (DBMS)—A data base, in contrast to a file, brings together a collection of interrelated data and integrates it in such a way that it is accessible by many users, with each user possibly having a different logical view. A DBMS manages this information by defining the structure of a relational data base in terms of tables, columns, and various constraints.

Data processing—As the transition from punch card technology to digital computer technology was taking place in the business world, data processing (DP), or administrative data processing (ADP), became the new term to distinguish this type of work from number-crunching computing, which researchers dominated.

Emulation—the ability of one digital computer to interpret and execute the instruction set of another computer. Generally speaking, the emulator program running on, for example, computer A would read the machine code of the computer being emulated, computer B, and create an equivalent set of instructions and data storage for computer A, then execute the program on computer A. Emulation was often used to eliminate the need for computer B, especially when the continuing costs were high and the workload was small. Depending on the relative speeds of the two computers and the differences in machine code, it was possible for the program to run more slowly under an emulator than on the original machine.

K—an abbreviation for 1,000. While technically incorrect—the real number was 1,024 (2 to the 10th power)—K was generally accepted as referring to 1,000.

Modem (modulator/demodulator)—In its early formulation, a modem allowed digital information (in bit form) to be transmitted over analog (in wave form) transmissions of telephone technology. Information at the terminal end was converted from digital to analog for transmission and, at the computer end, converted back to digital for the computer to use as required.

There were dial-up and direct-connected modems; dial-up modems (acoustic couplers) were used so that one telephone line could be shared between voice and computer communications. Speeds varied from 110 bits per second for typewriter-like devices to 4,800 or 9,600 baud for higher-speed devices.

Multiprogramming—the overlapping and interweaving of the execution of more than one program over a period of time; also known as multiprocessing, which allows several applications to run concurrently. For example, the control unit, instead of waiting for a long input/output (i/o) operation to be completed, goes ahead and executes another program until it is notified that the i/o operation has ended. At that point, the control program can go back to executing the original program or another task in the queue it maintains. Multiprogramming depends heavily on system interrupts that alert the operating system of events.

Object code—the basic machine language of a particular computer that is needed to run a program on that computer (often referred to as the object computer).

Source code—the high-level language used to develop a program or application—for example, a program in Fortran. In addition to the advantage of writing a program in some combination of English and mathematics, writing in such higher-level languages often meant that the program could run with no or few changes on other computers that had such compilers/translators to generate their own “object” code.

Subroutine—a subprogram of a program written in a higher-level language that performs a specific task. Subroutines were often used for mathematical functions but were also used so that a coded routine could be called whenever needed. Each such call would be given parameters or arguments on which to operate. One of the real advantages of languages such as Fortran was that many mathematical subroutines were built into the compiler and did not have to be coded by the programmer.

Throughput—the time it took a computer system to process a known amount of work. Generally speaking, the more powerful the central computer, the larger the throughput. But throughput could be influenced by the amount of memory and the speed of the supporting components, so that the goal was to have a reasonable balance in the performance of the interacting components. For example, a fast processor with small

memory might have lower throughput than a slower processor with large memory.

Time sharing—organizing a computer so that a number of users could interact with it simultaneously, each user operating independently and often at locations remote from the computer itself. Sometimes known as a multi-user system. This was in part achieved by time slicing, or giving each user a predetermined amount of time (time slice) for executing their program before switching to the next program in a rotation.

Turnaround time—the time it took for a single batch job to be completed. This was typically the time elapsed between reading the cards in the card reader and printing the results, but as operating modes changed, it became the time period between the job entering the system queue and finishing execution. Typically, small jobs were expected to be turned around in minutes, long jobs in hours, and very long jobs overnight.

Virtual memory—the ability of a digital computer to simulate a main memory larger than was actually available. It often provided a storage space so large that programmers did not have to worry about space allocation and did not have to reprogram or recompile their programs when the amount of memory changed.

1960 to 1969 at Cornell

The 1960 decade at Cornell was noted mainly for the bold step taken to combine the academic and administrative computing centers and to place all computing on a shared machine located off-campus. This was accomplished while coping with a new computer architecture and very immature operating systems and communications and applications development software.

1960 to 1965

The Development of CORC, the Cornell Compiler

Dick Lesser continues his narrative on the Burroughs 220 system.⁵

The next programming advance at Cornell came in the early 1960s, when Professors Richard Conway and William Maxwell of the Department of Industrial Engineering and Robert J. Walker of the Department of Mathematics wrote an interpretive compiler called CORC, the Cornell Compiler, similar to the widely used BASIC, which was being developed at Dartmouth. Whereas BASIC had a highly structured, formal syntax, CORC was programmed in English statements and was easy to learn and use. Initially written for the 220, and later for the Control Data 1604, it increased student usage greatly. Programs submitted before five o'clock could be compiled or run during the night and results made available to the student the next morning. By 1963 approximately 1,000 students were using the computers for course work. In 1962 the College of Engineering incorporated computer programming as part of the third semester of calculus, and in the academic year 1963–1964, 14 credit courses related to computing were offered on campus.

The staff continued to grow. Management additions included the appointment of Professor Seymour Parter, Mathematics, as associate director in charge of research, and John W. Rudan, as assistant to the director. In 1960 the staff consisted of 18, including 4 programmers.

Control Data 1604 Installed; Burroughs 220 Sold

Lesser provides the following information about the installation of the Control Data 1604 system.

The three-year cycle continued as Cornell ordered a Control Data 1604-160A com-

puter system for installation in the fall of 1962. The 1604 was 50,000 times the speed of the original CPC. This machine, valued at well over a million dollars, was again financed by a National Science Foundation grant⁶ as well as a substantial educational grant by Control Data.

The 1604 was installed in a new computer room contiguous to the 220, also air-conditioned with raised flooring. Student assistants, both undergraduate and graduate, continued to provide much of what we now call systems programming. Some of these were Philip Kiviat, George Petznick, John Behrenberg, and David Bessel. With the acquisition of the 1604, the center expanded into the second floor of Rand Hall to provide storage space and offices for additional staff. Assistant Professors Shayle Searle, Department of Plant Breeding, and Sidney Saltzman, Department of Industrial Engineering and Administration, joined the staff as research associates. By the time I left in 1964, the Cornell Computing Center full- and part-time staff numbered about 50. One of my last efforts at Cornell was to effect the sale of the Burroughs 220 to a former customer, John Middlebrook, a consultant to General Foods. I believe the sale price was \$75,000, not bad for an obsolete, tube-type machine!^b

By 1964, many Cornell departments were beginning to acquire computers to serve their own research and administrative needs, as were many of our former outside-Cornell customers. Some of these were the University of Rochester Institute of Optics, Syracuse University, GLF, Inc. (now Agway),^c Corning Glass Works, the GE Advanced Electronics Center, and optical companies in Rochester.

Expanding on Lesser's comments on the grant proposal, the following additional comments provide further information about the configuration, the grant, and some of the faculty who supported the acquisition of the 1604. The Control Data 1604 computer and equipment, including the 160A support computer,

⁵ Richard C. Lesser, "Richard C. Lesser's Recollections: The Cornell Computing Center, the Early Years, 1953 to 1964," Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

⁶ Cornell University, "Proposal to the National Science Foundation for Support of High Speed Computing Facilities," undated but presumably 1961.



Figure 1. Control Data 1604 console and magnetic tape units

eight magnetic tape units, a high-speed printer, an X-Y plotter card, and paper tape input/output units, had a list price of \$1,518,000. Control Data provided an educational grant of \$453,300 and NSF provided a grant of \$700,000 to offset these costs, leaving the university to pay \$364,700. An additional \$40,000 was spent for site preparation and other building facilities. Hans Bethe, Robert E. Bechofer, J. L. Hoard, and Ta-Chung Liu are some of the Cornell faculty who provided summaries of their proposed research projects that would benefit from the installation of the 1604.

It is worth noting in this context that in 1961 a technical committee of the Computing Center recommended the installation of the 1604 computer.⁷ In their investigation of alternatives, they considered the IBM 7090, the Philco 2000, the RCA 604, the Univac 1107, and the Burroughs B5000 and came down to the IBM 7090 and Control Data 1604 before recommending the 1604 for both technical and operating concerns. IBM made a last-minute proposal to Cornell after they learned of the university's decision to acquire the 1604. They proposed the installation of an IBM 7090 at Cornell and offered to contribute \$250,000 to help form a bio-mathematical center in combination with the Cornell Medical School–Sloan Kettering complex in New York City. This offer was declined.

⁷ R. W. Conway, R. C. Lesser, Seymour Parter, R. J. Walker, "Report of the Technical Committee on the Purchase of a Large Computing Machine," Cornell University, July 1961.

Cornell Computing Center—John W. Rudan, Director

When Dick Lesser left for SUNY central in 1964, John W. Rudan was appointed director of the center, reporting to Franklin A. Long, vice president for research and advanced studies. In July 1965 a number of changes were made to improve operations and services to users. The practice of allowing staff and graduate students to operate the computers by themselves and sign up for blocks of time ended. The second operator, hired earlier, was deployed to extend supervised computer operations to two shifts so that more runs could be put through the 1604-160A complex and additional services such as plotting could be done under staff supervision. A cadre of students was hired and trained as operators for staffing operations during weekday third shifts, weekends, and occasionally, any shift during periods of high activity, such as the end of the semester instructional crunch. David H. Bessel, who for years was one of the top technical staff members of the center, was promoted to assistant director of the center.

Changes in Rates and Operating Practices

Differential computing rates, distinguishing the limited- and high-demand weekday working hours (first shift) from the lower-demand weekday nights and weekends, were implemented at this time also. For projects supported with university funds, the weekday first shift rate was \$108 per hour and for all other shifts, \$60 per hour. The commercial rate was \$180 per hour. In all cases now, hours were obtained from an internal clock in the 1604 computer and not from a time clock at the console. In addition, restrictions were placed on the type of work processed during these different shifts. For the day shift, preference was given to short, limited-resource jobs, such as for student programs, to ensure quick turnaround. Longer jobs requiring more resources, in particular magnetic tapes, were restricted to the other off shifts.

Another major policy and operation change during this period was to make the Computing Center programming staff no-charge consultants instead of fee-for-service programmers. Since the beginning days of the center, sponsored research projects were treated on a "closed shop" basis. The assigned programmer would translate the client's problem into a computing language, run that program on the computer, and when satisfied with the results, return them to the client. As the base of programmers on campus continued to increase, more and more questions were being asked of the Computing Center staff, who then had less time for programming for hire. It had been a practice for some years prior to this time to encourage users to develop and provide their own programming staff as a

way of expanding the use of computers. This change in policy formalized the practice and was the first deliberate step to increase the base of computer technologists on campus beyond those in the Computing Center.

Finally, and perhaps most important for its long-lasting effect, the use of the computer for instructional courses and nonsponsored research was changed from a free (no charge) to a billed service, so that computer use became part of the base for allocating costs and resources. The Computing Center budgeted for the support of instruction while the deans of the schools and colleges were given allocated funds to support use of the computer by research staff in their respective schools or colleges. To provide for a smooth transition, appeals could be made to the Computing Center in the event there were problems in the coming year as a result of the new policies and procedures. This was the beginning of allocation funding, or more cynically referred to as funny-money funding—that is, general-purpose university funds restricted to pay for the use of computing resources.

The CO-OP Monitor system developed by Control Data and the users group was used to control system operations. The basic operations style was to use the 160A for all input/output operations, in particular, the reading of program and data from cards onto magnetic tape and printing output from magnetic tape onto paper. All input/output of the 1604 was done with magnetic tape to improve the efficiency and effective use of the 1604 as the computing engine. Program libraries were maintained on a permanently mounted library tape for easy and direct access by user programs. It was in this period that Control Data developed an innovative card reader that could read up to 1,200 cards per minute. This was a major improvement in speed, and because the device moved the cards with a vacuum drum instead of mechanical pickers, it not only operated faster but was more tolerant of bent cards and bent corners of cards. The major drawback was that when the machine malfunctioned, and it surely did, quite a few cards had to be fixed!

By this time, the earlier initiative to create professorial partnerships between the Computing Center and academic units came to an end. Searle and Saltzman were no longer affiliated with the Computing Center and Parter had left Cornell. Searle joined the Biometrics Unit in Plant Breeding as an assistant professor, and Saltzman returned to Industrial Engineering and Operations Research as a lecturer, although he came back in 1968 as the assistant director for academic computing. Both went on to have very successful professorial careers at Cornell.

The financial policy of balancing the budget with earned income from research grants and other sources continued. Processing administrative work for GLF

(Agway) and research affiliations with the Cornell Aeronautical Laboratory were key parts of this discipline. The issue became a little more complicated when rates had to be recomputed at the end of each fiscal year and the difference between the billing rate and the actual rate had to be refunded to each agency. To avoid such complications, an agreement was made to roll forward any such differential into future rate computations as a way of avoiding the transfer of funds after the fact. This helped the budget process in future years.

Instructional Computing

CORTRAN

CORTRAN was developed to provide a complement to CORC for student use as well as for researchers who were interested in low-cost, fast turnarounds for their Fortran jobs. A team of staff and student programmers under the leadership of Bessel developed CORTRAN^d as an in-core version of the Fortran compiler provided by Control Data. The availability of CORTRAN not only improved services at the center but conserved resources and improved the performance of the 1604 system.

More about the Use of CORC

Although Lesser provides some interesting information about the use of the Burroughs 220 and the Control Data 1604 computer for course work by students, additional information may provide a better sense of the practices at the time and the load volume generated. When the CORC process started, it could be considered rather archaic—students did not prepare their own input but wrote their programs on specially prepared coding sheets that facilitated both the writing and the transfer to punch cards by professional keypunch operators. Figure 2 gives a sense of the situation.

Table 1 presents the number of CORC programs from the initial start of the program in 1962 until 1966.

At the time it was the practice to record initial runs, that is the first run after the program deck was punched by keypunch staff, and the reruns, which were subsequent runs after corrections were made by students and until the program produced the expected results. One measure of the growth is that initial runs increased by 2.5 times from 1962–63 to 1965–66. Reruns showed a four-fold increase over this same period. These increases give witness to the increasing use of the computer in the instructional programs at the university.

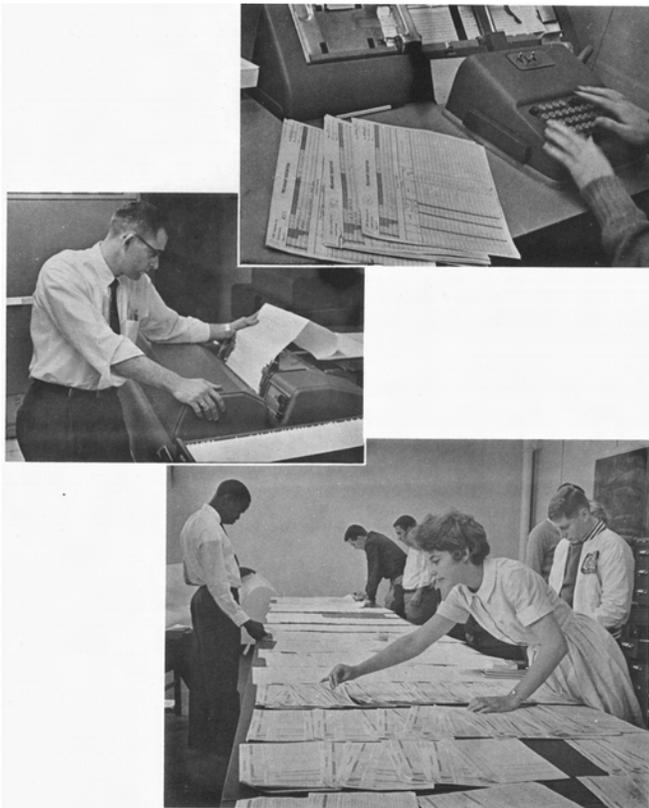


Figure 2. A brief pictorial of the CORC process for student programs. The original caption for this pictorial in a Cornell Computing Center publication circa 1963 was: "With the CORC simplified computing language, all Cornell students can make use of the Computing Center. After a student has completed his program descriptions, it is punched on cards by the center staff; run through the computer; then returned to the center laboratory for correction and resubmission by the students when necessary."

Instructional Statistics

The computer hours billed to instructional use from 1962 to 1965 attests to the increased use of computers during this period. From 1961 to 1962, when CORC processing first started, the recorded use on the Burroughs 220 computer rose from 268 hours to 453 hours. From 1963–64 to 1964–65, instructional use of the Control Data 1604 computer rose from 349 to 577 billed hours. The availability of both computers in the 1962–63 period and the differences between these computers makes direct comparisons over the full five-year period difficult, but it is obvious that the trend was toward increased use for instructional computing.

The courses Mathematics 293/294 and EPM^e (Engineering Problems and Methods) had the largest enrollments—an estimated 500 students each—while Industrial Engineering 9381 (formerly 3281, renumbered), the basic introduction to digital computing and programming, had an enrollment of 100 students. Other courses in Engineering and Arts and Sciences and Agriculture (about seven) had enrollments of 10 to 40 students.

Other statistics from this early period show that it took about nine minutes of keypunch operator time to prepare the typical program of 30 to 35 cards, and it took the Burroughs 220 computer one to one and a half minutes of cpu time to complete a problem. Comparison times were not recorded for the Control Data 1604 due to the method of operation (cards were first read to magnetic tape, and a collection of programs were run through CORC as a single job).

Table 1. Cornell Computing Center: Number of CORC Programs Run, 1962 to 1966

Year/Semester	Keypunch Initial Run	Rerun	Total Semester	Total Year
1962/fall (Control Data 220)	2,900	4,100	7,000	
1963/spring (Control Data 220 and 1604)	2,300	2,800	5,100	
Total				12,100
1963/ fall (Control Data 1604)	5,900	9,500	15,400	
1964/spring	6,000	9,100	15,100	
Total				30,600
1964/fall	6,900	10,900	17,800	
1965/spring	4,300	8,800	13,100	
Total				30,900
1965/fall	7,800	15,900	23,700	
1966/spring	5,300	13,500	18,800	
Total				42,500

Research Computing

The system of recording computer use with time cards and a console-mounted time clock, supplemented by computer logs, continued into the 1960s and worked quite well when the typical time increment for computer scheduling was an hour. The key information obtained was the name of the individual, his or her department, and the billing account number if the use was sponsored research. The few records that survive from 1960–61 are not complete but provide a starting basis for presenting the amount of research use of the computing facilities at the Computing Center. Note that in Table 2 there are only two years where the data are complete, including the total available hours of billed use; the difference between this figure and the total research hours is largely for hours of instructional use.

The chart shows that research use of the computers at the Computing Center exceeded 85 percent for the years for which the data are complete. It is a safe assumption that this same statement could be made for the other years given the conditions at the time.

Over this period, two colleges accounted for more 80 percent of the research use—Arts and Sciences and Engineering. Within Arts and Sciences, the Chemistry and Physics Departments and the research centers closely associated with Physics were the major users. Within Engineering, Industrial Engineering and Administration, Electrical Engineering, Thermal Engineering, and Civil Engineering were the major users. Toward the end of the period, computer use in the College of Agriculture was large enough to place it among the more frequent users.

Business Systems Data Processing

Machine Records Installs IBM 1401 Computer

In 1960 Machine Records started a major transition from tab shop operations to computer operations when it ordered and installed the first computer dedicated to administrative applications, an IBM 1401 card system. This represented a radical change in the way

new transaction processing applications and reporting systems would be developed. Charlie Evans, one of the senior staff members at the time, provides a historically fascinating insight into this transition:⁸

They gave me the choice of a 1401 with 1,400 positions of storage with multiply and divide or 2,000 positions of storage without multiply and divide. I took the 2K of storage and wrote my own multiply and divide routines, which were very efficient. Besides, you divided so little. When you did, you tailored the code to the application and played strange games.

This decision had significant future ramifications, but those were the kind of tradeoffs being made at the time. Evans also recalls that IBM really didn't understand how to position the 1401 computer, because all their programming literature made comparisons to 407 tab operations.

In 1961 magnetic tapes were added, mainly to store up to 100,000 mailing labels for alumni and replace the mechanical addressograph-based mailing system. Other upgrades were made each year, but the system's expansion had reached its limit by 1966. An IBM 360/40⁹ was ordered, but the order was canceled as the situation changed.

Business systems were developed on the 1401 using SPS (symbolic programming system) and later the Autocoder language, both provided by IBM. Autocoder was an assembly programming language that relieved programmers of remembering operation codes and allowed them to use symbolic addresses and assign information to storage positions. However,

⁸ Charles V. Evans, "Oral History, 1957–1994," Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

⁹ Dominic Bordonaro, "Dominic Bordonaro, Oral History, 1947 to 1968," Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

Table 2. Cornell Computing Center: Billed Hours to Research Computing, 1960 to 1965

Year	Computers	Nonsponsored Hours	Sponsored Hours	Total Research Hours	Total Available Hours
1960–61	B220	2,889	584	3,473	
1961–62	B220	3,871	1,063	4,934	5,212
1963–64	Control Data 1604	3,379	538	3,917	
1964–65	Control Data 1604	2,902	992	3,894	4,471

given the shortage of storage positions and the relative slowness of the assembler itself, once the program was in machine language, its use of storage, as well as run time, was improved by manipulating the object code. This itself was not a problem because the basic machine program deck could be punched onto a set of master cards, which became the production program. However, often these programming changes were done by changing the master cards themselves using stickies (a tape that covered the hole punched out in a card) to cover some holes and punching new holes to reflect the changes. As a result, there was no real documentation for the program other than personally reading the deck of cards to discern the entities involved and the operations performed. Although this came to be a major problem later in the decade, at the time, given the slow machines and limited staff, shortcuts such as these were considered the only way to keep up with the growing demands.

Machine Records Staff—Dominic Bordonaro, Director

By 1964 the Machine Records staff had grown to 26 people, up from 12 people in 1960. Correspondingly, between staff and equipment additions and inflation, the budget had grown from \$145,000 to \$275,000 in that same period. The key staff members of Machine Records during the early 1960s included Dominic Bordonaro, director; Fred Hoffman, operations manager; Smoky Stover, responsible for 1401 operations; and Bud Bradt, operator. Addison Locke was the tab room supervisor, and Jerry Buckland started as a tab operator. Other operators were Ruth Delong and Irene Van Zile. Charlie Evans was the programming supervisor, and his staff consisted of Danny Bahn, Al Seliga, Herm Kramer, Bruce Lloyd, and Carolyn Baker. Ruth Delong, and then later Irene Van Zile, supervised the keypunch section.

Statutory Finance and Business Office; IBM RAMAC Installed

In the Statutory Finance and Business Office (SF&BO), Robert Walsh replaced Lloyd Slater as director. In 1962 an IBM RAMAC 305 (Random Access Method of Accounting and Control) system was installed, a first step toward replacing the book-keeping machines used for financial records with a computer. Dan Argetsinger, who had joined the office in 1961, still recalls that the RAMAC had 5,000 electronic tubes and 8,000 characters of main memory.¹⁰ It was endearingly referred to as the “iron woodpecker,” because the read/write head moved vertically to find the appropriate stacked disk platter and then horizontally in and out as it read or wrote information on the platter. Given the slow speed of the device, it was still

necessary to presort transactions to minimize the head movement. In support of the RAMAC there was an attached 407 printer for different reports. The IBM RAMAC lasted until later in the decade when it was replaced by an IBM 1440 computer system.

Argetsinger also recalls how he became a computer programmer. Walsh requested that he take an IBM programmers’ aptitude test, and as the highest scorer of the staff members who took the test, Argetsinger was at programmers’ school within a week. In those days the programmers’ aptitude test was a common way for selecting and developing programmers. Bob Mack, who later worked with Argetsinger on developing the university accounting system, recalls that he also became a programmer by scoring high on the aptitude test.¹¹

1965 Business Systems Studies

The year 1965 was one of studies and reports on the future directions of computing, both with respect to equipment for academic and administrative use, as will be discussed, and to the organization and deployment of staff resources for business data processing. A number of position papers were prepared for consideration and discussion on this matter. Conway, Rudan, and Saltzman prepared a comprehensive plan for data processing at Cornell¹² that called for the creation of a consolidated administrative data processing unit. This unit would have responsibility for the work being done by Machine Records and the Statutory Finance and Business Office and would have its own machine (an IBM 360/40 was recommended), which would replace the 1401 and 1440 being rented from IBM. The plan recommended that the group and equipment be housed in Rand Hall and operate under the name Cornell Information Processing Center, or CIPC.

While this activity to develop the CIPC concept was going on, IBM proposed a joint study to examine the administrative information systems at Cornell. Cornell accepted the proposal, and a team of IBM and Cornell staff was assigned. The Cornell staff consisted of Bordonaro, Evans, Rudan, Saltzman, Wadell from Dairy Records, and Walsh. This was a comprehensive effort involving interviews with more than 50 administrative offices and operating units and subsequent meetings to make sure the problems and issues were

¹⁰ Danny Argetsinger, “Oral History, 1961–1996,” Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

¹¹ Bob Mack, “Oral History, 1965–Late 1970s,” Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

¹² R. W. Conway, J. W. Rudan, S. Saltzman, “A Proposal for the Organization of Data Processing Facilities at Cornell,” Cornell University, July 1965.

correctly understood. Later, 41 offices were requested to provide more comprehensive information about their information system needs. The final report made the following comments and conclusions:¹³

The academic theme of Cornell, “freedom with responsibility”—the long and proven valuable attribute of faculty relations—appears to have influenced the administrative hierarchy of the university in its management philosophy. The continuance of independence of action by administrative units and the single purpose and separate job responsibility is the apparent method of management as it progresses toward stated goals.

From this followed the conclusion:

To summarize, many departments have individual files which supply enough static information to get today’s job done. However, this information is suspect as to timeliness, may be redundant, is difficult to acquire by other departments, and generally is not useful for university planning purposes. There is an abundance of information available today, but scattered, unorganized, undefined, and impossible to integrate.

And further:

The committee estimates that not more than 10 percent of the current information-handling process is automated. The estimate of 10 percent is subject to a downward revision approximating zero when evaluated for the ability to generate comparative reports of historical data, projections for planning decisions, and any other task that requires a level of sophistication beyond simple addition and subtraction.

The report also made the following recommendations:

The establishment of a totally integrated information system. The university take the necessary actions to establish a totally integrated university-wide information system.

Management services group. A management services group should be established immediately. It will be composed of a director, system analysts, operations research

specialists, and programmers. To ensure the recognition of the broad responsibilities of the group, the director should report directly to the provost of the university. It should have two immediate functions: (a) plan and implement the total information system; (b) advise the administrative areas of the university on data processing-oriented applications.

The organization of computing resources. The committee recommends that there should be a single organization with responsibility for all data-processing applications. This does not call for the immediate reorganization of present computing installations but a planned conversion period.... The committee makes no specific recommendation on what computing capacity will be needed by this data-processing organization.

Fiscal policies for administrative computing. The committee recommends that the present policy of financial support for Machine Records undergo substantial revision and a new concept of cost sharing of computer facilities be developed. As a service function, the data-processing organization should not operate as an income-producing office, where its budget must be balanced by income from other university departments. This is not to say that there should not be rates.

Consideration of costs and means of implementation. The university administration should not require a definite assessment of potential dollar cost for development and maintenance of a total information system. The ultimate costs cannot be determined until *all* applications and the related elements of support are identified. However, the system can be developed in segments along department lines and/or functional areas with fairly accurate estimates of cost determined for individual segments.

A supplement to the report was issued later dealing primarily with recommendations for organizing and staffing the management services group.

No specific actions were taken to implement those recommendations, but they influenced future decisions on the key issues for improving campuswide business systems.

¹³ “Report of the Cornell-IBM Information Study Team,” Cornell University, November 1965.

User Groups—CUMREC and CAUSE

Cornell was an early and active participant in CUMREC, the College and University Machine Records Conference, a group that was supported by IBM and fostered the best use of card-based data processing practices in higher education.¹⁴ In 1956 Frank B. Martin, director of data processing at Michigan State University, called for a meeting of higher education professionals to share their expertise and experience with computerized systems. The group formed CUMREC, and each year a conference was hosted by a member institution. Bordonaro and the staff in Machine Records attended these conferences, both contributing papers and getting insights from others.

In 1962, 22 data-processing directors organized an IBM 1401 users group at a meeting in Chicago, for the transition had started from the unit record card equipment to the 1401 computer. Bordonaro, representing Cornell, was in this first group of organizers, which called themselves the College and University Systems Exchange, or CAUSE. Their objective was to share information about the new administrative information systems that they were beginning to develop. Later in the 1970s, with the help of a \$10,000 grant from IBM, CAUSE was formally organized and Bordonaro served on the first board of directors, although at that time he was with Ithaca College. In a way, that continued the tradition of Cornell starting user groups.

1965 to 1967

The Office of Computer Services (OCS)—R. W. Conway, Director

By the mid-1960s, when the 1604 system was approaching saturation after about two years, it was not clear which technology direction to pursue. The prevailing attitude of getting the most “bang for the buck” again brought into question whether the different organizations running their own computers might be brought into various consolidation configurations to further improve efficiency and reduce costs. At the same time, Vice President Long and the executive committee of the Computing Center were concerned about the current and future status of computing services at Cornell and took steps to evaluate the situation. They retained Dr. Bruce W. Arden to review the status of computing services at Cornell and make recommendations for improvement. Arden was a faculty member in the Computing Center at the University of Michigan and, at this time, was quite well known in academic computing circles.

Arden’s primary recommendation was that the Computing Center and Machine Records be combined into a single computing organization and that a single machine be used as their primary computing engine.^f

Although this was an innovative step at the time, it was not recognized for how forward looking it was. When put into practice, Cornell was one of the first, if not the first, major universities to combine these two different operating groups into a single organization using the same machine. Other universities came to this same arrangement 20 years later for the same reasons—cost savings and staff synergy!

In late 1965 a committee representing the three different campus computing centers (Cornell Computing Center, Machine Records, and Dairy Records) prepared a report¹⁵ on computing equipment at Cornell with the following summary:

It is this committee’s recommendation that Cornell take the necessary steps to establish a single, consolidated data-processing facility by the end of 1967.

1. To ensure the availability of suitable equipment for such a facility, a letter-of-intent for a 360/67 should be placed with IBM as soon as possible.

To provide necessary interim capacity and to facilitate transition to consolidated operation, a 360/40 should be ordered from IBM as soon as possible for installation next summer in Rand Hall.

2. Work should begin immediately to detail the organizational structure, operating procedures, and equipment configuration for a consolidated facility that would be entirely satisfactory to the Department of Animal Husbandry (and through the department, to DHIA), the major administration users of data processing, and the users represented by the Cornell Computing Center.

In the course of their evaluations, the committee members had considered equipment from six vendors. Burroughs was proposing to build their B8500 system, an outgrowth of their military work, which would provide fast computation and an advanced virtual memory feature. Control Data was already selling their 3600 system, successor to the 1604, and was proposing their 6400 system, which was possibly the first of the supercomputers^g of the era. None of these alternatives seemed attractive in features or price. General Electric was proposing their model 645 system. GE was an early leader in time-sharing systems and had secured orders from Project MAC at MIT and Bell Labs.

¹⁴ Taken from the web site www.cumrec.org.

¹⁵ D. Bessel, D. Bordonaro, R. Conway, C. Henderson, J. Rudan, and S. Saltzman, “Recommendations on Computing Equipment at Cornell,” Cornell University, October 8, 1965.

Honeywell and Univac were not serious contenders but offered equipment that had to be considered. IBM had already announced their new 360 system, but the initial announcements largely continued the batch processing style of computing. Responding to market pressures for time sharing in higher education, largely coming from GE, in 1965 IBM announced their Model 360/67 time-sharing system. IBM put together a convincing road-show^h about the 360/67,ⁱ and some 25 orders had been placed at the time the report was written. Conway¹⁶ recalls that in the end, IBM convinced over 150 universities, including Cornell, to place orders for this machine.

The executive committee of the Computing Center accepted Arden's recommendation and the equipment recommendation and moved to put them into effect with a complete new management and director. They offered Arden the position as director of this new combined organization, but he declined. Given the large number of commitments that had been made about the new organization and hardware, the executive committee prevailed upon Conway to assume the directorship. He agreed to a two-year term effective July 1, 1966.^j

Conway set about building his management team. He recruited Jeremy (Jay) Johnson from outside Cornell as associate director to help build the new organization and services. Bordonaro continued as assistant director for administrative services, and Saltzman returned as assistant director for academic services with oversight over the Finger Lakes Colleges Computing Center (FLCCC), a regional initiative with other schools. Bessel, who had been appointed director of the Computing Center, continued in the position until the center at Rand Hall ceased to exist.^k Pulleyn continued as manager of operations, responsible for all activity in the computer room and support functions and the campus terminals. Robert R. (Bob) Blackmun was named the first business manager for the computing organization. William S. Worley was recruited from the University of Chicago to head up the systems programming staff. This last recruitment was essential for overcoming some of the operating system shortcomings resulting from IBM's overselling of the 360/67 system.

A grant proposal for \$1,500,000 was submitted to the National Science Foundation,¹⁷ for funding for the

expansion and building of computing services in support of instruction and research. The proposal called for installing a simplex 360/67 at Langmuir Laboratory with an initial complement of 60 typewriter terminals and a satellite IBM 360/20 on campus for remote batch entry and for establishing a special rapid-cycle courier service between campus and Langmuir Lab. A second 360/20 was to be installed at Langmuir to handle support services such as printing, card reading, etc., mainly for administrative systems. It was estimated that the total cost of the project would be \$5,540,000 over the first three years. Contract research income was estimated to be \$1,225,000 over this period, with Cornell providing \$2,145,000 and other grant requests providing \$670,000. The grant was for a two-year period ending June 1969, and the funds were critical to support the increased expenses of OCS. However, NSF only approved a grant of \$650,000, which became available in mid-1967, and budget adjustments were required.

In a significant departure from past practice, it was decided that only a portion of the computing system was to be purchased, and the remainder was to be leased. The 2067-1 cpu, core storage, channels and drum storage (for performance), and controller and remote communication controllers and teletypes were to be rented, based on the supposition that this equipment was most likely to be replaced in a shorter time frame. Magnetic disk and tape units, card reading/punching and line printing units, and the 360/20 unit were to be purchased, based on the assumption they had a longer lifetime. The 512K-byte core storage, the drum, and the disks (in IBM terms, DASD—direct access storage device) were deemed to be a minimally viable time sharing system. The data cell drive (the first mass storage system at Cornell) was to be purchased to support the work of Professor Gerard Salton in Computer Science in information retrieval.

The Founding of the Department of Computer Science in 1965

During this intense period of computing activity on the campus, Conway also was involved with the formation of the Department of Computer Science. Conway and R. J. Walker, with the strong support of Andrew S. Schultz Jr., who was dean of the College of Engineering at the time, wrote a proposal to the Sloan Foundation and received a grant of \$1 million to help form the department. This graduate department was established jointly in the Colleges of Engineering and Arts and Sciences. By mid-1966, the annual report of the department, written by Juris Hartmanis, the first chairman, notes that there were three full professors, two associate professors, three assistant professors, and

¹⁶ R. W. Conway, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

¹⁷ Cornell University, "A Proposal to the National Science Foundation for a Facility Grant for the Expansion of Computing Facilities," Principal Investigator, R. W. Conway, July 1966.

one visiting professor. While listed as a full professor, Hartmanis noted that Conway “will have to dedicate his time to his new and difficult job as director of the Office of Computer Services, which is responsible for the new IBM 360/67 computing system. This will be a very severe loss to the department; on the other hand, we fully realize the importance of Mr. Conway’s new position, since there cannot be a first-rate Department of Computer Science without a first-class computing facility.”

Central Computing Moves Off-Campus to Langmuir Laboratory

The major issue that had to be quickly resolved was the location of a central computing facility and the campus locations for computer access laboratories for researchers and students, referred to as campus terminals. Up until this time, almost all input card preparation, processing, and printing of results was done at the computer site. Departments had a few more card punching machines than were recorded in the 1957 report on data processing and computing at Cornell University¹⁸, but program and data cards still had to be carried to the computer site for processing. A key feature of the 360/67 was the use of remote terminals, teletype-like devices that served as data entry and output devices and which were connected to the computer by telephone lines. This technology made it possible for users to access the computer from their own offices or from equipment clusters in campus terminals, most of which were also expected to include a high-speed printer.

Although Rand Hall had the space for siting a computer center, Thomas W. Mackesey, then vice president of planning, quickly ruled it out as an option. Mackesey informed the group researching options that Rand Hall was to be torn down in 1969 and a new, better-looking campus gateway¹ built on that site. As chance would have it, at about this time the General Electric Company was closing down its research operations in Ithaca and offered the university its buildings on Brown Road near the Tompkins County Airport for the proverbial one-dollar selling price. The university accepted the offer and Langmuir Laboratory became the prime new site for central computing. Given the amount and quality of space and the critical fact that principal access to the computer could be over telephone lines, the decision was made to locate the equipment and staff of OCS at Langmuir. A budget of \$100,000 was authorized to build the new computer

room and refurbish the space for ancillary equipment and staff. One feature that favored the Langmuir site was the availability of many parking spaces for employees at no cost.

Coping with the Failure of the IBM 360/67; 360/65 Installed

As OCS was taking shape in 1966, rumors began circulating that TSS, the time-sharing system for IBM’s 360/67 system, was not performing up to expectations. That is, instead of being able to simultaneously handle hundreds of users doing different kinds of tasks, it could barely handle 10 users. In January 1967, IBM confirmed these rumors. This failure of the 360/67^m precipitated a crisis in OCS, because the whole program for computing at Cornell had been built on the premise that the 360/67 would be primarily a time-shared system and permit what was referred to at the time as conversational computing. Batch jobs could be executed from the time-sharing component or by direct submission to the batch processing subsystem. Although OS/360, the operating system for the 360 line of computers, was capable of handling routine batch jobs, it was deficient in its ability to support batch jobs submitted from remote sites or handle the large volume of small batch jobs submitted by students.

In early 1967 a variety of new IBM equipment was installed at Langmuir. A 360/40 computer was installed for training and preparation for delivery of the 360/67 later in the year. Consideration was given to taking delivery of a 360/65, which was the underlying hardware of the 360/67 but without the hardware and software for TSS. There was some early expectation that the 360/40, installed with a special option, could be used to run 1401 programs directly and so permit the 1401 system to be removed. A 360/20 was installed to act in several capacities, one of which was for card operations, line printing, or magnetic tape support of the 360/65. A new multi-function card (reader/punch) machine (MFCM) came with the 360/20 to carry out many of the traditional card-processing operations on individual card machines. However, an important reason for installing this system was to evaluate its capabilities for a remote batch terminal on campus. Soon after all this equipment was in place, the decision was made to order a 360/65,ⁿ and attention shifted to providing the best possible services within the constraints of the budget and available technology. Plans were announced that in the fall, the 360/65 would be run as a multicomputer system, with the 360/65 operating as the computing engine and the 360/40 as an attached support processor, or ASP, for input/output operations and control

¹⁸ W. S. Gere, B. McK. Johnson, R. C. Morris, “An Analysis of Data Processing and Computing at Cornell University” (Cornell University, November 14, 1957, copy number 10).



Figure 3. The IBM 360/65 computer at Langmuir Laboratory

of remote devices. Limitations of the ASP configuration shifted the emphasis from time sharing to remote batch operations, and plans were put in place to have two satellite terminal locations on campus, in Upson and Clark Halls.

Announcements at the time estimated that the rough charge for the 360/65 would be \$350 per hour, comparable to an estimated rate of \$140 per hour for the 1604 time, a favorable comparison for a system that also was estimated to be five times faster. However, to take advantage of the operating system features, a more sophisticated priority-charging system was introduced. The intention was to make sharing of the computing resources more equitable and for the cost of the job to reflect the expected job turnaround time. Six priority levels were considered, with the highest level being for fast turnaround for small/short jobs and decreasing in cost and turnaround for overnight processing; fill-in work was at the lowest level, where turnaround could be five days. An important new element was introduced by charging for main memory—that is, the larger the program and data regions, the larger the cost. Additional charges also were proposed for all the ancillary services, such as line printing, card reading, tape mounting, and online storage of information on disks.

Operating System Changes and Modifications

A completely new job control language (JCL) had to be learned in order to submit work to the system. Classes were scheduled for users to learn this

new JCL and to learn about IBM Fortran and other language processors. CUPL, the Cornell University Programming Language, replaced CORC for small jobs from students. CUPL was developed by Conway, Maxwell, and Walker and built on their earlier work and experience and the improvements in technology.

Conway quickly organized what became known as the OCS skunk works to redress the deficiencies of the 360/65 computing system, which was installed instead of the 360/67. This skunk works included Worley, professors Maxwell and Howard Morgan from Computer Science, and a number of graduate students. Plans were changing rapidly, and by the time the 360/65 was installed in September 1967, the 360/65-360/40 ASP configuration was out of the picture, and the 360/65 was operating alone. The 360/65 memory was increased to accommodate this change to OS/360 augmented with HASP (Houston automatic spooling system)¹⁹ to “spool” jobs. HASP, which was cleanly interfaced to OS/360, was also an efficient way for handling online peripheral operations for remotely located units and for better operator control procedures. To address some of the deficiencies, Conway and Worley developed a plan that spelled out the work that needed to be done to make the system really use-

¹⁹ Initial references to HASP are Houston Automatic Spooling Priority system, and later on as Houston Automatic Spooling Program. The reasons for this change are not clear. It suffices to say that HASP itself was a “skunk works” project done by the IBM lab in Houston without corporate approval or involvement.

ful at Cornell.²⁰ The major deficiencies noted were the lack of multiprogramming during the execution phase, the inability to handle short, undemanding jobs, the lack of a simplified JCL, the lack of an accounting and access control system, and the lack of a convenient and efficient system for routine use of online storage. Of particular interest for administrative applications was the inability of HASP to anticipate the advance mounting of tapes and disks and the lack of a practical method for permitting the use of special forms, special alignment, and carriage control for special printing. This list ended up being the work plan for the skunk works to develop and deploy what became known as COOL/HASP (Cornell online) and later just COOL.

A special accommodation was provided to give fast turnaround to student jobs in CUPL, WATFOR, and ZAP (assembly), all in-core compilers. Part of the job-scheduling system was designed to accumulate programs for these special compilers and then invoke them to process the entire number that had accumulated. Using this discipline, almost all of the between-job overhead was eliminated, and this saving could itself be used to process many such small programs. A simplified Cornell JCL was developed to provide for the advance mounting of tapes and disks and for controlling special print jobs. Supporting the JCL but taking advantage of some of the facilities of OS/360, access control was provided by developing an accounting scheme for users to provide for both access control and subsequent billing for services.

One important development was known as the card data system (CDS), which addressed the online storage and file management of card images on the disks at Langmuir. Such virtual card decks could be edited separately as needed, and when a job was ready to be executed, the cards in the CDS could be inserted into the job stream as if they were part of a job totally on cards. There were significant operational advantages to the use of CDSs, the major one being the ability to reduce the traffic on the limited speed of the communication lines to send card images from campus to Langmuir. Last, a billing system was designed and written to take all the accumulated records from access control and individual job accounting information to create bills for services.

The deployment of COOL brought with it serious performance problems due to the single-threaded system design; that is, COOL was only able to execute one task at a time. The preference given to student jobs resulted in long turnaround times for other batch jobs. Performance overall was not very good, especially

for longer-running administrative production jobs and for similar jobs from the research community. It was noted at the time that almost always 500 jobs were in the job queue waiting for execution. Other solutions were considered.

Fortuitously, IBM introduced the LASP system, which appeared to provide relief for the single-thread bottleneck. LASP (local attached support processor) effectively provided a multiprogramming system; one part spooled work in and out of the system and the other part executed the jobs. A major project to install LASP on a fast track was started by the Systems Programming group in 1968. With the COOL modifications carried forward, the system was renamed CLASP (Cornell LASP), and the expectation was that installing this system would improve the turnaround problems as well as system stability by reducing the number of outages per day.

Campus Terminal Facilities for Accessing the 360/65

To support batch computer access, satellite terminal facilities were located in several different campus locations. As noted earlier, one facility was located in Clark Hall to support the Physics and Chemistry users and users at the north end of the campus. Another was located in Upson Hall to support Engineering and the south end of the campus. A terminal facility also was located in Warren Hall to support the statutory Colleges of Agriculture and Home Economics and the ILR School. Later a card-input-only station was located in Riley-Robb for Agricultural Engineering and the east end of campus.

When first in operation in 1967, these terminals did not have card reading or printing capability, and cards were taken by truck courier to Langmuir and returned, along with any output cards or printer listings, to the facilities when the job was completed.

In 1968 card reading and printing equipment under the control of IBM 360/20 computers (referred to as RJE for remote job entry stations) was installed so that the input/output operations could take place directly at the terminal sites. Each of these sites had a key-punch room for 10 to 20 keypunchers, an area where returned cards were filed, racks for holding computer printouts for retrieval by users, consulting rooms where advice could be sought for problems, and rooms to house staff on a permanent or rotating basis. The courier service still continued to return special print jobs, principally to Day Hall administrative offices, and served other terminal facilities as well, returning materials such as magnetic tapes and plotter output.

Upson Hall was noted for its unique "cage," which was a full floor-to-ceiling partition of steel bars and chainlink fencing that surrounded the RJE station.

²⁰ R. W. Conway and W. S. Worley Jr., "The Cornell-HASP System for the IBM 360," Office of Computer Services, Cornell University, February 1968.

The cage was built to protect the operator and the equipment from harm from overzealous and anxious students. Sandy Pastore was the first full-time operator at Upson, transferring directly from the keypunch section. Engineering students, who composed the majority of users, would hand in their program decks through a small window and return sometime later to find their output—or listings—filed in alphabetical order on the output racks that formed part of the outside wall of the cage. The anxiety level rose as assignment due dates came up, and mid- and end-of-semester crunches prevailed. During the campus disruptions of the late 1960s, this feature of the Upson Hall terminal prevented rioters from trashing the facility.

Academic Support Staff Move to Campus

At about this same time, academic support staff previously located at Langmuir, with only scheduled hours at campus locations, were permanently moved to Upson and Clark Halls on campus. Upson was the nominal headquarters for the staff as well as for the staff supporting instructional use, while the staff at Clark focused on the research users. Jim Manning, who became the supervisor of campus operations, recalls²¹ that Dan Bartholomew and Tom Jones^o were among his first hires as student operators.

Formation of Computer Activities Group—CAG

The Warren Hall terminal, although similar to those in Clark and Upson Halls, had a broader mission than just serving as a computer input/output station. It was to support computing activity in the statutory colleges. In many ways it was the successor to the previous tab shop that had supported the punch card-based technology for cooperating academic units in the statutory colleges. The new organization was named the Computer Activities Group and was referred to as CAG. Shayle R. Searle, then a professor in the Biometrics Unit of Plant Breeding, was the key person who organized this unit. He remembers (personal communication) that in 1966 Charles Palm, then dean of the College of Agriculture, came back from a meeting of his peers from other colleges with a real concern that the college was behind in computing. Dean Palm thereupon convened a committee chaired by Charles R. Henderson of Animal Husbandry, with Searle as a member. This led to Searle being the “grunt man” (as he describes it) for Nyle Brady, director of research in the college at the time, to develop an organization dedicated to helping faculty make better use of computers in their research and instruction. Searle prepared a report in late 1966 laying out the mission, responsibilities, and organization for what was then termed a Computer Service Group but which shortly

after became CAG. Searle, in turn, recruited Errol W. Jones to head up the facility, and Jones became director of CAG in 1968. Previously, Jones had been manager of the computing facility for the New Zealand Department of Scientific and Industrial Research, which started out as a biometrics computing facility, specializing in statistical computing.

When Jones came in 1968, the staff consisted of Joanne Haviland, who oversaw the operations of the 360/20-based input/output facility and the tab shop.²² Alma Coles, who ran the tab shop for many years, had already retired. In addition, Sandy Seafuse supervised a staff of six data entry clerks. Half of those clerks were CAG employees doing work for hire, while the remainder were dedicated to the farm accounting project under the direction of Austin Lowry. Later on, additional programming staff members were hired. These were Betsy Keokosky, Lois White, and Mabel Jennings. Although CAG was officially part of the Department of Plant Breeding and Biometry, practically it reported to the director of research in the College of Agriculture. Overall policy guidance for CAG was provided by a faculty advisory committee representing the different statutory colleges or college departments that had the most interest in taking advantage of computing technology. Shortly after Jones’s arrival, the statutory colleges agreed to fund five positions for CAG—the director, one operator, and three programmers. Other expenses had to be covered by earned income.

In typical Cornell fashion at the time, the funding of CAG followed a complicated cost recovery model. Because the statutory colleges were directly paying some of the costs of CAG, an agreement was made that OCS would reimburse CAG for all income generated by charges for cards read and lines printed at its Warren Hall facility. Reimbursement was straightforward for research users paying for computer use with contract or state funds. It became more complicated when use was funded by general purpose university funds allocated to OCS and not otherwise assignable except for use of computing resources at OCS. This was the so-called funny money that supported endowed instruction and research use. The computer accounting system was designed to track cards read and lines printed at all the terminal facilities by the home college of the user and the type of money used. At the end of every month, CAG would receive a credit of nonbillable funds, which could be used to

²¹ Jim Manning, “Oral History, 1965–1995,” Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

²² Errol W. Jones, “Cornell Computing Oral History,” Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, February 1998.

purchase computer resources only. CAG, in turn, could sell these computer credits for real dollars to generate income to support its costs. According to Jones, this system worked reasonably well for a few years when the amount of money was quite small. However as CAG's business increased and the credits grew to \$35,000 per year, the College of Agriculture positioned itself as the broker for those credits, and it collected the funds for selling the credits instead of CAG. The administrators of OCS, the College of Agriculture, and the central university were often at odds over this agreement, but no side was willing to give up its position. In fact, the credits continued to be a thorny issue well into the 1970s.

Finger Lakes Colleges Computing Center

One important project that was seriously impacted by the failure of IBM to meet the specifications it set out for the 360/67 computer was the Finger Lakes Colleges Computing Center (FLCCC). At the invitation of the National Science Foundation and on the strength of the 360/67 time-sharing capabilities, Rudan wrote a proposal that OCS submitted to the NSF to create an upstate New York computing center to serve the smaller colleges and some of the more progressive high schools.²³ The FLCCC grant was one of a number of such national grants made to gain experience with regional computing centers. NSF approved a grant of \$174,528 for a two-year period. It was a natural step for Cornell to make this proposal, because for many years the Computing Center had sponsored formal and informal training classes, seminars, and demonstrations for some of these educational institutions and occasionally had hired their students as summer assistants.^P The colleges included Ithaca College, Keuka College, Wells College, Elmira College, and Eisenhower College.^Q There were also 11 high schools, including Ithaca, Trumansburg, Interlaken, and Candor, and others that were part of the regional Board of Cooperating Educational Services (BOCES). Rudan was the principal investigator for the grant and the first manager of the project. Saltzman, who had returned to OCS in about 1968, later replaced Rudan.^r

It was not until 1968 that serious activity formally got under way using the name FLARCO, the Finger Lakes Area Computing Center, to represent the broader mission. Colgate University, for example, had joined the group. By that time, NSF had been informed of the probable changes due to the 360/67 difficulties, and additional funds were granted to supplement the budget. Also in that interval, OCS had developed a rudimentary time-sharing system, CTS (Cornell time sharing), and by using IBM 2741 typewriter/terminals and modems, students from the remote sites could access the 360/65. The CTS sys-

tem was designed for data and program entry and job submission to batch for execution, but it sufficed for these remote users. The cooperating high schools also ran a courier service by which coding forms were transported to OCS, punched onto cards, and processed and the cards and results delivered back to the students. Between 1968 and 1970 the Cornell contacts for FLARCO were Saltzman as the principal investigator and Bob Blackmun as the Cornell coordinator, followed by Erik McWilliams and Arlene Larsen, respectively, in those same roles when Saltzman left OCS in 1969 for the last time.

Some of the schools also went ahead and got their own small computers, IBM 360/20s and IBM 1130s, which were used cooperatively with the FLARCO services. By the time the second annual report covering the period June 1969 to July 1970²⁴ was written, it was clear that some inroads had been made in all the schools. In addition to the teaching of programming, use of the computer had expanded to other disciplines such as chemistry, mathematics, and statistics. While the operating problems at OCS did not help the situation, it is also clear that the cost and the unreliable nature of the telecommunications services provided by the small local telephone companies were a definite hindrance. At the time, for example, some of the local companies were still using operators to monitor calls.²⁵ Nonetheless, FLARCO initiated all the schools to the use of computer technology in the classroom, and they all chose different future paths.^s

Machine Records Moves to Langmuir

In January 1967 the entire Machine Records operation left Day Hall and moved to Langmuir. As part of this move, the operations, production control, and data entry staff were merged with the former Computing Center operations and data entry staff. This merger, along with the physical relocation, produced some early problems in running production work. Customers from administrative units had been accustomed to a close hand-holding relationship with Machine Record programmers and operators, who had served also as

²³ John W. Rudan, "A Proposal to the National Science Foundation by Cornell University on Behalf of the Finger Lakes Colleges Computing Center to Establish a Regional Computing Center," Cornell University, July 27, 1966.

²⁴ Erik D. McWilliams and Arlene S. Larsen, "Second Annual Report of the Finger Lakes Regional Computing Organization," July 1970.

²⁵ Mike Oltz, "Computing at Candor High School—Personal History, Regional Computing," Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

production controllers. There were some written procedures for executing the run of a program (with its associated data) to create a data file or end product report, but these procedures were not of the quality needed for routine repetitive running of jobs without user and programmer involvement. There were serious problems getting jobs processed correctly and on time.

Descriptions of several notable problems have survived since that time. In one incident, 10,000 student semester grade reports came out blank when printed on four-part forms because the important step of converting mark-sense forms to punching the grades into the grade cards had been omitted. In another incident a 1401 emulation run went awry due to a programming glitch, and OCS staff had to sort all the student grade slips by walking around tables in the computer room, then under construction, because there were no more pre-printed forms in the inventory. Because of such problems, the actual integration of the production and operations staff was suspended for almost a year while the former Machine Records staff prepared more complete documentation of procedures and protocols for interacting with customers and initiating production jobs. The final transfer of responsibility for production control to computer operations took place in October 1968.

Plan to Remove 1401 Computer

The basic plan for campuswide business systems from the time of the creation of OCS was to eliminate the 1401 system costing over \$40,000 per year and redirect the savings to other priorities. The first approach, as Bordonaro describes it, was to convert the applications from Autocoder to COBOL and then process the COBOL programs on the 360/67.²⁶ Once this was completed, analyst and programming resources would be redirected to developing new applications. The new business systems of greatest priority were for Public Affairs and Library Acquisitions. However, the conversion tools available at the time were not very good; they provided well below the 90 percent compatibility expected, and a considerable amount of programmer rework was required to get an application program running in COBOL. As a result, a new strategy was developed for using a 1401 emulator that ran on 360 hardware. Initially, this was to be done on the 360/40 that was installed in 1967. However, later that year, the 360/40 was removed for financial reasons, and the emphasis shifted to using software emulators on the 360/65. Given that much of the 1401 production code

was in the form of patched object program cards and not actual Autocoder source code, the commercially available emulators could not be used.

As a result, the OCS skunk works mentioned earlier took on the task of building a 1401 emulator that emulated the logic of the 1401 system. When first built, this object-code emulator running on the 360/65 executed code slower than on the 1401 itself. Thomas A. Dimock, who was a rookie programmer at the time, recalls that because the 1401 at Cornell had no “divide” instruction, simulation of a divide took over 250,000 instructions to complete (personal communication). The effort was terminated in 1968 when the total work of maintaining and upgrading the emulator and dealing with the problems encountered with its use seemed to be more work than the value gained. As a result, the 1401 system itself continued to be used for several more years. During the summer of 1969, an attempt was made to convert the remaining 1401 programs and some still-existing tab operations. When it was clear that the project would not be completed over the summer, an assessment of the need for the programs and new proposed systems for payroll and accounting led to a smaller project that focused on the systems that were likely to be needed in the future. It was expected that a successful outcome would lead to releasing the 1401 in early 1970.

Accounting and Billing for Computer Resources

Tracking of the type of resource use, type of client, type of funding, and department or college affiliation was typical of the times. It was almost the case that if you could track or count some aspect of computer resources or its use, you did. In part, this concern to account for everything could be blamed on the computer industry, the directors of computing centers, and the university administrators. When the industry moved from a single-user wall clock measurement of computer time to multiple users sharing a computer's resources under the control of an operating system, it also moved to a synthetic, contrived measurement of computer use by a single user. Coincidentally, because resources were being shared there was valid concern that various facets of resource use needed to be measured to be able to ascertain their effective use. For example, too little available computer memory at a point in time would require a job to be swapped out to disk for a temporary period, and as a result, computer cycles were lost both on the swap out and the later swap back in. Elaborate resource tracking and recording mechanisms were developed to enable systems programming experts to monitor individual component resource use and develop hardware or software solutions to improve system performance.

²⁶ Dominic Bordonaro, “Dominic Bordonaro, Oral History, 1947 to 1968,” Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

Computer center directors, always starved for adequate resources to meet the growing demand, used this detailed information in two ways. First, based on the fairness principle that users should be billed for resources they consumed, they constructed very complicated billing algorithms and systems to bill users for cost recovery. It was not uncommon at the time to have up to six pages of computer rates, full of mathematical formulae and explanations for each component of billing. The situation became even more complex when, in responding to different client interests, a priority system of job scheduling was superimposed on resource tracking, and various multipliers were used to reduce or increase selective charges depending on the priority stated when the job was submitted.

Second, computer directors relied heavily on these statistics to create elaborate justifications for capacity improvements. It was not enough for the central administration to hear the loud and repetitive complaints of key users, even those such as the payroll office or the registrar's office, who had very strict and time-critical pressures to meet their customers' needs. Nor did they listen to faculty and students who had to compete for resources with administrative departments and blamed them for their delays. University administrators, who were generally unfamiliar with the technology and who didn't trust their instincts about these operations, asked for seemingly endless accountability and usage reports, and so reports were developed to provide this information.

In the midst of all the hectic activity at the startup of OCS, the development of a billing system was critical to maintaining cost recovery from contract research. Up to the Control Data 1604 period, bills were only produced for contract research use. Computer time was recorded on log sheets, subsequently on time cards, and transcribed daily to billing sheets that were manually totaled at the end of the month and extended to compute a billing amount. Later these data were recorded on punched cards and the amounts computed and accumulated for the month and year. In either case, printed sheets were sent to the enterprise bookkeeper at the university for collection and credit of income to the Computing Center. It was not possible to continue this practice with the 360/65.

First, because the 360/65 computer was to be shared with different users with different funding and different affiliations with the university, it was necessary to create some way to track this information. A much more complicated billing system had to be designed and built to accommodate this need for tracking as well as the more complicated billing for different computer resources. Further, because of the complex rate structure and because the very same job might cost a

different amount depending on the priority requested when the job was submitted, detailed information on job resources and cost was requested each time the job was executed. To eliminate manual operations and transfers, the billed amounts were submitted to the rebilling operation on cards or magnetic tape.

Trying to do all this with limited and strained resources in a disruptive period caused protracted delays and aggravation. Delays meant that funds were not available when the bills were submitted and contracts had already terminated and closed accounts. In the case of CAG, delays meant that funds expected to be transferred at a certain time were not. Aggravation resulted because users were having amounts subtracted from their accounts without being able to verify the amounts in advance; in addition, technical difficulties and accuracy problems early on did not give them much confidence in the system. The whole accounting/billing system was made more complicated by the need/demand to have each job fully priced out when the job was completed or terminated and to print this detailed cost information on the trailing page of the job. This meant the operating system had to be modified to collect all this information from different sources and do all the necessary cost computations on the fly. While worthwhile in many aspects, this one feature alone set a tone for many years of having to move the code through all the numerous future upgrades of hardware and systems software.

1968 to 1969

OCS Transitions 1968 to 1969—Erik D. McWilliams, Director

In 1968 Conway had completed his two-year term and expected to return to his faculty position. Erik D. McWilliams was appointed director of OCS to replace Conway. McWilliams had earlier come to Cornell from the University of Chicago to replace Worley as head of systems programming when Worley went on to graduate school at Cornell. With the departure of Worley, and of Jay Johnson to the University of Maine at Orono (director of computing and data processing services) and Bordonaro to Ithaca College (director of systems and data processing), McWilliams set about rebuilding his management staff. He appointed Richard C. Cogger as assistant director for systems programming and promoted David W. Pulleyn to assistant director for operations. Cogger had earlier come from the University of Chicago as technical coordinator in Administrative Services. Peter Shames was appointed manager of academic services, and Arlene Larsen, who had done a number of turns in interim positions, as coordinator of off-campus services. Dave Jennings, who had been hired as assistant director for

administrative systems earlier, completed the team. In less than a year, Jennings left in 1970 and was replaced by Jerry Tucker, who continued to head the group for several years.

McWilliams's main priorities continued to be the performance of the 360/65 system and the integration of the staffs from the Computing Center and Machine Records. There still existed communications and service gaps between the administrative users in Day Hall and the Machine Records and production control staff at Langmuir—gaps that had been created earlier when the administrative support staff was moved from Day Hall. Further, there had been very little progress on upgrading administrative business systems. The lack of sufficient funds to deal with any of the above problems made the situation even worse. The NSF grant of \$650,000, which had been extended for another year with an additional grant of \$215,000 for a total of \$865,000 over the three years, was ending in 1969. The loss of these funds made the financial situation worse. Service problems, both in system availability and in executing jobs, caused further financial problems due to income lost from not running jobs for paying customers. Early optimism that administrative systems would be moved from the 1401 and replaced with new systems had resulted in a commitment to remove the 1401 in mid-1968 in order to reduce costs. This was proving to be impossible.

CLASP Deployed; CLAPTRAP Report

The project to deploy CLASP on the 360/65 was already under way when McWilliams was appointed director. One of his early actions pending the move to CLASP was to install additional memory and disk hardware to try to squeeze yet more computing cycles out of the 360/65 processor. Finally, CLASP^t was installed in September 1969, but the result was not quite what had been expected, and instead there were significant problems. The system was not quite ready, and when problems did occur, significant machine time was taken out of the schedule in order for the systems staff to troubleshoot the problem and find a fix. Late 1969 brought the outages of the 360/65 system to an acceptable level, although throughput was still not acceptable. Cogger recalls²⁷ that the job queue had shrunk to 300 jobs instead of 500 jobs and that at its best, there were perhaps 12 outages a day with the CLASP system. This, of course, exacerbated the turnaround for batch jobs as system restarts, and investigations of problems and fixes took time out of the schedule.

The communications problems between the campus facilities and Langmuir about the status of jobs, and other issues such as lost jobs or turnaround time, led

to the creation of a long-lasting programmed report dubbed CLAPTRAP: CLASP Transaction Analysis Program. According to Cogger, the objective was to eliminate the arguments and discussions based on anecdotal evidence or statements by programmers on campus. The program for the report was written by Walt Haas, who also gave it its unique name, and it provided a definitive and timely summary and analysis of daily system performance based on accounting data generated by the batch system. Copies were printed on the higher-speed printers at Langmuir, at times three inches thick, and delivered to the campus facilities with the first courier in the morning. This way the campus staff could review the report with campus programmers who had any questions about their batch job or other performance issues over the previous 24-hour period. Debate about its costs, besides that of the cost of paper, and its value continued well into the 1970 period.

Although the amount of systems programming work to build a stable, efficient, and well-performing operating system infrastructure tended to dominate the years from 1966 to 1969, it is important to point out this work also was done to provide time sharing and fast batch services. For time sharing there was first CTS, as mentioned earlier, followed by CRBE (conversational remote batch entry), which was installed as part of the CLASP system. CRBE was an IBM product that had operating advantages, not the least of which was support from IBM in the event of problems. A number of problems, stability and availability of computing cycles being the most important, prevented the system from being fully exploited. For fast batch, a one-step monitor, along the lines of that existing in COOL, was provided. However, with all the other problems experienced with CLASP, the monitor did not totally provide the fast turnaround expected. Despite the increased throughput for CLASP and the shorter job queues, it was not uncommon at times to find the job status sheets at the various campus terminal locations stretching from six feet off the floor to six feet along the floor.

Tension continued between the academic users, the administrative users, and the OCS staff, who all were in need of more resources. Not unexpectedly, each community felt it suffered from unfulfilled promises and searched for ways to improve its situation.

Instructional Computing

An indication of the growth in student interest in computing and computer science can be seen in Table 3,

²⁷ Dick Cogger, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

which records statistics from the Department of Computer Science:

In their first annual report in 1965–66, the Department of Computer Science noted that there were 241 students enrolled in both IE9381/CS301 and IE9481/CS401 for a total of 964 credit hours of instruction. (IE9381 was the successor course to IE3281.) Conway and Chris Pottle, who had a joint appointment with Computer Science and Electrical Engineering, taught the courses.

Table 3. Department of Computer Science Courses, Enrollments, Credit Hours; 1965 to 1969

	1965–66	1968–69
Number of Courses Offered	9	33
Number of Students	370	1,267
Number of Credit Hours	732	3,788

The Development of PL/C

In 1968 Dick Conway and Prof. R. J. Walker received a grant of \$109,270 from the National Science Foundation for “A Computer System for Introductory Instruction.” This work led to the development of CUPL for the 360/65 system, replacing versions on the previous computers. In 1968–69 Conway and Howard Morgan from Computer Science, with the assistance of Tom Wilcox, then a graduate student, developed PL/C at the request of IBM, which also provided financial support. PL/C was an ultra-high-performance compiler for a subset of the PL/I language and was intended for high-volume instruction. PL/C went into production use in September 1970. By 1970–71 more than 100 copies of PL/C had been distributed and 60 copies had been purchased for use at other institutions.

Research Computing

Almost no statistics about the number of jobs and billed use of time and very little information about instructional and research use of the OCS facilities between 1965 and 1969 were available for inclusion in this history. However, the general trend was to increasing use of the 360/65 by instructional computing and decreasing use by research. The operating problems with the 360/65 during this period may have contributed to this condition. However, as researchers were continually searching for less costly alternatives, those with grant funds were using them to acquire minicomputers, for example, from Digital Equipment Corporation, and satisfied their computing needs this way rather than pay for OCS computing services.

Business Systems Data Processing

In 1969 McWilliams made the decision that all new programming work was to be done in PL/I, with COBOL a second choice for strong and valid reasons, and so new systems were to be built using PL/I and conventional flat file (sequential) records. Investigations also began into the use of data base management systems (DBMS), for serious consideration was being given to starting a project to develop a new payroll/personnel system using this new technology. The Information Management System (IMS) from IBM was selected, and Evans, the lead analyst in the group, did some initial work on the payroll/personnel system using IMS.

New Public Affairs and Library Systems

Toward the end of the 1960s, staff resources could finally be assigned to the priority projects. A new system for Public Affairs was developed; Jim Brinkerhoff was the key contact from Public Affairs and Bruce Lloyd was the programmer analyst. For the library system, Ryburn Ross was the key library staff member and Edmund V. Hollenbeck was the programmer analyst. This new library acquisitions system was developed to automate backroom operations, such as the ordering and receipt of books and periodicals, and vendor tracking but not circulation control. Budget problems prevented the system from being completed until the early 1970 decade. Both systems were written using PL/I.

Administrative Systems Planning and Control Board

Also at the end of the 1960s, an Administrative Systems Planning and Control Board was organized to plan for new systems and allocate resources to projects. The chairman in 1969 was Wallace B. Rogers. Very little is known about this group, its membership, and its actions. No information was found in the archives, but numerous references to the committee are made in the correspondence between OCS and the Statutory Business Office (later the Statutory Finance and Business Office, or SF&BO) regarding the computing system for the proposed new university accounting system. Memos from McWilliams, Comber, Long, and Peterson refer to this committee, but it is not clear what role it played.

New Financial Accounting and Payroll Systems

As the end of the 1960s approached, serious discussions took place with SF&BO, the university controller’s office, the endowed accounting office, and the endowed payroll office as to how the accounting and payroll systems should be upgraded. Controller

Peterson felt that the payroll office in particular, but other offices as well, had not been well served since the merger of Machine Records with OCS. He proposed developing new payroll and accounting systems with the staff in SF&BO and upgrading their computer to accommodate this load. A compromise decision was finally reached for SF&BO to develop and run a single new accounting system for both the statutory and the endowed colleges, and for OCS to develop and run a new comprehensive payroll/personnel system using its staff and computing resources. An accompanying financial arrangement was made for appropriate cost sharing. For the payroll/personnel systems, cost sharing was 55 percent endowed and 45 percent statutory, based on the number of paychecks cut for each sector. OCS only received the 45 percent statutory payment as new income because the endowed share was to come from the general university allocation for computing services.

To support the new accounting system, SF&BO installed an IBM 360/25^u computer in 1969 to replace their 1440 system. Bob Mack²⁸ recalls that the 360/25 system had three disk drives (IBM 2311 drives with 7.2 megabytes per disk), two tape drives, a card reader and cardpunch, and a printer. The system was supported by a number of card-processing machines such as a sorter, collator, etc. According to Mack, cost sharing for the accounting system started at 50 percent endowed and 50 percent statutory and was adjusted each year. As experience was accumulated, the percentage ended up being 60 percent endowed and 40 percent statutory, largely as a result of there being more endowed accounts and transactions.^v

Dealing with Campus Disruptions

Unfortunately, the campus student disturbances during the latter part of April 1969, and the weeks and months preceding and following, caused significant concern at OCS.^w Other computing centers at Berkeley, the University of Wisconsin, NYU, and Columbia had been attacked, firebombed, and ransacked by the disaffected students protesting against the Vietnam War and other real or imagined injustices. Computing centers were particularly vulnerable, because the destruction of the computer supporting administrative applications would create considerable havoc in a university's operations. The large computer room windows, originally designed for easy viewing of the facilities for formal tours or curious visitors, now turned out to be a major liability as they provided very little security. A saving grace for OCS was the Langmuir location, where it was much harder for students to mobilize a sneak attack. While OCS was always conscious of site security, nonetheless it took extra precautions to protect the computers and

equipment at Langmuir and all the administrative systems and procedures. Blackmun²⁹ recalls that during the height of the crisis, OCS management staff took four-hour turns during the second and third shifts and weekends to act as a security patrol by walking around the roof at Langmuir.

Protecting the university data and business systems resources with backups was a routine process, except that all such backups were stored at Langmuir and not at some offsite location. Until secure access to an offsite vault could be arranged, most of the administrative system master card decks from the 1401 computer were temporarily stored in the car trunks of trusted employees. Blackmun recalls riding around with the payroll system cards in his trunk, and Bruce Lloyd³⁰ recalls carrying six to eight trays of program decks in his car trunk for 10 days. To obscure this storage of key system card decks in the trunks of employee's cars, the procedure was referred to as "the vault in the sky." When arrangements were made to secure vault space in the First National Bank (which became Fleet Bank in 2000), it became routine weekly practice to store a large volume of key administrative and computer backup information at this site in downtown Ithaca. This weekly practice was continued for over 20 years before different arrangements were made.

Closing Out the 1960s

If it isn't clear by now, 1967 and 1968 were very hectic and tense years for all the staff in OCS as well as for the users. The continued changing of system software to increase throughput and improve services, numerous daily outages, and the time it took from the workday to deal with problems all drained the OCS staff and users. McWilliams started publication of the OCS *Grapevine*, an internal OCS newsletter, in 1969 to communicate important issues directly to the staff, to boost staff morale, and to deal with all sorts of rumors circulating on the campus.

External factors did not help the situation. In 1969, for example, IBM unbundled its services and changed its pricing and educational allowance policy. McWilliams estimated the cost to OCS to be \$50,000 per year to functionally replace IBM systems engineering support staff previously available at no cost and

²⁸ Bob Mack, "Oral History, 1965–Late 1970s," Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

²⁹ Bob Blackmun, "Bob Blackmun's Recollections, 1962–1982," Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

³⁰ Bruce Lloyd, "Oral History, 1963–1989," Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

whose expertise was critical in dealing with operating system problems. He estimated the net cost to be \$35,000 per year after accounting for reductions in rental costs that were part of the IBM policy change. The change in the systems engineering policy took several man-years from the experienced staff dealing with CLASP issues and forced OCS to consider whether (and how) to pay IBM or develop its own talent.

There also was much uncertainty about the availability and costs of programs and program products previously available at no cost. On top of all this were financial problems from the ending of the NSF grant and the drop in earnings from research users. Reacting to this last condition, OCS reversed the policy announced earlier to rent the major components of the 360/65 and purchased them in March 1969 for a saving of \$52,000 per year and a total net saving of \$325,000 by the end of the five-year amortization period. Counteracting these savings were the anticipated budget increases due to the reduction in the IBM education allowance policy ranging from an effective discount ranging from 20 percent to 30 percent, depending on the product and its time of acquisition, to 10 percent on all new products.

The situation at OCS and for all users looked like it would improve as 1970 approached and the stormy 1960s came to a close. The university had earlier retained McKinsey and Company to study and recommend an organization for the overall administration of the university. In 1969, following the student disruptions on campus, President James Perkins resigned and was replaced by Dale Corson, who was provost at the time. Following the recommendation from the McKinsey study, President Corson created the new position of vice president for administration. The significance for OCS was that all directors of service organizations, including OCS, were to report to this new position.

Other Computing Organizations on Campus

Dairy Records Processing Laboratory

We continue with the developments in the Dairy Records Processing Lab (DRPL) taken from Wadell's account.³¹

During 1960 the IBM 650 became overloaded, and an IBM 1620 was added to help by removing some of the research from the IBM 650. Twenty-three employees were now handling 193,000 cows on central processing.

In 1961 a staff of 36 employees managed by Lyle Wadell moved into new quarters as the first employees to use the new Morrison Hall building. The number of cows had now grown to 272,000.

In 1962 an IBM 1410 computing system was installed. At the same time we went from a card system to a card-magnetic tape system. Forty-three employees were now handling 323,000 centrally processed cows, and by 1963 this had grown to 386,000 cows and 45 employees. The growth in volume and growth in applications with a magnetic tape system rapidly overloaded the IBM 1410, and in 1963 an IBM 1401 computer was installed to mainly handle the printing of reports and small card jobs.

By 1964 the centrally processed cows had grown to 428,000. The number of workers had grown to 46. In anticipation of further growth, an order was placed for an IBM 360-40. Central processing cows had reached 450,000 cows by 1965 and the number of employees had grown to 48. The new IBM 360-40 was installed in late 1966 and along with it were our first disk drives (IBM 2311s). The addition of disk drives opened up new applications doors.

³¹ Lyle Wadell, "History of the Northeast Dairy Records Processing Laboratory, 1948–1985, with Additional Comments by J. D. Burke and H. Wilmot Carter," Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

Endnotes

^a It is interesting to note that the term “computer utility” was being recycled in 2003. IBM uses the term e-commerce for their offering. However, the meaning is somewhat different in that the “utility” in this case provides complete applications rather than computer cycles.

^b Sometime in the middle of the 1980s, John Middlebrook called me from his office in New Jersey, and he was still using the Burroughs 220 system (which he had bought from Cornell in 1964) for his consulting business. At that time, more powerful and much less expensive minicomputers or desktop computers were available, but he had such an investment in the programs and data that he continued to avoid converting. I could just imagine his costs for electricity to power the computers and peripherals and air conditioning. I believe he purchased another Burroughs 220 computer system just for spare parts! He was fortunate to have hired a competent Burroughs serviceperson, who kept the system going all those years.

^c When I came to the Computing Center as a graduate assistant in September 1959, Tom Kemp and I were assigned the specific task of writing the monthly sales report for GLF (Grange League Federation, a farmers’ cooperative that later became Agway). Income from running this report was an important part of the budget of the center in order to cover the increased costs of the new Burroughs 220 computer. GLF had a RAMAC 305 computer (IBM recycled the name Ramac in 1996 for its high-volume disk storage), which could not accumulate data over their entire system of stores and compute incentive rewards for exceeding sales quotas. Every month we would receive five to seven trays of IBM cards that were the monthly sales statistics for each store. These statistics were organized by product lines and included last year’s and this year’s to-date sales dollars. Our job was to read in the information for each store, accumulate the data by region, district, and corporation, and for each such entity compute the awards and then print a report for each entity. It sounded simple enough until we had to deal with all the events that caused problems—a misread card, a misprinted line (on seven-part preprinted report forms), a computer fault, etc. We ended up building a very complex restart mechanism, storing all kinds of interim information on magnetic tape so we could avoid starting from the very beginning every time the inevitable error would occur during each run. Besides learning quite a bit about “data processing,” I also learned quite a few names of towns in New York State and northern Pennsylvania, where many of the GLF stores were located. Because the 407 printer operated at 150 lines a minute, it was possible to keep up a reasonably constant quality check on the printing by hovering over the printer when not feeding cards and watching the console. We continued to run that report well into the 1960s before GLF got its own equipment or abandoned the program!

^d The developers of CORTRAN listed on the surviving operations manual were Richard A. Stone, David Bessel, Tatjana Grenewitz, Frank Long, John Edgerton, John Emler, Eric Mintz, and Alan Goodman.

^e EPM—Engineering Problems and Methods—was a freshman course taught for the first time in the fall of 1961, Bob Blackmun’s freshman year. Blackmun offers the following explanation: “I don’t remember the names, but it was taught by several of the very senior Engineering faculty and was intended to get us ‘doing’ engineering work (to the extent that we could do so without having had all of the ‘basics’), and one of the ways that was done was to learn CORC and write some computer programs. One of the reasons that I remember this is that my first job working for the Computing Center (in my sophomore year, I think) was to be a student ‘consultant’ for the freshmen who were taking the course the following fall. We worked in the room on the fourth floor of Rand Hall where the printouts were put out on big tables, and there were some keypunches, drawers for submitting programs to be run, etc.”

^f No copy of Arden’s report was found in the files accessible to me. As a result, the statements made about the conclusions of the report are my recollections. However, they are supported by all the future actions that took place.

^g The 1604, 3600, and 6600 computers from Control Data Corporation (CDC) were designed by Seymour Cray, who later founded his own company. Cray Computer became the premier builder of more modern-era supercomputers and was a leader in this field well after 1990, some 30 years later. In a private communication, Dick Lesser mentions that he met Cray on a visit to Minneapolis with Gerald H. Larsen, the CDC sales rep for Cornell, to check out the CDC equipment. At that time, CDC was in makeshift quarters in an old commercial building and Cray in a no-frills cubicle.

^h A full set of slides of the 360/67 IBM “road show” was retained by Conway and is now in the archives of Cornell. The slides portray a system that was everybody’s answer to all the problems of providing contemporary computing services to different users at a very low cost! According to the *Encyclopedia of Computer Science*, 4th edition, only 10 of the 360/67 computers were built!

ⁱ The significant new component of the 360/67 was the Blauw box, named after its developer. This box contained relocation registers for programs that allowed the program and data to be located anywhere in memory each time a program was executed by defining all addresses as relative to a base address. The base address could be changed at each execution, depending on the amount and location of available memory at the time the program was executed.

^j Credit for keeping the newly formed combined academic and administrative computing organizations together and using the same single computer belongs to Conway. He firmly believed that such an organization would produce operating economies and synergistic effects through staff interactions. Although seemingly obvious, it took 20 or more years for other large universities to come to this same conclusion!

^k Frank A. Long, vice president for research and advanced studies, issued an official memorandum, “Death of the Cornell Computing Center on February 5, 1968.” In that memo he refers to the CDC 1604 as “retired from use and will apparently soon be shipped to Mexico.” In 1969 the 1604 was sold to a Mr. Beltran, who was suspected, but never proven, to represent Cuban interests while arguing that the machine would be shipped to Mexico. It was always assumed that the 1604 was eventually installed in Cuba despite the embargo against such transactions. Ralph Barnard from the Cornell legal office, who participated in drafting the agreement, believes the 1604 never got out of Texas. Nevertheless, he takes pride that the contract was so designed that Cornell received its funds at the time the computer was loaded on the truck in Ithaca! However, by the time the 1604 was sold, the price was \$100,000 instead of the \$150,000 first expected. According to a recent note from Dave Bessel (in 2001), he suspects that the 1604 might have ended up in East Germany, based on a brief news item stating that a restored system that was in full operation had been placed in a German museum. If the serial number of the system was 13, then that would be the Cornell system!

^l It is interesting to note that with all the discussions taking place in 2001 about building new facilities for the College of Architecture, Art, and Planning and tearing down Rand Hall, this concept of a new campus gateway was yet again recycled!

^m Dale R. Corson, then dean of the College of Engineering, and Frank Long made an appeal to graduates of Cornell in high positions in IBM seeking support to redress this situation. As Corson recalls, there was sympathy for Cornell’s position but no funding was provided. Also according to Corson, only Princeton University

obtained support from IBM to fix the problems for the 360/67 that they had ordered.

ⁿ There is some interesting correspondence in the IBM files from the period; depending on the final date of the order, the discount on the 360/65 could change from 40 percent to 30 percent, a substantial difference. Further, it seemed that IBM was not uniformly treating other universities, notably Princeton, in the same manner it was treating Cornell with regard to “free” rentals of the 360/67 for testing purposes. Long and Corson made appeals to IBM to provide some form of financial relief to Cornell due to the failure of the 360/67 product, but nothing materialized.

^o Tom Jones later became one of the leaders of the takeover of Willard Straight Hall by black students in 1969 and appeared in the famous photograph of their exit. He eventually went on to become president of TIAA-CREF, giving Manning bragging rights that he knew how to choose the good ones!

^p See the “Report on an Inservice Course in Computer Mathematics for High School Mathematics Teachers, Spring 1964, Given at Cornell University, Sponsored by the State Education Department.” The course was taught by Dick Lesser and John Rudan.

^q Eisenhower College folded in 1980. Before that, the Rochester Institute of Technology had purchased the facilities and tried to run the college as a satellite campus. The startup of the college, located near Seneca Falls, N.Y., was funded in part by proceeds from the sale of Eisenhower dollar coins. In 1988 the campus was acquired by the New York Chiropractic College, which moved their activities from Long Island in 1990 after renovating and refurbishing the site.

^r Because there did not appear to be a position for me in the new organization, I looked for other opportunities. With the help of Professors W. T. Federer and Shayle Searle and the generous support of NIH, I became a graduate student in statistics in the Biometrics Unit.

^s At some point in the 1970s when I was director of OCS, I recall going to Chicago for a wrap-up conference with the three other centers. Gerard (Gerry) P. Weeg from Iowa was the convenor of that meeting, and its purpose was to review the accomplishments and problems and prepare a final report for NSF. The report, “A Study of Regional Computer Networks,” was issued in February 1973. Other authors besides Weeg were Fred W. Weingarten, Claremont College; Norman R. Neilsen, Stanford University; and James R. Whiteley, University of Iowa.

^t Some disgruntled users were said to have sarcastically called CLASP “Cornell’s Last Attempt at Systems Programming!”

^u The staff in the statutory business office had placed an advance order for a 360/25, but before taking delivery they also considered the competitive NCR 200 system. They conducted performance tests and site visits to NCR installed systems and concluded the IBM system was a better deal. It is also worth noting that part of the motivation for continuing to have a separate machine for the university accounting system was the cost increases and unbundling of software services that IBM was putting into effect at the end of December 1969!

^v According to the oral history interview with Bob Mack in 1999, 60 percent endowed and 40 percent statutory was the cost-sharing formula being used for the accounting system in the year 2000!

^w During this period, the mayor of Ithaca was my neighbor Jack Kiely, and so I had an insider’s view of the seriousness of the concerns of the different police agencies about a full-scale riot and the alerts and preparations that were taking place. The height of the crisis was on April 19 and 20, when black students occupied Willard Straight Hall. I saw some of the comings and goings around noon on Sunday. I had gone to 11:00 a.m. Mass at Anabel Taylor and while walking out was alerted to the activity taking place at Willard Straight. I walked over and stood on the grassy knoll, where the Campus Store is now but wasn’t then, and watched events unfold. On Monday, April 21, 21,000 people jammed Barton Hall to discuss what was going on and how to deal with the actions of the black students. At about this time the city of Ithaca had converted several of the buildings at Stewart Park into temporary jails and holding facilities. Further, there were over 200 New York State police milling around in the Woolworth store (now the Tompkins County Public Library at the corner of Cayuga and Green Streets) parking lot, waiting for the call to deal with any disturbance. Thankfully, peace prevailed and no call was made. It goes without saying that these were very tense times at Cornell. The book *Cornell ’69 Liberalism and the Crisis of the American University* by Donald Alexander Downs, published by Cornell University Press, gives an excellent summary of the period.

1970 to 1979

The Time-Sharing Decade

Contents

1970 to 1979—Industry Overview	56
Hardware	56
Software	57
Services	57
Vendors	58
Technology Terms, Acronyms, and Buzz Words	59
1970 to 1979 at Cornell	60
The Early 1970s	60
<i>Conditions on Campus</i>	60
<i>Studies of Computing on Campus</i>	60
<i>Proposed New Campus Computing Leadership and Organization</i>	60
<i>Formation of the University Computing Board (UCB)</i>	61
<i>OCS Priorities</i>	61
<i>Campus Financial Issues and Plans</i>	62
<i>Campus Computing and OCS Leadership—John W. Rudan, Director</i>	62
<i>Cost Cutting at OCS</i>	62
<i>Off-Campus Computing</i>	63
<i>Communications with Users—Newsletters and the OCS Bulletin</i>	63
<i>Squeezing More Out of the 360/65</i>	64
Instructional Computing	64
<i>Time-Sharing Services for Instruction</i>	65
<i>Introduction of PL/C</i>	65
<i>Instant Turnaround (IT) for Student Batch Computing</i>	65
<i>Free IT—A Critical Milestone</i>	66
<i>Biological Sciences Initiative</i>	67
<i>PLATO for Instructional Use</i>	67
Research Computing	67
<i>Advances in Statistical Computing—COSSIS Proposal</i>	67
<i>New Policy for Distributing Allocated Research Funds</i>	68
<i>Research Advisory Subcommittee Projects—Media Conversion, ARPANet</i>	68
<i>Surveys of Computers on Campus</i>	68
Business Systems	69
<i>Formation of MSA (Management Systems and Analysis)</i>	69
<i>New Payroll/Personnel System</i>	70
<i>New Student Information System</i>	70
<i>New Financial Accounting System</i>	71
<i>Mark IV Query Language Installed</i>	72
<i>1401 Computer Finally Removed</i>	72
<i>Improving Budget Systems</i>	72
<i>MSA Staff and Business Systems Support Services Move Back to Day Hall</i>	72

OCS Initiatives and Activity in the Mid-1970s	72
<i>Recommendations to Increase Computing Capacity at OCS</i>	72
<i>Faculty Reports on Campus Computing—Chester and deBoer</i>	73
<i>An Early Computer Security Incident</i>	75
<i>Regional and National Computing Initiatives</i>	75
<i>NBER and TYMNET National Networking</i>	76
<i>PBM/STIRS, EFPM, and SENTRY (CBORD) Network-Based Initiatives</i>	76
<i>New Leadership for Academic Computing</i>	77
OCS Seminars	77
Computer Science	78
Business Systems	78
<i>The Financial Accounting System Moves to the 370/168</i>	78
<i>McNeil Task Force and the Demise of MSA</i>	78
Campus Transitions—1975 to 1977.	79
<i>W. G. Herbster Appointed Senior Vice President</i>	79
<i>A. H. Peterson's Laws on Computing and Information Systems</i>	79
<i>New President of Cornell—Frank H. T. Rhodes and Other Executive Changes</i>	79
<i>New Management Team at OCS—Robert R. Blackmun, Director</i>	80
<i>A Malicious Computer Incident</i>	80
<i>Institutional Planning Services</i>	80
Instructional Computing	81
<i>The Push for More Time-Sharing Services</i>	81
<i>Subset Conversational Monitor System (SCMS) Developed</i>	81
<i>Experience with First Use of SCMS</i>	81
<i>PLATO Use Continued</i>	82
Research Computing	82
<i>Acquisition of the Array Processor</i>	82
<i>Mass Storage System</i>	82
New Initiatives—1978 to 1979	82
<i>Microcomputers and OCS Support Services</i>	83
<i>Text Processing with SCRIPT and the Start of Electronic Mail</i>	83
<i>Network Services; Terminal Leasing Enterprise</i>	83
<i>OCS and the Cornell University Medical College</i>	83
<i>New Plans for Campus Computing Services; DEC20 Favored along with Use of Terak Microcomputers</i>	85
Instructional Computing	86
<i>The Cornell Program Synthesizer Using Teraks</i>	87
<i>DEC20 Computer Ordered for Installation at OCS</i>	87
<i>IT (Instant Turnaround) Batch Use Declines</i>	87
Business Information Data Processing	87
<i>Mark Sense Cards for Student Registration; Centralized Registration</i>	87
<i>CADE Data Entry System</i>	87
Network Plans for the Campus.	87
Academic Computing.	89
<i>CAG Merges with OCS as Part of User Services</i>	89
<i>Formation of DACS (Decentralized Academic Computing Support)</i>	90

New Leadership for OCS	90
<i>New Reporting Relationship: OCS to Report to Provost</i>	90
<i>Cooke and Rudan Direct OCS Operations</i>	90
<i>Search for New OCS Director</i>	90
Summary Comments	91
<i>Growth in Time-Sharing Services</i>	91
<i>Systems Software Development Accomplishments</i>	91
<i>OCS Space on Campus</i>	92
Other Campus Computing Activities	92
<i>Dairy Records Processing Laboratory</i>	92
<i>IBM System 7 for Monitoring/Controlling Energy Use</i>	93
<i>WATSBOX for Long-Distance Telephone Services</i>	93
<i>Laboratory of Nuclear Studies (LNS)</i>	93
<i>Department of Physical Biology in the Veterinary College</i>	94
<i>Program of Computer Graphics</i>	94
<i>College of Veterinary Medicine</i>	94
<i>Department of Computer Science</i>	94

Figures

Figure 1. IBM 370/168 in Langmuir Laboratory, 1974	73
Figure 2. Existing Cornell campus network in 1978	88
Figure 3. Proposed Cornell network, July 1978	87

Tables

Table 1. Number of Student Batch Jobs, 1965 to 1975	67
Table 2. IBM 360/65 Research Computing Use, 1968 to 1971	67
Table 3. Department of Computer Science Course Statistics, 1970–71 and 1979–80	86
Table 4. Growth of Time Sharing, 1967 to 1979	92

1970 to 1979—Industry Overview

The following is a brief description of the major hardware, software, and service innovations in the 1970–79 decade as well as how different vendors came into and left the market in these areas. Time-sharing use was now a mainstream service available on systems from most vendors, especially on systems in the rapidly growing minicomputer market. But by the end of the decade, microcomputers that gave the user full control over all resources started to have an impact on the computer service delivery and charging model. The spectrum of machines and their uses became very broad in this decade, capitalizing on the improvements in chip technology and software. At the low end were hand-held calculators with limited programming capability, while at the high end the first supercomputers were built. In between, special-purpose systems came in different varieties, from word processors for document processing to array processors for vector arithmetic and matrix calculations. The range of applications of computer technology was expanding rapidly, including the early development of new network technologies.

Hardware

The computing engine continued to shrink in size and increase in speed as the industry started aggressively to develop silicon chip technology. The first use of chip technology came in the form of increased size of memory. Intel Corporation is credited with the invention of the DRAM chip (dynamic random-access memory) in 1970.¹ Later, Intel developed the first microprocessor, the Intel 4004, which was a computer on a chip. Very large-scale integration (VLSI), in which more functionality was placed on a single silicon chip, became the driving force to increase speed and reduce the size and power consumption of computers. It was in this period that Gordon E. Moore, one of the founders of Intel Corporation, formulated Moore's Law. First stated a decade earlier in terms of transistors per chip, Moore's Law now said that the computing power of a silicon chip would double every two years; this version is better remembered. As with Grosch's Law stated two decades earlier about the relationship of price and performance of IBM computers, it is not clear if Moore's Law was the result of independent actions by Intel or if Intel used Moore's Law to drive its product development plans. Accompanying this dramatic increase in chip power were dramatic reductions in the price of chips.

During the 1970s Intel developed a series of chips, each improving speed and increasing the number of transistors per chip. The 8080 chips, first developed in mid-decade, incorporated 6,000 transistors running at

2 Mhertz, while the 8088 chip developed in 1979 contained 29,000 transistors and was rated at 4.77 Mhertz. Addressing space was increased from 64 Kbytes to over 1 Mbyte during this same period. Chips from other vendors such as AMD (Advanced Micro Devices), Texas Instruments, and Motorola were comparable and were used by different vendors in their systems. Chip technology was increasingly used in larger central systems, as in IBM mainframes, and this led to smaller computer boxes, reduced power consumption, and air-conditioning requirements. However, there were still "hot spots" in some components owing to the way they were packaged. These areas had to be cooled directly with chilled water. The raised floor cavity had to be made taller to accommodate the piping and all the other cabling that interconnected units for performance and redundancy.

Similar advances in technology occurred for peripheral equipment, but the offerings were different for microcomputer systems and the large central systems. For the microcomputer market, IBM introduced in this decade the 8-inch floppy disk drive. Initially such a disk could store 400 Kbytes, but before the decade ended that expanded to 800 Kbytes and compatible versions were being offered by several vendors. Later in the decade, Shugart announced the 5 1/4-inch "minifloppy" that was capable of storing 110 Kbytes. Different vendors used either type of these floppies in their system offerings. Individuals building their own systems could make this choice as well. It is estimated that 2.5 million floppies were sold over the decade.²

For the large mainframe computers, IBM introduced a new line of 370 systems early in the decade. The 370 line maintained the same architecture as the 360 line with some additions. However, it used the VLSI technology so that there were notable computing capacity increases and cost reductions compared with the 360. IBM also made significant advancements in magnetic disk and tape technologies over the decade. Using a combination of increased recording densities and double-sided recording on disks, the amount of storage increased from 26 Mbytes per unit to 1,260 Mbytes per unit for the IBM 3380 Direct Access Storage Device (DASD) introduced in 1980.³ At the same time, data transfer rates more than doubled, while costs (dollar/Mbyte/month) dropped by a factor of 10. Contributing to some of these advancements was the change from

¹ Throughout this document, technical terms are based on definitions from the *Encyclopedia of Computer Sciences*, 4th ed. or from Webopedia at www.webopedia.com, or from Polsson, Ken., "Chronology of Personal Computers," www.islandnet.com/~kpolsson/comphist/, 2003.

² Ken Polsson, "Technology of Personal Computers," www.islandnet.com, copyright, 1995–2003.

³ "History of IBM," www-1.ibm.com/ibm/history/.

removable to fixed disk packs.

Complementary activity was taking place with magnetic tapes. Using IBM as the example again, the IBM 3420 magnetic tape systems introduced in the early 1970s increased the recording density to 6,250 characters/inch, up from 1,600 in the earlier models, a major breakthrough. Combining this increased density with increased movement of the tape itself to 200 inches/second from 112.5, increased the throughput rate by a factor of four, from 320,000 characters/second to 1,250,000 characters/second. The new systems maintained the use of the 10.5-inch reel holding 2,400 feet of tape. The improvements in storage capacity and data delivery for both disks and tapes kept up with the improvements in computing throughput being generated by the new computers.

Software

Developments in software continued during the 1970s and, as with hardware, on several different levels. The most significant development was that software no longer was provided as part of the purchase or rental of a computer. While this change improved the bottom-line profitability of IBM and other companies, it introduced another level of expense for computing centers. On the positive side, however, a whole new industry grew from providing special-purpose systems, or applications programs, for specific business functions, such as accounting, or higher education functions, such as student registration and financial aid. There were now competitive offerings for all types of software.

With respect to computer languages, PL/I was never totally accepted by the scientific community, which was wedded to Fortran, or by the business community, which was wedded to COBOL, and it faded into obscurity. Nicklaus Wirth from Stanford defined the Pascal language that became popular for introductory computer science courses. Ken Thompson of Bell Labs invented a language he called B, which Dennis Ritchie extended to form C. Much of this activity took place on DEC computers running the Unix operating system developed and implemented in the earlier decade. It was in this context that the first e-mail application was developed by Ray Tomlinson, who is credited with first using the @ sign in the structure of e-mail addresses, a development that has continued to the present time, much to the consternation of many who opposed his choice.⁵

⁴ R. Bradshaw and C. Schroeder, "Fifty Years of IBM Innovation with Information Storage on Magnetic Tape," *IBM Journal of Research and Development* 47(4), July 2003.

⁵ Bruno Giussani, "A Brief History of @," *The Industry Standard*, May 14, 2001.

Very early in the 1970s, a lot of effort went into creating BASIC for the different microcomputers being built. Paul Allen and Bill Gates, who went on to later fame by forming Micro-soft (original spelling), started their activity by writing the BASIC compiler for the Altair system, available at that same time. They built a series of BASIC compilers for different computer chips and computers built by different vendors, making BASIC the computer language of choice for these systems. The first operating system for microcomputers, CP/M (control program/monitor), was developed by Gary Kildall, who formed Digital Research Corporation to market the system. CP/M was extended to different microcomputers and was the ranking operating system during this period. The most notable application program developed during this period was VisiCalc, the spreadsheet program first proposed by Dan Bricklin in the middle of the decade. In 1979, Dan Bricklin and Bob Frankston, operating under the name Software Arts, demonstrated the use of VisiCalc running on Apple II computers, and soon afterward VisiCalc became commercially available.

Services

Time sharing continued to be the focus of the service improvements in computing access and use. Most computer use continued to be for program and data input or for access to specific information depositories (databases) because general-purpose applications had not yet been developed. Improvement in this service required companion improvements in input devices (terminals), output devices (printers, mainly), connectivity (networking), and central computer operating systems. The computer utility model of the large central computer accepting input from local or remote devices of different types was still the basic operating model.

While Teletype terminals (combined keyboard and printer) continued to be popular, new terminals based on IBM Selectric typewriters were introduced. DEC brought out a line of its own typewriter-like terminals called Decwriters. Increasingly, though, CRT (cathode ray tube) video terminals became available. These offered the visual display of the input and output on a screen, and by the end of the decade these video terminals had become the most popular input device for programming and online data entry. They acquired the moniker "dumb terminals" because they simply provided or accepted characters and had no local intelligence. IBM had its own brand of terminal, called the 3270, which was somewhat smart in that if it was in the IBM-defined systems network architecture (SNA) environment, it could do some functions locally.

Connectivity between computers and individual terminals and other low-speed input-output devices

was achieved by the use of acoustic couplers that converted terminal signals to telephone signals, which could then be transmitted to a controller at the computer. These couplers were very slow, operating at 110 bits per second, but their speed increased to 300 bits per second by the end of the decade. Late in the decade, the Hayes modem (modulator/demodulator) was developed, which not only increased the speed of transmission but provided improved reliability as well.

This decade saw the introduction of several networking technologies that came to play important and significant roles in the decades ahead. In 1973 Robert Metcalfe developed the Ethernet protocols at the Xerox Palo Alto Research Center to be able to connect Xerox printers to computers. During these early years, Ethernet networks operated at 10 Mbits per second and were limited in their total length to less than 500 meters. As a result, networks using Ethernet protocols were limited to use within buildings or ones in close proximity to each other. At the end of the decade, Ethernet had become a standard, referred to as 802.3, defined by the Institute of Electrical and Electronic Engineers (IEEE), which created a committee to standardize network technologies. In 1974, IBM developed and implemented its system network architecture (SNA) in response to customers' complaints that they could not properly automate transactions because of unreliable communications networks. When first conceived, the SNA network was built out of expensive dedicated minicomputers managed by a central mainframe. These minis ran a special program that managed communications with the terminals, while the mainframe managed the total network and communicated with these minis. Protocols were defined and developed for all this to work harmoniously.

TCP (transmission control protocol) was first proposed by Vint Cerf in the paper "A Protocol for Packet Network Interconnection," co-authored with Bob Kahn.⁶ In 1974 Cerf, with the participation of Yogel Dalal and Carl Sunshine, formally described TCP. The objective was to define a standard information packet transmission protocol so that it was possible to build a network of cooperating computers. The protocol was designed to be flexible enough to handle the physical differences in host computers, routers, and networks and still provide for such physically different entities to transmit data between themselves despite differences in packet sizes. Later in the decade TCP split into TCP/IP (Internet protocol) where IP was responsible for routing packets of information and TCP was responsible for creating packets, error control, retransmission, and reassembly of packets.

A whole new technology and industry was started to develop the software and hardware to produce fast and inexpensive gateways to route information over networks stretching around the world. By the end of the decade, TCP/IP was reasonably mature and had become the standard military protocol for networking.

Vendors

IBM continued to be the dominant computer manufacturer of large computers and their supporting hardware during the 1970s. The B.U.N.C.H (Burroughs, Univac, NCR, Control Data, and Honeywell) had shrunk so that only Control Data remained a solid competitor. However, IBM's dominance was challenged on several fronts. In 1970, Gene Amdahl, who was instrumental in the design of the IBM 360 systems, left IBM and started his own company, Amdahl, to make and market IBM-compatible systems. Amdahl thus began a whole new subindustry with computers designed to run IBM operating system and applications software but on non-IBM hardware. Amdahl was able to take advantage of the decision made in 1969 for IBM to unbundle its software and hardware offerings. Others, such as Ampex, were less adventuresome and simply built plug-compatible components such as memory and tape and disk drives, which could replace the IBM equivalent at a much lower cost. These suppliers became known as third-party vendors, being the third party to IBM and the customer working together to configure and operate a computing system and center.

Seymour Cray, who had left Control Data, founded Cray Research to build his own scientific computers, the first of the so-called supercomputers—extremely fast computers running Fortran and supported by large, super-fast peripherals. The CRAY-1 supercomputer was available in 1976. Control Data countered in part by introducing the STAR-100, the first of the computers that could perform vector computations. The University of Illinois entered the supercomputer market by designing and building the Illiac IV system.

Digital Equipment Corporation (DEC) continued to produce new and improved versions of its PDP equipment line. It brought the PDP-11 to market in 1970 with the Unibus, the first multivendor computer bus (a pathway between parts of the computer), allowing other equipment to be connected to the system. Later in the decade it introduced the DEC VAX-11 systems and architecture.

A whole new breed of vendors started to take advantage of improving chip technology to build different kinds of computer engines. Hewlett-Packard (HP) and Texas Instruments built the first hand-held calculators, which in the case of HP was rudimentarily programmable. Nolan Bushnell founded Atari Corporation, which used the chip technology to develop the

⁶ *IEEE Transactions on Communications*, vol. COM-22, no. 5, pp. 637-48, May 1974.

Pong video game. Xerox Palo Alto Research Center (PARC) built the Alto workstation. In 1975, MITS (Micro Instrumentation and Telemetry Systems) built the Altair 8800, which many consider to be the first microcomputer. The Altair 8800 used the Intel 8080 chip. MITS also produced the first personal computer kit that sold directly to individuals. By mid-1970 Steve Wozniak and Steve Jobs had founded Apple Corporation and were selling the Apple I personal computer, and by 1978, the Apple II, which had debuted in April 1977 at the Consumer Electronics Show in Las Vegas. Their competition, Radio Shack, was selling the TRS-80 microcomputer. These two companies started the personal computer era.

The one vendor that combined hardware and applications software was Wang Corporation, which brought to market its first word-processing system in 1972. In 1976 Wang came up with an advanced system that within a year included a CRT display, a large disk for storage, and a fast letter-quality printer. Given the appeal and utility of the system, Wang was able to charge a much higher price than the competition at the time.

During this era the hobbyist represented an important element of activity in building microcomputers and writing software. Clubs such as the Bay Area Amateur Computer Users Group—Homebrew Computer Club provided forums for demonstrations and exchange of ideas and information. In 1977 the First West Coast Computer Faire was held in San Francisco's Brooks Civic Auditorium, and almost 13,000 attended the weekend event. Apple Computer, for example, sold its Apple I computer board in kit form. New magazines such as *BYTE Magazine* and the *Computer Hobbyist* were founded to appeal to this market. The ComputerLand stores were organized in 1977 to provide a sales outlet for microcomputers and components to the individual and vendors as well. By the end of the 1970s there was intense competition to build microcomputers to achieve breakthroughs in capabilities and cost. This activity repeated the 1960s cycle of vendors trying to capture the imagination of users as well as market share. It was also during this decade that microcomputers started to be called personal computers, before IBM copyrighted the term in the 1980s.

Technology Terms, Acronyms, and Buzz Words

Freeware—one of the terms used to describe software that is available at no cost. Freeware may or may not have restrictions on its use, the most likely being about making changes to the software and distributing it to others. In general, the author of public domain software gives away the rights to the software, and modifications are generally encouraged and are expected to be shared with others. The term has been copyrighted, so the generic terms “public domain software” and “shareware” have partially replaced it. Shareware is often free when distributed for evaluation, but payment is expected if the software is put into use.

Hackers—used in the 1970s to describe programmers who were obsessed with computers and whose main focus was to master the machine. A hacker had no apparent organized method of writing programs but was looking for elegant and tight code. The term evolved to describe programmers who were dedicated to computerized vandalism and who had little respect for personal property, security, and privacy of networked systems.

Megahertz (MHz)—one million cycles per second, a measure of the clock speed of a microprocessor that in turn determines the number of instructions that can be executed per second. As such, the MHz rating in large measure determines the power of the microprocessor.

MIPS—millions of instructions per second, used as a measure to rate the computing speed and compare the computing power of different machines, as in the IBM line of systems, and different vendors with the same computing architecture.

Turnkey—a term taken from the construction industry, meaning a computer system that could be installed, powered up, and put to its designed use without any further work. Generally first used in reference to minicomputer systems.

1970 to 1979 at Cornell

Several bold initiatives to bring overall campuswide leadership to computing and to improve computing services on campus were started early in the decade. The one objective that was achieved was the formation of a new administrative business systems group in 1971. Financial conditions severely limited the changes and improvements after that, in part forcing the computing center to cut its operating budget and look off-campus for customers and earnings to sustain the operation. Several studies of computing services on campus were undertaken during this decade. One led to positive results, the acquisition of the 370/168 computer, while another led to negative results, a significant cutback in staff supporting business systems.

The Early 1970s

Conditions on Campus

The year 1970 did not start out well for Cornell as a whole. Besides having to take steps to calm the campus after the 1969 disruptions and deal with a new and increased role of the faculty in campus governance, Cornell administrators faced a projected deficit of \$1.6 million, following a deficit of \$93,000 the previous year. Terms such as deficit, deferred maintenance, and other such connotations of doom and gloom became a large part of the administrative vocabulary for most of the decade. In July Samuel A. Lawrence was appointed vice president for administration, and early in 1971 he announced a 10 percent cutback in budgets to be achieved with 4 percent, 4 percent, and 3 percent reductions in the next three fiscal years, respectively. At the same time, a 7.5 percent salary improvement program was announced for the 1971–72 fiscal year. Given the dominance of salaries in most budgets, there was continual pressure to balance budgets by reducing costs through staff reductions and any other possible means or by finding additional sources of funds.

Studies of Computing on Campus

Nineteen seventy also was a year of considerable activity in examining the current and possible future state of computing on campus as well as more immediate changes for OCS. Early that year, Provost Plane announced the formation of the Academic Computing Advisory Board (ACAB) to advise him on all matters concerning research and instructional computing. Geoffrey V. Chester was appointed chairman. In the summer of that year Plane also appointed a group “to conduct a brief survey of the future developments of computing at Cornell and to give some thought to the organization of computing facilities.” In six weeks

during August and September, the group produced a report that became known as the Chester-Thomas report.⁷ (D. A. Thomas was associate dean of the Graduate School of Business at the time.) That report made tentative recommendations that the committee hoped would be subject to further discussion and debate. Their principal recommendation became known as the semi-autonomous model for organizing computing.

Our tentative recommendation is that a decentralized organization, with decentralized hardware, will serve the university best in the long run. However, we recommend that the decentralized units be subject to a sufficient degree of control so that suitable computing can, under careful priorities, flow between them. Hence our name “the semi-autonomous” model. We realize that this recommendation runs counter to much expert advice, and we have therefore set out the arguments for it in some detail.

Given the short span of time and the time of year that the report was written, the committee noted that “the replies from the deans and directors did not reveal the same sense of high priority as is felt by the committee.” Some deans called for at least a doubling of the use of computing over the 1970–75 period, while others were much more modest in their responses, which were contained in the appendix to the report. Taken in total, the committee recommended a significant increase in spending on computing of \$1.4 to \$1.8 million over the five years. Of this amount, \$300,000 to \$400,000 was recommended for educational computing, roughly \$300,000 for administrative computing, and \$700,000 to \$1 million for research computing contract funds. For research computing, the committee recommended that a subsidy of 20 percent to 30 percent be granted to research users to supplement the real dollars and stretch their use of funds coming mostly from grants and contracts.

Proposed New Campus Computing Leadership and Organization

With a mandate to improve the business systems of the university, and perhaps recognizing the tension between the different computer user constituencies, Vice President Lawrence decided soon after he came to Cornell to split OCS into two organizations. The administrative systems analysts and programmers would become part of a new organization,

⁷ J. T. Anderson, G. V. Chester, D. A. Thomas, R. McGinnis, H. L. Morgan, A. H. Peterson, and W. B. Rogers, “The Use of the Computer at Cornell 1970–75,” October 1970.

Management Systems and Analysis (MSA). The remainder would continue to be called OCS and would comprise the computing center, including academic computing and all support services.

Serious consideration was given to creating a new position, coordinator of computing, reporting to both Provost Plane and Lawrence. The directors of OCS and MSA would report to this new position. While this issue was being debated, recruiting began for the director of MSA, and McWilliams continued as the director of OCS. In October, McWilliams resigned to take a position with the National Science Foundation, and John W. Rudan, who had returned to OCS in September as a statistical consultant, was appointed acting director. At this time, the OCS management group consisted of Dick Cogger, assistant director for systems programming; Dave Pulleyn, assistant director for operations (Langmuir and campus); Peter Shames, manager of academic computing; and Bob Blackmun, business manager.

Formation of the University Computing Board (UCB)

In 1971 the University Computing Board (UCB) was established by Provost Plane to deal with all computing policies and practices on campus. Lawrence was the chair of the seven-person UCB. The initial meeting was held on April 30, 1971. The first members of this group were Nyle Brady (director of research, College of Agriculture and Life Sciences), Conway, Justin Davidson (dean of the Graduate School of Business and Public Administration), Peterson, Chester, and Wally Rogers. ACAB was replaced by a subcommittee, the Advisory Committee on Instructional Computing (ACIC), although the same membership continued. Saltzman was the first chairman of ACIC. A Research Advisory Committee also was formed as a subcommittee of the UCB and was first chaired by Harold Scheraga from the Department of Chemistry. The existing Administrative Systems Planning and Control Board was to become a subcommittee of the UCB, with Wally Rogers as chair.

At the time the UCB was formed, a Computer Advisory Council⁸ was also formed. This council, consisting of seven knowledgeable and experienced computing practitioners from outside Cornell, was to advise the provost and the vice president for administration on planning and use of campus computing resources in service to the different constituencies. The initial appointees to the council were Milton Rose (Colorado State), Charles E. Dykes (U.S. Gypsum Company), William A. Vickroy (McDonnell

Automation), Donald R. Wood (Touche Ross and Co.), Ronald Brady (Syracuse University), Gordon Runner (J. W. Thompson), and David Freeman (University of Pennsylvania). The council met several times and in particular met in March 1973 to consider the issue of replacing the aging 360/65. Norman Zachary (previously at Harvard) joined the four members who attended this meeting: Vickroy, Wood, Runner, and Freeman (now at Rutgers). While expressing some skepticism about the materials presented to justify replacing the 360/65, they proposed a review of the load statistics and policy assumptions and suggested different approaches to the proposal to continue with a larger single computer. They expressed strong opinions on the need to be realistic about growth projections and to have a strong planning function to stay ahead of developments in this fast-moving field.

OCS Priorities

Two major priorities dominated the agenda for OCS during the early 1970s. The first was to deal with a host of financial issues, some coming from the budget crunch and others from the general financial underpinnings of the OCS budget. The second was to provide better services to all users: time-sharing services for instructional computing, large-scale computing for research users, and reliable systems and services for administrators. These all naturally became intertwined with the upward creep in the need for computing resources, which continued to grow at a rate faster than money and capacity could be provided.

With regard to the OCS budget itself, the problems essentially boiled down to the distinction between “hard” money—that is, the billed and recoverable income from different sources such as research grants, statutory funds, and external customers—and “soft” or “funny” money, coming from general-purpose university funds. The research community was finding alternatives to using OCS facilities by buying their own equipment, a step favored by the short-term duration of a grant and the increasing availability of minicomputers at affordable prices. In addition, minicomputers could be placed in laboratories to record and store data that could be analyzed later. Once purchased, such computing equipment could be used at a rather modest continuing cost after the grant expired. This was a much better approach for many groups, compared with buying and paying for services at the computer center while funds lasted and then potentially being denied access when the funds ran out.

Soft money, or allocation dollars, was diminishing despite growing needs, especially in administrative computing. The approach that finally evolved was to

⁸ Cornell University Computing Advisory Council, charter and minutes of meeting of March 2, 1973.

consider allocation funds as real money and to hold users supported by those funds accountable for their spending. That is, if a user was granted \$100 to fund services, the user could spend \$100 only, no more or less. Significant overspending could be covered by pooling funds over larger units such as colleges or major administrative organizations, but failing that, had to be covered by paying out of college or department budgets. While such a system satisfied the auditors and brought some discipline to the users, its major advantage was that it tempered demand a bit; its major disadvantage was that it inhibited actions to use computers more effectively in all areas.

Campus Financial Issues and Plans

In 1971 President Corson appointed an advisory committee on long-range financial planning to take a broader view of the financial situation at Cornell and make recommendations. This 14-member committee was chaired by College of Engineering dean Edmund Cranch and became known as the Cranch Committee.⁹ The appointment of this committee started off a round of self-study, introspection, and evaluation directed toward improving the financial situation and maintaining the scholarly work of the university. While this committee did not deal directly with computing on campus, it set a tone for the coming years on the priority issues for the university and how those should be approached.

The following quotes from the report prepared by the committee give a sense of the times:

Section 11.3: At the present time, there is considerable interest in the use of modern audiovisual aids—tapes, slides, films, cassettes—as well as in the use of interactive computer systems. Use of these expensive tools must be developed with discrimination. Controlled empirical studies, either at Cornell or elsewhere, which compare new teaching methods with the traditional and seek to find the relative advantages of each, are required.

Section 16.2 (part of Section 16, Support Functions): For example, there are at least 110 Xerox machines with an annual expense of \$500,000. It is difficult to believe we are making the most efficient possible use of this equipment.

The sentiments above were not unique to Cornell. A 1971 report from the Ford Foundation considered the question, “Does Every Campus Need a Computer?”¹⁰ The report, based on a survey of liberal arts colleges of different sizes, assessed the situation at the time and

attempted to provide guidance to college presidents on how they should address this increasingly important issue.

In 1973, President Corson responded to the Cranch Committee report with his own plan, “Cornell in the Seventies.”¹¹ In that report the only reference to computing dealt with the broad topic of information systems, which was included in the recommendations related to academic support services:

4. That expenditures be made to improve the university’s information systems to supply basic data about the institution, for planning and operations, and to satisfy anticipated government requirements.

The capital project appendix to the report mentions the need for increased computer capacity and upgrading of the computer electrical system, alluding to the activity going on to replace the 360/65 computer in OCS.

Campus Computing and OCS Leadership—John W. Rudan, Director

During 1972 the process to recruit the computing executive was moving forward. UCB members, college deans, and members of the Computing Advisory Council were asked for recommendations, and some of those individuals were contacted. The UCB and Lawrence expressed concerns about the authority that would be granted to such a position, the amount of money it would take to attract someone to Cornell, and the reporting relationship of the position. Although not all those issues were resolved during the year, two individuals—Richard Mills and Robert Woodruff—came to the campus to be considered for the position. However, in August 1972, Lawrence, citing the support of the UCB, suspended recruiting for the position pending clarification of the university’s willingness to step up its investment in computing by a substantial amount. The outcome was an indefinite postponement. Rudan continued as director of OCS.

Cost Cutting at OCS

As part of the budget accommodations in the early 1970s, there were numerous opportunities for cutting expenses by changing equipment configurations and vendors. At this time, third-party vendors began to supply CPU components, peripherals, and controllers,

⁹ Edmund T. Cranch et al., “Report of the Advisory Committee on Financial Planning to the President,” September 1972.

¹⁰ Does Every Campus Need a Computer? The Ford Foundation, Association of American Colleges, Robert Heller Associations, Inc., October 1971.

¹¹ Dale R. Corson, “Cornell in the Seventies; Goals, Priorities and Plans,” April 1973.

making it possible to achieve significant cost savings by replacing IBM equipment. IBM first met this challenge by raising the specter of the increased risk of experiencing both more down times, as a result of having less reliable components, and extended down times, as a result of not having onsite staff and having multiple vendors solving a problem. After evaluating different vendors, OCS selected Ampex Corporation to provide extended memory (bigger, faster, and cheaper), magnetic tapes, and communications controllers for the 360/65 system. Installing Ampex equipment saved upwards of \$50,000 per year. In general, all these moves were very successful and there were few serious performance and repair problems. The professionalism and concern of the repair staff from IBM and Ampex forced them to cooperate, for they all understood that they would collectively be blamed for any serious problems, regardless of the source.

The success of the installation of Ampex equipment was offset in part by an extended series of problems with the Ampex tape drives—magnetic tapes written at Cornell could not be read at other locations. It was really embarrassing when the PL/C project in the computer science department shipped out a new release of its software and few sites could read the tape. A major external client, National Planning Data, which was using OCS facilities to provide census data to clients, had the same experience. After some quick engineering work to correct the “skew” adjustments on all the tape drives, the problems were overcome, and increased vigilance kept the problems under control for the next several years. Somewhat ironically, by the mid-1970s IBM recovered from this experience and adjusted its product offerings and prices, making it possible for OCS to replace all the Ampex equipment with IBM equipment and save an additional \$50,000 per year over the already lower base budget.

In addition to these changes at Langmuir, Unitech printers replaced the higher-cost IBM equipment at all of the campus terminal facilities. Credit for this change goes to Errol Jones from CAG, the Computer Activities Group in Warren Hall.¹² Because CAG was motivated to increase its print volume, and hence its income, it could increase its profit by lowering costs. As a result, Jones took particular interest in looking for acceptable lower-cost options. Granted, the IBM equipment was of higher speed and quality, but the line-print quality from Unitech chain printers was judged acceptable for most of the “throwaway” type of work conducted at the campus terminals. At the same time, higher-quality printers were still available

at Langmuir when quality was required. In addition, Unitech was willing to place a technician^a in the Ithaca area to provide rapid response time for repairs and maintenance. Working in cooperation with CAG, Unitech printers were installed in 1972 for a total annual cost saving of \$50,000 to OCS. Soon after, Unitech-supplied card readers were installed at most of the campus terminal facilities, and self-service operations began so that users could now punch and load their own card decks and retrieve their own printouts.

Off-Campus Computing

The other major initiative to deal with lack of university funding for OCS was to accelerate efforts to find new external customers. Owing to concerns about Cornell's nonprofit status from federal auditors, OCS was restricted to having no more than 15 percent of its activity come from non-Cornell sources. To further inhibit external customers, they were charged 150 percent of the standard computing rates. Despite these higher costs, during this period Cornell was the only large-scale computing center in the Finger Lakes area that was accessible to commercial firms. By this time, the Finger Lakes Computing Center/FLARCO, the regional small college and high school computing initiative initially underwritten by NSF in the 1960s, had more or less folded, although individual colleges still used Cornell facilities. The largest external client at this time was National Planning Data, a newly formed census data processing and information company, founded when the U.S. Census department farmed out access to its data resources to private companies specializing in this activity. This arrangement was serendipitous, because the founder of NDP was a Cornell graduate, Peter Francese, who lived in Ithaca. Cooperating with OCS was to his advantage, and to Cornell's as well, because faculty could get no-cost access to census data for research and educational purposes.

To maintain and improve external income, John Aikin, who had earlier been hired as manager of the Uris Interactive Facility, was named regional computing coordinator to tour the area and see what new relationships could be established. It was at this time that a long and fruitful relationship started with Hamilton College. As the decade moved on, other attempts to secure additional customers and income were made.

Communications with Users—Newsletters and the *OCS Bulletin*

Erik McWilliams first published the *OCS Newsletter* in 1970 to improve communications with all users. Before that time, issues of concern to users were communicated by memoranda from the director. Rudan

¹² Errol W. Jones, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, February 1998.

continued publication of the *OCS Newsletter* after McWilliams left, and in 1972 a monthly schedule was adopted when Cecilia Uren (later Cowles) became editor. The newsletters were a combination of policy communications, descriptions of current and future issues, and technical hints and resolutions of problem reports.

In January 1974 Cecilia started a new weekly publication, the *OCS Bulletin* (the “yellow sheet”), to provide more timely reporting of problems, systems events and changes, and other technical items. As this publication proved to be more useful than the *OCS Newsletter*, the latter soon faded away and important items of policy or procedure were reported as the occasion demanded. Continuing the practice that was established in the 1960s, more care was taken to inform users of daily events by printing descriptions of problems or reminders of upcoming events on the header sheet of all printouts.

Squeezing More Out of the 360/65

One of the key priorities for OCS was to resolve some of its differences with ACAB, the Academic Computing Advisory Board. The hot issue was the yet-again recommended installation of a new operating system on the 360/65 system. Given the past history of systems projects, ACAB was concerned and doubtful of the success of OCS systems programming projects. As part of its program to raise the consciousness of users about computer issues, ACAB published its own newsletter, which was widely circulated on the campus and to the academic users. Relations between ACAB and OCS were not the friendliest during those times, although both professed to have the same objectives—to improve services to academic users.

The 360/65 had been augmented with additional memory, additional slower-speed memory, and larger and faster disks and still was not performing up to its potential. Dick Cogger, head of systems programming, recommended the reinstallation of HASP (Houston Automatic Spooling System) to gain efficiencies to carry the system for another year or two. ACAB was opposed. After discussions with ACAB, the University Computing Board, and interested users, the recommendation was approved and accepted. In 1972, CLASP 360/OS MVT was replaced by HASP 360/OS MVT in an installation that went smoothly and produced the anticipated improvements in service. At this time the campus terminals at Clark Hall and Upson Hall were staffed 24 hours a day to provide increased access to the computer.

The several positive incidents that the chairman of ACAB, Geoffrey Chester (personal communication, February 1998), recalls from this period were the

ways in which major systems programming projects were approached. Sunday afternoon “test times” were scheduled for major system changes. During these times users were encouraged to use the computer at no charge and see if such real jobs would break the new system. Initially, system outages occurred, but as the testing progressed, users were often able to get some relevant work accomplished at no cost in exchange for acting as testers. Another change that was adopted was skipping system software upgrades. When IBM issued an upgrade to correct errors or to make improvements, that upgrade was skeptically examined to see if the improvements were important or significant enough to warrant changing the system. Later on this review evolved to the point where selective beneficial changes were introduced into the current system to gain any advantage and to postpone a major change.

In discussing these system software changes, it is important to remember that in the 1960s OCS had introduced several major local modifications, or mods, including the online system for accounting for computer resources and the card file system (CFS), the replacement for the earlier card data system (CDS). These and other mods had to be integrated into the new system, which often involved extra work and creativity to avoid stressing the users—one of the constraints resulting from an increased customer base. Only much later could these mods be eliminated when IBM started to provide equivalent or near-equivalent features or a feature was no longer needed.

The other significant operations change made at this time was to reduce and reschedule the “test time” taken by the systems programmers to develop and test the system. In early 1971 there were three test times a day every weekday: from 8:00 to 9:00 a.m., from 12:00 to 1:00 p.m., and from 5:00 to 6:00 p.m. These test times created a conflict with the schedules of users, who, despite coming in early, staying over lunch, or staying late in the afternoon, did not get the turnaround they needed to keep their work up to date. Test times were rescheduled for early in the morning and for late evenings, when they were shared with hardware maintenance activities, if possible. This action alone reduced some of the tension between OCS and the different user groups.

Instructional Computing

Very little progress was made in providing time-sharing services in the early 1970s. In 1969 APL was ordered and installed to provide a degree of time-shared services. APL, while a very powerful language and subsystem, was used mainly by those who found it interesting and fascinating or for some particular feature such as matrix algebra calculations. By 1971 a few

IBM 2741 and Decwriter printer-based terminals were in Upson and Clark Halls to supplement the batch remote job-entry terminals.

Time-Sharing Services for Instruction

In July 1972 a modest-sized interactive terminal facility was opened in Uris Hall as a cooperative venture between OCS and the Advisory Committee on Instructional Computing (ACIC). ACIC submitted a proposal to Provost Plane for \$50,000 over two years to upgrade interactive computing by adding 12 more campus-based slow-speed terminals and additional full-time and part-time technical consulting staff. Only \$2,000 was granted, and as a result, all that could be accomplished was to increase the number of slow-speed terminals (SSTs) in Uris Hall from four to nine. These were printer-based IBM 2741s and Decwriters that had access to the 360/65 computer. The services included APL and CPS (an IBM offering, Conversational Programming System) for BASIC and PL/I and CRBE (Cornell Remote Batch Entry), which replaced the earlier Cornell Terminal System (CTS), to create and submit batch jobs.

However, between the lack of good software and all the other pressing priorities mentioned earlier, batch continued to be the primary way that students used computing, and OCS was under pressure to supply better service for the small student jobs. A semi-automated procedure called the student batch monitor was put in place as a result of significant pressure from ACAB. With this monitor, operators would assemble 20 student programs as a batch and submit it as a single run instead of 20 separate runs. The single run was given priority processing with the in-core compilers, and turnaround was made more acceptable without excessive use of scarce computer resources. Later the system was modified to load jobs every hour to further improve the turnaround of the small jobs.

Introduction of PL/C

One of the primary languages for introductory computing instruction changed from CUPL to PL/C. Conway and his various associates, who had developed CORC in the early 1960s then CUPL in the late 1960s, now followed with PL/C as the language of choice for this purpose. (CORC was used from 1962 to 1966 on both the Burroughs 220 and the Control Data 1604, while CUPL was used from 1965 to 1969 on both the Control Data 1604 and IBM 360/65.) Adopting PL/C was in keeping with the computing industry trend to adopt the PL/I language for more applications. IBM commissioned Conway¹³ and his associates to develop a student version of PL/I, which was named PL/C, and provided the initial support that was also

supplemented by support from Seimens. The development of PL/C was built on the previous experience noted above as well as new technologies, so when completed it was the first high-performance compiler with advanced error-correcting techniques for a subset of the PL/I language. Once under way, the project was partially funded by income from sales. In 1970–71 more than 100 copies of PL/C had been distributed, and 60 sales had been made. Later in the decade, at its peak, PL/C was used at 250 universities around the world, according to Conway. Members of the team that produced the first PL/C release were Conway, Howard Morgan, R. Wagner, and five graduate students in Computer Science, the principal one being Tom Wilcox.

Instant Turnaround (IT) for Student Batch Computing

From the very early 1960s, students could access the computer for class work by having their instructor obtain an appropriate number of computer accounts for the class. This account had to be backed up by funds: the statutory colleges used their own departmental funds, which counted as “real money” (income) to OCS; the deans of the endowed colleges could use “soft money”—allocated funds. Such student batch jobs entered the job stream and were executed according to the priority and job selection criteria built into the system. However, exceptions were made for the efficient in-core compilers such as WATFIV for Fortran and PL/C for PL/I that executed typical student programming assignments with a minimum of overhead and quick turnaround. This was accomplished in the early days with the locally written COOL system but migrated to the semi-automated student batch monitor after CLASP was installed. Professor John Hopcroft, from Computer Science, championed again building a “fast batch” system for these compilers in the HASP environment. He had returned from a sabbatical at Stanford University, where the students had used a fast batch system that provided essentially instant turnaround for small student jobs with very short job execution times, very small memory requirements, and only card input and line print output options.

In 1971 Dick Cogger and the Systems Programming staff developed the instantly successful instant turnaround processor, which became popularly known as “IT.” Students could hand in their program decks for the operator to read them into the system, and by the time they were handed back, the job would have started printing. Later when the Unitech equipment

¹³ R. W. Conway, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, November 2002.

was installed at the campus terminal locations, students were able to read in their own decks and tear off their own output listings at the self-service stations. As important as this service improvement was, future refinements led to the fundamental changes in the financing of computer access to students.

When IT was first introduced as a special service, students continued to use regular computing accounts, requiring faculty members to create course accounts and OCS to use computer resources to track and bill for such use. The first simplification to the IT process was made in 1972 when IT tickets were introduced. Basically, a course instructor was given a number of tickets to distribute to his class. The instructor could distribute these as he chose—all at one time or throughout the semester as assignments dictated. The computer science department, in particular, wanted to control the number of tickets per problem exercise. Such tight controls led to a flourishing black market in IT tickets that benefited those students with more on hand than they needed. To diminish the potential for theft of services and to provide an alternative to the black market, the UCB authorized OCS to sell IT tickets and to amend the service so that two IT tickets could gain the student twice the resource limits.

The UCB had no problem with students paying for instructional computing services with their own funds even if they had to go outside the course limits. The argument was that if the university did not cover the cost of as much computing as some students desired and those students were keen enough to want to use the computer more, then they could pay for the extra use with their personal funds. Jim Manning¹⁴ recalls that operators at the campus terminals sold batch tickets, one for a dollar or five for five dollars. Some class instructors gave students 5 to 10 such tickets for an assignment. A student who was not careful, for example, by not checking what was punched on the program cards, could easily run through his or her quota before completing the assignment. In such cases it was considered a fair penalty for the student to use his or her own funds to complete the homework assignment. It was not uncommon for graduate students, particularly in the humanities, to spend their own funds for using the computer for their thesis work if the department or chairperson had no access to computer funds. These students did all of their computing after midnight, during the third shift when the rates were the lowest and turnaround of jobs was acceptable for their work.

Free IT—A Critical Milestone

A year later, in the spring of 1973, under the auspices of ACIC, an experimental program was undertaken to simplify the IT process. Any student, in fact any employee, could get a \$50 restricted account in his or her name simply by going to the terminal facilities at Upson Hall or Uris Hall. In March 1973 more than 2,200 IT accounts were distributed on the campus, and as a result, 140,000 jobs were processed over the spring semester, compared with 55,000 the previous fall semester! Resources were tracked and accumulated by college or department. Use by endowed units was billed against a central allocation fund, and use by statutory units was billed directly to them. The \$50 per semester, later \$50 per year, limit was simply there to place a restriction on use, although, practically, it could be renewed with little effort.

The following year, the process was simplified even further when free IT was introduced. All a user had to do was punch his or her name on the IT batch card, and the system could be used any time and any place. OCS continued to track resources and accumulate them by college or department, but the information was used to distribute an estimated fixed annual cost of the IT program to the endowed and statutory units. This major change in the funding for instructional computing can be credited to controller Arthur H. Peterson, who persuaded the statutory college deans that the savings in OCS accounting staff and computer resources and from billing by prorated share were worthwhile. Free IT and this special IT billing arrangement continued throughout the decade. Although in itself fast batch was not that innovative, the trend started with free IT, that is, funding instructional computing at the college or university level instead of direct billing to these units, was a critical milestone. This more rational approach became known as the library model, implying that students should not be charged for routine educational computing services just as they were not charged for routine library services.

Although the available statistics are fragmentary, the effect of making student computing more generally available—and eventually free and unlimited—is demonstrated by the information in Table 1 on student jobs.

The other part of the ACIC experimental program in the spring of 1973 was to offer free interactive computing accounts to students. This proved very popular, and more than 400 such accounts were distributed. Regrettably, the distribution stopped because it created overload conditions on the 360/65 computer and it wasn't possible to maintain adequate services.

¹⁴ James F. Manning, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, December 1998.

Table 1. Number of Student Batch Jobs, 1965 to 1975

Year	Period	Student Batch Jobs	Comments
1965	spring semester	13,059	
1969	fall semester	53,257	
1971	academic year	178,000	first year of IT
1973	academic year	200,000	
1974	academic year	370,000	first year of free IT
1975	academic year	510,000	

Biological Sciences Initiative

Elsewhere on the campus other departments were creating their own decentralized computing facilities for student work. In March 1971, Howard C. Howland announced the opening of the Biological Sciences interactive computer facility in Stimson Hall with a DEC PDP-8/L with four teletype terminals. The facility was intended for Bio Sci 106 students, but others could make reservations at \$2.50 per half hour. Mini-grants to support other use could be applied for. Shortly after the opening, it was announced that unre-served time was available at \$0.50 per half hour.

PLATO for Instructional Use

One of the interesting experiments in instructional computing was the installation of several PLATO terminals as a joint venture between OCS and several of the larger Cornell colleges. PLATO (programmed logic for automatic teaching operations) was a system developed at the University of Illinois for use as a programmed learning device for students up to the college level. It was noted for its tutorial capabilities and an extensive library of course materials in a number of disciplines. It used an innovative, and for its time advanced, monitor that could display graphics as well as text to support class exercises. For example, the screen could display images of fruit flies for succeeding generations as the population expanded over time. Control Data Corporation provided the commercial service, which could be accessed by telephone connections to one of its computing systems. One terminal was installed in Uris Hall in 1974, and Aikin coordinated the writing of a proposal to the NSF to expand the program. When that was not funded, the colleges funded the acquisition of a second terminal in 1975. The estimated annual cost of the program was \$14,500, which was shared by Arts and Sciences (\$3,000); Agriculture, Engineering, Law, Academic Funding, and OCS (\$2,000 each); the Business School

(\$1,000); and the Veterinary College (\$500). While these dollar amounts roughly represented the actual use, the Department of Music was expanding its use rapidly after purchasing an audio output device. PLATO was thus another instructional service offered in Uris Hall, which was becoming the central location for OCS special services and staff to support the faculty in their use of the 360/65 computer for instruction.

Research Computing

The changing nature of research computing on the campus is best reflected in the statistics from the 1973 UCB computer planning policy statement:¹⁵

Table 2. IBM 360/65 Research Computing Use, 1968 to 1971, by Percent

	1968–69	1970–71
Research	60.0	40.0
Administrative	30.0	47.5
Instructional	10.0	12.5

These statistics also offer some insight into the concern at that time that the central facility was losing its academic orientation and becoming more of an administrative support service. Had all the applications from Machine Records and the Statutory Finance and Business Office been on the 360/65 during this time, the figures would have been even more skewed to administrative use. Nonetheless, in many cases the availability of special software for performing statistical analyses, or the lack of funds for other alternatives, resulted in the 360/65 at OCS being heavily used by researchers on campus despite the competition from minicomputers being installed at a rapid pace in laboratories.

Advances in Statistical Computing—COSSIS Proposal

In 1970, OCS, in cooperation with CAG, published a statistical computing newsletter for wide distribution on campus to inform the research community about statistical analysis offerings on the 360/65. The library of statistical programs included the Bio-Med series (from UCLA), CAG Express (which CAG wrote for small data sets), and other user-contributed and -supported analytic routines. The newsletter, published for a number of years to keep users informed of changes,

¹⁵“Recommendation to Increase Computing Capacity at the Office of Computer Services, Cornell University,” University Computing Board, July 1973.

in 1971 announced the availability of several new packages: SPSS (Statistical Package for the Social Sciences), OSIRIS, and SSP. OSIRIS, developed by the Michigan Institute of Social Research, replaced the earlier ISR package on the system. SSP was IBM's Scientific Subroutine Package, offering programs that supplemented other statistical programs or alternate choices. These powerful new programs enhanced the ability of researchers with extensive data collections to perform complex statistical analyses at reasonable cost and convenience.

In 1971 when ACIC submitted its proposal to the provost to expand interactive computing, it also submitted a proposal to establish a computer-oriented social science information service (COSSIS) built around the already available census tapes. While there was some concern that this service would overlap with the responsibilities of the Cornell Library, the library supported the proposal. The UCB also supported the proposal, which requested \$45,000 largely for staff. It was suggested that this service center be located in Uris as a part of the Center for Urban Development research. The proposal was not funded and the initiative was put on hold.

New Policy for Distributing Allocated Research Funds

Some units of the research community at Cornell were using their grant funds both at OCS and in acquiring their own computers to enjoy the benefits of access to computing beyond the life of the grant. Soon after it was formed, ACAB took a bold step in July 1970 and announced a new policy for distributing allocated computing funds to the campus. They recommended that allocated computing dollars available to researchers for the use of OCS facilities be distributed by the Cornell Research Grants Committee. Provost Plane accepted this recommendation. The first solicitation was to take place in July with awards to be made in October. The total of such awards was to be \$60,000 for the 1970–71 year. These funds were obtained by reducing the computing allocations to deans. Later, as the UCB got more experience and became more knowledgeable about computing affairs on campus, they took over this responsibility from the Research Grants Committee. However, after a few years and given the complexity of the issues and the need to make hard decisions, the UCB moved this process back into the colleges for the deans to make the decisions.

Research Advisory Subcommittee Projects—Media Conversion, ARPANet

The Research Advisory Subcommittee was involved with several projects in its first years ranging from the

quite small to the very large, all directed to improving services to research users.

The subcommittee was an ardent advocate for the OCS proposal to install a generally available “media conversion” terminal in Clark Hall and argued for the funding to make this facility possible. This terminal would not only replace aging equipment at Langmuir but also provide local campus services for plotting information and reading paper tapes. To meet this requirement OCS installed a small PDP-8 computer in Clark that was connected to the 360/65 at Langmuir. Plotting information was downloaded to the Clark facility while paper tape, still then being used to record research data, was read and transmitted for storage on the 360/65 or copied onto magnetic tapes.

At the very large end, the ARPANet project generated considerable interest to the members of the Research Advisory Subcommittee. In 1971 and 1972, Educom was attempting to generate interest in network and computing plans being developed by the Advanced Research Projects Agency (ARPA) of the Department of Defense. Although the network was of interest itself, a second item of importance was the building of the Illiac computer, which would be accessible from the network and capable of large-scale computations. While the UCB and the subcommittee were interested, they were apprehensive about the costs involved: \$50,000 to \$100,000 one-time costs for hardware and up to a man-year in technical support plus operating costs of at least \$20,000 per year plus network usage charges.

OCS, on the other hand, was more supportive of regional networking. The request also came at a time when the financial pressures on OCS were such that it had lapsed paying Educom fees and only a letter from the president of Educom to Cornell President Corson restored the membership. The Research Advisory Subcommittee was more interested in access to computers that could be immediately used, like the CDC 6600, and not in a proposed machine. In fact, a small number of research users (led by Chester and Scheraga), working with OCS installed a single terminal in Clark Hall so that they could access these large remote systems at major national research laboratories. Consideration of ARPANet was dropped at this time.

Surveys of Computers on Campus

The University Computing Board (UCB) concerned itself with broad university issues, the most important of which was the overall status of computers on campus, and with policies governing OCS operations and financial stability, including off-campus use. In this role the UCB acted very much like a board of directors concerned with OCS rates, services, income and

expense budgets, and results. With respect to the overall status of computers on campus, preparations were made in early 1971 to canvas the campus to determine the number and type of use for all the computers that could be identified.

Based on the information known before the survey, it was estimated that there were 21 computers in research labs on campus. Of these 21, two were IBM 1800 systems and the rest were a combination of PDP systems from DEC. They included PDP-8s (8), PDP-9s (1), PDP-10s (1), PDP-11s (7), PDP-12s (1), and PDP-15s (1). The survey¹⁶ confirmed the presence of these machines and identified a whole range of other “computers,” most of which were programmable calculators from Wang, Hewlett-Packard, and other vendors. These systems could have a variety of other equipment connected such as keyboards and printers and so offered a powerful tool in the classroom or the lab for carrying out computations.

The survey was repeated the following year, 1972, and at that time identified 25 laboratory computers dedicated to data acquisition and analysis, with occasional use by students for homework assignments. An increasing number of programmable calculators was also reported. Nine other “computer centers” were identified. These centers served one or more departments (in one or more buildings, as in the case of the Materials Science Center) and a broad spectrum of users and uses. Included were the Dairy Records Lab in Animal Science and computers in the Hotel School, Arecibo Observatory (CDC 3300), Chemistry, Physical Biology, Geneva Agricultural Experimental Station, Materials Science Center, Laboratory of Nuclear Studies (PDP-11s connected to a PDP-10), and the Division of Biological Sciences in Stimson Hall. This was the last formal survey and compilation of computers on campus.

To keep track of how these smaller computers were changing over the years, the UCB intended that this survey be kept current. A procedure was put in place that required Central Purchasing to route all computer or identifiable computer component purchase orders to OCS for compilation and presentation to the UCB for approval. Because the UCB as a whole met only once a month, this procedure resulted in delays and complaints from researchers. Accordingly, the procedure was modified so that the UCB would concern itself with larger systems—those that could be considered “computing centers”—while Rudan could directly approve changes to existing laboratory-based systems or new smaller such systems unless serious issues needed to be discussed with the UCB. In this review Fred Hiltz ably assisted Rudan in evaluating those requests. Since the requests were seldom turned down, unless university funds were involved, the pro-

cedure diluted over time and OCS was removed from the loop unless called for advice. The UCB continued to exercise review and approval of the installation and upgrade of the larger multidepartment, multiuser systems throughout the decade.

OCS recognized that it had to broaden its mission and campus services and had earlier hired Fred Hiltz to be a consultant to the campus on the acquisition and use of minicomputers. For many years, Hiltz had been a technical specialist in the College of Veterinary Medicine and had helped the college acquire and develop computers and data acquisition systems and had a particular interest in these smaller systems.

Business Systems

Earlier recommendations in reports to form a separate organization for administrative data processing and the intentions expressed by the UCB resulted in the formation of a new unit.

Formation of MSA (Management Systems and Analysis)

In January 1971, Henry (Hank) G. Vaughan was hired as director of Management Systems and Analysis (MSA), which was charged not only with supporting the university-wide business systems but with the broader mission of examining university management activities. The staff from the former OCS administrative programming group was combined with the staff from Institutional Research to carry out this larger mission. At the formation of MSA, William Tetlow was director of Institutional Research and Jerry Tucker was director of the administrative programming section in OCS. Within a short while, Edmund V. Hollenbeck was appointed to this leadership position. Senior staff at the start of MSA were Libby Gruppuso, John Joubert, Dolly Sewell, and Charlie Evans. Shortly after, a small number of the administrative programming staff moved from Langmuir to Day Hall to improve contact and communications with the campus-based administrative offices.

Within three months of his arrival, Vaughan wrote a report critiquing the state of administrative systems at Cornell.¹⁷ In the summary of that report Vaughan states:

The challenge of the future is immense, but not overwhelming. The first priority activity should be to improve customer relations.

¹⁶ John W. Rudan, “1971–72 and 1972–73 Surveys of Computers on Campus,” Summer 1971 and 1972.

¹⁷ Henry G. Vaughan, “The Crisis and Challenge of Administrative Data Processing at Cornell University,” March 1971.

Second priority is the development of new administrative processing software in a controlled environment. The third priority is to catch up with technology by integrating individual software requirements using the facilities of a teleprocessing-oriented data management system. The ultimate objective, of course, is to develop an integrated and cost-effective administrative data processing system at Cornell that can serve as the nucleus of an integrated university management information system.^b

The opportunity to work toward these objectives came quickly. Continuing with the initiatives started in the late 1960s, there were two major systems under development in 1970: a new university-wide Payroll/Personnel (P/P) system being developed by the staff then in MSA and a new university accounting (general ledger) system being developed by the Statutory Finance and Business Office. The P/P system was to bring Cornell into compliance with new federal guidelines, particularly biweekly instead of semimonthly paychecks and overtime pay for hourly paid staff. In addition the new system would incorporate a large number of policy exceptions being done off-line with other systems or by hand and reflect major changes in related policies. The accounting system would totally revise the department account numbers and the tracking of all the different expense categories. Essentially, these were systems that would affect the whole campus and change the way departments and central offices conducted business on these two important systems.

New Payroll/Personnel System

The first new system developed soon after the formation of MSA was the Payroll/Personnel (P/P) system. This system was the first at Cornell to use new database technology, the IMS (Information Management System) from IBM. MSA staff reviewed the earlier process to select a database system and confirmed the decision to adopt IMS as the standard for new systems development and acquired the fee-based programming product instead of using the freeware system for the P/P project. The project leader for the P/P system was Charlie Evans, who was granted a temporary leave from the Systems Programming group to take on this assignment. Other members of the project team were Bruce Lloyd, Al Seliga, Alan Doniger (who left part-way through the project and was replaced by a new hire, Thomas A. Dimock), and later Libby Gruppuso.

In May 1971, Controller Peterson announced that the new system would be installed on September 15, 1971, so that the old and new pay periods would coincide. Programming and testing proceeded on a pace

to achieve this goal, although the usual policy loose ends existed in which decisions had to be made by the systems analysts and programmers for the programming to move ahead. Installation took place as scheduled, but there were problems described as a “computer data problem,” which delayed 4,300 paychecks by a day. Vaughan summarized the main problem as the data volume being greater than expected and the quality of the data being poorer than expected. The statutory and endowed payroll offices had to close because the requests were coming in faster than they could be dealt with. Nonetheless, the heroic effort of the programmers kept the problems from getting worse. Dimock¹⁸ recalls working 32 hours straight without a break on resolving problems as they appeared. Some policy peculiarities made repeat runs take a considerable amount of time. For example, the paycheck number assigned by the program had to match the preprinted number on the actual paycheck, and if something went wrong, one essentially had to start all over again, recomputing all the checks. The P/P problems lingered for a while, and it was reported in late October¹⁹ that 1,000 employees were still having problems getting their checks. Of these 1,000 problems, 400 resulted from late submission of information, 200 had no appointment forms, 100 had no Social Security number, and 300 had incorrect details on the voucher. The new, unfortunately sophisticated, system with tighter controls was bringing to the surface a variety of problems that had been tolerated by the older system. Campus offices viewed this new system as a failure of central computing, not the success anticipated when the project started.

New Student Information System

A second opportunity for MSA arose in late 1971 when the Office of Student Records and Finance was established to consolidate and improve student services for admissions records, student biographic and academic records, and financial services. R. Peter Jackson was appointed director while continuing as registrar, and Garry Lee, Ralph Miller, and Byron McCalmon were the key department heads representing Scholarships and Financial Aid, Bursar, and Systems Development, respectively. Within a short time, McCalmon was appointed university registrar. This new organization took a bold step forward to create a “one-stop shop” (although that term was not

¹⁸ Thomas A. Dimock, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, July 2001.

¹⁹ “Payroll Problems Not Solved Yet,” *Cornell Chronicle*, October 28, 1971.

in vogue at the time) for almost all student-related business activity. But no business system integrated all the different activities of the combined units and their independent systems. After a long, complicated, and intensive process of preparing specifications, soliciting vendors, and reviewing proposals, the decision was made in early 1973 to spend a total of \$418,000 to build a new integrated system, the Student Information System (SIS), to replace the 15-year-old system for student records. The implementation, with a primary undergraduate focus, was to take place in three parts: (1) selections—admissions and financial aid; (2) active—biographical and academic records; (3) financial—student payment and financial aid records. The university took another bold step in buying software from a reputable vendor, Systems and Computer Technology Corporation (SCT),^c instead of continuing to develop its own administrative systems. The \$148,000 contract price with SCT was included in the overall project cost.

The design of the new SIS system began well, with expectations and announcements by McCalmon that the system would be operational by fall 1973. However, things seemed to fall apart when it came to actually programming and implementing the system. The SCT project leader who had done such a fine job during the design phase, and who was asked by Cornell to be the leader of the implementation phase, was not able to deliver the software modules on time and with agreed-upon functionality and performance. He was replaced. Even then the programmers who made their appearance at Cornell to test modules looked very much like recently recruited junior staff who wrote the code on the way to Ithaca instead of coming to Ithaca with already tested materials. Despite contract stipulations and countless meetings between SCT and Cornell principals, it appeared that project deadlines would not be met. In September 1973, McCalmon was making statements that the plans to implement SIS by spring 1974 would be delayed and “we would be extremely imprudent to wait any longer to reinstitute the old system as a backup.” Basically, the old system was put back into production and negotiations for terminating the SCT contract started and were completed a year later in May 1975.

Once the SCT contract was terminated, serious consideration was given to importing the Student Records and Registration System from the University of Pennsylvania. However, because of incompatibilities, costs of acquisition, and costs to create a better fit for Cornell, this approach was abandoned. Taking a cue from the sentiments being promoted at the time to consider minicomputer-based commercial systems, the WISE student system being marketed by DEC was examined. However, it had a hard limit on the num-

ber of students that was far less than the number of Cornell students and so this approach was not feasible. As a result, the Student Records and Finance Office limped along with no feasible alternative to the old system. In time, many of the principals took other positions at Cornell or elsewhere, and reorganizations eventually brought back the separate and original departmental structure.

New Financial Accounting System

The new Accounting/General Ledger system went into production on schedule on July 1, 1971, despite having its own set of problems and setbacks according to Bob Mack.²⁰ Les McKee was slated to head up the new effort of this joint endowed/statutory system. McKee, who had joined the university in 1965, was the designer of the first computer-based accounting system for the endowed colleges that was run on the 1401 in Machine Records. It was decided to staff this project with existing employees even though they had no programming experience, so it was at this time that Dan Argetsinger and Bob Mack became programmers by scoring high on the IBM Programmers Aptitude test. Mack recalls that someone from IBM came to Cornell and gave the group, which also included Matt Hayes, one week of training to learn COBOL, a first computer language for most of the group. About six weeks into the project, Les decided that he had had enough and announced he was retiring and leaving the university in January 1971. That left Mack and Argetsinger as co-designers and co-project leaders. They split the system responsibilities so that Mack would be responsible for the general ledger portion and Argetsinger for accounts payable and labor distribution. To shield them from distractions, a room was rented in the stacks of Mann Library where the two of them worked, as Mack recalls, “horrendous hours” to make the whole system come together.

It was also about this time that Stew Comber had become director of the Statutory Finance and Business Office, moving from Endowed Accounting, and that Joe Bates became head of Endowed Accounting. Much of the final design and policy came from these two directors.

According to Mack, it was this “small committee” of Bates and Comber who decided on the account numbering schema of letters and digits for departments, the practice still in place in the year 2003. Mack recalls an interesting sidenote about the peculiar way in which negative numbers were captured in the system. Largely because Comber didn’t feel confident

²⁰ Bob Mack, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, June 1998.

that negative transactions would be entered into the system correctly, he insisted on having a self-checking mechanism. The convention at the time, continuing from the “tab” machine days, was that a negative number was entered as an “11” overpunch in the digit column of the number field. Correspondingly, because the alphanumeric character codes were made up from combinations of the 10 digits with the “11” and “12” overpunches, the “11” punch over a 1, for example, would be read as an “A.” This was much less straightforward than punching a “minus,” for example, in the leftmost column of the number field, as one had to wait to punch the minus in the digits field, and thus the perceived need for added redundancy in the object code field. As a result, the practice, which continues today, is that using a letter for the first character of the object code indicates negative numbers. Mack takes responsibility only for writing the small subroutine that examines the first character of the object code and made the appropriate sign change. The university accounting system, with the usual changes resulting from policy and technology changes, continued to be run on the 360/25 until the mid-1970s.

Mark IV Query Language Installed

One of the significant improvements made in 1972, in keeping with Henry Vaughan’s priority for developing management information, was acquiring the Mark IV product from Informatics at a cost of \$22,000, roughly 40 percent of list price. This was a high-level query-language/report generator for easily extracting information from business system data files into special reports. It also had the advantage of being inherently self-documenting. Although the system ran only in batch mode, it greatly improved the ability of owners or users of a major system to compile special reports. Before Mark IV, special reports were customized programs and, unless the programmer could “tweak” an existing program, they had to be written from scratch each time. When first brought on campus, Mark IV was used exclusively by programmers, but in time and with training sessions, departments were able to write their own requests as needed.

1401 Computer Finally Removed

Not to be lost among all these other priorities in OCS and MSA, the 1401 system was removed in January 1972, roughly four years after the first plan for its removal. A champagne party to celebrate this achievement was held on February 4.

Improving Budget Systems

To expand the initiative to improve Cornell’s business systems, in 1973 Lawrence recruited John (Jack) S. Ostrom from Princeton University as the new director of financial systems development. Ostrom’s challenge was to find expeditious ways to move the endowed budgeting and expense tracking system from a manual-based ledger system to a more contemporary computer-based system.

MSA Staff and Business Systems Support Services Move Back to Day Hall

Other improvements continued to be made. In 1974, the production control staff, the data entry staff, the 360/20 printer/card reading system, along with a variety of still-required card-handling equipment, was moved back to Day Hall from Langmuir. To accomplish the move, the MSA staff was temporarily housed in Rand Hall before moving back when the Day Hall basement space had been minimally refurbished.

So, in roughly a 10-year time span, the administrative computing programming and support staff were back where they started in Day Hall. The return of the operations and production staff to the campus helped to restore, and shore up, the working relations between OCS and the administrative departments in Day Hall. With this move, the basement of Day Hall once again was the location of all the computing staff supporting the campuswide university business systems, with the exception of the Accounting/General Ledger system. That situation was continued for another 12 years, until the mid-1980s.

OCS Initiatives and Activity in the Mid-1970s

Deteriorating computing services, the negative feedback from new business systems, and the failure of the SIS project created a tense situation on campus regarding computing services. But unless computing capacity was expanded, it would be difficult to accommodate the modestly increasing use from the current instructional and administrative users. In contrast, other factions on campus felt the need for a closer examination of some of the problems before any action was taken or funds expended during these tight budget times.

Recommendations to Increase Computing Capacity at OCS

By 1972, it was clear that despite all the improvements and changes the 360/65 was at the end of its useful life and needed to be replaced. The UCB dealt with this issue by first developing a Computer Planning Policy Statement that was issued to the campus in May



Figure 1. IBM 370/168 in Langmuir Laboratory, 1974

1973.²¹ The UCB did an excellent job in documenting the growth in central computing capacity in the different segments of instruction, research, and administrative computing as well as the ways the future needs for computing might be addressed in each of these segments. It concluded that OCS needed to upgrade its capacity by mid-1974. No recommendation was made in the report but some alternatives were described.

During this time, and forming part of the deliberations on alternative plans to increase computing capacity, Judy Campbell, associate director of MSA, wrote a report on the future of administrative computing at Cornell.²² Picking up on the “semi-autonomous” model described in the earlier Chester-Thomas report, Campbell argued strongly against a single shared system and for separate systems for academic and administrative work but as part of a larger organization that shared resources. Campbell also referred back to the 1957 report,²³ which argued for centralization but not necessarily one computer to carry out all the work. The Campbell report reinforced the presentation made by OCS in a 1971 report²⁴ prepared for the UCB by Rudan that also argued for separate computers to satisfy the different user requirements and improving services.

After several open campus meetings in June 1973, despite some objections of continuing with a single large IBM computer and not seriously considering other vendors, the UCB recommended the installation of an IBM 370/168 to replace the IBM 360/65. In the end, expenses played the decisive role in putting the “semi-autonomous” alternatives on the back burner, as the alternatives with more than one computer were more expensive. The total cost of the 370/168 was estimated at \$3,728,000 including some site preparation costs. (By this time the 360/65 had been in use

for seven years.)^d After a lot of ritual posturing—first getting Board of Trustee approval for sending a letter of intent to IBM, then their approval for the transaction and then for the method of financing—the system was installed in August 1974. The financial plan called for the above amount, less estimated resale value, to be amortized over six years with income from user charges! The board also authorized the sale of the 360/65 cpu for a total of \$271,125.^e

Faculty Reports on Campus Computing—Chester and deBoer

In response to a request from the Executive Committee of the Faculty Council of Representatives, the dean of the faculty was asked to form a small ad hoc committee²⁵ to report to the faculty on the recommendation of the UCB to install the IBM 370/168 computer. The committee, which included Conway and Saltzman from the UCB, supported the UCB recommendation and stated:

²¹ University Computing Board, “Recommendation to Increase Computing Capacity at the Office of Computer Services, Cornell University,” July 1973.

²² J. A. Campbell, “The Future of Administrative Computing at Cornell,” *Management Systems and Analysis*, February 1973.

²³ W. S. Gere, B. McK. Johnson, and R. C. Morris, “An Analysis of Data Processing and Computing at Cornell University,” November 14, 1957, Copy Number 10.

²⁴ J. W. Rudan, “Preliminary Plans for Computer Development,” Prepared for the University Computing Board, October 28, 1971.

²⁵ G. V. Chester (chairman), R. W. Conway, N. R. Lyons, S. R. Searle, S. Saltzman, and D. E. VanHouweling, “Report on the University Computing Board’s Recommendation to Increase Computing Capacity of the Office of Computer Services,” September 1973.

They (the UCB) estimate that the acquisition of this machine will increase our computing capacity by a factor of between three and four and will require an increase in expenditure on computing in the range of \$500,000 –\$700,000 per annum. This increase in expenditure would raise our total expenditures on computing from \$1.9 million to between \$2.4 and \$2.6 million, an increase of about 25 percent. While the committee strongly supports this recommendation, it does so with one major reservation. We believe it is essential for the university to greatly improve the range and quality of its computing services in the next two or three years. To increase our computing capacity is a vital step toward this goal. The committee's position is that the recommendation from the UCB should be regarded as the first step in a two- or three-year program to improve our computing services until they are of the same quality as those currently provided by the library system. If our library services were reduced to the level of those provided to computer users, the campus would be in an uproar and many of our faculty would depart to other institutions.

The committee presented some comparative statistics with other institutions showing Cornell on the very low end and argued strongly for improving services in concert with capacity improvements by further budget increases and sharing resources on campus and regionally to accomplish these service improvements.

In 1973 the Research Policies Committee of the University Faculty Senate formed a subcommittee to consider the following resolution:

Be it resolved that the Dean of Faculty will appoint a committee to provide an outline of the various uses of the Cornell University computer.

This resolution arose as a side issue while the Research Policies Committee was considering another resolution of December 3, 1973, on the "governance structure of computer operations."

A subcommittee chaired by P. C. T. deBoer was appointed. It took as its charge:

Specifically, the need for a new computer has arisen from extensive increases in usage. There appears to be a general lack of information on the scope and character of these increases. It is important that such information be made available.

Their report issued in April 1975 became known as the deBoer report.²⁶ It dealt very briefly with research and instructional computing. It noted that the increase in instructional computing expenditures from \$257,000 in 1971–72 to \$377,000 in 1973–74, coming largely as a result of the increases in IT jobs, was in line with the expansion of capacity. The remainder of the report dealt with administrative computing in general and expressed concerns on the projected growth of administrative computing based on information from the earlier report by Judy Campbell. Sections were devoted to the newer systems such as Payroll, Admissions, Student Information System, and the Library. While the report made no hard recommendations, it questioned the need for more complicated and complex administrative systems that stored more and more data but did not appear to serve the needs of the faculty. It also questioned the need for more government reporting requirements and how faculty at Cornell in concert with faculty at other institutions might work to temper this demand.

The general tenor of the report is demonstrated by the following excerpt from the section "Directions of Increasing Computer Usage":

Increased use is being made of the computer for storing and processing various academic records. We have not been aware of any great benefits to either students or faculty from this practice. Grades are not processed any more rapidly, possibly less so. For example, some years ago, mid-term grades were available for advisors at pre-registration time; this has not been true for some years now. It is our subjective impression that the various computer outputs we are provided from time to time have not been particularly useful to us.

As mentioned before, the new direction in administrative computing envisions, eventually, a complicated, expensive on-line system with capabilities which we believe to be of dubious usefulness. The system has rather disturbing implications with regard to privacy and confidentiality of records. It imposes on future generations of students the burden of frequent updating of records in order to keep the system operating.

VP Lawrence, chairman of the UCB and the vice president responsible for both MSA and OCS, wrote a brief response to the report, basically clarifying some

²⁶ P. C. T. deBoer (chairman), N. R. Lyons, W. L. Maxwell, A. Silverman, P. C. Stein, "Report of Subcommittee on Computer Usage," February 1975.

of the comments about rates and charges but primarily disagreeing that new or proposed systems were unnecessarily complex. He emphasized that such systems were necessary as the primary source of operating and reporting data. Further, these issues could not be dismissed without some additional discussion.

An Early Computer Security Incident

An incident in early 1974 made the campus, and administrative computer users in particular, feel acutely aware of the potential insecurity of their information on the OCS computer. Several student operators obtained OCS computer accounting information. This raised some concerns about unauthorized access to sensitive and private personal information stored on the central system. Errol Jones, director of the Computer Activities Group (CAG), recalls hiring the five part-time student operators who were studying computer science. In their spare time during idle late night or weekend shifts they would experiment and discover system “holes,” deficiencies in the 370/168 operating system software. They discovered that the computer accounting file was encrypted but not password protected. They experimented and found the encryption key and were able to read the file. Jones recalls²⁷ walking into the CAG terminal room on a late Friday afternoon and finding one of the operators with a six-inch stack of paper containing the entire records of OCS computer accounts, passwords, access codes, and funds available. He quickly confiscated the materials and notified OCS, who took appropriate precautions to disable the perpetrators’ access, but on Monday morning the *Cornell Daily Sun* had banner headlines about the incident. Rudan and Vaughan thoroughly investigated the matter and were able to satisfy themselves that the worst that could have happened was the theft of computer time. The only files of consequence that could have been read were academic and not administrative information. Access to sensitive information from student records or payroll, for example, was under the strict control of production staff. Nonetheless, the incident caused a furor on campus.

The University Computing Board issued a strong statement to the campus about this issue, saying that the university needed to strive to achieve a balanced approach to protecting information while working to achieve more general and easy access to computing:

Unfortunately, the very complexity of computers and of protective systems that are

used to defend them appears to challenge the virtuosity of some thoughtless computer users.

Although this incident was closed, it was but one of a continuing sequence of incidents to steal computer time or otherwise achieve no-cost access to the computer or try to access administrative records. It was just a forerunner of more serious attempts to compromise the integrity of computer systems at Cornell.

Regional and National Computing Initiatives

Before the installation of the 370/168, Cornell was part of a SUNY (State University of New York) Region II initiative to consider regional computing for that area. The schools in Region II included all public and private higher educational institutions in the New York State Region II area and ranged from community colleges to research universities, of which Cornell was the largest and most comprehensive institution. A Coordinating Computer Committee was organized as part of the broader initiative to consider cooperative interests that involved a total of 12 committees such as course coordination, library resources, and continuing education as examples. Robert E. Roberson of SUNY Binghamton was the first chairman of the Computing Committee, who served until 1973 when Rudan became chairman. The Region II Computing Committee, and all the other regional computing committees, worked with Harold Wakefield, assistant vice chancellor for computer system development in SUNY. (Wakefield was the SUNY Central Computing “czar” who had replaced Richard C. Lesser in that position when Lesser, who had been the first director of the Cornell Computing Center, took another position.)

By this time most of the SUNY institutions had a combination of local computing on their campus and/or were linked to a central system in SUNY. As a result, they had varying degrees of interest in a regional computing center to serve all or part of their needs. Several proposals were developed, but it was difficult to get agreement that would involve all the institutions in the area. The final agreement that was forwarded to SUNY Central Computing called for Cornell to upgrade and install an IBM 370/168 and for Binghamton to upgrade its computer and install an IBM 370/58 but to link both machines with a high-speed network connection. The network connection would allow some common operations that could serve the region and yet allow independent operations. Rudan and Lawrence from Cornell and their counterparts, Edward Demske and David Stonehill from Binghamton, developed that proposal. When the proposal was forwarded to SUNY Central Computing

²⁷ Errol W. Jones, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, February 1998.

under the signatures of President Corson from Cornell and Peter Magrath from Binghamton, funds were requested from SUNY to install and support the network between the computers. When there was essentially no response to the proposal, regional computing more or less ended, although Cornell held independent meetings with Ithaca College and the New York State College of Forestry in Syracuse to explore their interests in a cooperative agreement. Neither of those succeeded.

NBER and Tymnet National Networking

Continuing with the renewed sense of serving off-campus customers to sustain the increased expenses at OCS, discussions were initiated with the National Bureau of Economic Research (NBER), which was seeking additional computing capacity to support its activities. In addition to the pure monetary advantage, NBER provided additional advantages to Cornell in that it made available its TROLL system for economic modeling, specialized statistical software, and some 30 years of economic data to the campus at no extra charge.

To do this effectively while not significantly affecting Cornell work, it was recommended that OCS install VM (Virtual Machine) as the highest level of the operating system on the 370/168. Other systems such as HASP for batch would then operate under VM. An additional and significant benefit would accrue as a result of this change, owing to the installation of CMS (originally Cambridge Monitor System but later Conversational Monitor System), a functionally lean but efficient time-sharing system. After appropriate review by the UCB and other technical discussions with users, VM was installed in November 1974. A further advantage, which came to be significant some time later, was that the late weekday evening systems test time (and as needed the early morning test times) were converted to virtual test times. This took advantage of the ability to bring up different virtual machines under VM, and so a test VM system could be brought up by a programmer to write, debug, and test the operating system changes while other operations continued. While this late evening procedure could affect performance without care being exercised, it was preferable to shutting down the system completely.

Coincident with accepting the NBER work, OCS installed Tymnet, a commercially available networking service that allowed NBER and its associates and customers to access the Cornell computer by a local telephone call to a Tymnet node. In 1974 Tymnet had nodes in some 60 cities in the United States and Europe. A nice feature of this arrangement was that

Tymnet would be responsible for establishing the connection and transporting the information at acceptable speeds and would levy time-use charges to the calling customers.

PBM/STIRS, EFPM, and SENTRY (CBORD) Network-Based Initiatives

OCS succeeded in various other ways to piggyback on the *Tymnet* connection by offering computing services to users outside of Cornell. The most successful was PBM/STIRS developed by Prof. Fred McLafferty and his associates and grad students in Chemistry, which was installed on the 370/168 in May 1975. The PBM (Probability Based Matching)/STIRS (Self-Training Interpretative and Retrieval System) was designed for analyzing mass spectrometry data. PBM compared the mass spectrum submitted by the user with those stored in the Cornell database. If the "fit" of a submitted spectrum was not close enough, the STIRS program could be used. STIRS provided the identification of some or all of the unknown spectra components, which could assist in the identification of the spectra. Charges were \$10 for a PBM run and \$15 for one STIRS use plus the *Tymnet* charges. Part of the charge had to be refunded to John Wiley and Sons for use of their Registry of Mass Spectral Data. When developed initially, there were 40,000 compounds (30,000 unique); after being "ported" to the 370/168, this grew to 80,000 compounds (70,000 unique). This program ran on OCS facilities for the remainder of the decade.^f

Another successful piggyback on the *Tymnet* connection was the EFPM project initiated by Educom. EFPM, the Educom Financial Planning Model, was a model and modeling program for university budget planning. Dan Updegrave,^g who earlier had been a student at Cornell but was now with Educom, developed the model. The availability of networking services served both to help Updegrave develop and maintain EFPM and to allow Educom to make it accessible to the higher education community. Because Educom favored the use of Telenet, a competing national networking service to *Tymnet*, a connection was also made to Telenet at this time. Within a short while after EFPM was made available, over 75 customers scattered from Belgium to Australia, but largely in the United States, were using the model for planning purposes.^h

Yet another such application accessible on the 370/168 using *Tymnet* was the SENTRY system developed in 1976 by Cornell Dining for menu planning, cost estimating, and recipe development. Art Jaeger, then director of University Dining Services, and John Alexander, a graduate student in the School of Business and Public Administration (now the Johnson

Graduate School of Management), were responsible for developing this system specifically tailored to the Cornell environment. Jaeger is quoted as saying, "We used to try doing this by hand. But what used to take 20 hours per dining manager each week can now be accomplished in about 20 minutes using the computer." Within a short while Jaeger and Alexander made an agreement with Cornell to secure the rights to the software in exchange for providing no-cost support to Cornell, and formed the CBORD corporation to market their services to other food preparation agencies. Initially CBORD customers used the IBM 370/168 at Cornell, but as the customer base grew and other computing and networking options became available, CBORD took on this responsibility.

New Leadership for Academic Computing

To improve its offerings to the instructional and research constituencies on campus, OCS initiated several efforts to improve services to the academic users, including changing the leadership of the staff on campus. In 1972, the terminal operations staff was separated from the Langmuir operations staff and became part of the academic support group on campus. It was expected that this would provide a fuller range of services and assistance to academic users, most of whom came to the campus terminals to do their computing work. A year later, the campus User Services professional staff was all located in Uris Hall ("Old Rusty"). This was an ideal central location and the first time for many years all such staff were in a single location. Before that, the User Services staff was located in Upson and Clark and at Langmuir. John Aikin was manager of the whole campus staff for several years as part of the larger Technical Services group under Dick Cogger. Technical Services combined the Systems Programming Group at Langmuir and User Services on campus to improve communications and focus on academic users' requirements.

In late 1973 Rudan asked the FCR (Faculty Council of Representatives) Committee on Computing, originally appointed to consider the issue of the 370/168 computer, to consider how to improve academic computing at Cornell. The committee, chaired by Chester, issued a report recommending the creation of an Academic Computing unit independent of OCS with the director reporting directly to the provost.²⁸ The UCB did not accept this proposal, as OCS had already decided to hire John Williams, who had been

an assistant professor in Computer Science, as assistant director for academic computingⁱ in 1974. His principal responsibilities were to work with the college deans and with department heads and individual faculty to determine their aggregate or individual needs for computer services and then work with OCS staff to see how these needs could be met.

The combination of new leadership and location was expected to produce improved relations and services in support of instruction and research use of the 370/168 computer. When Williams left Cornell in 1976, Douglas E. Van Houweling was appointed assistant director of academic computing. Van Houweling, as a professor in the government department, had been an active user of computing in the classroom and a member of the University Computing Board for several years. He was well known on the campus and in the Ithaca area for promoting the use of "Metro-Apex," a computer simulation of the operations of a city. Van Houweling actively promoted new initiatives in expanding the use of time-sharing in instructional computing and for support of non-mainframe computing. In these last issues he was supported by Fred Hiltz, who continued to be a consultant to the campus on the acquisition and use of minicomputers, and later by Alison Brown who took on this responsibility. Toward the end of the decade Brown began to focus her attention on newly available microcomputers (now called personal computers, or PCs).

OCS Seminars

Throughout the decade, but accelerating with the installation of the 370/168 and as new software and services became available, OCS held seminars and short courses to make known its offerings and to train Cornell staff and students in the use of its equipment and applications. Typically these seminars were held at the beginning of semesters or at major events such as the acquisition of the 370/168. In 1973 and 1974 the topics included the use of open access computing (Free IT) and the newly installed 370/168. This was followed by the introduction of CMS and interactive computing; the use of Fortran, PL/1, APL, and SAS (for statistical computing); and the use of SCMS, a simplified and controlled environment student version of CMS. By the end of the decade the seminars were about Terak microcomputers and Pascal. It is interesting to note that as technology changed during the 1970s, only Fortran continued to be offered throughout the decade. Other topics came and went as technologies and interests changed. The OCS staff responsible for organizing and promoting these seminars included Michael Steinberg, Mariann Carpenter, James F. (Jim) Manning, Joan Winters, Chuck Boeheim, and Alison Brown.

²⁸ FCR Ad Hoc Committee on Computing, "A Proposal to Form a New Organizational Unit: Academic Computing Services," December 1973.

Computer Science

Computer Science continued to increase its course offerings for instructional computing. Early in the decade it took on the responsibility for Engineering 105, Introduction to Computer Programming for Freshman Engineers. To improve the presentation of courses when APL was taught in 1973–74, a video camera was used to project what was happening on terminal on a large screen in the classroom. The APL section of the course was oversubscribed one day after it was announced.²⁹

To create more options for learning different computer languages, the department made arrangements for OCS to offer one-credit sessions on Fortran and APL. The plan was for CS 100 to be a two-credit course on elementary programming, which would be followed by the OCS part on a particular language. This was a nice fit with the offerings in the OCS Seminar Series that were extended for this purpose. Following the success of the video experiment mentioned above, the department started to make videotapes of the large elementary programming courses. These tapes were to be held in the Cornell Library for self-study by students or others. By 1976–77 videotaping of courses had expanded to six undergraduate programming and systems courses, and the videotapes were being used by outside industrial organizations.

That same year the “Programming for Poets” course was taught by Charles van Loan and attracted 60 students for credit. The course was intended as a “soft” introduction to this increasingly important technology. Despite these innovations and the continued reliance on batch processing for such course work, the department was also noting the lack of good interactive computing facilities for its teaching and research needs.

Business Systems

The Financial Accounting System Moves to the 370/168

During the discussions about the acquisition of the 370/168, one opportunity for increasing income to support the added costs was moving the processing of the university accounting system and other related work from the 360/25 in the Statutory Finance and Business Office to the 370/168. Serious negotiations got under way in 1974. Several test trials and processing of job streams on the 370/168 showed that moving the work would provide adequate service at a cost not greater than the current cost of the 360/25.

After addressing concerns about service levels and cost escalators, the work was moved to the 370/168 in the fall of 1974, and the 360/25 was returned to IBM, providing an additional income of \$120,000 per year to OCS. This move proved to be correct in the long term. Having all business systems resident on the 370/168 not only improved the ability to transfer information between such systems, a step toward integrating systems, but also led in the direction of using more contemporary technology for existing or proposed business systems.

McNeil Task Force and the Demise of MSA

In 1975 President Corson announced another round of cost reductions to save a total of \$2.3 million over the next three years. The process for dealing with these cost reductions was to appoint a large number of task forces to examine different facets of the university’s operations. Ian R. McNeil, professor in the Law School, was appointed leader of this effort, which thus became known as the McNeil Study. One of the task forces, Task Force 16, was directed to examine administrative data systems and administrative computing. This examination of business systems was motivated by the earlier failure to develop a new student information system and continued the pattern of having faculty-led committees and task forces examine the state of computing on campus and make recommendations for changes.

In September 1975, McNeil Task Force 16, Administrative Data Systems and Administrative Computing, issued its final report. The summary recommendation reads:

Administrative data systems and computing must be handled with care and consideration of the public which such systems are designed to serve. If the servant becomes the master, a mounting bureaucracy tends to obscure human relationships and is detrimental to the learning environment. A reduction in MSA (Management Systems and Analysis) and SIS (Student Information System) would not be detrimental to but rather an improvement on our present situation.

The more detailed recommendations (with minor editing) were as follows:

1. We recommend that no new large administrative computing systems nor large modifications to existing systems be undertaken. By large we mean a system that costs \$10,000 installed or \$10,000 a year to operate.

²⁹ Department of Computer Science, Annual Report, 1973–74.

2. We recommend that MSA staff be substantially reduced. . . something like a 50 percent reduction might be possible from the 1974–75 staff.
3. We recommend that the planned lifetime of the 370/168 computer be extended by three to five years.
4. There should be a change in attitude toward administrative computing. . . They (systems) should only be instituted after the need for them has been clearly established by the final user; they should be kept as simple as possible; costs of running them should be regarded as real costs; changes in programs should be infrequent.
5. In principle, administrative users should write computer programs themselves, in the sense that researchers should write computer programs themselves.
6. Before deciding on a future student information system, the possibility should be investigated that a less expensive or complex alternative, such as a minicomputer or a simpler, already existing 370/168–based package, might satisfy Cornell's needs.

Within a short while, the MSA staff was reduced by 40 percent, from 30 to 18, and administrative computing entered a maintenance mode for existing systems. There was a minor exception in mid-1976 when the Board of Trustees approved an expenditure of \$60,287 to rehabilitate the outdated 15-year-old student records system. The funds were those remaining from the original 1972 SIS plan. The work was done not to achieve the functional improvements of this earlier plan but to improve the reliability and efficiency of the current systems at minimal expense. Two important elements of the work were to eliminate both the processing of 1401 programs under the 1401 simulator on the 370/168 and the punch card processing still being done on the 360/20 in Day Hall. When completed it would be possible to remove the 360/20 being used for these card-processing functions and save \$50,000 per year in rental costs. The project was going to be accomplished through reassignment of existing personnel and incur only minor incremental expenses.

By early 1977, MSA ceased to exist. The administrative programming staff was reorganized as Administrative Programming Services (APS) and became a division of OCS under the leadership of Edmund V. Hollenbeck as assistant director. The Institutional Studies staff of MSA was transferred to the new institutional planning organization under Lawrence.

Campus Transitions—1975 to 1977

Significant changes were made in staff at the very highest levels of the university, including the president and provost, and the addition of a new top-level position, senior vice president responsible for administrative operations. The responsibility for OCS remained with Lawrence, but a new director was appointed for OCS.

W. G. Herbster Appointed Senior Vice President

In 1976 there was a major reorganization of the executive staff of the university. William G. Herbster was appointed senior vice president and was to be one of the three top-ranking executives along with the president and the provost. The provost would continue as the chief academic officer, while the senior VP would be responsible for overall support services. There were to be nine VPs reporting to Herbster as new vice presidential positions were created to broaden the responsibilities of previous directors. As part of these changes, Lawrence became vice president for financial and planning services but continued as chairman of the UCB and responsible for OCS operations.

A. H. Peterson's Laws on Computing and Information Systems

Also in 1976 Arthur H. Peterson retired as university controller and treasurer and his many responsibilities were assigned to other VPs. In paying tribute to Peterson on his departure, the UCB formally stated Peterson's Laws on Computing and Information Systems, gems of wisdom that he postulated back in 1962:

It's going to cost more than you say.

It's not going to work nearly as well as you say.

Having seen these, Bob Blackmun believes there was a third law that said:

It's going to take longer than you say.

It should be clear that there was ample historical proof of the validity of all these statements and, as will be seen, additional proof that these laws would continue to be relevant in the future.

Shortly after Peterson's retirement, John S. Ostrom was appointed university controller.

New President of Cornell—Frank H. T. Rhodes and Other Executive Changes

In 1977, President Corson resigned his position and was replaced by Frank H. T. Rhodes. David C. Knapp left the provost position to assume the presidency of the University of Massachusetts.

W. Keith Kennedy became provost in June 1978 and started to build his management team. In September,

Don Randel and Alison Casarett were appointed part-time vice provosts, and Jim Spencer was later appointed vice provost for planning and budgeting.

In January 1979 Thomas E. Everhart was appointed dean of engineering. All these changes had significant bearing on the developments in computing at the end of the 1970s.

New Management Team at OCS—Robert R. Blackmun, Director

Bob Blackmun was appointed acting director of OCS in 1976 when Rudan joined Institutional Planning and Analysis. Taking advantage of this interim period, Lawrence and the UCB invited several consultants to review the computing situation at the university, and Lawrence started the process to recruit a permanent director. One of the visitors was Bruce Arden, now at Princeton, who had come on a similar visit in 1966 before the formation of OCS. The other visitor was James Emery, president of Educom. Both Arden and Emery expressed their concerns about the future of the 370/168 computer, in view of IBM's announcements of the newer systems that offered the same computing cycles at lower cost. Because the remaining amortization of the 370/168 was the same as the cost of the newer systems, there was no clear advantage to change. However, they advised that the UCB and the university administration be cognizant of the increasing support costs for staff, software, etc., and the decreasing costs and packaging of computing cycles in planning for the future.

In late 1977, after interviewing several external and internal candidates, Lawrence appointed Blackmun as director of OCS and at the same time appointed Van Houweling as associate director, computer services, and director of academic computing. He articulated that this two-man team was going to be better than any single individual who could be recruited. When these appointments were made the following staff rounded out the management team: Cogger as assistant director for system programming, Pulleyn as assistant director of operations, and Hollenbeck as assistant director of administrative programming services.

A Malicious Computer Incident

In early 1977, the 370/168 at OCS experienced a series of random and unexplainable outages that continued for over two months. Without warning the computer would simply shut down. Despite the best efforts of IBM to track down the problem and the almost complete replacement of sections of the computer, the problem could not be overcome. Over the two months approximately 100 hours of computer time were lost as downtime went from the normal 2–3

percent of the schedule to 6–10 percent of the schedule. The problem was finally resolved when it was traced to tampering by a computer operator.^j Again the likelihood of lost or stolen records and potential security violations was raised even though such issues were inconsequential in this particular case. It was simply a case of months of dealing with unpredictable downtimes and the serious inconvenience this caused everybody.

It did not help that these incidents were followed by another sequence of computer outages caused by small pin leaks in the water cooling system of the 370/168. For that six-month period, it seemed that the 370/168 was jinxed, and water cooling of computers took some ridicule from those who were using air-cooled systems from other vendors. Blackmun was commended for the way he handled these two incidents.

Institutional Planning Services

In 1977, Lawrence organized a new central planning group. He recruited three senior-level staff—Ian Stewart, assistant professor on leave from City and Regional Planning, Jack Lowe from the Sponsored Research Office, and Rudan^k on leave from OCS—and with the staff from Institutional Studies formed Institutional Planning and Analysis (IPA). The other new staff member was Simeon Slovacek as a research analyst. Rudan was nominal director of this new unit. The charge for this group is best stated in a quote from Provost Knapp:³⁰

To improve the information and analysis necessary to make decisions about Cornell's future we need to improve the process by which problems are anticipated and defined, alternatives analyzed, likely consequences weighed, and consensus developed around the proper courses of action.

Henry Vaughan, who had been heading up the Institutional Studies group, left Cornell and opened Computerland, the first computer store in Ithaca and one of several such stores in New York State.

Although the IPA group did not have anything to do directly with administrative or academic computing, for the record it produced two successful long-term projects during this period. One was the creation and production of the first *University Fact Book* under the leadership of Ian Stewart. This book was crammed full of historical multiyear statistics and derived information to create a consistent set of Cornell facts to

³⁰ "Planning Services Group Staffed," *Cornell Chronicle*, September 22, 1977.

be used for reference and to assist the administration in planning activities. Since that time it has been more or less kept up to date. There was a time in the 1980s when such information was considered to be more confidential than in the 1970s, but the *Fact Book* has matured into a web page that makes most of such information widely available. The other project headed by Rudan was to compile the first-ever university-wide capital budget to improve the short- and long-term planning for financing such projects, including the solicitation of support from alumni, friends, and foundations.

Instructional Computing

During the mid-1970s, the faculty at Cornell continued to lobby for more and improved time-sharing services as this mode of computing was becoming more popular and available on other computing systems.

The Push for More Time-Sharing Services

Responding to the continued need to improve interactive computing services, in 1976, ACIC (the Advisory Committee on Instructional Computing) again made a proposal to the University Computing Board for increased use, access, and support. ACIC, then under the chairmanship of Van Houweling, made four recommendations:³¹

1. A new interactive service be developed consisting of a carefully limited subset of facilities.
2. This service be priced to recover incremental costs involved in providing it, but well below the full average cost of all services.
3. Two additional staff members be hired for the purpose of supporting the users of this facility.
4. A new cluster of interactive terminals be installed.

The total cost for the program was estimated at \$111,000. The UCB endorsed the report and forwarded it to the president and provost strongly supporting the implementation of the program but changing the funds requested. The UCB requested \$50,000 for three years to cover the costs of new staff, recommended funding computer use from the current allocations, and recommended that the college deans be

asked to cover the costs of new terminals and connecting hardware. The program got under way in January 1977. In anticipation of this new program, but also recognizing the increasing use of the 370/168, earlier in 1976 OCS installed VMA (Virtual Memory Assist). This hardware feature increased the computing capacity by 10 to 15 percent.

Subset Conversational Monitor System (SCMS) Developed

OCS developed a new interactive operating system called the Subset Conversational Monitor System (SCMS) with a design criterion that it be easy to learn and easy to use. When put into use, it provided for access to PL/C, WATFIV, APL, SPITBOL, the MINITAB statistical package, and the ability to submit small batch jobs to the IT system. It was so configured that instructors could develop course materials using standard CMS and then provide these to students on a read-only basis in SCMS. It was further priced at a rate substantially below those for standard interactive services with an attempt to make the cost of a SCMS session roughly comparable to an IT job. Two new staff were hired and three new clusters of terminals were located in ILR (four terminals), Agriculture (three terminals), and Engineering (five terminals). These terminals augmented the eight OCS terminals in Uris Hall and six terminals in the Business School.

Experience with First Use of SCMS

Experience in the spring semester of 1977 was very encouraging. A total of 1,150 accounts were distributed: 600 accounts (\$50 each) through OCS and 550 accounts set up by a faculty member for a specific course. In addition, another 220 standard accounts used SCMS for their work. Faculty reactions were very positive toward this new service, both from an evaluation questionnaire and solicited comments, with reservations expressed that further investment to develop course materials would depend on the long-term continued availability of the service and the funds for its use.

OCS prepared an evaluation of the semester's experience along with several recommendations. The recommendations were to continue SCMS as the primary vehicle for interactive instructional computing at Cornell, to continue the same lower rates, to fund the support for the installation of an additional 15 terminals, and to continue the two-channel distribution of SCMS accounts—through OCS and through departments. While the issue of additional terminals remained unresolved, the other recommendations were put into place after review by the UCB.

³¹ Advisory Committee on Instructional Computing, "A Report on Interactive Computing for Instruction to the University Computing Board," May 4, 1976.

PLATO Use Continued

PLATO continued to be of particular interest to the Music Department but also to biologists for the teaching of biology. A proposal was advanced to increase the number of terminals to the hardware limit of four, but this failed to gather support. As a result, the PLATO initiative continued with two terminals and more or less ran by itself for the rest of the decade.

Research Computing

Several new initiatives to improve computing and data storage and retention services were started in 1977.

Acquisition of the Array Processor

On the research computing front, Alec Grimison, recruited in 1976 to serve as OCS liaison with research users, was able to form a consortium with researchers from Chemistry and Physics and several of the grant-funded research centers and laboratories to acquire a Floating Point Systems (FPS) 190L Array Processor (AP).^l Before this decision discussions were held within the research community and the UCB as to whether the AP should be connected to a separate minicomputer or the 370/168. On the basis of the capabilities of VM, large and easily expandable disk and tape storage facilities, network connectivity, and the ability for many users to access the system—along with strong support from OCS—the decision was made to connect this processor to the 370/168. The belief at the time was that this was just the first of many such specialized minicomputers to be placed on campus.

It was also the case that installing the AP was the first cooperative project of OCS with a campus constituency. After discussions with the UCB, OCS agreed to pay for one-third of the \$114,000 acquisition cost, allowing it to sell this share of its use to non-consortium members as a way of recovering costs. By paying a portion of the annual costs, consortium members were guaranteed a low rate for use of the AP in addition to paying for 370/168 charges at prevailing rates. Other users paid a higher rate for the AP in addition to the usual 370/168 charges.

The FPS unit executed floating-point arithmetic instructions very efficiently for vector calculations. In effect, it executed the equivalent of a Fortran “DO” statement, to say add two vectors, in about three add instruction times instead of “n” add instruction times where “n” was the length of the array. The Cornell box was serial number 2 and the first to be connected to an IBM system. The array processor provided a significant improvement in execution time for complex matrix operations and was rated at three times

the speed of the 370/168 and almost as fast as a CDC 7600 system. Installation of the AP was the first step in bringing “supercomputing” to the campus, and Ken Wilson, the future Nobel laureate, played a key role in bringing about this improvement for larger-scale research computing on campus. Other original consortium members were Harold Scheraga (Chemistry), Keith Gubbins (Chemical Engineering), and Geoffrey Chester (Nuclear Studies).

The FPS 190-L array processor was installed in March 1978 after delays resulting from problems with the interface hardware. In the intervening period between the placement of the order in 1977 and the installation, OCS and consortium staff were developing Fortran and Fortran linkages to make more effective and early use of the AP using a simulator provided by FPS. To show their interest, Floating Point Systems contracted to pay Cornell staff to write a special Fortran compiler (APTRAN) for use with the AP. The relationship with FPS was fruitful for both parties. FPS was able to penetrate the IBM scientific market, and Cornell and OCS were now in a position to provide cost-effective, large-scale computing to Cornell and other interested users.^m

Mass Storage System

At about the same time that the AP was installed, the T. J. Watson Laboratory of IBM invited OCS to participate in a joint study to evaluate the use of mass storage in a VM environment. This project fit nicely with the AP project as it was likely that once the AP started generating computations there would be a need to store massive amounts of data for future retrieval. At the same time, other research at Cornell and elsewhere was generating a large amount of data that were being stored on magnetic tapes because of cost considerations, and that later required much more computer time and processing to retrieve. An IBM 3850 MSS (Mass Storage System) capable of storing 35.5 billion bytes (gigabytes) was installed and put into use for a short time before the effort was abandoned as not sufficiently productive to continue.

New Initiatives—1978 to 1979

Given the changing technologies as microcomputers started to become a discussion item and an emerging technology for those on the leading edge of development and computer use tilted toward information processing instead of computation per se, OCS responded with several new initiatives.

Microcomputers and OCS Support Services

Recognizing the future importance of the microcomputer in research and instruction, Van Houweling promoted the formation of a new unit in OCS to focus on this technology. In the summer of 1977, for example, OCS held a three-day workshop on microcomputers. The workshop covered the technology and projects for using the 8080 chip-based computers in the classroom and in research applications. About this time, the Small Systems unit was started in the User Services group in Academic Computing. The initial staff who had been assigned or more or less drifted to this technology were Alison Brown, Dan Bartholomew, and Gary Buhrmaster.

Their concern was about not only operating system and applications software for the new microcomputers but also the small PDP-11 systems being offered by DEC. Pascal and BASIC compilers were being offered for all these systems, which included the TERA (using the LSI-11 chip from DEC), and for microcomputers built with 8080 or Z-80 chips and running CP/M or Cromemco CDOS, the Apple II as well as for TI and General Automation machines. The list is indicative of the vendors coming into the market toward the end of the decade.

Text Processing with SCRIPT and the Start of Electronic Mail

In about the mid-1970s, greater interest began to develop in new aspects of computing. First there was the growing number of minicomputers and the very new microcomputers. This interest was seen in not only the increasing number of such systems coming on the market but also the steadily declining price/performance that favored these systems for certain applications. Different forms of networking became of interest, and forward-looking pundits were predicting a totally interconnected world where all systems could communicate with each other and users could seamlessly communicate with them and one another.

A growing development was the increased use of computers for information and text processing and not just for number crunching as in the past. Early in the decade, OCS had promoted the use of SCRIPT, a text-processing system that ran on the 360/65 and the 370/168. When first introduced, SCRIPT could print only upper-case text until printers were able to print both upper- and lower-case letters. Cecilia Cowles once commented about the finger contortions she had to use when trying to input lower-case character equivalents on the keypunch by striking several keys at the same time! As the shift started from punched card input to the use of typewriter or CRT terminals, which easily permitted lower- and upper-case alpha-

bet characters, there was an added convenience and incentive to use such systems. By mid-decade almost all OCS publications were composed in SCRIPT, which made it simpler and more efficient to produce copies and manage updates and make the publications retrievable by users from online files.

One of the most interesting developments in text processing was the first electronic mail system using the 370/168 computer. Steve Worona³² designed and wrote what he refers to as “Mail 1” in CMS.¹¹ Basically, this system transferred files between users’ accounts. These were text files that conveyed a memo or a letter using the computer’s facilities to transfer bits instead of the surface movement of paper. In conjunction with development of Mail 1, Worona produced “Talk-to,” a system of interactive communications in which one or more users in real time could send messages to each other.

Network Services; Terminal Leasing Enterprise

Also being explored were new networking options beyond point-to-point communications, i.e., from the terminal to the computer using modems at the terminal end and a communications controller at the computer end in use at the time. Improvements had been made raising the transfer rate from 300 baud to 1200 baud, half duplex to full duplex (being able to transmit and receive at the same time). Other options were being explored, particularly the emerging Ethernet technology. Campus meetings were held to inform the community of developments and solicit their input and needs.

To promote the increasing use of time-shared computing, OCS organized a “Terminal Enterprise” to lease Datamedia 1521 crts and model 087/089 Decwriters at low rates. For some years, IBM model 2741 and Datamedia 1521 terminals had been placed at different campus locations, including department offices, as the primary means to access the 370/168 to input programs and data, and to initiate the execution of programs. The new lower rates were \$45 per month for a one-year lease, a considerable reduction from the previous rate of \$85 per month, as a result of volume purchasing.

OCS and the Cornell University Medical College

Initial discussions to form a relationship between OCS and the Cornell University Medical College (CUMC) in New York City started in 1975 and involved Jim Peters, associate dean for business affairs, John

³² Steve L. Worona, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, April 1999.

Daniels, director of computing at CUMC, and Rudan and Aikin from OCS. At that time CUMC had to decide on continuing its relationship with SECOS (Shared Education Computer Systems), its service provider, or find another alternative for computing services. SECOS was experiencing financial problems and was negotiating for CUMC to take over the lease for its 370/145 computer until the lease expired in two more years. The 370/145 computer leased by SECOS was providing academic and administrative computing services for CUMC.

The way CUMC arrived at its situation is relevant to the steps that led to CUMC's relationship with OCS. SECOS³³ was started in 1971 by Stephen W. Dunwell to provide APL time-sharing services to a broad spectrum of education and not-for-profit institutions in the area around Poughkeepsie, N.Y. Dunwell was an IBM fellow and received initial start-up support from that company. CUMC was one of SECOS's early customers because several of the medical staff at the college were interested in using APL for their academic work. When IBM withdrew its support from SECOS, creating the financial problem, it was then that they asked if CUMC would be willing to take over the computer lease. As one alternative CUMC contacted OCS about providing services. OCS made a formal proposal to provide computing services to CUMC in late 1975, but this was rejected on the basis of cost and possible problems with providing the same level of service. The SECOS 370/145 was a souped-up system with a special feature designed specifically to speed up APL processing, and the 370/168 configuration at the time was considered to be noncompetitive for this work.

CUMC decided to continue its relationship with SECOS, took over the hardware leases, moved the hardware to its buildings in New York City, and became the operations manager for the computer while SECOS became the marketing agent to sell APL services to both educational and other agencies, including commercial companies. In 1976 CUMC was experiencing difficulties with its administrative computing systems and operations and made inquiries as to whether the Ithaca campus could be helpful. One of the concerns was that several doctors and their staff were writing and supporting APL programs for selected CUMC business systems. OCS also became more interested in developing a relationship with CUMC in 1976 when it seemed that NBER was going to end its relationship with OCS and move its computing activities to MIT. There would be spare cycles on the 370/168 to sell to support the OCS operations at their

then-current level. This request from CUMC also fit nicely with the opportunity to start a path toward better integration of administrative systems between the Ithaca campus and CUMC. There was also some desire for CUMC to take better advantage of the computing expertise and services available on the Ithaca campus that could be used by academic staff at the medical college.

Despite the interests of all parties, the relationship between CUMC and SECOS and OCS was somewhat in flux during the interim 1976-77 period. For example, CUMC held talks with the New York Hospital about cooperating on administrative systems and examining the installation of purchased systems for selected applications. These talks failed. So, despite being next door to the hospital, CUMC was forced to run its own independent business systems and provide its own academic computing services. Also by this time, CUMC was experiencing maintenance problems with the special microcode of the 370/145 and faced the prospect of IBM withdrawing support for this feature. The special APL used at CUMC was becoming dated and not compatible with other such systems so that it was difficult to export programs to other sites. Converting to another in-house system was going to be a long process, requiring another computer and additional support staff. With a hard deadline of December 1977 when the 370/145 lease expired, a decision about future computing services had to be made before then. SECOS submitted a proposal for continuation of services and OCS developed its own project plan by organizing a joint CUMC/OCS task force to consider how OCS could service the CUMC load.

After extensive explorations during the summer of 1977, the task force concluded that a merger of the two operations at OCS and CUMC was not only feasible but beneficial to both parties. A key part of the proposed agreement was for OCS to hire several of the CUMC computing staff as OCS employees on site at CUMC to improve user support, which had not been always available and helpful. Also as part of this agreement, OCS and SECOS would make an independent agreement for SECOS to act as the marketing and billing agent for its customers, who would also be serviced by the 370/168 at OCS.

In early 1978, the high-speed memory of the 370/168 was expanded in anticipation of the additional CUMC load. The transfer of work from CUMC started early in 1978 with program testing and benchmarking activities in anticipation of a coordinated transfer of activity from the 370/145 to the 370/168 later that year. For reasons that are not fully understood almost 25 years later, the newly hired director of computing at CUMC, Susan Schwimmer, decided to transfer the

³³ R. R. Blackmun and John R. Daniels, "Draft Proposal OCS and CUMC," 1977.

full load to the 370/168 in July and at the same time removed the 370/145 from the CUMC computing room. This full load created a tremendous overload on the 370/168 and created unacceptable conditions for all users. The Ithaca campus and CUMC and the SECOS customers were not well served. OCS reacted by taking steps to better manage the 370/168 load and make technical corrections, but also aggressively moved to place the 370/145 back into service and transfer the CUMC APL load back to this system. OCS stationed George Cameron as on-site manager at CUMC and rotated different staff on temporary assignments as the situation demanded. In the midst of this turmoil, CUMC decided that it would alone take the responsibility for its future computing services and retained the consulting firm of Haskins and Sells to support some of its activities and give advice on future actions.

There was much angst on all sides about this failure, but in the end it seemed to come down to insufficient staff and unwise decisions made on the part of CUMC. By January 1978, Peters and Daniels, who had been involved with all the negotiations, had left CUMC and there was no director of computing there during most of the year. However, in the interval between Daniels' leaving and the hiring of a new director, the CUMC staff installed the payroll package from Information Associates (IA)^o to replace the general accounting system. After Susan Schwimmer was hired, for one reason or another almost all the staff left CUMC, taking with them much undocumented technical expertise.

Shortly after the unsuccessful transfer of work, Schwimmer was dismissed and CUMC was supported by Cameron and other OCS staff and by staff from Haskins and Sells. A new director, Phil Ackerman, was hired in the fall of 1978 and took responsibility for CUMC operations after that time. OCS continued to process the SECOS load and the CUMC administrative batch computing on the 370/168 throughout the 1978 fall semester and into 1979 while other arrangements were being pursued. The conditions for all users of the 370/168 were only marginally acceptable and OCS was subjected to criticism for "creating" these conditions. SECOS had made other arrangements to obtain computing services and were no longer using OCS facilities by the end of 1979. CUMC continued to run its administrative systems that relied on IBM mainframe services, payroll, and accounting for several more years.

The failure of this joint effort, which resulted in a \$1.4 million cost overrun in the CUMC budget for 1978-79, was publicly attributed to several factors. According to a press release from Sr. VP Herbster,³⁴ the reasons were the complexity and inadequate docu-

mentation of the CUMC programs, the lack of trained computer personnel at CUMC, and an unrealistic estimate of the time needed to run CUMC's programs on the 370/168. An important component of this cost overrun was the significant loss of revenue from outside users of the CUMC computer who could not be served properly during the attempted shift to the 370/168. Some OCS staff who were involved with this project to this day take exception to the statement of "an unrealistic estimate of the time needed to run CUMC's programs on the 370/168." What is not arguable is that the combination of unpaid use of computer resources and additional expenditures for consultants and added hardware to cope with the crisis produced the large discrepancy between expense and income at CUMC.^P

New Plans for Campus Computing Services; DEC20 Favored along with Use of Terak Microcomputers

The experience with CUMC and the effect of poor service on the campus during late 1978 and 1979 forced a consideration of the short- and long-term future of computing at Cornell. The UCB members were concerned about users' increasing call for improved and different services as well as by the rapidly changing technology. The discussions were a continuation of the process started when Arden and Emery came in 1976-77.

During this period there were many crosscurrents and opportunities. For example, IBM approached OCS to participate in a joint study to build an SCMS-like time-sharing system on its newly announced 4341 series of 370-compatible but low-cost entry systems. OCS was very much interested in leveraging its skills and success to create what could become a high-performing, entry-level time-sharing system product. However, this request was turned down to focus on other more pressing issues. The positions that the UCB took and the items they pursued were the following:

- There would be a need to continue IBM 370-compatible processing past 1980.
- Improvements in time-sharing services should come from other vendors, in particular, from DEC.
- Instructional computing, particularly the teaching of introductory programming, might best be done using individual microcomputers per student.

³⁴ "Computer Overrun Costs \$1.4 Million," *Cornell Chronicle*, October 18, 1979.

- The current scheme of providing direct funding to OCS, the so-called “allocation funds,” would be phased out and shifted to college/department budgets where the funds would be theirs to use as they desired. Exceptions would be made for SCMS and IT (Instant Turnaround for small batch jobs) services, which would continue to be centrally funded.

These positions contributed to several actions and developments. In view of the need to bring the OCS income and expense into better balance and to continue 370 capacity, OCS started negotiations to enter into a lease/purchase contract to reduce the annual costs of the 370/168 computing system. To try to contain the deficit, a major cost reduction was achieved by negotiating a lease/purchase contract for the 370/168 and other selected components with CIS^a systems of Syracuse, N.Y. Such lease/purchase contracts were reasonably new during this period, but they involved CIS purchasing the system and then being guaranteed a multiyear lease contract for the same system. At the end of the contract period, CIS would own the equipment and would be betting that its market value was greater than its depreciated cost and they would make a profit as a result. It is also important to note that extending the 370/168 lease for three additional years was in keeping with the recommendations of Task Force 16 of the McNeil Study to extend the planned lifetime of the 370/168 for three to five years.

With regard to improving time-sharing services, there was considerable discussion on whether this should be done on IBM or DEC platforms. Given the mixed success using IBM hardware and software, the UCB was favoring DEC. In this case, however, it was not clear whether the specific platform should be the newly announced VAX systems or the older DEC 10 and 20 series. Consensus eventually favored the DEC 20 computer.

With regard to individual microcomputers for teaching introductory computer programming, the UCB was very interested in the innovative and interesting development of the Cornell Program Synthesizer by Tim Teitelbaum, professor in Computer Science. Using an LSI-11 chip, Teitelbaum developed this self-contained system for writing structured programs from a video display tube using a subset of the PL/CS processor running in a “syntax-directed synthesizer.” As it turned out, the Terak company was building microcomputers that contained the LSI-11 chip, and so several of these computers were purchased for use as development and test machines as a joint OCS and Computer Science project.

Instructional Computing

Along with promoting the use of more time-sharing services for instructional use and for casual student use, experimentation continued with alternative technologies. In 1979 seven small IBM 5100 computers were placed in Uris Hall and charged at a rate of \$3 per hour. These stand-alone computers used APL and Basic and were considered powerful calculators. The experiment only lasted a short while owing to lack of interest.

Another experiment in 1979 was Computer Science’s new course, CS 103, Introduction to Pascal, which was taught by Hal Perkins and used a Pascal compiler brought from Australia and running on the 370/168.

One indicator of the growth of instructional computing can be seen from the increase in the number of courses in the Department of Computer Science over the decade (Table 3). The growth factors ranged from a 45 percent increase in the number of courses and a 104 percent increase in the number of students enrolled. Because the credit hours for CS 100, one of the larger courses, changed from 2 to 3, the combination of the increases in courses and enrollment resulted in an 80 percent increase in credit hours.

Table 3. Department of Computer Science Course Statistics 1970–71 and 1979–80

	1970–71	1979–80
Number of Courses Offered	38	55
Number of Students	1,589	3,236
Number of Credit Hours	5,117	9,178

Source: Department of Computer Science annual report

The Cornell Program Synthesizer Using Teraks

The most innovative and interesting development in the fall of 1979 was setting up a pilot project to evaluate the use of the Cornell Program Synthesizer with a small number of students in CS 100, the introductory programming course. This action was in keeping with the earlier sentiments of the UCB that introductory courses should use microcomputers. The pilot was done in two parts, during summer school and then during the fall semester. In the fall semester 75 students from CS 100 and 25 students from CS 211 took part in the pilot project to gain experience and thus be able to decide in time to make plans for the spring semester. Six Terak microcomputers were installed in Upson Hall at a cost of \$33,000 (\$5,500 per machine). OCS assumed responsibility for operations and maintenance costs while Computer Science continued to

develop and support the Synthesizer software. Because the Synthesizer could compose only PL/C programs, OCS built the necessary software and hardware connections so that the Synthesizer could send completed programs to the IT batch processor on the 370/168 for execution and printing of results. The pilot proved successful enough to make plans to expand the number of Teraks to handle classes of over 550 students in the spring of 1980.

DEC20 Computer Ordered for Installation at OCS

Following up on the earlier UCB action favoring a DEC20 to improve time-sharing services, the issue came to the fore in 1979. Two members of the UCB, chairman Everhart and Doug Reece, combined interests to make a strong case for the acquisition and installation of a DEC20 computer. Besides the goal to make a significant improvement in time-shared interactive computing services for instruction and research, Reece was interested in bringing a complete suite of DEC20 applications developed at Carnegie-Mellon for use by all students in the Graduate School of Business and Public Administration.

In addition, to demonstrate strongly a complete new direction and theme for instructional computing, this new computer had to be on campus and easily available. After a search for suitable campus space, room G20 in Uris Hall was selected. Building a new raised-floor computer room, and providing additional air-conditioning requiring the extension of chilled water lines, raised the costs from an original estimate of \$500,000 for the project to \$1,000,000.¹ An order was placed in 1979 and delivery and installation was planned for the following summer.

IT (Instant Turnaround) Batch Use Declines

By the end of the decade IT use had pretty much run its course. The number of jobs run per year peaked at 595,000 in 1978–79 and then started to decline, numbering 493,000 in 1979–80. Some of the decline was a result of the success of SCMS and the shift to more interactive use including the shift to other technologies such as the Synthesizer.

Business Information Data Processing

Activity in administrative computing was modest toward the end of the decade. With a significantly reduced staff resulting from the cutbacks in 1977, only maintenance and minor new work were possible. Despite this situation, several new technologies were introduced late in the decade.

Mark Sense Cards for Student Registration; Centralized Registration

After McCalmon left the registrar position, Eleanor Lundy Rice was appointed registrar in early 1977 and initiated a series of changes to improve the operations of her office and services to students. The first major change was to make registration faster by having students fill in “mark sense” forms to produce initial registration requests so that college registrars could more quickly confirm registration. According to Lynne Personius, associate registrar, this new method eliminated a great deal of manual processing of cards. The following year the registration process was centralized in Barton Hall. Much of this activity and experience led to a new student information system project in the next decade.

CADE Data Entry System

Reflecting the shift away from punched cards, the keypunch operation made a partial move toward online data entry by acquiring a CADE (Computer-Assisted Data Entry) system from Unisys in 1979. This was a joint project of Production Services and the Endowed Accounting Office to move away from the mechanical keypunch stations. Jim Doolittle and Irene Van Zile were part of the OCS staff that recommended the installation of the system, took responsibility for creating the data entry programs for all the applications, and provided continuing support. The system was installed in the production area in the basement of Day Hall with several entry stations installed in the accounting offices on the first floor. The CADE system was considered as an RJE station so that when data entry was completed the information could be transferred to the 370/168 to create OS data sets to be called by the appropriate application program (personal communication, Jim Doolittle).

During the data entry operation the programming support made it possible to eliminate and automate keystrokes along with some editing of the information keyed into the system. The combination of the programming support and the elimination of the mechanical operations of the former keypunches resulted in a 20 to 30 percent gain in productivity. As a result, the gain in efficiency allowed CIT to provide additional data entry services without increasing staff. At this time there was also a name change: keypunch operators were now known as data entry operators or clerks.

Network Plans for the Campus

In the summer of 1978, Blackmun, Cogger, and Worona from OCS made a proposal to modernize

and upgrade Cornell's data network.³⁵ One part of the argument was increasing load. At this time approximately 150 CRT and hard-copy terminals were connected to the 370/168 (about half being installed since the first of the year) and the expectation was for another 100 to be connected before the end of the year. Another part was that the old installed technology was expensive, costing about \$100 per month for 120 character-per-second service speed while newer technology could reduce the cost to about \$30 per month for 960 character-per-second speed. Service

improvements could also be expected by being able to conveniently access other computers on campus, to connect to national networks, routinely exchange electronic mail, and connect to special processors. The proposal envisioned moving from point-to-point communications to a packet-switching network and is best captured by Figures 2 and 3.

The proposed plan called for the evaluation of two competing technologies. One was *TYMNET*, installed in 1974 with the NBER project, a leading commercial supplier of packet transmission services. The other was Telenet, a commercial version of the ARPANet built by Bolt, Beranek, and Newman. CUMC had experience with Telenet, which was also favored by the

³⁵ R. Blackmun, R. Cogger, and S. Worona, "Modernization and Upgrade of Cornell's Data Network; Preliminary Exploration," memo to S. A. Lawrence, July 1978.

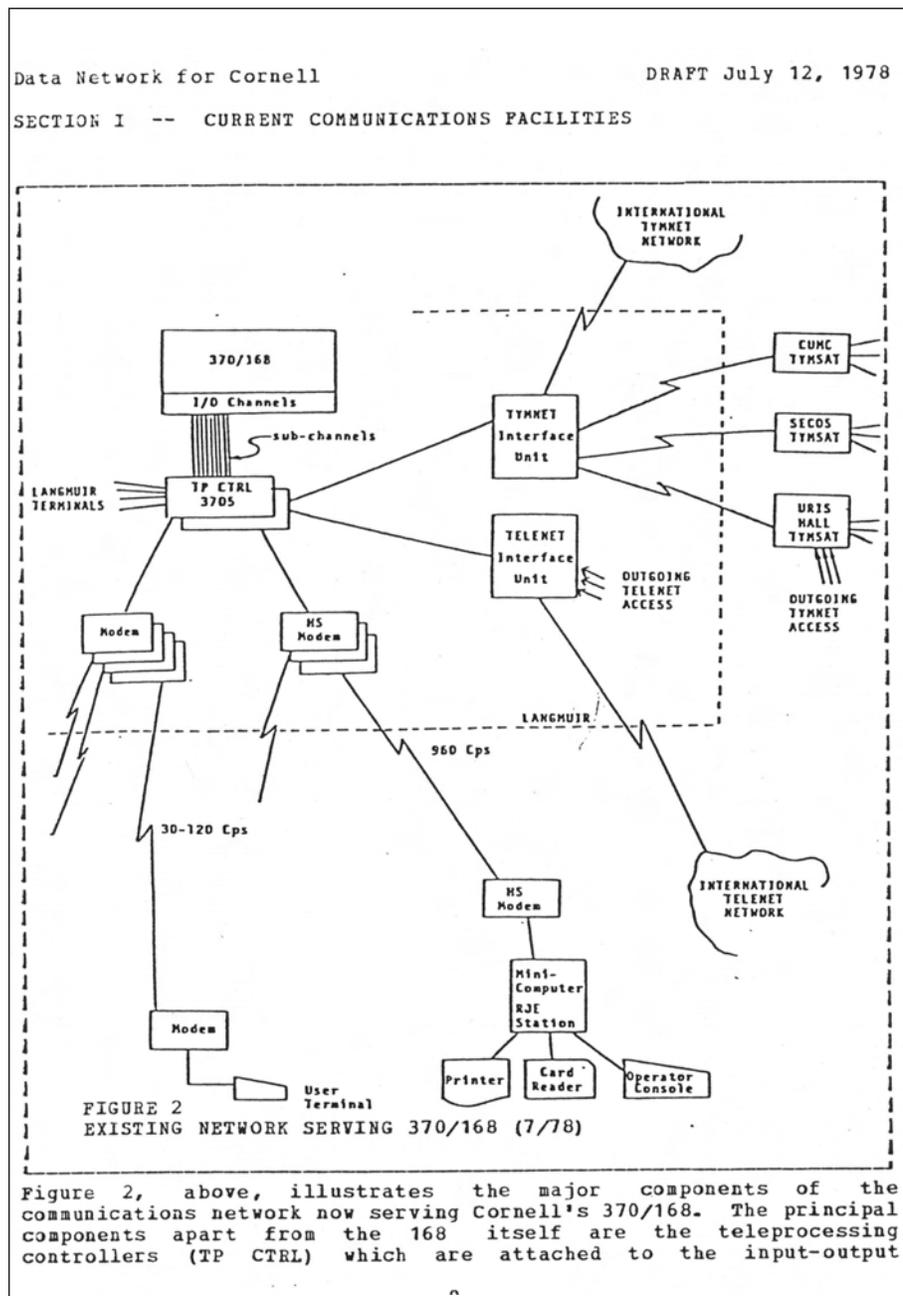


Figure 2. Existing Cornell campus network in 1978

Educom clients using Cornell facilities. No action was taken on the proposal, given all the other priorities in 1978-79, although some low-level work continued to be done in preparation for moving ahead at some future date.

Academic Computing

Several organizational changes were made in the Academic Computing Group. The merger of the Computer Activities Group with OCS created a single entity supporting the instructional and research users. The formation of a small group to focus on microcomputers came from the increasing interest and use of the small systems.

CAG Merges with OCS as Part of User Services

OCS and CAG (the Computer Activities Group supported by several of the statutory colleges) merged in 1979, and the CAG operations in Warren Hall became part of OCS User Services under the direction of Steve Worona, newly appointed assistant director of user services.³⁶ J. Robert Cooke and Van Houweling, to whom Worona reported, were the prime movers behind this merger. Cooke, in his role as director of instruction in the College of Agriculture, expected several benefits to flow from having one instead of

³⁶ "Computer Merger Expected to Improve Service," *Cornell Chronicle*, March 1979.

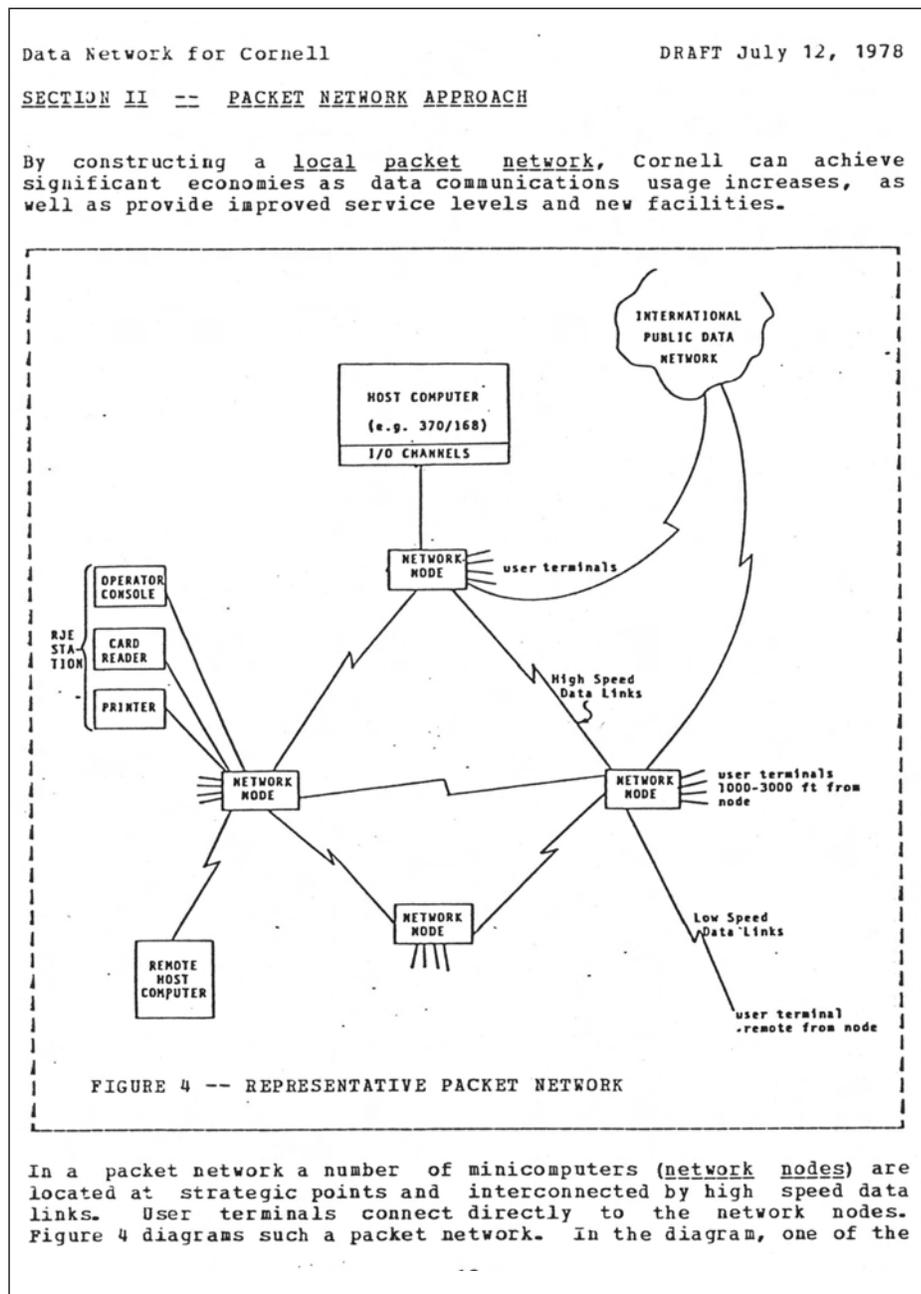


Figure 3. Proposed Cornell network, July 1978

two organizational bodies dealing with faculty and students. One was improved coordination and efficiency in the delivery of services. Another was a more aggressive use of newer technologies. Part of the agreement was the installation of more interactive terminals in Riley-Robb Hall and in Ives Hall for the School of Industrial and Labor Relations. In addition, a Terak microcomputer was to be installed in Riley-Robb. Errol Jones, the first and last director of CAG, took a technical position with the Small Computing group in OCS and later transferred to a position with Agricultural Engineering as a microcomputer specialist.

Formation of DACS (Decentralized Academic Computing Support)

The effort to give increasing support to microcomputers and computing activity in academic units led to the creation of the Academic Small Computing Group as part of Academic Computing in 1979. Van Houweling appointed Alison Brown acting assistant director of this unit and lobbied for funds to create a permanent position. Later that year, with budget approval and the support of Cooke and Rudan, recruiting began for the head of this new unit, now named Decentralized Academic Computing Support, or DACS. Brown led the recruiting effort as one of her principal short-term responsibilities.

New Leadership for OCS

As part of the outcome of the discussions on the future of computing at the University Computing Board and the senior levels of the university administration, the reporting relationship was changed and new leadership was appointed for OCS.

New Reporting Relationship: OCS to Report to Provost

One of the positive outcomes of the budget and service crises in OCS was a change in the reporting relationship of OCS. In the spring of 1979, Lawrence announced his intention to leave Cornell. At this time, the responsibility for computing on campus shifted to the provost's office. Don Randel became the responsible senior executive and also chairman of the UCB. With this change, Provost Kennedy announced that OCS had accumulated a total deficit of \$800,000 to \$900,000. There was no single reason to explain this shortfall, but the major reason seemed to be unrealistic budgeting on the part of the university. According to Kennedy, instructional use was projected to increase by 12 to 15 percent and had actually increased by more than 20 percent! Similarly, administrative computing use had increased beyond

that expected when budgets were prepared. The loss of income from NBER's departure and the failure of the CUMC project made the situation worse and showed the extent to which outside income was carrying the expenses of OCS. Kennedy committed to bringing the OCS budget into balance by eliminating the deficit with a special appropriation.

Cooke and Rudan Direct OCS Operations

In the summer of 1979, Kennedy made further changes in the OCS reporting structure. Randel had returned to his professorial position, and so Kennedy appointed J. Robert Cooke as part-time assistant to the provost to oversee OCS operations. Rudan was appointed acting director of OCS reporting to Cooke.

There was a wholesale change in the UCB with all new members and a new chairman. Everhart, dean of engineering, was appointed chair along with G. V. Chester, Nuclear Studies; Juris Hartmanis, Computer Science; Timothy D. Mount, Agricultural Economics; J. S. Ostrom, controller; Don M. Randel, Music; Douglas K. Reece, Business and Public Administration; Paul F. Velleman, Industrial and Labor Relations; and Richard N. White, Civil Engineering.

Although the budget overruns and subsequent organizational change affected OCS as an organization, demoralized the staff, produced tensions within and outside the organization, and affected the careers of competent and dedicated individuals, what cannot be overlooked are all the OCS successes that have been mentioned.

Search for New OCS Director

Given the conditions in campus computing and OCS at this time, it seemed appropriate to look to experienced and reputable non-Cornell academic computing practitioners for advice before undertaking the recruitment of a new director for OCS. Cooke approached this task with his usual gusto. In 1979, working in conjunction with the UCB and Everhart, a visiting committee was appointed to advise Cornell about several aspects of campus computing. They were to consider the quality of academic, administrative, and research computing in comparison to peer institutions, the effectiveness of service delivery given the resources being expended, and how quality and effectiveness could be improved. In summary, they were asked to help Cornell understand "how modern computing can benefit education, research, and efficient administration and how much such computing should cost." The visiting committee was made up of Fred Harris from the University of Chicago, John McCredie from Carnegie-Mellon (chair), and Saul Rosen from

Purdue. They were able to visit the campus for several days in January 1980, and after meetings with representatives from different constituencies, they came to the following conclusions:³⁷

OCS lacks cohesion. It gives the impression of a number of groups loosely coupled together. This is at least partially due to the history of OCS in which groups have been moved in and out. It is partially due to the geographical separation of the Langmuir Center from the rest of the campus.

OCS needs strong management to pull the organization together. It needs strong management to decide what belongs together and what doesn't. It needs a center on campus in which it can operate as a unit. Perhaps most of all it needs a well-defined mission and a high level of support from the Cornell administration.

How these recommendations came to influence the future will be discussed in the next decade. The search for a new leadership in campus computing also did not conclude until 1980, when Kenneth M. King was hired as the vice provost for computing. His significant impact on computing at Cornell will also be discussed in that decade.

Summary Comments

The 1970s was surely a decade of dashed hopes and expectations that started with promises of improved services to the three constituent groups—student users, researchers, and administrative users—and ended up with very mixed accomplishments. The lack of adequate funding forced OCS to seek funding from selling services to outside agencies to maintain more or less adequate services and cope with budget cut-backs. The several well-documented and noticeable failures should not mask the significant accomplishments that started during this decade. In particular, there was the beginning of supercomputing services for research users, the availability of entry-level services for students without the need of advance approvals, and the deployment of microcomputers on campus. Only administrative computing failed to move ahead significantly during the decade.

We close the 1970s with a summary showing the progress made in providing time-sharing services and some of the significant achievements in improving the VM operating systems software. Augmenting and improving the systems software, while seemingly a

potential waste of talented resources, should properly be considered a necessary step to getting a lot more out of the computing resources as an alternative to increasing them.

Growth in Time-Sharing Services

Because this decade was titled the Time-Sharing Decade it's appropriate to present some data to see how these services developed. The data were compiled from several different sources and represent the information that could be found. Even so, the growth over the decade is evident. Also noted are the various changes in hardware, operating systems, and time-sharing systems as they changed over the decade. (See Table 4.)

Systems Software Development Accomplishments

Starting in 1975–76, after a year's experience with VM, the System Programming staff began an extended series of improvements in various facets of the product being shipped by IBM. This multiyear effort, which extended into the 1980s, positioned Cornell as a leading world-renowned expert in VM systems software. The different improvements, with their author(s), were as follows:

- Cornell VM/370 Scheduler was implemented to improve control over resource utilization, to offer different levels of service to defined classes or users, to improve control over contention for resources, and to keep response times short and consistent for trivial commands. The majority of the work was done by Bob Cowles, hence the alternate name used was the Cowles Scheduler; it was built on the Favored Machine Virtual code developed earlier by Bob Lent.

- Page Migration improved the availability of needed pages in the highest-speed memory by migrating the less frequently used pages, which tended to slowly fill high-speed memory, to lower-speed storage devices. This improved the response time for trivial transactions severalfold. Larry Brenner developed this facility.

- OS Format Minidisks made VM/CMS minidisks available in both VM and OS so that the same data could be used in both the interactive and batch environment. Writeable OS format minidisk facilities provided by IBM had read-only access from CMS. Initial development was started by Herb Weiner and completed by Larry Chace.

- Minidisk Manager automated the crude manual tools provided for allocating and de-allocating minidisks available for each user's virtual machine and added tools for installations to manage and suballocate space in storage subpools. Andy Hanushevsky did the initial work and subsequent improvements. Although not

³⁷ Report of the Cornell University Visiting Committee on Computing, "Summary of Key Issues and Recommendations," January 21–22, 1980.

Table 4. Growth of Time Sharing 1967 to 1979

Year	Machine	System Control Program	Interactive Systems	Average Daily Users
1967	360/65	HASP/MFT	CTS	15
1969	360/65	CLASP	CRBE	10
1971	360/65	CLASP	CRBE	10
			APL	15
			CPS	10
1972	360/65	HASP	CRBE	10
			APL	15
			CPS	10
1973	360/65	HASP	TSO	10
			APL	15
1974	370/168	HASP	TSO	10
			APL	15
1975	370/168	VM/370	CMS	10 (limit 35)
1976	370/168 with VMA	VM/370	CMS & SCMS	25 (limit 65)
1977	370/168 with VMA	VM/370	CMS & SCMS	45 (limit 90)
1979	370/168 with VMA	VM/370	CMS & SCMS	71 (limit 120)

directly part of the VM environment the daily backup of online files was improved at about this same time.

- Incremental Data Set Backup selected only the changed datasets to be copied on the backup rather than the entire file space. As the amount of total file space increased, it became increasingly harder to execute a full daily backup at the end of the third shift to try to capture as much of the daily activity as possible. This was a joint effort with contributions from Peter Shames, Paul Whitted, Steve Stein, and Barbara Skoblick.

Several of these products—the VM Scheduler, the OS Format Minidisks, and the Minidisk Manager—were developed into program products and marketed by the Adesse Corporation to sites all over the world. While this enhanced the reputation of the involved individuals and Cornell, and brought financial rewards to OCS, these creations improved the computing environment at Cornell. For example, the development of a restricted interactive environment for student work relied on the VM Scheduler to deliver effective services without significant impact on the system. These improvements improved the performance of the 370/168 system well beyond the rated capabilities of the hardware.

OCS Space on Campus

A distracting issue during this whole decade was that of OCS space on campus. Lawrence as chairman of the UCB and as the executive responsible for OCS

championed this issue at all occasions. Earlier in the decade, it seemed that OCS could be assigned some space in the new Biological Science building, later named Corson-Mudd Hall, but that fell through. Later, Facilities (Buildings and Grounds) was going to move its staff from the Humphries Building on the edge of campus, and that seemed like a close enough site to move all of OCS to “campus.” However, that opportunity also fell through, and so OCS continued to have the computer center at Langmuir, academic support staff in Uris, and facilities scattered over the campus and administrative computing staff in Day Hall.

Other Campus Computing Activities

Dairy Records Processing Laboratory

We continue with the history of this organization taken from the document prepared by Lyle Wadell.³⁸

From 1966 through 1973 the cow numbers fluctuated up and down but stayed within the 440–458,000 range. Over this time period our applications grew, making use of the faster printers and the disk drives so that it was necessary to install a new computer in the spring of 1973 (system 370/135) along with larger and faster disk

³⁸ Lyle Wadell, “History of the Northeast Dairy Records Processing Laboratory, 1948–1985, with Additional Comments by J. D. Burke and H. Wilmot Carter,” Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, 2002.

drives (IBM 3330s). In late 1973 we took a step that has had a major impact on the efficiency of our operations in the years since. We installed the first group of online cathode ray tubes (IBM 3277s).

The number of cows being processed grew steadily, data files were expanded, and new applications were developed. (In the year 1975, there were 463,000 cows on record that grew to 548,000 cows in 1979 taking information from the table in the document.)

As this occurred appropriate equipment changes were made. In the fall of 1976 the disk drives were replaced by newer IBM 3344s. In the summer of 1977 a new computer, a system 370/138, was installed along with higher-speed printers (IBM 3203s). Our first teleprocessing applications were started in the summer of 1978 with a telephone connection to Eastern AI Cooperative. After a little over five years of use we replaced our cathode ray tubes with IBM 3278s in late 1978.

One of the more positive results of the availability of more and better records on milk production and artificial breeding, and the increased computer capabilities at Dairy Records, was the improvement in milk production in New York State. In the announcement that Charles R. Henderson was awarded the 1977 National Association of Animal Breeders Award from the American Dairy Association,³⁹ it was noted that Henderson's work helped increase milk production from 6,810 pounds per cow in 1950 to 13,612 in 1976. Henderson was noted for developing methods of estimating sources of variation among production records and for using these estimates in the evaluation of the genetic merit of animals. Using his methods, artificial insemination organizations had been able to select sires with the best genetic traits for breeding. It was estimated that 30 percent of the improvement in milk production came from the applications of these methods.

IBM System 7 for Monitoring/Controlling Energy Use

In 1975 a System 7 Demand Limiting Computer from IBM was installed for monitoring and controlling energy use. The system was to monitor electrical power used by numerous air-conditioning and environmental control units on campus and remotely shut off certain units to prevent excessive demand peaks and extra charges from NYSEG for peak demand power. This was a "turnkey" system, which IBM developed in response to the energy crisis in the early 1970s. Edward Hartz

was responsible for initiating the project and managing the installation. According to Mike Newman, attempts to manage Cornell's use of electrical power started much earlier in the 1960s with electrical relay systems that were designed to detect if a fan or pump was operating as scheduled.⁴⁰ As the technology developed, a central control room was built in the Chilled Water Plant building near Beebe Lake on Forest Home Drive to provide round-the-clock supervision and response. The System 7 was installed at the Chilled Water Plant. The System 7 and its upgrades continue to function to this day in a much-expanded role.

WATSBOX for Long-Distance Telephone Services

Another of the campuswide computer systems was the WATSBOX, which enabled Cornell to save money by managing its own long-distance dialing.⁴¹ Cornell telephones would dial into the "box," which would then route calls over leased WATS (Wide Area Telephone Service) lines. The WATSBOX was installed in January 1976 and Edgar A. "Ed" Swart directed the WATSBOX control center. It was estimated that the project would save Cornell over \$1,000,000 for the following five years. The system had built-in security protection based on "access codes" so that only authorized users could access the system and suspicious use could be reported to a control center. It also had the advantage that often-repeated numbers could be speed-dialed. At the same time, being a computer, it had the usual problems of programming errors causing service problems, downtime, and theft of services. In 1977, within a year of installation, the Cornell Safety Division was investigating the theft of \$1,000 of unauthorized telephone calls. In a problem familiar to most of the users of computers, the Safety Division was quoted as saying, "The problem is that many people assigned authorization cards leave them on their desks, paste them to their telephones, or give coworkers the number."

Laboratory of Nuclear Studies (LNS)

The Laboratory of Nuclear Studies installed a PDP-10 in 1971 at a total cost of \$600,000 provided by NSF.⁴² The installed IBM 1800 was too small, and the new system provided faster data handling and the ability to conduct more complex experiments.

³⁹ "Henderson Wins Genetic Prize," *Cornell Chronicle*, July 14, 1977.

⁴⁰ H. Michael Newman, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, September 1999.

⁴¹ "WATSBOX in Operation," *Cornell Chronicle*, January 22, 1976.

⁴² "New Computer Aids Synchrotron Physicists," *Cornell Chronicle*, February 11, 1971.

Department of Physical Biology in the Veterinary College

Also in 1971, the Department of Physical Biology installed a PDP-15 in its Radiation Biology Lab at Langmuir at a cost of \$90,000. Funds were provided by a grant from NIH. Howard Moraff was the director of this service facility for recording and processing experimental data.

Program of Computer Graphics

In late 1974 the National Science Foundation provided a grant of \$500,000 for Donald P. Greenberg, professor of architecture, to establish the Program of Computer Graphics.⁴³ Greenberg was in the forefront of the development of computer graphics. One of his statements at the time turned out to be prophetic: "It is entirely feasible with the current state of the art of computer graphics to develop programs for weather forecasting, develop detailed topographic maps of Mars and Venus using radar data, produce pictorial time displays of population growth and ethnic changes, to name a few."

One of the memorable productions was "Cornell in Perspective," which took a walk through the Arts Quadrangle as it appeared in 1900. The production included the trolley that could be seen moving past the now-demolished Boardman Hall. The computers and staff supporting this program were installed in Rand Hall, taking advantage of the computer space abandoned by OCS in 1967–68. In 1979, the program received an additional grant of \$1,000,000 from NSF to refine computer graphics techniques and to apply them to structural engineering and water resources problems. In addition to these projects the program was working on animation projects to eliminate the time-consuming hand painting of each frame of animated cartoons.

College of Veterinary Medicine

In late 1976 the College of Veterinary Medicine announced its use of the MUMPS system on its computer.⁴⁴ MUMPS, an acronym for Multi-User, Multi-Program System, was developed at Massachusetts General Hospital in Boston for human patient records and adapted by John Lewkowicz and his staff for animal patients. The system was flexible, allowing different and simple means to find the patient and its owner and other information, including historical treatment information. At installation the database contained more than 22,000 visits by some 15,000 patients. In addition to improving the operations of the large and small animal clinics at the college, the system was expected to assist faculty and students considering different aspects of animal treatment and diseases.

According to Lewkowicz the system had already been in operation for a year and a half and proven its worth.

Department of Computer Science

In 1977–78, the Department of Computer Science temporarily installed its first departmental computer, a DEC PDP-11/60, in Upson Hall. When the new fifth-floor addition to Upson was completed the following year, this system plus an additional PDP 11/60 was installed in the new computer room. In addition to these systems several LSI-11 processors were installed. These systems became the research and development machines for the department.⁴⁵

Endnotes

^a Dave Perkins was the first Unitech technician who stayed on at Cornell to become a member of the Networking group. He retired in 2004.

^b Thirty years later, Cornell still lacks an integrated university management system. One can argue about the reasons for this condition as attempts have been made over the years to redress this situation, the latest being Project 2000 in the late 1990s. However, if one considers where much of the activity was in the intervening years, one soon recognizes that a lot of effort was expended on building and maintaining transaction systems for different key operating areas, and integration was never accomplished as events forced time and attention on the transaction systems.

^c One of the principals in SCT at the time was Fred Gross, the president, who participated in the contract negotiations and signings and then in the contract terminations. He was well-known in the higher education community.

^d The Campbell report notes, "The extent of change (in the 360/65 over its lifetime) is equivalent to that of keeping an automobile's chassis frame and physical engine block, then modifying the engine for higher performance, and adding other new components including a new body." What also should have been said is doing all this while keeping the car on the road!

^e The records show that the sale was to be to L&A Computer Industries, Inc., of Overland Park, Kansas, but whose office and legal services were provided out of Atlanta, according to Ralph Barnard, who again was part of the Cornell team involved with the sale. This multiplicity of locations and possible sources for funding the acquisition caused some initial anxiety. However, on the direction of Controller Peterson, Barnard recalls personally flying to Kansas City to pick up the bank check and then flying to New York City to deposit the funds directly in the Cornell account with Morgan Stanley Bank.

^f PBM/STIRS ran on the Cornell computers until the late 1980s. During this interval many changes were made, including the reduction of charges to \$4.50 per compound for using overnight batch processing of runs. Eventually, the availability of less expensive minicomputers and direct connections to mass spectrometry equipment made that route more economical. Later still, the pack-

⁴³ "Computer Aids Visualization," *Cornell Chronicle*, October 31, 1974.

⁴⁴ "MUMPS at the Vet College," *Cornell Chronicle*, October 7, 1976.

⁴⁵ Department of Computer Science, "1977–78 Annual Report."

age was commercialized to run on the more generally available and powerful micro, or personal, computers.

^g Dan Updegrave was a student at Cornell in the early 1970s and was one of Van Houweling's student assistants for Metro-Apex. After graduating from Cornell he went on to a successful career, holding leadership positions in computing/information technology with Educom, the University of Pennsylvania, Yale University, and the University of Texas at Austin.

^h EFPM continued to run on the Cornell computers until December 1987 when the availability of desktop computers and spreadsheet systems allowed higher education institutions to do the same planning on their own local computers.

ⁱ John Williams approached his job with vigor and spent the summer meeting with a variety of administrators and faculty. This was at the time when the use of Free IT was rapidly increasing and all his information indicated this growth would continue. In anticipation of this situation, additional keypunches were ordered for Upson and Clark Halls. At this time OCS, at the suggestion of IBM, signed one-year leases for keypunches so they could be returned to IBM in the summer, after being beat up by students. New ones would be installed in August. Given IBM's lead time for ordering that year, quite a few keypunches arrived early in July. Deans Cranch and Levin were made aware of this storage of equipment in hallways and wrote a "nasty" letter to Lawrence chastising OCS for its wasteful ways. Had the deans taken the trouble to ask about this situation, they would have learned that IBM did not charge OCS until the machines were actually installed and so no extra expense was involved by having the machines arrive early. I had difficulty in composing a response to this unjustified criticism.

^j The problem of 370/168 outages in 1977 was finally resolved when IBM deduced that the incidents were caused by external actions. A hidden camera was used to record the activity in the console area and caught several occasions when an operator improperly pressed several console switches that resulted in the failure. Credit for finally solving this problem is due to the persistence of IBM and Bob Blackmun, then acting director of OCS. The individual was dismissed and no further such outages were experienced. An unfortunate side effect of this incident was that the person was the first affirmative action candidate hired by OCS and so these actions soured management for a time about hiring more such candidates.

^k The year 1976 was an exciting one for me. I was selected as the one Cornell administrator to attend the Executive Development Program held at the Graduate School of Business during the summer. I not only refreshed a lot of past education but felt it was time for a job change and that is when I moved to Institutional Planning and Analysis for several years.

^l Andy Hanushevsky and Donna Bergmark received the best paper award at the Fourth Annual FPS Array Processors Annual Meeting in 1980. Their paper, titled "Document Retrieval: A Novel Application for the AP," described how use of an AP can achieve a two-thirds reduction in the price of searching for documents in large databases. At the same conference, Ben Schwarz received critical acclaim for his paper, "A Dynamic Segment Loader for the AP."

^m There was a bit of a flap when Grimison was featured in an FPS advertisement in technical computer publications, even though all Cornell clearances were obtained!

ⁿ "Mailfiles," as they were called back then, limited formats. Nonetheless, in the 1970s some of the forward-looking staff in OCS started to use e-mail to their advantage in recording and transmitting textual information, forcing others to do the same or be left out of the "loop."

^o Information Associates was from Rochester, N.Y. One of the founders was David Duffield, who later also founded PeopleSoft Company and whose software was installed for Human Resource/Payroll in 1998. He has been a generous benefactor to Cornell; Duffield Hall is named in his honor. One of the other founders of IA was John Robinson, who later founded TRG, The Robinson Group, and whose systems were considered for installation in the 1990s.

^p I had the unwelcome task of negotiating a settlement of the payments that CUMC would make to OCS for the aborted transfer of its work to Ithaca. Phil Giuca was the principal at CUMC, and while admitting that they owed OCS some amount, they did not feel compelled to pay the amounts billed. We eventually settled for something like 10 cents on the dollar and had a memorable celebration lunch at the Tavern on the Green.

^q Ralph Barnard recalls that the first company contacted for this lease/purchase contract was OPM Systems. OPM had a strong reputation for entering into such contracts with the major airplane manufacturers and, although Ralph checked with these companies, he recalls getting very little good information. During this investigation some members of the Board of Trustees were assuring Cornell executives that lease/purchase agreements were the new way of doing business and there shouldn't be a problem. The agreement was not completed with OPM because they would not provide any guarantees at completion of the contract, and so CIS was given the contract. Barnard happily recalls that a few years later several of the OPM principals were indicted for keeping two sets of books for their activities and that OPM had bilked banks out of millions of dollars for fake leases.

^r It was later acknowledged by some of the executive staff involved in the decision that the DEC20 should have been installed at Langmuir, where there was quality computer room space, on-site operations and support staff, and high-speed connections to campus. However, the UCB, feeling that it was necessary to make a strong statement that computing was going to be done in a different way and on campus, forced the installation to be made in Uris Hall and at considerable extra cost. To be fair, the extra chilled water lines and the presence of a raised flooring in the computer room were very useful in future years.

1980 to 1989

The Decade of the Personal Computer

Contents

1980 to 1989—Industry Overview	101
Hardware	101
Software	103
Services	104
Vendors	104
Technology Terms, Acronyms, and Buzz Words	105
1980 to 1984 at Cornell	105
Campus Issues and Transitions	106
<i>Assessment of Computing Needs by the UCB</i>	106
<i>Ken King Becomes the First Vice President for Computing at Cornell</i>	106
<i>From OCS to CCS</i>	107
CCS and New Campus Initiatives	107
<i>More IBM Computers Installed: CornellA, CornellB, CornellC</i>	107
<i>Adabas Acquired for Business Systems</i>	108
<i>DEC20 Installed</i>	108
<i>CCS Organization and Organization Chart</i>	108
<i>A Building on Campus for CCS and Telecommunications Study</i>	110
<i>Supercomputers; Wilson Wins Nobel Prize</i>	111
<i>Explosion of Computers on Campus</i>	111
<i>New Publications and Outlets Developed</i>	111
CCS Financial Restructuring/Cost Recovery	111
<i>The Glob Model for Cost Recovery</i>	112
<i>CSARS (Computer Services Accounting and Reporting System)</i>	113
Instructional Computing	113
<i>Instant Turnaround (IT) Fast Batch and the DEC20</i>	113
<i>The Use of SCRIPT and the Development of CUPAPER and CUTHESES</i>	113
<i>Plato Service Ends</i>	114
<i>VMBatch Developed; Free Distribution Accounts</i>	114
Research Computing	114
<i>Windfall Computing</i>	114
<i>Statistical Computing</i>	115
<i>Large-Scale Computing—More Array Processors Acquired</i>	115
<i>A New Computer Billing Category: Flatrate</i>	115
Business Systems	116
<i>Policy Statements on Administrative Systems</i>	116
<i>Administrative Priorities Committee</i>	116
<i>Application Transfer Team Study—Registrar’s Office</i>	116
<i>Adabas Installed; Database Administration Started</i>	116
<i>New Systems Development: Student, Payroll, Financial Accounting</i>	117
<i>Installation of IBM 3270 Terminals</i>	117
<i>Improved Relations with CUMC</i>	117

Text Processing and Information Services	118
<i>Office Automation Group Formed</i>	118
<i>Word Processing Systems Support Defined</i>	118
<i>Electronic Mail</i>	118
<i>Advent of CUINFO</i>	118
Telecommunications and Networking	119
<i>SYTEK and the Start of High-Speed Networks on Campus</i>	119
<i>BITNET and Network Addresses in Telephone Directory</i>	120
1984 to 1986—Microcomputers Rise to Prominence at Cornell	120
Microcomputer Sales Start	121
CCS Seminars and Workshops	122
User Groups Form (and Dissolve)	122
Thefts of Computers	123
Software Piracy	123
Microcomputer Viruses, Trojan Horses, and Worms	123
Instructional Computing	124
<i>The First Use of Macs and IBM PCs</i>	124
<i>Concerns about Microcomputers and Course Instruction</i>	124
<i>Project Ezra</i>	125
<i>Apple Mac•Ed Center: MacAdemia Conferences</i>	128
<i>Intellectual Property Rights</i>	129
Business Systems	129
<i>UCB Recommendations for Decentralized Systems and Support</i>	129
<i>Microcomputer Support Group; Apple Loaner Program</i>	130
<i>Development of IRIS (Inquiry and Reporting Information System)</i>	130
<i>First High-Speed Laser Printer Installed</i>	130
<i>Automating Production Control</i>	131
Telecommunications and Networking	131
<i>The New Cornell Telecommunications/Telephone System</i>	131
<i>The Computing and Communications Center (CCC)</i>	132
<i>Telecommunications Equipment Installed</i>	133
<i>Pronet Installed to Use TCP/IP Protocols for Campus Network</i>	133
<i>CCSNet and TheoryNet</i>	134
<i>Omninet for Low-Speed Local Area Networks</i>	134
<i>Terminal Emulator Development</i>	134
<i>NSFNet</i>	135
<i>NYSERNet</i>	135
<i>GateD</i>	136
Research Computing and Super Computing	136
<i>IBM 3081D Floating Point Errors</i>	136
<i>Supercomputing at Cornell</i>	136
<i>The Theory Center Is Formed</i>	137
<i>The Production Supercomputer Facility (PSF)</i>	137
<i>Cornell National Supercomputer Facility (CNSF) Computers in CCC</i>	138
<i>CISER</i>	138
Information Services	138
<i>Word Processing Systems for Microcomputers—WordPerfect</i>	139
<i>Project Milestone—The Electronic Age Comes to Day Hall</i>	139

CUINFO Grows, “Dear Uncle Ezra” Started	139
Desktop Publishing	140
Cornell Library Goes Online	140
Mid-Decade Transitions	141
The End of the Punch Card Era.	141
The End of Mainframe Plotting Services.	141
DEC20 Removed—Vax Systems Installed.	141
New Deans and Executive Staff	141
Dim Financial Outlook.	142
CCC Computer Room Opens	142
CCS Staff Move into CCC	142
Office Equipment Center Merges with CCS Microcomputer Sales	143
Ken King leaves Cornell.	143
Formation of User Groups CAUCUS and SUPER!	144
1987 to 1989 at Cornell	144
CCS and Campus Networking Activities	144
<i>Status in 1987</i>	<i>144</i>
<i>AT-Gateways for LAN Connections; Network Operations Center Started</i>	<i>146</i>
<i>The Call for Pincnet</i>	<i>146</i>
<i>NSFNet, NYSERNet, and GateD</i>	<i>146</i>
<i>Networking Study</i>	<i>146</i>
Business Systems	147
<i>Table of Central Administrative Information Systems</i>	<i>147</i>
<i>CUDA (Cornell University Distributed Accounting) System</i>	<i>147</i>
<i>High-Speed Laser Printing Systems</i>	<i>149</i>
Research Computing	149
<i>The Push to Parallel Computing at CNSF</i>	<i>149</i>
<i>High-Performance Computing—A National Concern</i>	<i>150</i>
<i>Mass Storage Systems</i>	<i>150</i>
<i>The Theory Center Building</i>	<i>151</i>
The Movement from Computing to Information Technology	152
<i>M. Stuart Lynn Appointed Vice President for Computing.</i>	<i>152</i>
<i>Cornell Computing Services to Cornell Information Technologies</i>	<i>152</i>
<i>Mission of Information Technologies</i>	<i>152</i>
<i>New CIT Organization</i>	<i>153</i>
<i>CIT Internal Advisory Committees: Technologies, Applications, Services</i>	<i>154</i>
<i>CIT Organization Chart.</i>	<i>154</i>
<i>New Publications: CIT News, CIT Briefly Replaces the Bulletin</i>	<i>154</i>
<i>CIT Logo</i>	<i>155</i>
<i>Training and Workshops</i>	<i>156</i>
<i>Help Desk</i>	<i>156</i>
Microcomputers/Personal Computers Update	156
<i>MOS Becomes Part of the Services Division</i>	<i>156</i>
<i>Product Fairs Evolve into Educational Distribution Events</i>	<i>156</i>
<i>NeXT Workstations.</i>	<i>157</i>
<i>Cornell Faculty Continue to Win Awards for Software</i>	<i>157</i>
<i>Microcomputers/Personal Computers on Campus</i>	<i>158</i>
Networking and Telecommunications	158
<i>Networking Task Force Report.</i>	<i>158</i>
<i>Other Communication Services</i>	<i>159</i>

Information Services	160
<i>Electronic Mail—New System for IBM Computers</i>	160
CUINFO	160
<i>Dear Uncle Ezra</i>	160
<i>Bulletin Boards and Electronic Conferencing</i>	161
<i>Library System Goes Online</i>	161
<i>Libraries Provide More Information Resources</i>	161
Business Systems	162
<i>Administrative Systems and Data Policy</i>	162
<i>High-Speed Laser Printing</i>	162
Status of Mainframe Computers at CIT	162
<i>Mainframe Computers Struggle to Keep Up with the Load</i>	163
<i>Statistical Computing</i>	163
The Cornell Worm Incident	164
Other Campus Computing Activities	165
<i>Dairy Records Processing Lab</i>	165
<i>CADIF (Computer-Aided Design and Instructional Facility)</i>	165
<i>CISER (Cornell Institute for Social and Economic Research)</i>	165
<i>Department of Computer Science</i>	165

Figures

Figure 1. VP Ken King at the DEC 20 opening ceremonies	109
Figure 2. Floating Point Systems, FPS164 Array Processor (Ken Wilson and Alec Grimison)	115
Figure 3. CUINFO terminal (kiosk) in the Day Hall lobby	118
Figure 4. CCS Microcomputer Sales Office in Uris Hall—taking orders	121
Figure 5. The first Macintosh facility at Cornell (Uris Library)	124
Figure 6. Statue of Cornell founder Ezra Cornell, with IBM PC keyboard	125
Figure 7. Unloading an early shipment of Project Ezra equipment	126
Figure 8. 1986 Project Ezra fair “Tools for Learning” in Statler Hall	127
Figure 9. Telecommunications cable being laid on the Arts Quad	133
Figure 10. Computing and Communications Center (new addition to left)	142
Figure 11. The campus network in 1986–1987	145
Figure 12. GE robot breaks ground for new Theory Center building, on command by the center’s founding director, Ken Wilson	151
Figure 13. M. Stuart Lynn, vice president of information technologies	152
Figure 14. CIT organization model	153
Figure 15. CIT staff taking orders for personal computers in Lynah Rink	157
Figure 16. NeXT Workstation—printer, keyboard, monitor, CPU	157
Figure 17. Cornell networks in 1989	159

Tables

Table 1. Cornell Computer Services Organization Chart, 1983 to 1987	110
Table 2. Administrative Information Systems Status	148
Table 3. Cornell Information Technologies (CIT) Organization Chart, 1988	155
Table 4. Statistical Software Available on CIT Mainframes	163

1980 to 1989—Industry Overview

In the 1980s two major innovations radically changed the use of computers and the delivery of computer services. The first was the commercial introduction of microcomputers, or personal computers, on a large scale so that millions were sold in the latter half of the decade. The second was the development of the so-called “killer applications” of word processing and spreadsheets. Combined with the introduction of low-cost laser printers, these applications changed the way computers could be used in all facets of activity. These new developments moved computers from computational devices to information-processing devices of much greater use and appeal to individuals and industry alike¹.

This decade also saw the creation of local area networks that linked personal computers with each other for sharing expensive resources such as printers and file servers for common or private data storage and for providing pathways to other more far-reaching networks and computer systems. Electronic mail started to be used extensively once the networks enabled data transfer between computers. Further, during this decade there was agreement on network standards that permitted all varieties of computer systems to be interconnected locally, nationally, and around the world.

Hardware

In overall terms this decade saw rapid growth in the number of installed computers in the United States, from an estimated 1 million units in 1980, to 5 million in 1983, to 17 million in 1987.² This growth alone is a good indicator of the activity in computing during the 1980s.

First, in chip technology there were advances in speed and in the number of bits stored on a chip. By the early 1980s the Intel 8088 chip, which was rated at 4.77 MHz, was succeeded by a sequence of chips: the 80286, the 80386, and the 80486, which, by the end of the decade, were rated at around 30 MHz. Not only was the speed increased but the design went from 8- to 16- to 32-bit bytes, which increased the size of memory and also the instruction-addressing space. Channels, which moved the information inside the computer, also went from 8- to 16- to 32-bit channels, further increasing the speed of systems. The main competitor of Intel was Motorola, which produced a similar set of chips, the 88000 series, running at similar speeds. Memory size, which at the beginning of the 1980s was in the order of kilobytes, increased to multiple megabytes by the end of the decade. Most vendors selling microcomputers used either Motorola or Intel chips in their systems.

Although microcomputers were generally referred to as “desktop” machines, for some models it was more convenient to put the processing unit on the floor. From that reference, microcomputer offerings went in two different directions during this decade—portable computers and, in contrast, workstation computers that approached the power of large mainframe (centralized) computers in a much smaller physical space. It is generally acknowledged that in 1981 Osborne Computer produced the first portable PC, although it was more often referred to as “luggable,” owing to its size and its 24-pound weight. Kaypro, Compaq, and IBM also developed portable systems. At the larger end, companies such as Sun Microsystems, Apollo, and DEC were offering powerful microcomputers that became known as workstations. In 1985, Apollo was offering a workstation running at 1 MIP, which was in the same range as an IBM 4341 mainframe. In 1989, NeXT Corporation, formed by Steve Jobs after he was ousted from Apple Corporation, Inc. (we will use just “Apple” in most future references), brought out the NeXT Workstation, which was rated at 5 MIPS, with 8 Mb of memory, and was offered at a selling price of \$9,995. NeXT used a variant of the Unix operating system and its own graphical user interface (GUI). It was noted for having a large black cube for the CPU and an accompanying black case for the CRT monitor.

IBM continued to dominate the mainframe market for large central computers. Amdahl had earlier entered this market, as had Fujitsu. They both offered 370-compatible machines, which could run with the IBM operating systems, along with application packages from IBM or other vendors. IBM moved to offering small and air-cooled 370-compatible computers such as the 43xx series that used the same peripherals and operating systems. These computers did not require chilled water for cooling and had a large range of operating temperatures so they could be placed into use outside special computer rooms. In effect, IBM created new entry-level systems to continue to expand its base. Other vendors followed suit, although longtime vendors such as Burroughs continued to offer machines and operating systems of their own design.

Cray Research and Control Data continued to offer the powerful fast “supercomputers” of the times. During the decade, Cray came out with a succession of systems—the Cray-2, Cray X-MP, and Cray Y-MP—that increased the Cray’s supercomputer ratings from 1 gigaflop to over 2 gigaflops (1 billion floating

¹ Throughout this document, technical terms are based on definitions from the Encyclopedia of Computer Sciences, 4th ed. or from Webopedia at www.webopedia.com, or from Polsson, Ken, “Chronology of Personal Computers,”

² Polsson, Ken, “Chronology of Personal Computers,” www.islandnet.com/~kpolsson/comphist/, 2003.

point operations per second).³ ETA systems, formed out of the Control Data supercomputing operations, announced its ETA-10 supercomputer to compete with Cray systems.

RISC (Reduced Instruction Set Computer) architecture computers began to be mentioned in the literature based on developments in IBM in the late 1970s. RISC started to be used to build servers and computers for scientific computation. RISC machines were so named as a contrast to mainframes that were built on CISC (Complex Instruction Set Computer) processors. Instructions in CISC machines were not real hardware instructions but relied on microprograms written in the real instructions. On RISC machines there were no microprograms, and machine instructions were implemented directly in hardware and executed in a single machine cycle. More complex tasks requiring more than a single computing cycle were done by a sequence of basic instructions or by the use of a subroutine. Executing instructions in a single cycle was a key part of the performance improvement for RISC machines, since instructions could be “pipelined” so that independent tasks could be executed in parallel instruction streams with predictable execution times.

In the 1980s there were several decade-defining incidents and hardware innovations. One was the commercial by Apple that announced the Macintosh system during the 1984 Superbowl game of the National Football League. This commercial played upon the book *1984* written by George Orwell and emphasized how the Macintosh liberated people from the “big brother” mainframe computers. The Macintosh, affectionately known as the Mac, which sold for \$2,495, introduced and popularized the use of a “mouse”-driven GUI, which was inspired by previous work done at the Xerox Palo Alto Research Center (PARC). In 1981 Xerox actually sold a system, the Star 8010, that had similar capabilities to those of the Mac, but its price of \$16,000 did not generate many sales. Apple captured the imagination and market with the Mac even though by this time it had sold over 2 million Apple II systems. In 1984, *Time* magazine also named the personal computer as the 1983 “Man of the Year.”

The other ground-breaking hardware development came from Hewlett Packard (HP), which introduced several new printing technologies. In 1984 HP came out with the Thinkjet printer, which introduced inkjet printing technology and shortly followed this by the introducing the Laserjet printer, which used an even more innovative technology: laser printing. These major changes quickly made obsolete the fixed head (and fixed font) and dot matrix (variable font but poor quality) printing mechanisms.

Further, when the Laserjet printer was combined with Postscript (see definitions below), a whole new industry called “desktop publishing” was created. The Mac was the first computer to take advantage of these offerings, became the standard for desktop publishing, and dominated the field for many years.

Apple was the first vendor to introduce the 3.5-inch diskette as a standard offering on its Macintosh computer. Such mountable and replaceable disks were referred to as “floppy” disks when they first came out in the 1970s. Originally these disks were 8.5 inches in diameter. They were so floppy that they had to be stored in a stiff paper folder and carefully inserted into the vertical reader slot. The next generation of floppy disks was 5.25 inches in diameter, which IBM first introduced and then adopted as its standard. The 3.5-inch diskette was hardly floppy, and its size gave it great portability because it could fit into a shirt pocket. It was used for storing both programs and data. The 3.5-inch diskette was used to boot up the first Macs because it contained the operating system, there being no internal hard drive. When first introduced, the diskette had a capacity to store 340 kilobytes (Kb), but improvements during the decade increased this storage to over 1 megabyte (Mb). By the end of the decade the 3.5-inch diskette had replaced the 5.25-inch real floppy as the diskette of choice for micro-computer data and program information storage.

The CD-ROM (Compact Disc-Read Only Memory) was a new addition using optical disks for the offline storage of computer information. CD-ROM drives were first popularized in 1984 by Phillips, one of the developers of the technology. This medium, which most commonly could store 650 Mb (the equivalent of several hundred floppies), became one of the principal means for the distribution of software and data. Because the computer-generated CD-ROMs shared the same technology as audio CDs, the opportunity to merge these technologies became a reality.

Once IBM introduced its personal computer (PC) in 1981, it became a best-selling system. When IBM introduced a series of new improved models, the PC-XT and the PC-AT, sales continued to boom. In short order, the term “PC” became an interchangeable term with “microcomputer,” even though IBM owned the copyright. To capitalize on this emerging market for computers, other vendors started to produce IBM-compatible systems. In 1982, Compaq Computer, which had been formed a year earlier, spent over \$1 million to reengineer the ROM BIOS (basic input/output system). This allowed Compaq to avoid IBM copyrights and enabled the company to offer an IBM-compatible

³ www.cray.com/company/history.html

PC at a lower price than IBM's.⁴ Phoenix Technology reverse-engineered a compatible chip to the Intel 8088 PC chip and sold it to other clone makers. At about this time, IBM published the PC specifications in an "open architecture" statement that gave a great boost to the PC clone or "compatibles PC" market, as any machine conforming to those standards could run the PC-based operating systems and applications. In time, Intel started selling the IBM chip to others and the market grew. IBM tried to counter this activity by coming out with its PS/2 line of computers, but by this time it was too late; the PC-compatibles market was too well established.

Software

Activity and innovation in software during this decade was no less frantic than in hardware. New vendors and new offerings came on the market constantly, often in combination with hardware but in most cases independent of hardware vendors.

DOS (Disk Operating System) from Microsoft became the most used operating system on microcomputers. It became the de facto standard for the IBM PC and compatibles market. About mid-decade Microsoft brought out the first versions of its Windows software, which many claimed was a feeble attempt to duplicate the GUI introduced by the Mac. As the first of many attempts to deter Microsoft's dominance in this area, even though IBM had given Microsoft the rights to DOS, IBM developed OS/2 (Operating System 2) as its own offering for its new line of PS/2 computers. Interestingly, IBM and Microsoft signed a joint agreement to develop this new operating system! Numerous legal actions were taken during this decade to preserve copyright protection of such software. Apple sued Franklin computer and won the suit; then Apple sued Microsoft for using Apple's icon-driven user interface for its Windows system. In turn, Xerox sued Apple for "borrowing" the Star graphical user interface for its Macintosh system. Although such actions provide some sense of the competition, it is clear that during the decade the two most-used operating systems were DOS and the one for the Mac.

The decade-defining software—the so-called "killer-applications" or "killer-apps"—were the applications for spreadsheets and word processing, with desktop publishing a close third. At the beginning of the decade, the Visicalc spreadsheet was the dominant software available on several microcomputers, including Osborne, the Apple II, and others. However, when Lotus 1-2-3 was introduced in 1983, it became the primary spreadsheet used on PCs and compatibles, and soon after that Visicalc faded into obscurity. At about this same time, Microsoft brought out the first versions of its spreadsheet software called Multiplan—first

for the PC and later for the Mac. By the end of the decade Multiplan was almost obsolete, having been replaced by Excel, which was first introduced in 1987 by Microsoft. During this same year, Borland introduced its Quatro spreadsheet as a competitor for the IBM and IBM PC-compatible microcomputers.

At the beginning of the decade Wang provided the dominant word processing software on its proprietary CRT-based computers with more than 50,000 installations around the world. However, products such as Wordstar, Displaywriter, WordPerfect for the PC, and MacWrite for the Mac very quickly took over the market. The great advantage of these text editor offerings was that they were used on general-purpose microcomputers that were more and more commonly attached to the new inkjet and laser printers. The marked improvement in appearance of the resulting documents, printed with a much larger selection of fonts, made this very appealing.

Unix was the operating system of choice for workstations and the minicomputer market during this decade. The Open Software Foundation⁵ was formed by IBM, HP, DEC, Apollo, and others to counter AT&T's ownership and control of Unix and to define an offering that could compete.

New software references were invented in this decade, including "vaporware," coined to describe nonexistent software that was talked about and promoted by vendors to manipulate the market. There was also "dribble-ware," when features came out one by one, and "idea-ware" to respond to a perceived threat to an existing product.

For larger-scale business applications, but applicable to other activities, new database management technologies were developed in the 1980s. The most important were called relational database management systems (RDBMS), the first one of which was developed by Oracle, based on earlier work by IBM and military contracts. RDBMSs store data in the form of related tables. Their power comes from specifying how data are related to each other and how they will be extracted from the database. The companion power of these RDBMS came from the query languages that were used to request information from the database. The original language was called SEQUEL (Structured English Query Language) and was designed by IBM in the mid-1970s. The name was shorted to SQL and was first introduced as a commercial product by Oracle in 1979. By the end of the 1980s, ANSI (the American National Standards Institute) had approved a rudimentary version of an SQL standard although there continued to be different "dialects."

⁴ "Plugged in to a New Millennium," *Inforworld*, October 26, 1998.

⁵ Open Software Foundation from FOLDOC, <http://wombat.doc.ic.ac.uk.foldoc>.

Services

With the emergence of microcomputers, support services changed to meet the demands of this new technology. At the beginning of the decade, putting together a microcomputer configuration was a technical issue: what CPU to select, what CRT to use as the monitor, how much memory to add over and above what the vendor provided, and how, and what, other peripherals to obtain, especially printers. Once the hardware was chosen, a choice had to be made about software—what operating system and which applications. Often, these items had to be ordered through the mail because there were few retail outlets. Building a system was a craft that required experienced technologists. Specialists were developed who could advise prospective clients on their choices, their relevant advantages and disadvantages, and costs. Seminars and classes were developed to extend and share this knowledge base. Soon consultants had to be trained to deal with clients who already had systems and who were experiencing difficulties. By the end of the decade there were many computer stores providing advice and sales as microcomputers moved to be more of a commodity, off-the-shelf item.

Although some of the excitement about microcomputers was driven by the appeal of having one's own system, once this experience subsided, the need to connect with other computers or devices became apparent. There would be advantages for individual PCs or workstations to communicate with each other and in particular in work environments to be able to share expensive resources like printers and large storage devices. Further, since PCs had so much more capability than “dumb” terminals used to communicate with mainframe computers, why couldn't they be used as terminal devices? In effect, one personal computer on a desk not only could now serve the needs of a single user but could as needed be connected with other users and other computers in different operating modes.

Connecting personal computers to each other led to the development of local area networks (LANs), and as was typical for the time, equipment vendors drove developments. Macintosh computers from Apple had their own built-in technology, AppleTalk, which made it easy to string telephone cable between Macs and create a LAN. In mid-decade IBM introduced its own Token-ring LAN intended for sharing printers and files with personal computers that were in close proximity. Ethernet was available for some systems from independent vendors. Novell became a leading company in this area offering Netware, a LAN operating system that could run on different network technologies. Netware provided users and programmers with a consistent interface, independent of the hardware

involved. LANs could be distinguished from each other by several characteristics: the topology (arrangement of devices in a ring or line), protocols for sending data (peer-to-peer or client-server), and media (the interconnecting cables that could be twisted pair wire, coaxial cables, or fiber optic cable).

To have the workstation or PC function as a terminal to a mainframe required special software to be written that would emulate the interactions between the PC and the mainframe or host computer. This was a complicated task because these programs had to contend with a variety of different technologies, one simple example being character representations. One of the most popular video terminals at the time was the VT100 from DEC that, like the Macs and PCs, used the ASCII format for representing characters. (ASCII, the American Standard Code for Information Exchange, represents the English characters in 7-bit format or in 8-bit extended format.) IBM, on the other hand, used EBCDIC (Extended Binary-Coded Decimal Interchange Code) on the mainframe computers to which these ASCII devices were connected. As a result, any terminal emulator had to ensure the mainframe received what the user intended.

The use of e-mail accelerated as the decade progressed. More workstations connected to LANs using terminal emulators could take advantage of sending mainframe-based e-mail. However, because of new products that enabled e-mail to be sent without going to the mainframe, the presence of LANs now made it possible to send LAN-based e-mail within a LAN itself.

The facilities for e-mail expanded rapidly after that with the creation of WANs (wide area networks), which were a combination of LANs linked by gateways. The gateways were so designed that they could accommodate networking schemes from different vendors. In short order the e-mail amoebae came into being, and it was possible for mail to be sent and received between all these different sources.

Vendors

As can be gathered from the preceding account of developments in hardware, software, and services, many new vendors entered the market during this decade. Vendors from the 1970s continued to be active; they kept up with the times through new product offerings. There were also some reorganizations and combinations. Sperry Rand became just Sperry early in the decade, and both the Sperry and Burroughs names disappeared when they merged to form Unisys Corporation. HP acquired Apollo before the decade was over. Compaq, formed at the beginning of the decade, acquired Wang Corporation before the decade ended. ETA Systems, formed out of CDC

Supercomputers earlier in the decade, closed out its operations in 1989.

Many new companies, dealing with microcomputer hardware and software, were formed during this decade, notably Compaq, Lotus, Sun Microsystems, NEC, and NeXT. Osborne, an innovator and early entry in the microcomputer market, declared bankruptcy in 1983 and went out of business as did many others, too numerous to name.

Technology Terms, Acronyms, and Buzz Words

Cyberspace—Term coined by William Gibson.

GUI—Mouse-driven graphical user interface popularized by the Macintosh that frees the user from learning the possibly complex command language of the system.

Inkjet printers—Nonimpact printers that operate by using a series of nozzles to deposit extremely small droplets of ink onto paper. There are different technologies to create the spray of droplets of fast-drying ink to produce quality output.

Laser printers—Printers that operate much like (Xerox) copiers but use a laser to create an image of the page on a copier-type drum. The technology then uses fine particles, the toner, which are deposited on the drum and transferred to the paper and heat-treated or fused to adhere to the paper, much like the copiers.

Mainframe—An industry term for a large centralized computer for business data processing for big commercial firms or for large-scale research computing.

Megaflop—One million floating-point operations per second.

Mouse—A device that controls the movement of the pointer or cursor on a computer display screen. Typically the size of a pack of cigarettes and looking like a real mouse with its wire-like tail leading to the computer, it is used by rolling it around a flat, hard surface. The device was pioneered at Xerox in the 1970s with its experimental and commercial systems.

Pagemaker—A page layout language to create publications. The output is postscript to drive a variety of different printers.

Postscript—A page description language whose main purpose was to provide a convenient language to describe images in a device-independent manner so that the same description could be used on postscript-enabled printers without modification. First defined by Adobe Systems, Incorporated, in 1985 and first implemented in the Apple LaserWriter.

Trojan Horse—A program often mistaken as a virus that appears to be a legitimate program but has hidden agendas. Trojan Horses can put up messages, bomb programs, erase information, and do other harm. A distinguishing feature of Trojan Horses is that the damage is localized and they do not contaminate other programs.

Virus—A computer program designed to spread by infecting executable files and then to replicate itself by copying itself into a valid program or attaching itself to other applications and so move to other platforms as information is exchanged between systems. Viruses typically operate without the knowledge or desire of the user of the computer and, depending on their objective, can destroy files or interfere with operations or just spread themselves around without causing any damage.

Word processor—A single-purpose microcomputer linked to a fancy typewriter-like console and CRT screen packaged together and sold as dedicated word processing devices to improve the typing of manuscripts. Typically linked to a high-quality printer for professional-looking results and also including some form of storage, initially “floppy disks,” to retain archival copies of letters and manuscripts.

Worm—An intruder, similar to a virus in that it is self-replicating, that tunnels its way into a computer but remains a self-contained entity. It does not need to be part of another program to spread itself to other systems and, again, may be designed to do some harm to the resident system.

1980 to 1984 at Cornell

The appointment of Kenneth M. King as the first executive-level vice provost (vice president)⁶ for computing at Cornell in 1980 started an exciting and busy period for computing activities at Cornell. VP King greatly increased the computing resources and staff support for instruction, research, and administrative computing with new programs in these three service areas. Microcomputers and/or personal computers increasingly were placed on the campus in cooperative ventures with both Apple and IBM through joint studies, and a retail outlet was started to sell these systems. The use of these computers started to change radically the delivery of computing services for instruction and administration. For research, the significant achievement was that Cornell was chosen by the National Science Foundation (NSF) as one of the five new national supercomputer centers, and the only one partnering with IBM. In 1984, the Production Supercomputing Facility created jointly by the newly founded Cornell Theory Center, IBM, and Cornell Computer Services was the first of the five centers in operation.

⁶ Later in the 1980s, all vice provost positions were relabeled as vice presidents so that this entire level of peer positions would be consistent in name regardless of whether the position reported to the provost or senior vice president. After this, all references to these positions will use the vice president title or the designation VP.

During this decade the former Comstock Hall was renovated and refurbished as the Computing and Communication Center (CCC), the new home for Cornell Computing Services and Cornell Telecommunications. Most of the staff from all CCS divisions, with the exception of the satellite operations, moved to CCC and were together in one building for the first time in 20 years.

Campus Issues and Transitions

The 1980s started out well for the university as a whole. The administration of President Rhodes, Provost Keith Kennedy, and Senior Vice President Herbster had been in place for over two years and was setting an aggressive tone for the university. Perhaps the most pressing issue left over from the 1970s was the very high interest rates, greater than 10 percent, which held back investing in facilities and other capital projects. Nonetheless, it was quite clear that President Rhodes's goal was to put Cornell into the premier ranks of major research universities where it enjoyed such standing in only a few disciplines.

In comparison to the modest tuition increases 10 years earlier, the university was bullish in setting tuition rates in the early 1980s. Tuition increases exceeded 10 percent per year and went as high as 13.6 percent in 1982.⁷ They abated a bit in mid-decade but continued at a high level for the rest of the decade. Such increases were justified by several arguments, but the main one in the early 1980s was the increase in spending on computing. By 1983, the university was reported as enjoying budget "surpluses," which Controller Ostrom was explaining as being the result of the intricacies of the university's use of fund accounting but nevertheless were an indication of "a healthy—but not really wealthy university." Early in the decade several new buildings were planned, and by mid-decade construction activity was occurring all over the campus.

Assessment of Computing Needs by the UCB

In February 1980, under the leadership of Dean Everhart of the Engineering College, the University Computing Board (UCB) wrote a report to the provost assessing the critical computing needs on campus. Taking advantage of the period during the recruitment of new leadership for computing at Cornell, and the potential at the time for rationing of computer resources, the board attempted to assess the immediate and longer-term needs of the Cornell campus. Going back to the 1973 "Computer Planning Policy Statement" the board believed the principles articulated in that document continued to apply in the 1980s. Two critical needs were identified:

- The campus computing facility is now seriously overloaded, and there is an immediate need for increased computing capacity.
- Student Computing has suffered disproportionately from the current overload of the system, and there is also a need to improve the badly outmoded instructional facilities.

The UCB's principal recommendation was

that a Digital Equipment Corporation DecSystem 2060 be ordered for installation during the summer of 1980. This machine has been widely used at other universities to support instructional computing and can meet this critical need without a major development effort.⁸

The UCB deferred any recommendation about increasing IBM-compatible capacity until new leadership was appointed. They also noted that not only were Cornell's expenditures on computing lower than the average of peer institutions (2.5 percent of education and general expenses), but the portion derived from Cornell general funds was also low. They further recommended that Cornell tune its computer acquisition strategy to developments in the marketplace to take advantage of declining costs for hardware.

Ken King Becomes the First Vice President for Computing at Cornell

As noted earlier in this chapter, Kenneth M. King was appointed vice president for computing in July 1980 after an extensive search.^a King's last position before coming to Cornell was as vice chancellor for university systems at the City University of New York, but earlier he had held leadership positions in computing for the City of New York, Columbia University, and in IBM computing laboratories. He was well known in higher-education computing circles. Soon after coming to Cornell, King used the goals articulated by the University Computing Board in the 1980–81 period for developing his own plans:

Goal 1: To provide ready access to modern computing facilities and services for all students who require such access for purposes of: acquiring computing skills and a fundamental understanding of computing necessary to achieve their career objectives and to fulfill their societal roles; developing such skills to a high degree of proficiency as appropriate to their interest and disciplines.

⁷ "Endowed Tuition Will Increase by 13.6 Percent," *Cornell Chronicle*, March 4, 1982.

⁸ "University Computing Board Recommendations to the Provost on Critical Computing Needs," February 22, 1980.

Goal 2: To provide access to essential computing facilities and services to support faculty, students and staff research and extension services, and to achieve and maintain quality consistent with the high academic goals of Cornell University.

Goal 3: To provide the facilities needed for the operation of a modern, efficient administrative support system.

Goal 4: To further the growth of computing awareness throughout the university to the point where faculty, students and staff determine the extent of their involvement by informed choice.

Goal 5: To progress toward the above goals in a fiscally sound manner consistent with overall institutional requirements, and in a fully accountable manner.

King talked about these goals in more familiar terms in the public announcements⁹ he made at the time. For students his priority was to expand computer literacy—the ability to create and understand computer programs:

In the 1980s, every student who graduates from college should have a basic knowledge of computing.

King was proposing a five-year buildup to achieve this goal with expanded equipment, staff, and classroom use. He estimated that while 100 percent of undergraduate Engineering students could be considered computer literate, only about 20 percent of the undergraduates in Cornell's largest college, Arts and Sciences, could "talk" to computers. The same percentage appeared to apply to undergraduates in the College of Agriculture and Life Sciences and the School of Industrial and Labor Relations, while King estimated 10 percent for the School of Hotel Administration and the College of Architecture, Art, and Planning. He further stated that computer literacy at Cornell was "about average" among comparable major universities, although far behind Dartmouth's 90 percent. He planned to expand the number of interactive terminals for students from the 77 then on hand to more than 650 interactive terminals and microcomputers in five years. King predicted that within a short while students would be bringing microcomputers to campus as now they were bringing hand calculators and stereos.

King expected research computing use to increase by 15 to 20 percent per year, and while some special equipment and needs would be funded by research grants, additional university funds would be needed to support this increased use.

⁹ "Upgrading Computer Literacy King's Goal, New Vice Provost Proposes Expanded Program," *Cornell Chronicle*, October 23, 1980.

For administrative computing he proposed incrementally to improve existing systems in a manner that did not exceed the offices' ability to absorb and direct the changes to improving services to students, to enhance planning, and to improve resource allocation.

In summary, King stated:

All the deans recognize the extent to which computer knowledge is necessary for their students. Gradually, all faculty will come to realize the place of computing in the courses they teach. The administration realizes that computing is something Cornell has to excel at. This has the support of the president and the provost. If it didn't, I wouldn't be here. The climate exists for major improvement.

From OCS to CCS

One of VP King's first acts was to change the name of the organization from Office of Computer Services (OCS) to Cornell Computing Services (CCS), a name that stayed for most of the decade. The environment at CCS improved considerably after King came to Cornell. He was immediately liked by all of the staff, not only for his easy-going manner but also for his statements of direction and his commitment to improving the state of computing at Cornell. The morale of the staff improved. During the time he spent on campus and his conversations with a large number of concerned individuals on campus before accepting the position, he was able to get a measure of the need for additional funding. With his acceptance of the position the budget of OCS went up \$2 million a year, from \$4.5 to \$6.5 million, in addition to several \$600,000+ funding allocations to address onetime expenses for equipment and other such needs. In early planning for the 1981–82 budget Provost Kennedy was projecting a tuition increase of 16.7 percent, one quarter of which was a result of increased costs of computing. Later on this increase was reduced to 13.2 percent, which was still a substantial increase for a single year. More appreciated by the staff were salary increases of 7 to 10 percent during this period.

CCS and New Campus Initiatives

More IBM Computers Installed: CornellA, CornellB, CornellC

The availability of increased funds spurred a flurry of activity on several different fronts. The computer room at Langmuir Lab was expanded into space previously occupied by card processing equipment, which was now out of use, to accommodate the addition

of two IBM 4341 computers. One was dedicated to student computing and the other to administrative computing as a test machine; each 4341 was rated at having about one-third the computing power of the 370/168. The extended storage system on the 370/168 was replaced with a faster and less expensive Intel system. Most of the peripherals, such as disks, tapes, and printers, were purchased by taking advantage of purchase credits that had accumulated. It was at this time that the IBM machines were named CornellA, B, and C for student, administrative development, and research/administrative production computing, respectively. (CornellC, at the time the IBM 370/168, was shared by these two large users.)

CornellB served as the development platform for the new administrative applications for only a short while. With the various shifts in hardware that were taking place, the IBM 4341 CornellB proved too small for serious work and was thus removed, leaving only CornellA and CornellC to support the three different kinds of work. In 1983 a new IBM 3081D computer was installed at a cost of \$2.5 million and became the new CornellC, replacing the 370/168, which was sold to Tulane University.^b

The installation of the 3081 was accompanied by several other significant changes in the system software. The MVT/HASP batch system was replaced with MVS, ending an era at Cornell that had started in 1969 when these systems were first installed. While that change enhanced administrative use of the machine, it required considerable changes on the part of users to learn how to get along without features that had been customized for Cornell and use the new JCL (Job Control Language). As part of these software changes, SCMS was removed from CornellC in line with making CornellA the machine for student use. All administrative computing was moved to this new CornellC and remained there throughout the various upgrades to this system. The 3081D was expected to provide three times the computing power and cost \$400,000 less per year to operate than the 370/168.

Adabas Acquired for Business Systems

Within months of coming to Cornell, VP King supported the recommendation to acquire Adabas, a new and powerful database management system, for developing all new administrative computing applications. Tom Dimock and others had investigated the different systems available and concluded that Adabas was the best choice for Cornell.¹⁰ An attractive feature of Adabas was that it came with NATURAL, a programming language that facilitated the creation of

application programs. NATURAL, a 4th Generation Programming Language (4GL), featured the support of IBM 3270 terminals and Adabas access built into the syntax.

DEC20 Installed

The DEC20 was in the last phases of installation in G20 Uris Hall by the time King arrived in the summer of 1980. He enthusiastically supported the availability of the system for instructional computing and presided over the formal opening of operations in October 1980. To encourage use and experimentation with the DEC20, announcements of its availability emphasized that there would be no charges for its use until the spring semester. This prompted experimental use by students and faculty to consider its use for their classes as well as their research programs.

A DEC20 advisory committee was formed to recommend policies for its use. Members of the committee were Alan Demers, Computer Science; Ron Furry, Agricultural Engineering; Tony Ingraffea, Civil Engineering; Bob McGinnis, Sociology; and Doug Reece, chairman, from Business and Public Administration.

Credit for the successful installation is owed to a cross-organizational team from Computer Services: Alec Grimison, Cecilia Cowles, George Cameron, Richard Alexander, and Alan Personius. They attended not only to the hardware and software installation but to providing user documentation and technical reference manuals and developing the materials to announce and publicize the availability of this new computer. Larry Fresinski was appointed interim facility manager before the installation.

CCS Organization and Organization Chart

At the time King came to Cornell, the headquarters offices for CCS were on the fourth floor of Uris Hall, although some offices remained at Langmuir. After the move to the ground floor of Uris in 1974 with a terminal facility and modest staff space, the space for computing staff in Uris had expanded considerably. By 1980 CCS occupied about half the office space on the ground floor and had a suite of offices in a large central area on the fourth floor. Initially King's office was in this fourth-floor Uris space, but in keeping with his vice presidential role an office suite was developed on the second floor of Day Hall, on the west side overlooking Sage Chapel. Soon after this space was available, King started his soon-to-be well known Friday afternoon Wassails.^c Everyone—staff, users, faculty, and administrators from around campus—was invited to enjoy this end-of-week social. Many people took advantage of this hospitality to meet each other,

¹⁰ Tom Dimock and Jim Quiggle, "DBMS Report," October 5, 1980.

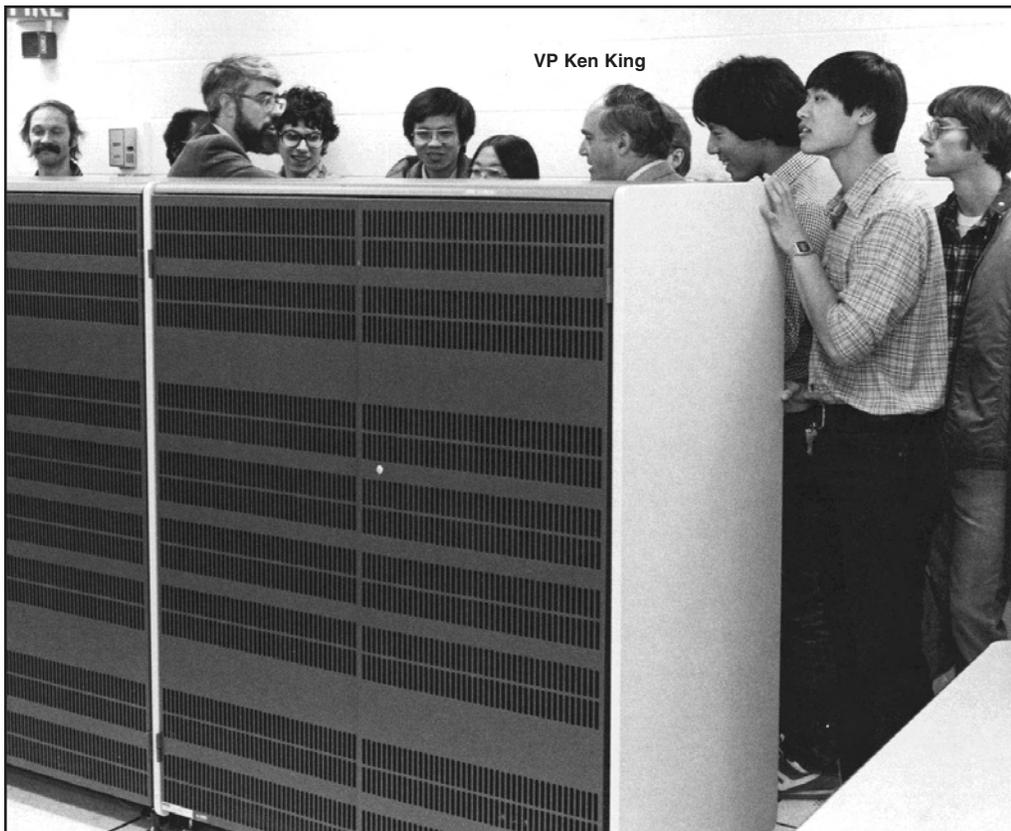


Figure 1. VP Ken King at the DEC20 opening ceremonies

exchange views, promote their agendas, and catch up on the latest jokes. King was an outstanding raconteur. His ability to recall jokes that captured the spirit of the moment is legendary, at Cornell and beyond.

At the time King came to Cornell his senior staff was as follows. Rudan was named assistant vice provost for computing and was mainly responsible for transition issues and the business and personnel functions assisted by Blackmun. All previous directors and assistant directors now reported directly to King until he decided on future steps, although routine operations continued under the previous management grouping. Accordingly, Van Houweling continued to be responsible for academic computing with assistant directors Steve Worona and Doug Gale. Gale was fairly new to Cornell, arriving just months before King as the assistant director responsible for Decentralized Academic Computing Support (DACS).

As discussed at the end of the 1970s chapter, steps had been taken to position Cornell at the forefront of using the emerging microcomputer technology for which OCS had formed this subunit of Academic Computing. In early 1980, Gale was recruited to be the head of DACS and he continued in this responsibility. Other senior staff were Dick Cogger, head of systems programming and network support; Dave Pulleyn, head of Langmuir operations and production

control; Ed Hollenbeck, director of Administrative Programming Services (APS). Lynne Personius and Libby Gruppuso were assistant directors in APS.

In early 1981, King created a new division, Communications and Text Systems, consisting of three operating units: network development, text systems, and data communication services. Cogger was appointed director while Cecilia Cowles was given the responsibility for text systems. Within a short time, Alan Personius, who had been central systems operations manager, joined the group as assistant director of data communication services. This was an experienced, solid, and versatile group and, given direction and support from King, was responsible for carrying out the early changes and improvements in facilities and services until he could build his own organization and management team.

That opportunity came in 1981 when Van Houweling^d left to become vice provost for computing and planning at Carnegie-Mellon University in Pittsburgh. Until a new director could be recruited, first Alec Grimison^e then Doug Gale served in an acting capacity during this interim period. The recruiting effort culminated with the hiring of Gordon L. Galloway as director of academic computing in 1982. Before coming to Cornell, Galloway had been a professor of chemistry at Dennison University, an active

and involved user of computing services, and a member of that university's computing board.

Under Galloway a new organization and management was put into place. Agelia Velleman became director of user services, Donna Bergmark directed software support, Larry Fresinski was manager of the DEC20 system, and Cecilia Cowles was responsible for publications. Doug Gale continued as head of Distributed Computing Services (the "Academic" was dropped so DACS became DCS to indicate a broader campus mission) in Academic Computing, until he left in 1984 to become director of computing at the University of Nebraska.^f Tom Hughes, the senior consultant in DCS, replaced Gale at that time. Blackmun^g left in 1982 to take the position of director of computing at the University of North Carolina at Charlotte. He was in time replaced by Tom Cardman in 1983. When Ed Hollenbeck left in 1981, recruiting started for a new director of administrative computing.

Russell S. Vaught, then director of computing at SUNY Binghamton, was appointed director of administrative computing in 1983. His senior staff was Lynne Personius, Libby Gruppuso, and Tom Dimock. Following Vaught's appointment, Rudan was appointed director of operations^h with Pulleyn as head of central operations and production control, Alan Personius as head of network services, and Bob Cowles

as head of systems programming reporting to Rudan. In 1984, Production Control was made a separate unit headed by Jim Doolittle. Cogger and Worona took staff positions reporting to King with Cogger focusing on network systems development and Worona on special projects.

The CCS organization chart from 1983 to 1987 is shown in Table 1.

The University Computing Board continued to provide policy direction to CCS and oversight of computing on the campus. Dean Everhart of Engineering continued as the chairman until 1983 when J. Robert Cooke succeeded him. In the early 1980s members of the UCB were Goeffrey Chester, Juris Hartmanis, Tim Mount, John (Jack) Ostrom, Don Randel, Doug Reece, Paul Velleman, and Richard White.

A Building on Campus for CCS and Telecommunications Study

Before 1984, in this first three hectic and busy years of King's tenure at Cornell, several initiatives were started that were to have a profound impact on the development of new computing services for the future. First was a grant of \$2 million from the J. N. Pew Jr. Charitable Trust for construction of a new computing center on campus.¹¹

Table 1. Cornell Computer Services Organization Chart, 1983 to 1987

Academic Computing	Gordon Galloway
• User Services	Agelia Velleman
• Distributed Computing Services	Doug Gale (Tom Hughes)
• DEC20 Services	Larry Fresinski
• Software Support	Donna Bergmark
• Publications and Special Assignments	Cecilia Cowles
Administrative Computing	Russ Vaught
• Payroll and Related Services	Libby Gruppuso
• Student and Related Services	Lynne Personius
• Technology Database Services	Tom Dimock
• Office Support Systems	Mark Mara
Operations Services	John Rudan
• Computer Operations	Dave Pulleyn
• Production Control (1984)	Jim Doolittle
• Network Services/Microcomputer Sales	Alan Personius
• Systems Programming	Bob Cowles
Vice President's Office	
• Network Planning	Dick Cogger
• Information Services	Steve Worona
• Business Office/Personnel	Tom Cardman

¹¹ "\$3 Million Pew Gift to Aid Computing," *Cornell Chronicle*, February 5, 1981. (\$1 million for Computer-Aided Design Instructional Facility)

The second was the authorization to spend \$100,000 for a telecommunications study and engineering design services to plan a new and independent telephone system for the Ithaca campus to replace the 10,000-line system being operated by the New York Telephone Company. Network Analysis Corporation was hired as consultants. This new system would include totally wiring the campus with Cornell-owned optical fiber and copper wire and connecting this to a commercially available switch that would provide both campus and external worldwide telephone connections and services. One of the key benefits of this new arrangement was to augment the data communications network being developed by CCS.

Supercomputers; Wilson Wins Nobel Prize

Another key initiative was on the research front, where Ken Wilson was promoting the need for supercomputing for research projects and developing a vision of how this would be accomplished. As one of the early participants and supporters of the installation of the Floating Point Systems array processor in the late 1970s, Wilson was able to envision how such devices could revolutionize the approach to large-scale computation.

In 1982 Wilson was awarded the Nobel Prize for Physics. Winning this prestigious award gave him even further platforms to advance his vision of supercomputing. He promoted the need for supercomputers to improve the models of phenomena and to complete calculations in a reasonable amount of time. In some cases, as in simulations, supercomputers could complete the calculations in almost real time. He postulated that a new approach involving thousands of smaller computers might be the answer to achieving parallel operations and hence both speed and price advantages beyond that available from new and faster chip technology.

Explosion of Computers on Campus

At the other end of the spectrum, microcomputers were becoming increasingly evident on campus, both in CCS and almost everywhere else the technology attracted interest. There were Apple II and Apple III systems along with systems from Commodore, Osborne, and Amiga, Tandy (Radio Shack), and others. It was also a time of “do-it-yourself” activity where the extremely active and interested “techno” types were building their own systems from parts ordered from a variety of outlets. These were exciting yet stressful times; computers now ranged from the desktop models in labs and offices to the most advanced supercomputers in large computer rooms, and new uses were being found every day in every discipline

and field. As one headline in the *Cornell Chronicle*¹² proclaimed: there is a “Virtual ‘Explosion’ in Use of Computers on Campus.” How true that was, and it was not virtual—it was *real*!

New Publications and Outlets Developed

New and innovative formats were developed by the Publications Group during the early 1980s to open up and increase the flow of information about technology. A computer page was introduced in the *Cornell Chronicle* so that readers of this weekly publication would be exposed to educational and informational computing activities on campus. Besides the continuing weekly *CCS Bulletin* (yellow, single sheet), additional focused publications were developed to reach certain audiences. The *Microcomputer Newsletter* dealt with sales and service issues; *Network News* covered word processing and text processing applications; *nibbles* dealt with microcomputer technology and developments; and *observations* covered statistical software. The *CCS Bulletin* changed to include fewer items about mainframes and more items about microcomputers and network technologies. About this time, CCS put into place a “hands-on” training facility in G25 Stimson Hall so that 20 people could be seated at Macintosh computers and learn their use by following an instructor who guided the group through elementary and more advanced operations.

CCS Financial Restructuring/Cost Recovery

A legacy of the 1970s was the continuing discussion about how the central computing organization should recover its costs and in particular how instructional computing services should become more like the university library. The cost recovery issue was seriously discussed starting in 1978–79 when the UCB and consultants began to be concerned with the predicted future in which hardware costs would rapidly decline while other costs, in particular staff salaries, would continue to increase. The fundamental problem was that the OCS/CCS financial model was still a full-cost recovering model that included a large number of overhead items in the cost base for computer rates. These overhead items included most of the management staff, all the campus user services, consulting and operations staff and their support costs, and all the campus terminal equipment and operating costs. In addition were the large and more direct computer costs such as hardware and software rentals, maintenance, power, and operations and systems staff, plus depreciation of the purchased components. Central computing rates were

¹² “Virtual ‘Explosion’ in Use of Computers on Campus,” *Cornell Chronicle*, May 12, 1983.

thus deemed to be too high and not competitive with alternate services available with minicomputer installations in labs or departments, so computer users avoided the services of central computing.

In pursuing the so-called library model of computing, of easy and simple access to computing, with some restrictions but without the sense of rigid clearance and tight rationing, OCS had made special “subsidized” cost arrangements for SCMS. Separate funding arrangements had also been made for DACS/DCS and early in the 1980s for the Teraks. Interested in further simplifications, King proposed charging colleges/academic units for IT (fast batch) on the basis of student enrollment rather than prorated use. This was put partially into practice in 1982.

Between 1980 and 1982 several reports were written by a subcommittee of the UCB appointed to study the issue of CCS rates, charges, and financing. A report authored by Controller Ostrom made the important point that when the rate base is increasingly dominated by overhead items “computing cycles from central facilities will not be competitive with computing cycles from decentralized facilities even in cases where centrally supplied cycles are in fact less expensive to produce.” The report recommended a number of different solutions.

Another subcommittee (On Financing Computing) issued a report authored by Chester in mid-1982. The committee recommended that only the direct costs of providing a computing service be charged to the users, arguing that this method would be inherently fairer to users and would lead to a significant reduction in computing rates. The committee argued that when all CCS income was generated on the basis of total costs, decisions were strongly influenced by how income would be affected and not by what was the best choice. They believed that it was a matter of rearranging the cash flows, which by themselves would not require an increase in budget. Echoing an earlier point, one of the recommendations made was that all computing should be purchased and budgeted in “fungible dollars” (university funds that could be used for any service at any agency inside or outside Cornell), but with an appropriate transition strategy.

The discussions continued without reaching a formal conclusion, although as the years rolled by more and more CCS resources were going into nonchargeable university-wide support services. These were services such as consulting and teaching, planning for new services and infrastructure, along with new business systems technology and infrastructure. In addition, there was increased support for microcomputers and the research users exploiting the Floating Point System array processors. By 1983 the charges for the central computer were increasingly considered to be out of

line with the alternative services from minicomputers and microcomputers, whose power, capability, and software were constantly improving. Central computing rates continued to be tweaked to be more in balance with capacity/cost/income as hardware changes were made, but also as a result of the annual rate review and computation demanded by federal government auditors.

The Glob Model for Cost Recovery

The last UCB subcommittee to deal with this issue finally came to develop a two-tiered model of indirect and direct costs in 1983. Members of this group were Controller Ostrom, who headed the group, and Juris Hartmanis from Computer Science, Paul Velleman from ILR, Mike Whelan from the Budget Office, and Rudan. The two-tiered model categorized CCS activities and the associated indirect costs into four cost pools: those that could be attributed to support of university-wide global activities, to research, instruction, and administrative computing activities. After redistributing the global cost to the other three costs, in proportion to their sum of costs, the three remaining costs were distributed to research overhead, instructional activity, and administrative overhead through the Responsibility Center Analysis (RCA) mechanism already in place for university cost-recovery policies and practices. The RCA distributed these total costs to colleges and administrative departments on the basis of mainframe use and to colleges and research centers on the basis of research spending. These indirect costs together became known as the Glob (Ostrom being prone to simple and descriptive labels), that is, indistinguishable as to direct allocation.

The other half of the model took the now smaller direct costs of providing services and used these to develop computing and personal service rates that would be charged for services. Most rates were substantially reduced once the Glob model was implemented. This use of the word “Glob” unfortunately underplayed the advance made to bring rates and cost allocations into better alignment with reality but caught the attention of the executive staff decision makers who accepted the model. Because of delays of one kind and another, new and significantly reduced rates did not go into effect until late 1986. It should be noted, however, that this model still continued the hard-soft (nonfungible) money model where CCS was still a fee-for-service operation, but now most of the billing was against accounts supported by university general-purpose funds and restricted to use for OCS services (nonfungible) as the amount of research and external (hard) funds had dwindled considerably.

CSARS (Computer Services Accounting and Reporting System)

The CCS billing system itself was designed and originally deployed in the late 1960s and had been carried forward with enhancements into the 1980s. The system was written in PL/I, a language that was growing increasingly obsolete and maintained by programmers from the CCS Business Office. Once Bob Blackmun left in 1982, the intimate knowledge of this old system left with him. VP King believed the system should be upgraded to use contemporary Adabas technology being used for all other university systems, which would then allow the programming staff in Administrative Computing to look after the system.

In 1983–84, a project was organized under the direction of Jim Doolittle to design and implement this new system called CSARS—Computer Services Accounting and Reporting System. With considerable effort from the project team of Alan Hubberman, Gary Buhmaster, George Cameron, Andy Hanushevsky, and Paul Zarnowski the system went into production on July 1, 1985. The installation was successful in meeting the requirements but brought out a serious problem. The sheer volume of transactions was far greater than Adabas had ever encountered, and so the new billing system ran more slowly than the old system. Even worse, when running, it degraded performance for other jobs. After some fine-tuning and changing the run schedule, these problems were overcome and the system was used well into the 1990s. Although the need for direct billing to accumulate real dollar income continued to decline, the data collected continued to be important for cost allocations and capacity and use studies.

Instructional Computing

Improving services to instructional computing was a top priority for CCS, so the building of new campus facilities, increasing the number of interactive computing terminals and Terak microcomputers, and providing additional computing capacity followed from that priority. In keeping with VP King's goal to increase the number of student-accessible terminals from 77 to 650 in five years, within a year, in the fall of 1981 the number of interactive computing terminals doubled. Fifteen new terminals and two Terak computers were installed in Baker Lab.

In 1981 the students and instructors had a choice of using the DEC 2060 or the IBM 4341 (CornellA), both available from interactive computing terminals at campus locations or by modem from other campus locations or off campus. Terak microcomputers were available at several campus locations for students taking courses in Computer Science.

In the spring of 1982 the number of public terminals had gone from 32 to 118 and the number of Teraks from 24 to 47 in the previous year.¹³ At this time the popular terminals were VT100 video terminals and the DECwriter printing terminals. New terminal facilities were opened in Martha Van Rensselaer Hall and in the basement of Carpenter Hall where new Teraks were installed in addition to those moved from Upson Hall. The Martha Van Rensselaer facility was a joint effort between CCS and CISER (Cornell Institute for Social and Economic Research).

Instant Turnaround (IT) Fast Batch and the DEC20

Free IT continued to be used by everyone on campus. CornellA instead of CornellC now hosted the service. The use of SCMS (Student CMS) continued to grow as the number of access terminals grew. Teraks were being used by students in CS100 and CS101, introductory courses in Computer Science, and by students in courses in the Colleges of Arts and Sciences, Engineering, and Agriculture and Life Sciences. As a result over 1,100 students moved from the use of IT batch services to the use of Teraks for instructional purposes.

Neither students nor faculty immediately flocked to the DEC20, as some had anticipated, because this required a whole new technology to be integrated into the teaching materials for each class. The Department of Computer Science did not switch any of its courses to this system. Even though Douglas Reece from the Business School (one of the original supporters for acquiring this system) left Cornell within a year, the school continued its strong commitment to the system by covering about one-third of the operating costs of the system. The DEC20 was used by students who were interested in the services provided and by courses where the instruction converted the class materials. In all cases they received strong support from Larry Fresinski, who took on the responsibility for keeping its software and systems current, supervising its operations from the very beginning, and promoting its use across the campus.

The Use of SCRIPT and the Development of CUPAPER and CUTHESIS

Other new and innovative uses of computing were introduced in instruction. One of the earliest was the use of SCRIPT for undergraduate writing workshops. SCRIPT was a text-processing language developed at the University of Waterloo in Ontario, Canada, for use on IBM computers. Between 1981 and 1983,

¹³ "Student Computing Terminals Tripled in Past Year," *Cornell Chronicle*, March 4, 1982.

Cecilia Cowles and Debra MacInnes, CCS staff, developed improvements using macro-commands to make this computer program much easier for students using the computer for the first time. The use of SCRIPT was viewed as a success because it resulted in better-crafted compositions by providing editing commands that could be used to change, move, insert, and rearrange text instead of retyping all or part of a paper. Shortly after SCRIPT started to be used in classrooms, the macro-enhanced system was packaged together as CUPAPER to provide features to facilitate the writing of term papers.

Key to the success of SCRIPT and to general time-shared interactive use of the computer was the increasing use of CRT display terminals and the movement away from punched card input. During the years 1981 to 1983, for example, many announcements were made concerning the removal of keypunch machines from campus locations as a means of warning users to consider moving to online services. Time-shared use allowed the person directly to see the data entered, which could be information or commands, and have the computer respond to an issue at the completion of entering a line and pressing the “enter” or “return” key.

CUPAPER was extended in 1983–84 to create CUTHESIS, which allowed graduate students to prepare their theses using the mainframe systems for text input and editing to produce output that satisfied the thesis rules of the Graduate School at Cornell. CUTHESIS accelerated the use of “automation” tools for text preparation.

The availability of laser printers at several campus terminal locations, and at the central site after 1984, contributed to this growth because high-quality text output could be produced. When laser printers were introduced on campus, charges were \$0.15 per page printed, and various schemes were used to control and pay for their use.

Plato Service Ends

With all the other services being offered for instructional use, interest in the use of PLATO, the computer-aided instructional system, continued to decline. The service was shut down in June 1982, and instructors were encouraged to contact CCS for alternative services.

VMBatch Developed; Free Distribution Accounts

A new service was introduced in the summer of 1982: VMBatch. Developed locally by the Systems Programming staff, VMBatch provided simpler and more effective commands to initiate batch jobs on either of the IBM systems. If a program took a “long” time to execute, students could submit the job to

VMBatch during the evening or weekend hours when the job would be completed. One additional advantage of VMBatch was the lower costs. Jobs were charged only 80 percent of the CMS rate, and because there was no terminal connect charge (typically \$1.80/hour) job costs could be reduced by as much as 40 percent.

Also, by this time, computer accounts for “free” access to the different CCS machines and services were called “Free Distribution Accounts” to continue the notion of no-cost access started back in the 1970s. Accounts, centrally supported out of university allocation funds, were initially set at \$100 and refreshed to that amount as requested. Later in the decade, these were simply called Distribution Accounts because they were typically “distributed” to new students after registration and remained active and were refreshed with funds during their entire stay at Cornell.

Research Computing

Several new initiatives were started in the early 1980s to foster the increased use of CCS computing resources in support of research programs at Cornell. These steps were in keeping with VP King’s stated goal of improving services and increasing access for these users. With additional money being released into departments as a result of the increased university allocation to CCS, researchers who previously had very limited funds were now able to use CCS computing resources.

The number of campus computers dedicated to research projects or research centers continued to grow rapidly. Data capture was made increasingly easy by the ability to connect different instruments directly to computers, which then filtered, edited, and stored the data on disk or tape.

Windfall Computing

In early 1981 VP King came up with an innovative program to make idle weekend computer capacity on CornellC available to researchers who otherwise did not have access to funding. The program was named Windfall Computing. On a quarterly basis, colleges and research centers were given an allocation of funds that could be used to run jobs at the lowest computer processing priority (4) during weekends. The funds could not be used for any long-term expenditure, such as the rental of online storage, but strictly for executing jobs. Because unexpended funds expired at the end of each quarter, individual users had to appeal to their local Windfall coordinator for the next quarterly round of funds. The Windfall program was a great success, allowing researchers who previously had no access to the computer to do serious work by using spare weekend capacity. It was another case of VP

King effectively using increased university funding to move toward achieving his stated goal of increasing research computing 15 to 20 percent per year.

Statistical Computing

Reflecting CCS's interest in data analysis tools, in 1982, the Statistical Computing User's Forum (SCOF) was organized by Agelia Velleman as the first of other such new users groups on campus. Membership was open to anyone interested in statistical computing activities. The purpose was to establish a forum conducive to the exchange of ideas among researchers who had projects requiring statistical computing applications. The first seminars sponsored by this group included such topics as using the Terak microcomputer to enter and edit data and later to transfer these data to the IBM systems for analysis.

Large-Scale Computing—More Array Processors Acquired

The major emphasis of CCS research activity centered on exploiting the Floating Point Systems (FPS) array processors (APs) for scientific computation in physics and chemistry. The FPS 190-L, installed in 1978, proved to be an overwhelming success in the early 1980s. It enabled researchers to consider much larger theoretical and simulation models for their research and made such computations very inexpensive and convenient to schedule as the AP relieved the 370/168 of doing this extensive work. A second array processor, the FPS-164, was installed in 1982. It was faster, provided 15-digit precision, and had the capacity to hold up to one-half million floating-point

numbers and many thousands of Fortran statements. Obtained at a cost of \$300,000, this unit was financed by the NSF and a consortium of university researchers including Ken Wilson, the 1982 Nobel laureate.

In late 1983, Wilson was one of the organizers of a conference, *Forefronts of Large-Scale Computational Problems*. The focus of this conference, which brought together computer researchers and planners and research users in university and government laboratories, was to emphasize partnerships of industry, government, and universities to address the need for such supercomputers.

A New Computer Billing Category: Flatrate

Extending the 1970s tradition of lower rates for use of CornellC during university holidays, in 1983 CCS created a new billing category for weekday evenings, weekends, and holidays: "Flatrate." Flatrate was introduced in October with a fixed charge of \$2.50 per hour for CPU, input/output, and connect time, plus additional charges for using public terminals or outside network services such as Tymnet. All other charges for services such as printing were charged at priority 4 rates. This practice continued throughout the decade.

One of the major beneficiaries of the Windfall program and the Flatrate charging schema was CISER (Cornell Institute for Social and Economic Research), formed in 1982 as the cooperative effort of several colleges to create a center for researchers interested in social and economic data and statistical and other analysis tools. Computing resources were critical to CISER's mission. These two schemes of charging low rates for off-shift use of CCS computer resources effectively provided startup funds for CISER members. In

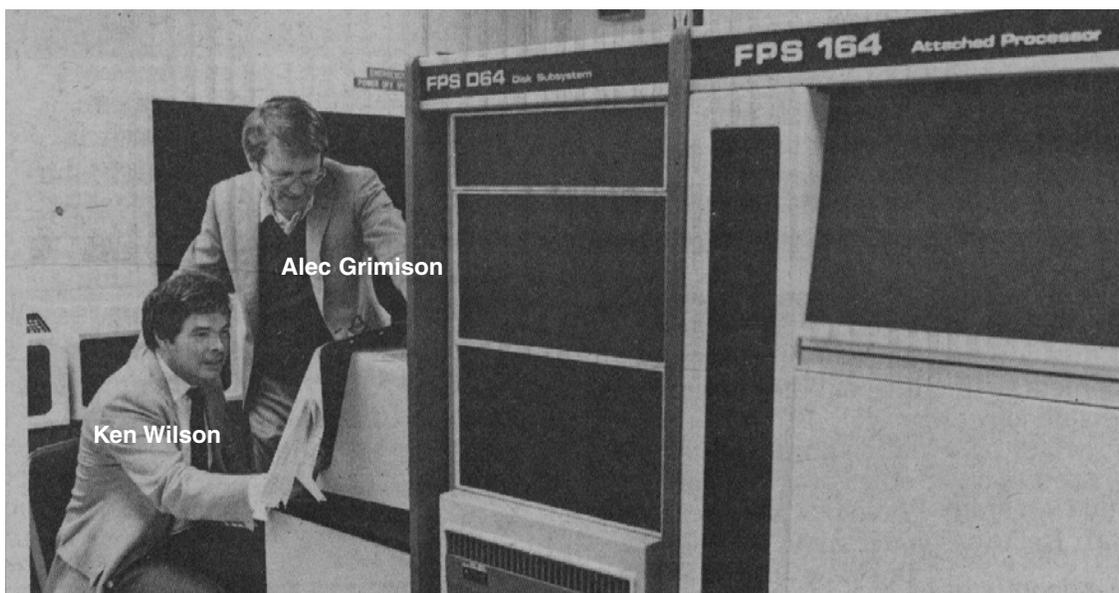


Figure 2. Floating Point Systems, FPS164 Array Processor (Ken Wilson and Alec Grimison)

combination with funds from CISER's own resources and the research grants of its members, CISER was able to get off to a fast start in the acquisition of data files and the use of the computers for analysis.

Business Systems

In conjunction with all the improvements taking place in computing facilities and in instructional and research computing, the activity in administrative computing was equally busy at the start of the decade. On one level, new policies were introduced to provide a framework for developing new systems and studies were started to determine requirements for new systems. On another level, a new environment was being built for new systems to use the latest in database technology for new online systems.

Policy Statements on Administrative Systems

In the summer of 1981, VP King issued an important memo, "Status of Administrative Systems Development," which was distributed widely around campus. The memo outlined a plan for systems development and provided status reports on the key systems in development. The most important introductory statement made was that efforts would be directed to

store and maintain the university's basic data in an electronic filing cabinet where it will be accessible from terminals in offices around the university (provided that the office is authorized to see the data). The impact of this change will be to reduce the manual effort of maintaining paper files with redundant and inconsistent data in many places, to improve communication between offices, and to make information retrieval easier.

This statement specified that data were a university resource that had to be not only well managed but also shared. The concept of a central "electronic filing cabinet" contrasted with data on paper then stored all over the university in filing cabinets. King's memo went on to outline the new team-based processes that would get new systems designed, approved, and implemented. The processes called for a cooperative effort of users and technologists to produce an acceptable design that was approved along the way by appropriate agents and committees.

In that letter King also enunciated an important principle that user departments would not be billed for design and programming services until they were realizing the benefits from new systems. This was done to overcome some of the resistance for creating new systems without having the funds on hand and the advantages realized.

Administrative Priorities Committee

One of the new committees announced in King's memo was the Administrative Priorities Committee consisting of Senior Vice President Herbster, Vice Provost James Spencer, Controller Ostrom, and King. Their responsibility was to allocate resources to competing projects.

Application Transfer Team Study—Registrar's Office

One major project, announced in early 1980, was a new student record system. The first phase, to study the functions of the various college registrars, started in 1980 as a joint project with a study team of university staff and IBM consultants using Application Transfer Teams. This was IBM's term for using its staff experienced in higher-education student systems to achieve good systems designs at institutions like Cornell. The team consisted of two IBM staff, three college registrars, University Registrar Eleanor Rice, and Jim Quiggle from Administrative Computing. Although it was expected that it would be several years before a new system would be in place, Vice Provost Larry Palmer stated that it would take time to study the problems and review the options for improving services and record keeping. One clear intention was to change from the present batch system to a more interactive and responsive system for students, faculty, and staff.

Adabas Installed; Database Administration Started

While the decision had been made to use Adabas as the primary database management system, it was taking time to develop staff expertise and production facilities. Cornell was the first installation that used Adabas with CMS as the time-sharing monitor for executing online, interactive transactions. The vendor of Adabas, Software AG, based in Germany, was interested in this new Adabas environment as all previous use had been executed under MVS using IBM 3270 interactive terminals connected with IBM's proprietary SNA networking technology. Tom Dimock was leading this effort for Cornell, building the new infrastructure for administrative systems.

Tom Marciniak was appointed as the first database administrator at Cornell. This happened during the time Rudan did a short stint as acting director of administrative computing, which in the early 1980s had been brought back to the previous level of more than 30 staff, after the cutbacks in the late 1970s. The increase in staff was in line with all the new systems development and support functions under way.

New Systems Development: Student, Payroll, Financial Accounting

The first new business system in the development pipeline was the student system, which included student registration, admissions, financial aid, and the bursar's office. After Eleanor Rice left her position, Keith Ickes was appointed registrar. Using the conceptual design produced by the Application Transfer Team working with IBM, Ickes engaged Peat, Marwick, and Mitchell to produce a refined design and implementation schedule, and work started on the project in early 1981 with some elements to be available later in the 1981–82 academic year. This was named the Student Registration and Records System (SRRS).

As it turned out, the first implementation of online business systems was done in the Admissions Office, where there was a more pressing (and simpler) need to deal with incoming admissions requests. The then-current process of using mark sense forms¹ filled out by prospective students was error prone and not the public face Cornell wished to present. New forms were designed for direct online operator entry of student admissions information. This not only improved the quality of the information but formed the basis for making it more accessible across the university. The system was designed to be consistent with the future larger system.

In early 1982, phase two of the student system was started with a study to review and develop “a secure financial aid system” using the same methodology. Later that year a third phase to review and develop a new admissions system was started, to complete the earlier work in admissions. Some of the new systems started to show up in production after the studies were completed in 1982. New registration procedures went into effect in the fall semester of 1982. When Russ Vaught became director of administrative computing in 1983 he conducted a review of the administrative systems activity since VP King's arrival and noted the many achievements accomplished by the staff in Administrative Computing without their having any continuous leadership. Besides the accomplishments mentioned above he added the success in implementing a new Public Affairs System with VIP designations, gift history, and other features; a new faculty personnel database; and the activity under way for redesigning the Financial Accounting system.¹³

In mid-October 1981, direct deposit of paychecks was put into effect (after being rejected in 1979 owing to the loss of interest earned on university funds on deposit while waiting for checks to clear). With this new system, paycheck amounts would be deposited directly to designated personal accounts in most of the local banks and financial institutions.

In 1982, staff from the Controller's Office and the Statutory Accounting Office formed a project team to upgrade the university's financial accounting system. Members of the project team were Controller Ostrom as chair with Stew Comber from Statutory Finance; Lee Cartmill, director of endowed accounting; Mary Jo Maydew from the Controller's Office; and Louise (JR) Schulden from Administrative Computing.¹⁵ According to Bob Mack, who was responsible for programming support for this system in the Controller's Office, this was to be a whole new system written in Adabas and using the most contemporary technologies. It was to have the capability of bringing about major changes, such as a whole new account numbering schema and structure. However, after a number of workshops in late 1983 and early 1984 and discussions with users of the system, it was decided to change only the externals by changing some of the reports and introducing new monthly reports requested by users. Maydew and Cartmill played a principal role in conducting workshops for training users and dealing with feedback.

Installation of IBM 3270 Terminals

At about this same time, CCS obtained a loan of \$300,000 to purchase IBM 3270 terminals and communication controllers to start a leasing program and make these units more broadly available to administrative systems users. As a result of standardizing on 3270s, all new business systems could use this technology as a design point for creating screens and to take advantage of the powerful built-in function keys.

Improved Relations with CUMC

Although he did not make any major pronouncements about the issue, early in his tenure VP King took steps to improve the working relationship between CCS and CUMC (the Cornell University Medical Center in New York City), and to treat them as another operating unit of the university. An agreement was made in 1980 to have CCS Production Control services take responsibility for ensuring the routine processing of some of their administrative systems and for central operations to keep the college up to date on current and proposed activities. One consequence of this new operating arrangement was that CUMC researchers were eligible for Windfall Computing to support their computer needs and were able to piggyback on any favorable volume or legal discounts extended to the

¹⁴ Russell S. Vaught, “Review of Administrative Computing,” September 16, 1983.

¹⁵ J. S. Ostrom, “Upgrade of Financial Accounting System,” memorandum to executive officers, deans, and directors, October 13, 1982.

Ithaca campus. CUMC continued to use CCS computers until late in the decade when it obtained its own machine. When microcomputers rose to prominence, CUMC was able to obtain these computers at the same prices as were available on campus.

Text Processing and Information Services

Office Automation Group Formed

One important change taking place in the use of computers at the beginning of the decade was the increasing noncomputational work. Evidence of this change was seen in the use of SCRIPT and its derivatives CUPAPER and CUTHESES, all of which enabled the use of the mainframe for text processing. The next major advance in text processing that came to the campus was the availability of stand-alone self-contained “word processors” that revolutionized the whole office process of preparing and storing typewritten materials. Word processing systems from Wang Corporation, the IBM Displaywriter, and a host of others were installed in the early 1980s. Some of the more adventuresome offices on campus, especially those interested in high quality and repetitive output, installed these machines as soon as they came on the market.

To deal with these different “text-processing” systems as well as with their increasing use in offices interested in obtaining data from central business systems, the Text Processing group was reorganized as the Office Automation Group with Cecilia Cowles as the manager. Other members of the group were Cynthia Frazier and Mike Oltz.

Word Processing Systems Support Defined

One of the group’s early studies was on the capabilities of different word processors so they could advise and support users who wanted to push the limits of the systems, such as sending or receiving data from related mainframe-based business systems. That study concluded that the Xerox 860, the Micom system, and the IBM Displaywriter would be the supported systems. These specialized word processors lasted two to four years before all the software makers produced programs such as WordPerfect, which could turn any microcomputer into a word processor in addition to all the other things it could do.

Electronic Mail

By 1982–83, Steve Worona had written a new version of mainframe mail that he called Mail2 with service and performance improvements. Although the population of users was quite small, considerable effort was

put forward to promote the service and point out its advantages for those who were computer connected. An interesting feature called “Talk” was also noted. Talk permitted a user logged on to CornellC, for example, and using Mail2 to send an instantaneous message to another user on the same machine and use the mail services. That was an advantage at times when one wanted to contact another person without using the telephone system.

Advent of CUINFO

In the early 1980s Worona was developing a forerunner to what became CUINFO for displaying information typically contained in Cornell catalogs and brochures. In early 1980 Worona started by using individual CMS commands to display the course roster and final exam schedule on CornellA and CornellC. When this approach proved too cumbersome to maintain and extend, the CUINFO umbrella structure was developed to accommodate many different information sources. Worona developed this improved system with the assistance of Cecilia Cowles (also from CCS) and Kathy Beauregard from the Information Referral Center in Day Hall.

The first use of CUINFO itself was in 1982. Since “free distribution” accounts to access CornellA and CornellC were available to all Cornell staff and students by this time, CUINFO was a free service to everyone at Cornell.

CUINFO was the first of what came to be called Campus Wide Information Systems (CWIS), computer-based systems providing general information to the campus. The first public CUINFO terminal was installed in the Information and Referral Service center in the Day Hall lobby in 1983. This video display terminal, which was always connected to CUINFO,



Figure 3. CUINFO terminal (kiosk) in the Day Hall lobby

had 17 categories of information; bus schedules, academic calendar, CCS information, dining and housing services, exam schedules, library schedules, movie schedules, religious services, and sports schedules are examples. All of these were text-only displays organized as lists or tables since at the time there were no graphics, hyperlinks, or color screens.

Telecommunications and Networking

Given all the developments in computing and communications on campus in the 1980s, networking took on increasing importance for CCS and Cornell. The plan announced in 1978 to build a packet network based on Tymnet technology did not get further than installing prototype networks from Uris Hall to the Johnson School of Management in Malott Hall. Other than this, the emphasis continued to be on installing technology that would provide more cost-effective transmission between computer terminals on campus and the computers at Langmuir Lab, five air miles away at the airport. Most electronic communication between the end user seated at a “dumb terminal” on campus and the computers at Langmuir took place over a single direct-link telephone line with a modem at the user end and a modem and controller at the computer end. This did not require high speed. Modems had achieved transmission speeds of 2400 baud (bits/sec), which seemed adequate and so much faster than the 100, 300, or 1200 baud available until then. When modems were able to transmit at 1200 or 2400 baud, it was the first time that the transmission could stay ahead of the typing and stay ahead of the user’s ability to read the text. High-speed (9600 bits/sec) network lines were available, but the low cost-benefit ratio made more moderate speeds appropriate for the need.

In 1980–81, ENA multiplexers were installed. Using X.25, a then-popular packet switching technology, the ENAs connected the computers at Langmuir through equipment in Uris Hall with a special link from Uris to Malott for access to the DEC20. According to Cogger, the ENAs permitted the use of a single bisynchronous (bisync) or alarm circuit-quality telephone line operating at 9600 bits/sec to carry the traffic from 12 to 24 hard-wired terminals on campus. Given this configuration, the actual speed at any terminal depended on the number of terminals active at the time, but at most times the speed would exceed the single line connections at the lower speeds.¹⁶

In 1982 VP King discussed a proposal for systems and communications architecture to support computing at Cornell. Very early in rebuilding the CCS

organization Dick Cogger was appointed director of network development to ensure that Cornell was carefully considering the options and design of a network architecture. That same year, King spoke about the central networking issue of having distributed computing elements being able to talk to each other as a means of improving the delivery of computer services to the campus. While the need was clearly recognized the delivery was slow in coming. Cogger and other CCS staff were part of the planning team for the new telecommunications infrastructure and so it seemed prudent for large-scale campuswide data networking plans to await the completion of this project. In the interim, CCS continued to improve and increase existing network offerings in different ways.

Sytek and the Start of High-Speed Networks on Campus

On campus CCS took advantage of coaxial cable that it had installed earlier as part of the Tymnet project, and in 1983 it put into operation a broadband network using technology from Sytek. This was done to implement interim but cost-effective improvement in connecting terminals to computers. As Cogger describes it, “Sytek used a combination of packet switching and broadband transmission technologies to implement network-based serial links.” The coaxial cable was extended to run from Uris Hall to Upson via a leg through Barton, from Uris to Malott to Warren Hall, and from Uris to Day Hall. The two-way communications link to Langmuir was implemented using coaxial cable leased from ACC (American Community Cablevision), the local cable TV provider at the time. The great advantage that Sytek brought was that boxes used to connect terminals to the network operated at 9600 bits/sec, and each box could connect two terminals at any location where they were installed. Further, connections to the network could be made anywhere the cable passed. Operating with one terminal per box, or with two terminals per box sharing the 9600-bits/sec speed, resulted in a significant increase in speed at the terminal.

Another advance came in 1983 with the installation of IBM 3270 terminals. These were connected over 9600-bits/sec bisync lines from the controller at Langmuir to controllers on the campus, which, using special IBM-supplied cables, connected the 3270 video display terminals in office work areas. While this was not the standard IBM way to connect these terminals, it sufficed to make the special function keys and other features of this “smarter” terminal available to the administrative applications being developed.

¹⁶ Dick Cogger, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, September 2000.

BITNET and Network Addresses in Telephone Directory

On the national and international scene, in 1982 Cornell became one of the founding members of BITNET, an intercampus network first proposed by Ira H. Fuchs at CUNY. BITNET was aptly named: Because It's Time Network (according to network lore this is attributed to Fuchs). The first link was implemented in 1981 between CUNY and Yale, where Greydon Freeman was the director of the computing center. The next expansion connected the computers at MIT, Harvard, Boston U., Brown, Yale, Columbia, City University of New York, Rockefeller University, Cornell and Cornell Medical College, Rutgers, Princeton, and Penn State. A major link to the West Coast was completed shortly afterward as soon as extensive routing tables were in operation so that files could be passed along to their ultimate destination.¹⁷

BITNET used the IBM NJE (Network Job Entry) communications protocol to send files from one IBM computer to another using leased telephone lines. The files transferred between the computers could be any information; they were just encoded bits, but the foremost use was for e-mail files. In August 1984 the *CCS Bulletin* announced that BITNET now connected 64 institutions and 200 nodes (computers).

Quickly BITNET spread to higher-education institutions and research organizations around the world. Scientists, students, and staff were soon sending e-mail to their colleagues at many institutions. In 1984, the first international link was to the University of Toronto with Cornell providing routing support. Shortly after, NetNorth was formed to connect most of the Canadian universities to BITNET. Connections to Israel and to EARN (European Academic and Research Network) soon followed. Seeing the advantage of these worldwide connections, vendors or users of other computing systems soon developed their version of the NJE protocol, and BITNET expanded rapidly after that.^j

At Cornell, most minicomputer installations soon had BITNET addresses permitting them to communicate with each other. It became almost a badge of honor at computer conferences to hand out a business card with a BITNET address, such as `jwr@cornella.bitnet`.

As another mark of progress, and a result of extensive lobbying by Worona, in 1985 individuals were invited to list their e-mail addresses in the Cornell telephone directory to facilitate the use of electronic mail. That practice has become commonplace today, but it was leading edge at the time, when Worona estimates there would have been only hundreds of e-mail users out of approximately 8,000 Cornell employees.

1984 to 1986—Microcomputers Rise to Prominence at Cornell

As noted earlier, microcomputers started to appear in campus laboratories and offices in the late 1970s. In the early 1980s, OCS, and later CCS and individual faculty and staff, began to acquire microcomputers based on their interests and needs and the possible exciting new uses that looked promising at the time. Many notable hardware and software developments took place in the early 1980s as discussed in the technology overview, but at that time microcomputer technology was very unstable. It was another case of vendors trying to outdo each other to capture the imagination of purchasers and obtain some significant market share. IBM, given its long history and market share in mainframe computing, had an edge. Early in the decade the terms microcomputer and personal computer, or PC, started to be used interchangeably even though IBM held the term PC as its own copyright.

By early 1983, Cornell Computer Services had concluded agreements with vendors such as Apple Computer, DEC, Hewlett-Packard, IBM, and Tandy Corporation to provide computers and accessories at highly discounted prices to the campus. Discounts ranged from 20 to 60 percent. Cornell was invited to join the Apple University Consortium, a select group of 24 colleges and universities that Apple expected to be pacesetters for their use of personal computers in higher education.¹⁸ Before concluding all these arrangements, Cornell's departments, faculty, and staff had to obtain systems from off-campus vendors. Usually, a new buyer would consult with DACS/DCS microcomputer support group in CCS and then make a decision about which system and components to acquire. In these early days there was typically only one local supplier that, for example, carried IBM equipment, owing to restrictive agreements that gave a supplier exclusive rights to market equipment in a specific region. The person could buy from this store or go further afield to other stores. The big changes with these new contracts were that Cornell users including students could now buy their choice of microcomputer directly from CCS at highly discounted prices. Nevertheless, there were certain restrictions, the most difficult being that each person could buy only one system; if that person wanted a second system, the first system had to be sold to another eligible user. In time these restrictions disappeared. People could later buy one system a year, but it took a while before an open market developed.^k

¹⁷ www.cren.net/cren/cren-hist-fut.html

¹⁸ "The Era of the Microcomputer on Hand on Campus," *Cornell Chronicle*, January 1984.

Microcomputer Sales Start

Encouraged by vendors such as Apple and IBM to create a better interface between them and all the possible clients on the campus, and to bring some order to this booming activity, VP King decided to form a microcomputer sales unit within CCS. He set two stringent parameters for this group: first, it would keep prices as low as possible and second, to help achieve that goal, it would offer no technical consulting. He envisioned a low-budget sales operation styled after the “bargain basement” 47th Street photo shops in New York City. The markup over cost was set at 5 percent, estimated to be sufficient to cover operating costs. In this model, consulting services would continue to be provided by the DCS organization. The microcomputer sales group was made part of the Network Communications division of CCS (NetComm), then headed by Alan Personius, on the assumption that there was a compatible relationship between the marketing/sales and technician staffs on these two fronts.

The process worked as follows. DCS would recommend a system for the customer’s stated needs, and the customer could then decide where to buy the system, possibly from CCS. This was not universally accepted as the best arrangement because customers wanted and needed serious hand holding when they were going to lay out \$2,000 to \$5,000 for a computer, a brand new technology to most of them. Nonetheless, by late 1983 microcomputer sales and service had been organized, staff were recruited, and orders were being taken, albeit in handwritten form. To further simplify

the situation for CCS, King made arrangements for software to be sold by the Campus Store because the store was much more capable of dealing with high-volume commodity items. The Campus Store created a “Computerware” department in 1983 and handled both software and supplies such as ribbons and paper.

The first PC sales office was located on the fourth floor in Day Hall in space that was available and assigned to CCS. This enabled a quick start but was not a good choice because it increased the foot traffic to Day Hall. Complaints were heard from the Cornell Legal Office, which had some of its staff on that floor, and from senior Cornell administrators whose offices were on the third floor.¹ It soon became evident that to provide decent delivery, and to meet the vendors’ ordering minimums and special package pricing and delivery restrictions as well as to install and test system components and upgrades, it would be necessary to create a warehouse and technical facility. As chance would have it, several biological sciences units were vacating space at Langmuir and a large amount of first-floor space convenient to the main door was available. In 1984, Personius quickly outfitted a secure warehouse facility^m and workbenches for equipment assembly and repair in the space, and business boomed.

Another of VP King’s objectives was to achieve a \$1,000 price for a system. This nearly happened in 1984 when Apple first introduced a 128K Macintosh system including Macwrite and MacPaint software, which could be purchased for a discounted price of \$1,125 compared with the list price of \$2,495. A year later the selling price for this same unit was \$1,050, as



Figure 4. CCS Microcomputer Sales Office in Uris Hall—taking orders

a result of the reduced list price of \$2,195. Given these low CCS prices, students were advised to purchase computers after they had come to campus and not before, to get the best price and, equally important, a configuration recommended and supported on the campus. While this recommendation went counter to King's earlier prediction of students bringing microcomputers to campus, it was a sound recommendation in those early days of rapidly changing technology and equally rapid changes in prices.

In 1984 the sales office was moved to the ground floor of Uris Hall (see Figure 4), and in 1985 it was moved to Caldwell Hall when this building was made available to CCS. Some of the key staff who were part of this new and busy group were Maureen Quartararo, Nancy Flynn, Karen Fromkes, Cindy Durbin, Chris Jones, and Rick Cortright. Some sense of the "ramp up" of this activity is conveyed in Personius's recollection that 536 micros were sold in the first six months of operations, 1,200 in the next six months, and then the volume jumped to over 3,000 systems and stayed over that amount every year thereafter.¹⁹ While this in itself was a significant increase, it didn't reflect the total volume of activity that increased commensurately with the sale of printers, external disk drives, and other peripherals and augmentations.

During 1984–85, the first full year of microcomputer sales in CCS, 3,000 microcomputers were sold. Of these, 78 percent were systems from Apple (a couple of Lisas but mostly 128K Macs and 512K "Fat Macs"), 17 percent from IBM (mostly PCs but some PC-XTs), 3 percent from DEC (mostly Rainbows but some Pro350s), and 1 percent from HP (Vectras). This share of market continued for the first several years of sales activity until the IBM and IBM compatibles, with improved software and operating systems from Microsoft, gained the larger share. The total dollar value of all sales in this first full year of operation was \$7.75 million, even with the large discounts on systems but including components like printers and external disk drives and modems. Cornell departments accounted for 50 percent of the dollar volume, undergraduate students for 21 percent, graduate students for 14 percent, and faculty and staff almost equally sharing the remaining 14 percent.

CCS Seminars and Workshops

The large influx of microcomputers on campus forced an expansion of the educational and informational activities of CCS. When the decade started, CCS continued to hold fall and spring workshops on new

and old computer applications such as Fortran and statistics packages, but now text editors and database systems were introduced. Special workshops were started in the winter intersession and over the summer in 1983 to introduce faculty to computing technology. Between 150 and 300 faculty and research staff took specialized workshops (over 30 in number) covering a wide variety of topics, from mainframes and their particular applications to microcomputers and their applications.

By 1985 the interest definitely switched to microcomputer topics to which over 75 percent of the workshops were directed. The most popular workshops dealt with selecting a microcomputer (not only IBM vs. Macintosh but also different models within each), selecting printers, and with using the IBM PC itself. A whole seminar was dedicated to discussing the different types of dot matrix printers or "daisy wheel" printers as preferred options to the more expensive laser printers then available.

In 1985–86 workshops were held on new products such as Lotus 1-2-3, Word for the Macintosh, and Desktop Publishing on the Mac. In fact, there was such unprecedented demand for microcomputer courses that in the summer of 1986 six courses were repeated. Only the tried and true Fortran workshops continued throughout the 1980s, with the emphasis on the use of vector arithmetic becoming the hot topic toward the end of the decade.

A software lending library was started in Uris Hall with over 450 different software packages that faculty or staff could borrow to test and evaluate for their purposes. Workshops continued to be a strong training, educational, and informational vehicle throughout the decade as the technology continued to change.

User Groups Form (and Dissolve)

Recognizing that CCS could not possibly cover all the topics of interest to the level of detail desired, interested parties started forming user groups. CCS fostered and supported these groups because they provided yet another vehicle for education and information and experience sharing. Some groups, independent of CCS, published their own newsletters and included members outside of the Cornell community. In 1984–85 there were user groups for Apple II, Radio Shack color computers, Commodore computers, CP/M (rival operating system to DOS), Displaywriter, IBM PCs, and WordPerfect. There were FLOG (the Finger Lakes Osborne Group) and RUG (the Rainbow Users Group). For the Macintosh computers there were the Mac Developers and Macintosh Users who under the name Mugwump (Macintosh Users Group for Writers and Users of Macintosh Programs) published the *Muggers Monthly*. For minicomputer users there was

¹⁹ Alan Personius, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, October 1999.

IMLUG (Ithaca Minicomputer Local Users Group), which met and communicated about issues relevant to those systems. Later in the decade there were SMUG (Student Mac Users Group) and SIGNUM (Special Interest Group for Novice Users of Macintoshes). The groups dissolved when the knowledge base was rich enough or the item of interest became more or less obsolete.

Thefts of Computers

An undesirable side effect of having many microcomputers on campus was theft. Instead of simply trying to steal computer time from a central system it was now possible to steal the whole system. Besides the first 128K Macs and later 512K Fat Macs, which were easily portable and even had a convenient built-in carrying handle on the top, most of the other systems were definitely “luggable,” even though the monitor and CPU box were separate units. The open-door policy prevalent on the Cornell campus aided robbers in walking off with systems. CCS and other public facilities used tie-down steel cables and special screwed-on and glued-down stands, but robbers used cable cutters or simply ripped the stand off the table to steal units. The *Barton Blotter*, the Cornell Safety Division report of incidents on the campus, would list one or two stolen systems a week, during 1985 and 1986, and started to enter these occurrences on a daily basis into CUINFO to spread the word more easily and encourage safekeeping practices. There was at least one major theft over the winter holiday break in December 1986 when an assumed gang of robbers stole over \$90,000 worth of equipment in 15 separate burglaries and even left behind \$10,000 of equipment when they may have been interrupted in their activities. As vigilance and security increased, thefts more or less stabilized. Still the increasing portability of lighter laptops made stealing systems easier and easier.

Software Piracy

A concomitant issue that became much more important in the latter half of the decade was “software piracy.” In the early days of microcomputers, software was pretty much “freeware,” and sharing with others was encouraged. As commercial firms entered this business activity, the practice of users sharing (copying) their software continued. The firms made some attempts at requiring master disks and password-type restrictions as ways to keep their rightful income from being eroded. However, eager users and other hackers would quickly find a work-around solution and the “game” continued. By the end of the decade, it was common for anyone buying software to agree to the terms of the contract printed on the box by simply

breaking the plastic wrapper on the box containing the software media. This became known as “shrink-wrapper licensing.”

Microcomputer Viruses, Trojan Horses, and Worms

A pervasive negative phenomenon that microcomputers introduced was software “viruses” on the campus. In the early days software was copied from one computer to another by the use of floppies or diskettes that were made on one machine and read on another. In fact, because there were no “hard disks” on the first systems, one had to use a systems diskette to boot up (start) the system when it was turned on to be able to start application programs. Some individuals would hide malicious programs on these system diskettes. At a public computer facility, an “infected” system disk would infect the users’ application disk and files so that when that disk was used with another computer, it would infect the system disks of that computer. Often viruses would simply destroy the application and systems software and not destroy the data directly. nVIR was one of the early viruses that acted this way. More esoteric ones like “Friday the 13th” would erase any infected program on any Friday the 13th. Notices were constantly being posted at all the public facilities about the latest virus. Viruses also became hardware specific for the Mac or IBM-compatible PCs and were more destructive.

The whole subject of hidden and malicious damage to microcomputer software and data was addressed by CCS in workshops titled “Microcomputer Viruses,” “Trojan Horses,” and “Worms” as the numbers of such undesirable programs multiplied and caused serious problems for a large number of users. It isn’t hard to imagine that the infection rate increased tremendously as microcomputers were networked together and could easily and quietly infect all the systems on a local network and then spread outward as the local nets were connected together to form larger networks. One of the favorite techniques was hiding the software in the electronic mail systems so that it could not be easily detected.

As was usual for computing technology, virus detection and eradication software was developed to try to anticipate and avoid whatever problems the virus presented. Protection software soon started to become available on the commercial market and through “freeware.” When this protection software was installed on a system, it inspected the incoming information from the disk for the presence of this rogue software. Cornell had to exercise vigilance in these matters and act quickly because a new virus started at Cornell could quickly become a national or international problem as networks started to connect the world at large.

Instructional Computing

The really big change in instructional computing got under way in 1984–85 when microcomputers started to be used as the primary vehicle for instructional computing in many departments but particularly in the Department of Computer Science. The use increased dramatically when IBM gave Cornell a grant (Project Ezra) to promote the use of PCs by courses, and Apple both provided computers for laboratories and encouraged the development of what they termed “courseware.” The combination of additional machines for faculty and laboratories led to significant changes in computing services for students.

The First Use of Macs and IBM PCs

While the Teraks had been successful for the teaching of PL/I programming in the constrained Cornell Synthesizer environment, they did not fully exploit the potential of these computers. The first breakthrough occurred in the spring of 1984 when the Department of Computer Science committed to using Macintoshes in CS 100 for the fall of 1984. It was expected that more than 1,500 students would be using these systems during the next academic year.

In 1984 a facility was opened in Uris Library with 24 Macintosh microcomputers and six printers. The following year a major step forward took place when additional facilities were opened in two dorms, the first in McFaddin Hall, followed by another in Clara Dickson Hall, making the latter the tenth public facility opened on the campus. In 1986 another facility was opened in McGraw Hall followed by the first facility in campus housing at Pleasant Grove apartments.

Mann Library, with the cooperation of CCS, established a microcomputer center²⁰ in the spring of 1984. The initial configuration had 33 microcomputers (29 IBM PCs, 2 DEC Rainbows, and 2 Macs with 9 dot-matrix printers shared between systems) with an arrangement so that 21 of the PCs could be separated for classroom teaching using a video projector. This Mann facility had a broader mission of providing access to library and information sources, but it supplemented the campus facilities.¹

Concerns about Microcomputers and Course Instruction

The increasing presence of microcomputers on the campus and their ever-increasing use by students and faculty in course instruction brought a number of important side effects to the campus. In late 1983, the Educational Policy Committee of the College of Arts and Sciences, while understanding the many benefits of computing, was concerned about the university entering into agreements with major computing companies. The committee expressed two major issues of principle:²¹

1. Should the university agree to terms with any company which would make the university an explicit and willing subject for commercial promotion? The scale of micro-

²⁰ Howard Curtis, ed., “Public Access Microcomputers in Academic Libraries,” *The Mann Library Model at Cornell University*.

²¹ “Gifts and Purchases of Microcomputers Reviewed,” *Cornell Chronicle*, October 13, 1983.



Figure 5. The first Macintosh facility at Cornell (Uris Library)

computers' presence on campus will be far greater than any technology in the recent past and hence makes the question of commercial cooptation serious.

2. Widespread availability of individual computers will necessarily redirect the resources of the university. To what extent should this be allowed to occur? In addition to the cost of machines themselves, factors which must be taken into account are

- (a) maintenance
- (b) time: for training, developing, and perfecting instructional materials
- (c) equal access to computers: will students who can't afford machines have easy access to university machines?
- (d) space: will computers be placed so they are accessible?

They went on to consider other related issues, such as the realism of the three- to five-year time frame for projects, the possible incompatibility problems if different systems were installed and supported on campus, and issues of academic integrity as computers were increasingly used for instruction. While understanding that some of the issues were "mechanical," the committee was most concerned with the fundamental changes being initiated by the microcomputer revolution happening on campus. VP King, J. Robert Cooke as chairman of the University Computing Board, and Richard W. Conway as acting chairman of the Department of Computer Science issued a statement on the issues raised. They clearly answered "no" to the question of commercializing Cornell's name and transfer of any intellectual property rights, but they could say only that the other issues would have to be addressed as both the colleges and service providers developed future plans for using computers in the curriculum and in research.

Project Ezra

In 1984 IBM made an \$8 million equipment grant to Cornell to promote the use of personal computers in instruction and research. Cornell was one of 19 universities chosen for such a grant under IBM's Advanced Educational Program (AEP) sponsored by the Academic Computing Information Systems division (ACIS) of IBM. The grant was organized as Project Ezra with the objective of permitting faculty in diverse disciplines to explore the utility of current and advanced microcomputers, now referred to as workstations, in classrooms and laboratories.



Figure 6. Statue of Cornell founder Ezra Cornell with IBM PC keyboard

Project Ezra²² had the following four specific goals:

1. To create and document innovative instructional software that will function on workstations.
2. To ascertain where, and with what facility, suitable configured workstations may be used to improve the efficacy of academic computing.

²² Gordon Galloway, Cecilia Cowles, and Carrie Regenstien, "A Report on Project Ezra IBM Advanced Education Project 1984-88, Developing Innovative Software to Improve Instruction and Research." April 1989.

3. To determine the potential and the limits of combining mainframes and workstations in delivering academic computing.

4. To investigate ways of integrating individual efforts by means of networking strategies.

The grant was for a period of three years. Gordon Galloway, director of academic computing, expected faculty members from the humanities, social sciences, professional schools, physical sciences, and engineering to be involved in creating new instructional software. It was estimated that 500 workstations would be installed from the grant, the first of which was to be in place for the summer of 1984.

To select proposals for funding, the “Ezra Coordinators” planning committee was established under the direction of VP King. This committee developed guidelines for the different resources provided by IBM to be equitably distributed among Cornell’s various schools and colleges. They also reviewed and monitored faculty proposals submitted to the project. A key decision was made that “each participating school or college would be allocated proportional shares of the equipment awarded, based 50 percent on the number of student contact hours taught and 50 percent on the number of fulltime faculty within that college.” A portion of the grant funds was reserved for the discretionary use of the provost.

A support structure was also put into place in CCS. Cecilia Cowles took on the position of director of Project Ezra, working in Academic Computing. Carrie Regenstein was appointed to coordinate activities among participants, their colleges, and CCS. The project got off to a fast start after the announcement, and more than 100 faculty members submitted proposals for developing new software for their courses. As expected, the number of requests far exceeded the amount of equipment and software that could be obtained with the funds.

To supplement these equipment grants, a grant of \$300,000 over two years was provided from anonymous donors to be used to hire student programmers. Faculty who hired students to support their software development projects using these funds were asked to get matching funds (two for one) from their college. The first year, 28 projects were granted funds to support student programmers working on both IBM and Macintosh projects. The donors also allowed a portion of the funds to be used to hire a full-time programmer/analyst dedicated to assisting the faculty programmers.

In 1985 additional support was provided to Project Ezra participants in several ways. In May, a support center was opened in Uris Hall with staff from other areas of Academic Computing as well as staff assigned to the project. A nearby room was equipped with workstations likely to be used by the project to assist evaluation and consulting. Later that year a monthly



Figure 7. Unloading an early shipment of Project Ezra equipment.

Project Ezra Newsletter was established and continued to be published throughout the life of the project. To publicize the many innovative uses fostered by Project Ezra, the “Lunchtime Bytes” seminar was organized and held in Caldwell Hall. It took advantage of the lunch break to offer presentations by developers of instructional and research software to demonstrate and discuss their applications software. This proved very successful as it involved actual “users,” who were not themselves technologists, talking about their own achievements with computers.

To provide broader exposure to the project, a Project Ezra fair called “Tools for Learning” was held in the spring of 1986, and again in 1987, inviting the campus to see demonstrations of all the projects in many disciplines.

By 1986, 372 microcomputer systems were distributed through Project Ezra to 252 faculty and staff for a total of 186 different projects. When the final report for Project Ezra was issued in 1989, it listed 638 machines distributed to 329 participants for 264 different projects. It is worth noting that the official total number of systems distributed exceeded the 500 expected at the start of Project Ezra. This was a result of both IBM’s making supplemental grants during and after the three-year period of the original grant and the lower prices of units over the lifetime of the project.

Given the rapidly changing technology during this period, just over one-third of the computers were PC-XTs, almost the same number were PC-ATs, and one-fifth were model PS/2s, with the residual labeled as “other” systems. The largest number of systems was distributed to the College of Arts and Sciences

(168), 121 to the Engineering College, and 117 to the College of Agriculture and Life Sciences. As a result these three colleges were given about two-thirds of the systems, and the other colleges and selected research centers received the remainder.

Halfway through Project Ezra, the colleges were invited to propose public facilities that would be equipped by using the remaining allocation for each college. The final Project Ezra report states that “approximately 200 workstations have been made available to Cornell students through Project Ezra. These workstations augment the approximately 400 workstations from other sources.” The report states that “planners expected 80 percent of the equipment received to be in the hands of students by the end of the project. Public workstation facilities were first placed in the Colleges of Engineering and Arts and Sciences and later in the Colleges of Agriculture and Life Sciences; Architecture, Art and Planning; and Veterinary Medicine. Facilities were also set up in the ILR School, the Hotel School, and the graduate business school. Project Ezra resources were also used to upgrade the CCS training and teaching facility in Stimson Hall. In 1985, Stimson was equipped with 21 IBM PCs networked with Omninet to a file server and 21 Macs awaiting the installation of a server.

One measure of the success of Project Ezra was the number of Cornell faculty presentations at the national IBM Advanced Educational Program (AEP) conference held each year. This national conference brought together faculty software developers from all 19 AEP schools to share their work and ideas. In May 1985, six Cornell faculty members were selected to



Figure 8. 1986 Project Ezra fair “Tools for Learning” in Statler Hall

present papers at the 1985 national AEP conference. Cornell faculty presented 10, 8, and 4 papers at this same conference in the years 1986 to 1988, respectively. Gordon Galloway chaired the program committee in 1985 and Cecilia Cowles assumed this responsibility in 1987.

Although the number of projects is too large to list in this summary, an earlier status report on Project Ezra²³ mentions four exemplary projects, two of which (numbers 1 and 4 below) went on to establish their own joint studies with IBM. The four projects were:

1. Education Applications of Videodisc Technology: Geri Gay, CCS, and Education (joint study with IBM)
2. Astronomy Undergraduate Laboratory Exercises: Martha Haynes, Astronomy
3. Dynamic Aspects of Structure and Reproduction: Karl Niklas, Plant Biology
4. Databased Writing Instruction in Foreign Languages: Jim Noblitt, Don Sola, Modern Languages and Linguistics (joint study with IBM)

In early 1988 an Interactive Media Center was formed as part of a joint study with IBM and Cornell Computing Services.²⁴ The center was a spin-off of Professor Geri Gay's videodisc project in Project Ezra. The center's aims were to tie together research, evaluation, and communications on the newest technologies such as CD-ROM, digital video, videodisc, and optical discs and find out how they could or could not be used to enhance higher education. Geri Gay and Cecilia Cowles were the directors of the center, and the staff included Mike Oltz and Margie Wilson as well as graduate students.

Apple Mac•Ed Center: MacAdemia Conferences

Although Apple Computer, Inc., did not have a formal program such as the IBM Project Ezra, it supported numerous initiatives for faculty to incorporate Macs in their classroom teaching activities. Apple worked with CCS on separate but similar initiatives in academic and administrative computing, and also made specific grants on direct appeal from faculty. The Academic Computing/Distributed Computing Support (DCS) staff were the first avid supporters of Macintosh equipment, having started with the first Lisa system acquired from Apple in 1983. They also fully supported faculty and staff with consulting activities before and after CCS started to sell systems.

In May 1985, Academic Computing/DCS and the Higher Education Marketing Division of Apple held the first "MacAdemia" conference in Ithaca. The principal organizers were Tom Hughes from Cornell and Mike Looney and Stacey Bressler from Apple. The conference brought together technical and sales/marketing staff from Apple and technical and support staff from the Apple University Consortium schools

to share their work and activities in the use of Macs in the curriculum. Well-known Apple visionary Alan Kay was the keynote speaker at the concluding lunch. The conference set a tone of excitement for use of the Mac, now present for just over a year on campus.

With a grant from Apple in 1986 the Mac•Ed Center was organized to assist faculty in the development of Macintosh-based "courseware," software specifically designed to assist in teaching students about an academic discipline.²⁵ The center was part of the Academic Computing section headed by Assistant Director Tom Hughes, and the initial staff were Chris Pelkie^o as manager and Jill Levien as the support programmer. The Macintosh *Monitor* publication was developed and edited by Cynthia Frazier and was distributed on the campus to keep interested parties informed about the center's program, new developments and practices, and changes in hardware and software technology. The package getting the most "press" in 1986–87 was the newly available "HyperCard" application, a simple database tool with an index card metaphor, for list and stack processing. This exciting new product—a revolutionary precursor of today's hyperlink-based web applications—was capturing the imagination of application developers. New uses and support features were being announced constantly.

The Mac•Ed Center followed a two-pronged approach. A demonstration facility was set up in Computing and Communications Center, on the third-floor balcony, to enable faculty and their support programmers to work with the center staff to familiarize themselves with Macintosh systems, get additional training, and make decisions concerning their own projects. To get faculty started on their project, the center was granted a pool of equipment that was loaned out to faculty to set up in their own space. Faculty members had to make a loan application that was reviewed by the center staff. Loans were made twice a year, fall and spring, and were generally for a period of six months. If the application was approved, a faculty member could keep the equipment for that period of time, after which it would be returned for loan to others. Because Apple continued to increase the number of computers, the loans sometimes were extended for longer periods. Not surprisingly, and to everyone's benefit, loaner systems were often either purchased directly or returned, and new systems were

²³ Carrie Regenstein, Cecilia Cowles, and Gordon Galloway, "Cornell University's Project Ezra; Background, Objectives, Infrastructure and Support." November 1987.

²⁴ "Center Created to Improve Interactive Video Instruction," *Cornell Chronicle*, February 11, 1988.

²⁵ "The Mac•Ed Center, Support for Faculty Developing Macintosh Courseware," *The Macintosh Monitor*, Fall 1987.

purchased to support the continuing activity once the application was fully developed or in active use.

A report on Macintosh courseware from the Mac•Ed Center (winter-spring 1987), after six months in operation, summarizes 24 courseware projects. Many of the projects exploited the graphics capability inherent in the use of a bit-mapped screen. Several of the projects were in the field of mathematics such as Function Probe: Teaching Calculus Constructively, MacMath: Graphic Exploratory Tools for Higher Mathematics, MacElastic: Instructional Finite Element Analysis, and MathWriter: The Master Key of Mathematical Typesetting. Other areas included Statistics: DataDesk: Statistics Both Powerful and Easy to Use, and Teaching Large Statistics Classes with Data Desk. Modeling, simulation, drawing, visualization, and subjects unique to certain disciplines, such as the Hotel School's "The Restaurant Game, Food, and Beverage Operations Management," cover the remainder of the applications.

During this period, from 1984 to 1987, a number of Macintosh-only facilities were created on campus. The most notable was the Cornell University Writing Program, which had 20 Macs in its Goldwin Smith facility. There was also a comparable facility in the ILR School.

Intellectual Property Rights

One of the less obvious effects of micros on the campus involved the copyrighting of locally written software. For faculty the major question was one of intellectual property rights: "Who owns the software we developed for our research or for classroom use?" Debate on the question started early in the 1980s. There was the academic tradition that creators of works owned the copyright to their creation and the long-standing tradition that faculty clearly owned the copyright to their books and papers. The university position was that software creations were the property of the university. There was no such debate in the case of staff employees who came under the "work for hire" clause where the university had clear ownership of their software products. Occasionally, exceptions were made where the university granted these rights to employees, but generally the question was cloudy only for faculty. The university argued that it owned the rights since most often the software was developed using university-provided computers. The counterargument was that often faculty used their personal home equipment for such developments, and so the equipment ownership criteria did not help clarify the situation. This issue was debated throughout the decade and was not resolved until the 1990s.

Business Systems

The concerns of how microcomputers would be deployed for administrative applications started to receive the attention of the University Computing Board, and initial steps were taken to improve microcomputer support for these users. There was a continuing discussion on data as a university resource, and several steps were taken to improve services and service efficiency for administrative applications.

UCB Recommendations for Decentralized Systems and Support

In 1984, the University Computing Board (UCB) made several recommendations to the provost about administrative computing. Earlier, the UCB had organized two subcommittees to respond to issues being raised by administrative users across campus. The first, chaired by Jack Lowe, concerned decentralized administrative computing and the need for support for microcomputers and related technologies. (Lowe was financial director in the College of Arts and Sciences.) The second concerned access to centralized databases and increased computing resources and was chaired by Controller Ostrom.

Based on reports issued by these two subcommittees, the UCB made the following recommendations:

In that our philosophy lacks the concept of a SHARED UNIVERSITY INFORMATION RESOURCE, and that our strategy for building information systems does not yet extend to decentralized activities or cope with decentralized information management needs. . .

Recommendation One

That the university provost establish a program committed to the orderly development of systems in a distributed/decentralized administrative computing environment. And

Recommendation Two

That the university provost establish as policy the concept of information as a shared university resource, and put in place a mechanism for university-wide implementation.

Given that there is not sufficient

- consultant help available to evaluate systems and determine computing resource needs.

- consultant help available to evaluate equipment choices and to facilitate purchases,
- technical help to interface equipment or information systems,
- information flow to help users to help themselves.
- training to allow decentralized units to develop their own support . . .

Recommendation 3

That the provost reconsider the scope and organization of decentralized computing support services to include all form of centralized and decentralized administrative computing.

Given that current levels of funding for administrative computing efforts may be unrealistically low in light of the university's perceived need to embrace automation.

Recommendation 4

That the provost reconsider funding levels for the support of decentralized administrative computing activities, and particularly, that centrally funded administrative consulting services be expanded.

There is no record that the provost explicitly accepted these recommendations,²⁶ but it is clear that they influenced the actions that followed.

Microcomputer Support Group; Apple Loaner Program

In 1985–86 Russ Vaught organized the Microcomputer Support Group in the Administrative Computing division. It addressed some of the issues raised by the UCB in an innovative way. The division was granted a pool of equipment by Apple to be loaned out to staff in administrative departments, for example in the registrar's office and admissions, to evaluate the systems for use in their office environment. This was comparable to the program that Apple sponsored in Academic Computing but with a twist: instead of simply training users in the use of the technology, the group initially consisting of Mark Mara and Doug Hornig would provide a complete service for the department. That service included installation of the computers and printers and interconnecting network as well as developing applications that integrated the computer with data and applications of the central administrative business systems, and of course, continuing support

and consulting services. This became known as the Administrative Loaner Program.

Generally speaking, an administrative department applied for a loaner system, but in some cases the group solicited departments to try a loaner system when they envisioned a particularly advantageous use of the technology. After a couple of months, when the loaner period had expired, the department was asked to either return or purchase the equipment. According to Mara, who headed up the group, about 90 percent of departments would opt to buy the equipment and keep it for their continuing use.²⁷ Administrative Computing would then use the funds to buy new equipment to put back into the program to loan to other users. Mara estimates within a few years that 90 to 95 percent of the administrative offices on campus were using Macs for their office use as well as for access to applications on mainframes. He notes that IBM may have regretted its decision not to support a requested loaner program after the initial success of Project Ezra. It is also worth noting that it was well into the 1990s before IBM and IBM-compatibles became the more prominent systems in Cornell offices.

Development of IRIS (Inquiry and Reporting Information System)

One of the interesting and innovative service improvements at the time was the development of IRIS, first proposed by Dave Koehler and implemented by him and Bill Borgida in 1983–84. IRIS was initially developed for online ad hoc inquiry into the Public Affairs data files but later extended as a general batch service as well. According to Koehler, the main feature of IRIS was that it incrementally applied Boolean logic to a set of data until it reached the population of interest. At that point information could be displayed or printed using standard outputs. Because IRIS was table driven, it could be applied to any data files, and its use spread accordingly for those special requests that often were not anticipated when a system was first developed for an operating unit such as the registrar or payroll.

First High-Speed Laser Printer Installed

In 1984 the first high-speed laser printer, a Xerox 8700, was installed in Day Hall to print business reports and other materials. CIT Operations and Production Services staff led by Jim Doolittle and

²⁶ Memorandum to University Provost from the University Computing Boards Re: Recommendations on Administrative computing, February 28, 1984.

²⁷ Mark Mara, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, July 1999.

Peter Baker had been investigating ways to start a transition from using line printers, forms decollators, and forms detachers/bursters with continuous forms to more contemporary printers using cut-sheet 8 1/2-by-11-inch paper. The recommended Xerox printer was capable of printing 70 pages a minute. Baker took on the technical support of this production printing system. His efforts eventually led to the elimination of all continuous-form printing, but more important, he was able to eliminate the use of very expensive preprinted custom forms for such items as grade reports or financial statements. One of the features of the 8700 system was the capability to store fixed formats for forms that could then be blended with the variable information to produce custom reports.

Automating Production Control

In 1984, and almost every other year thereafter, automating production control of batch jobs was an issue. Fundamentally these automated systems allowed one to build a computer-based master schedule of jobs, days, and times. The system, using this schedule, would automatically put the job into the batch queue for processing. One potential problem with automating production was the elimination of positions to pay for purchasing and maintaining the automated system. Another was that job scheduling would become tighter and more formalized. After experimenting with systems, and considering the pros and cons and costs and savings, such a system was never acquired, although the issue came up from time to time in the future.

Telecommunications and Networking

While studies were under way for a new Cornell telecommunications system, the allied concern was where this would be housed. When this issue was combined with the need for more computer room space for the prospective supercomputer program, and the need to meet the obligations of the Pew Grant for a campus computing center, it was clear that action had to be taken. The space needs were met by converting the former Comstock Hall to more contemporary space and using Caldwell for additional functions. Once the conversion of Comstock Hall to the Computing and Communications Center (CCC) was complete, the Telecommunications division (Telecomm) and CCS moved into CCC over the next few years.

The New Cornell Telecommunications/Telephone System

In 1984 the Cornell Board of Trustees authorized the spending of \$17.4 million for the installation of a telecommunications system. The project had been under way since 1981, when the Board of Trustees first approved it. In 1982 the university had hired Harold D. (Hal) Craft Jr. as director of telecommunications and Pat Nelson^P as manager of the telecommunications project. Craft had been director of observatory operations and last served as acting director of the National Astronomy and Ionosphere Center at the Arecibo Observatory in Puerto Rico, and Nelson had been a project manager with Corning Glass. Working with Network Analysis Corporation they conducted a major and extensive investigation of alternative vendors and costs to install a complete new wiring plant and telephone system on the campus.

When the project was authorized in 1984, it was estimated that savings from telephone services alone would amount to \$583,000 in 1986–87 besides providing the basis for a computer data and video network.²⁸ In October, Robert M. Matyas, vice president of facilities and business operations, signed a contract with AT&T to install the major portion of the equipment. VP King and Craft were making statements that while Cornell was not the first university to embark on such a project it was among the leaders as others were just starting their planning process.

The key expectations of the new system were listed as follows:

- a new multipurpose wiring plant including fiber optic lines, coaxial cable, and computer-controlled voice and data switches for increased speed of data transmission beyond present capabilities;
- a backbone network made up of primarily optical fiber and coaxial cable which can interconnect at high speeds with Cornell's mainframe computers, departmental minicomputers, and personal computers and will be accessible to any computer user on campus;
- small networks in the dormitories through which students with microcomputers in their rooms can take advantage of shared facilities such as printers and disk storage equipment;
- links from dormitory to larger campus networks to incorporate a student's micro-

²⁸ "Telecommunications System for Campus Gains Approval," *Cornell Chronicle*, March 29, 1984.

computer into the more comprehensive university computing system, which could allow communications by electronic mail and provide access to centralized facilities²⁹

All of these expectations were eventually achieved. Some came to pass before the decade ended; others took a little longer.

The Computing and Communications Center (CCC)

The need for locating the telecommunications switch in a central campus location, the need for larger space for the production supercomputers, and the need to meet the obligations of the Pew Foundation grant for new space for central computing services all came together in 1984. By this time, the Department of Entomology occupying Comstock Hall and Caldwell Hall, on the northwest corner of the “Ag Quad,” and other academic and administrative units were scheduled to be moved to the new Comstock Hall on Garden Avenue, across from Barton Hall. After some discussion on the pros and cons of making Comstock Hall into the new campus computing center or looking at other alternatives, the decision was made to undertake a major renovation of Comstock Hall for the telecommunications center and the computer center and to rename the building the Computing and Communications Center (CCC).

The planning for remodeling Comstock into the Computing and Communications Center put a large strain on CCS staff. Dealing with questions such as how much space to provide for which activity and where it should be located in the building and all the thousands of questions and concerns surrounding these simple questions took an enormous amount of CCS staff time.⁹ The plan was to execute the remodeling of the CCC building in three phases. Phase one was to complete the annex for the telecommunications equipment and staff and a section of the computer room. Phase two was to complete the CCS computer room in the old Comstock basement and join this with the new computer room in the annex to create one large room. The third and last phase was to provide for staff offices and other support space on the upper floors. As a result individual plans had to be made for the movement of equipment and/or staff in each phase.

There was a small one-story section of CCC, as it was in 1984, along the north wall facing Martha Van Rensselaer (MVR) Hall, and one key decision was whether this annex was needed or not. After evalu-

ating the space needs of CCS and Telecomm it was decided not only to retain this extension but to make it both larger and taller. There was some earlier blue-sky planning about a separate Telecommunications building between Caldwell and CCC, creating an archway into MVR, but this was rejected in favor of a more ordinary and less expensive addition, albeit with the promise to match the exterior brick and stone of the original Comstock-CCC building.^r The original project budget was estimated at \$6.5 million with the costs to be shared by Telecomm, the Theory Center, and CCS after using the funds from the Pew Grant, which had grown to over \$4 million by this time.

Because Telecomm was on a faster track than CCS and was targeted to go into the annex, the decision was made to build the new annex to meet their space needs. The original design called for a three-story addition, but this was shortly expanded to a four-story addition to accommodate the space needs of both Telecomm and CCS. The four-story layout consisted of a sub-basement, a basement (ground level in the annex), and first and second floors. The sub-basement was to house the utilities and equipment such as air handlers for the entire old and new buildings and a backup diesel generator for the telephone system. The basement would be part of the larger CCS computer room that would encompass both the new addition and the old original building basement when it was remodeled. The first floor was for the telephone-switch equipment and telecommunication terminations and the battery backup, while the second floor was for Telecomm staff offices. The old addition was demolished in 1984, while Entomology was still in Old Comstock proper, awaiting the completion of New Comstock, and building began, only to run into a major complication—an underground stream. To isolate the new annex from this underground water an outer shell of interlocked steel pilings was driven into the ground around the whole perimeter of the addition. This was one of many complications, such as adding an additional floor to the annex, that raised the building budget to over \$8 million before the project was completed. Nonetheless, the CCC annex was completed in 1985.

In concert with the CCC project, in 1985 additional space was made available to Cornell Computing Services (CCS) in the basement and first floors of next-door Caldwell Hall, including shared use of the Caldwell 101 lecture room so that CCS could have its own classroom space. The Caldwell remodeling was a comparatively low-budget operation with most of the space given only a minor repainting and upgrading of electrical services. The basement was made available to the fast-growing Theory Center and was wholly used by it for staff and its Graphics

²⁹ “New Telecommunications System to Be Installed,” *Cornell Chronicle*, October 11, 1984.

Computing Facility. Almost the whole basement was air-conditioned with chilled water systems, first requiring the installation of extensive piping for the chilled water itself. The first floor was used by Network Communications (NetComm) for its staff, for the microcomputer sales office, and for the CCS computer accounts office that also sold special computer supplies to departments.

Telecommunications Equipment Installed

The installation of the new telecommunications equipment and fiber/wire plant continued at a fast pace. Every building on the campus had to be totally rewired with new copper wire to every desk from a new building distribution frame that then connected the building wiring with the campus telephone backbone loop. Each telephone in use (by this time there were over 11,000) also had to be replaced by a new telephone. The campus loop itself consisted of five major node rooms, CCC being one of the nodes, con-



Figure 9. Telecommunications cable being laid on the Arts Quad

nected by optical fiber to form the telephone backbone. This design was chosen so that any current or future telephone would be less than 2,000 feet from a node room and so avoid having performance affected by distance problems. Connecting the buildings to the node rooms and the node rooms to each other involved the installation of miles of underground conduit and then the pulling of yet many more miles of optical fiber and copper wire cables. The campus was constantly being dug up for infrastructure trenches of one kind or another, e.g., steam or telecommunications cables (see Figure 9), and roadways were blocked as a result of construction.

The installation of the System 85 switch from AT&T was completed in late 1985 and well tested before going into production on March 3, 1986. A few operating problems were experienced. Some resulted from the new telephone numbers, which at cutover were all 255-nnnn or 253-nnnn local extension numbers, canceling the old 256-nnnn numbers and not necessarily keeping the same last digits. Some problems came from removing the special long-distance WATSBOX from service and being able to direct dial long-distance numbers. All problems were quickly resolved. By any measure, this long and complicated system installation and service changeover was a smashing success. Some of the staff responsible for this project were Pat Nelson, who was appointed director of telecommunications when Hal Craft became associate vice president for facilities and business operations, and staff members Jeff Wilber, Jan Brown, Kathy Parker, Ann Smith, and Holly Biglari.

Pronet Installed to Use TCP/IP Protocols for Campus Network

During 1984 to 1986, networking activity in CCS was taking place on three major and interrelated fronts. To be able to participate in the nationwide networks such as ARPANet and in the linkup of the NSF-supported national supercomputing centers, CCS had to come up to speed on network technologies involving TCP/IP protocols. This same protocol was also used for the campus network backbone (major traffic carrying core network) linking all computers and workstations or LANs on campus in a consistent and cost-effective way. Next, there was the need to install LANs that would accommodate different personal computers and differing LAN technologies and be connected to the campus backbone. Last, at the workstation level there was the need to provide terminal emulators to access different mainframe computers over this spectrum of different technologies. These were especially daunting challenges, given the immaturity of these new technologies in a rapidly changing environment.

When Sytek was installed in 1983, it was understood that this was a temporary move to improve the situation during the building of the telecommunications system to support the Cornell telephone system. The future direction was clearly TCP/IP-based networking, but there were several significant complications. One was that the main CCS computers (as well as those from the Production Supercomputing Facility as of 1984) were at Langmuir and far from the reach of optical fiber runs. Given that these computers would be moving to campus in a few years, installing a fiber link to them was not sensible. Other alternatives had to be found. Cogger and his staff found a workable solution that involved the vendor Proteon, which had a commercially available token-ring product, Pronet, that supported the TCP/IP protocols. By using special modems available from Fairchild, Pronet could be extended to Langmuir by using the ACC cable installed for Sytek.³⁰ While that was a pleasing scenario, two major hurdles had to be overcome.

Given that TCP/IP originated in the Unix/VAX world, IBM had no hardware or software on the market at the time that would connect its mainframes to a TCP/IP network. CCS had to use hardware and software components that were works in progress. IBM built a special PC-based box called the DACU (Direct Attached Connection Unit) that accommodated the physical attachment of Pronet and the IBM systems. This PC accommodated a Pronet card and connected to a mainframe channel via programs that could translate between the different protocols. For mainframe software CCS and IBM obtained WISCNet, developed at the University of Wisconsin, which accommodated the TCP/IP protocols under VM on the mainframe.³¹

To expedite this project, IBM and CCS participated in a joint study to install, test, and evaluate the feasibility and performance of what some called a “classic kludge,” and to make whatever software changes were necessary in VM for acceptable performance. IBM particularly emphasized that the focus of the study was to be on large-scale computing and large file transfers to support the supercomputing initiative and avoid dealing with large numbers of time-sharing users. CCS staff members participating in this joint study were Scott Brim, Mike Hojnowski, and Nick Gimbrone.³²

CCSNet and TheoryNet

Once a stable environment was achieved in 1985–86, in short order two Pronet backbone rings were installed. One was TheoryNet, which connected the large IBM supercomputers at Langmuir to the campus and the world at large via the developing NSFNet. The other was CCSNet, which connected the CCS computers at Langmuir with other computers or LANs on campus. By the time the supercomputers moved to campus in 1986, IBM offered formal products for both the DACU and WISCNet. The DACU, then using an IBM PC-AT, was called the 7170 Communications Controller; the successor to WISCNet was offered as the VM Interface Program for TCP/IP.

Omninet for Low-Speed Local Area Networks

With respect to LANs, the differing technologies posed particularly complex problems. Macs came with their own AppleTalk LAN networking built into the machine so that it was easy to create Mac-only LANs. For IBM-compatibles, there was a choice of the IBM token-ring technology or Ethernet protocols, and some installations used Novell Netware for supporting access to file servers. There were no popular and reasonably priced commercial offerings that accommodated a mix of workstations. In their investigations on how to resolve this issue, Cogger and his staff focused on Omnet from Corvus, which had the potential to provide the LAN technology for mixed workstation environments. In contrast to the 230-Kbit speed of AppleTalk, Omnet ran at 1 Mbit and so could offer faster speeds. The difficulty that presented itself was the high cost of using commercial products to connect Omnets to the TCP/IP-based backbones. Overcoming this deficiency led to a project to develop an Omnet-Pronet interface that in time became known as an AT-Gateway, as these were built on the then-popular and sufficiently robust IBM-PC/AT microcomputer. The AT-Gateway could support two to four LANs, depending on the type. By using AT-compatibles, it was expected that the cost would be substantially lower than commercially available gateways then coming on the market. During the 1984 to 1986 period, John Lynn, assisted by others, developed AT-Gateway programs to connect Ethernet, AppleTalk, and Omnet LANs to the Pronet backbones.

Terminal Emulator Development

The increasing need for terminal emulators required an increasing variety of software to accommodate the different types of personal computers (Macs and IBM-compatibles), the different network technologies, and the different host computers that needed to be accessed. There was a need for line-at-a-time termi-

³⁰ Dick Cogger, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, September 2000.

³¹ Peter Hoyt, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, July 2003.

³² IBM to Kenneth M. King, “Jointly Defined Effort on TCP/IP Connectivity,” November 5, 1986.

nals, such as the popular VT100 from DEC, and page-at-a-time terminals, such as the IBM 3270. To complicate matters, all these emulators had to carry the same “look and feel,” whether using the backbone network or Sytek, so that users could migrate to the new network technologies without major disruption. Further, there was increasing need for file transfers because users wanted to download or upload information files between workstations and host or server computers. Cogger and his group of developers (Peter Hoyt, John Lynn, Kevin Saunders, and Rich Kennerly) developed and supported these emulators for use at Cornell.

For the line-at-a-time terminals, several Telnet terminal emulators were readily available from other sources. TN could be used on the campus backbone network from Macs or PCs. NCSA Telnet, another version of TN, could also be used; its VT100 emulation allowed communication with VAX systems. The problem was that there was no easily available 3270 emulator. Taking a partial implementation from MIT for the IBM PC, called PC/IP, Cogger and his staff proceeded to develop what came to be known as TN3270 for both the Mac and the IBM PC-compatibles that worked with the TCP/IP backbone being implemented. For Sytek-connected terminals, the group provided a C19 emulator based on the Heath C19 terminal that provided equivalent “look and feel” so that users could easily transition to the newer technology. Cogger and Hoyt provide much more detail about these issues in their oral histories (footnoted earlier), but here it suffices to say that by early 1986 there were working versions of TN3270 for both the Mac and the PC-compatibles. For file transfers, Kermit, developed at Columbia University, was imported and made available and supported at Cornell for the Mac, IBM-compatibles, and different target hosts.

Of increasing importance was the ability to transfer files between systems using the FTP (File Transfer Protocol), available with the TCP/IP networking protocols. With the Sytek networking protocol there was a C19 emulator based on the Heath C19 terminal and MS-Kermit developed at Columbia University for emulating a VT100 terminal on an IBM-compatible PC and for executing file transfers between mainframes and workstations.

NSFNet

As noted earlier, driven by the need to link the five national supercomputer centers to each other as well as to provide high-speed access to these computers from researchers in New York State and nationally, there were two concurrent efforts to create high-speed networks. One came to be called NSFNet, the other NYSERNet.

Nationally, NSF created NSFNet, which initially connected the five centers and the National Center for Atmospheric Research in Boulder, Colorado, with a TCP/IP-based backbone operating at 56 K bits per second (Kbps) over leased telephone lines.³³ NSFNet had its roots in the ARPANET that developed the TCP/IP protocol suite and the use of packet switching technologies during the 1970s.³⁴

In July 1986, the Theory Center assumed operational responsibility for NSFNet for monitoring the backbone, installing new nodes, and distributing the router software. Before taking on this responsibility, in December 1985 the Theory Center had established its Network Information and Support Center (NISC) to improve and facilitate access to the CNSF supercomputers. Initially the NISC staff worked to develop TCP/IP-based networking expertise and worked with NSF and other institutions on the interim operation of NSFNet and Theorynet, the campus-based network Pronet backbone. To make it easier for off-campus researchers to access the supercomputers, in December 1986, the Theory Center installed 24 “800” telephone numbers that connected to the Sytek network on campus (at speeds of 300/1200/2400 baud). Off-campus users could also access the Production Supercomputer Facility (PSF) through ARPANET or BITNET. On campus, Theory Center members were encouraged to connect to the PSF using Theorynet, but they and others could come through the CCS Pronet backbone because the two were interconnected (see Figure 14).

NYSERNet

In 1985, Cornell took the lead in creating NYSERNet (New York Educational and Research Network) to connect 15 research institutions in the state to Cornell’s supercomputer and to Princeton University in New Jersey. It was expected this network would operate at 1.5 million bits per second, once optical fiber was in place, a speed fast enough that video and graphics input and output could operate successfully, but would initially operate at 56 Kbits per second using copper wire connections. VP King from Cornell and Richard Mandelbaum from the University of Rochester, who became co-chairs of the executive committee, organized the effort for NYSERNet.³⁵ Bill

³³ “National Network Update,” Center for Theory and Simulation in Science and Engineering, Cornell University, *Forefronts*, Vol. 2, No. 6, October 1986.

³⁴ NSFNET—The National Science Foundation Network, <http://moat.nlanr.net/INFRA/NSFNET.html>

³⁵ “NYSERnet Update: Looking Good,” Center for Theory and Simulation in Science and Engineering, Cornell University, *Forefronts*, Vol. 1, No. 9, January 1986.

Schrader, who was executive director of the Theory Center, was appointed president of NYSERNet.⁵

One of the innovative proposals of NYSERNet was to use the 4,489-mile Interstate Highway system as the route for laying the fiber. Governor Mario Cuomo, who dubbed NYSERNet the “digital thruway,” put this idea forth in a proposal to the Department of Transportation seeking an exception from Secretary of Transportation Elizabeth Dole. The first exception granted was for fiber to be laid in the Albany-Schenectady capital district corridor for 6 miles while awaiting permission (which was given soon after) to lay fiber cables along the 559 miles of the New York State Thruway to form the backbone of NYSERNet.

In late 1986, NSF granted Cornell \$1.2 million to begin implementing NYSERNet and to link up to NSFNet. As a result, NYSERNet was up and running in 1987. It was one of the first state and multistate networks, if not the first, to link up educational and research institutions to NSFNet and so to broaden access to the supercomputers as well as to allow for exchange of data and e-mail. The fiber routed along New York’s interstate highways, which gave rise to the term “information superhighway,” represented the beginning of the “Internet,” which in a few years linked the whole world.

GateD

One successful development that came out of the Theory Center’s involvement with NSFNET and NYSERNet was GateD (pronounced Gate Dee) routing daemon. Basically, GateD automatically coordinated the three major routing protocols used by Unix sites, the NSFNET backbone, and the ARPANET gateways, and chose the best route between them based on a set of rules applied to all the protocols. So, for example, if the best route became unavailable, GateD would switch to the next best route, maintaining connectivity and network availability, transparent to users, fundamentally providing a more stable routing environment. GateD was installed at all the NSFNET sites and at other sites that expressed interest. GateD was conceived by Scott Brim and initially developed by Mark Fedor, a systems programmer working for the Theory Center NISC. According to Brim (personal communication), the availability of GateD was crucial to the initial success of NSFNET. Over the next few years, Jeff Honig extended the initial GateD daemon to make it more robust and applicable to other protocols.

Research Computing and SuperComputing

CornellC, along with the FPS Array Processors, continued to be used as the high-performance computer

at Cornell, and was used extensively by researchers for their computations. However, there was a scare in the spring of 1984 when potential problems with floating point arithmetic were discovered.

Supercomputing at Cornell got off to a fast start using CCS facilities at Langmuir, but with the formation of the Cornell Theory Center, and its selection as one of the national supercomputer centers, the program moved into high gear when the computer room was completed in CCC.

IBM 3081D Floating Point Errors

In May 1984 it was discovered that floating point arithmetic on the 3081D was producing erroneous results. At first it was thought that the problem was contained between 12:00 noon, April 22, and Friday, May 11, but further investigation showed that the problem could have been occurring as early as December 29, 1983.³⁵ The December and April dates were connected with equipment changes. In December the High Performance Option (HPO) was installed; in April the Performance Machine Assist (PMA) option was installed. What made the situation worse was that the problem was intermittent and so not every program would be affected, but further, the condition applied equally to independently written programs and library programs. Users were advised to report immediately any anomalous results to CCS. IBM, once notified of the problem, issued an advisory that this problem was occurring on all 3081 systems with HPO under high load conditions.

A fix from IBM was installed on May 15 and normal operations resumed, although user advisories continued to be issued. When the problem was more or less officially closed in September after no further incidents were reported, Dan Barthomolew, who had been investigating the problem from the beginning, reported that there were only 14 problem reports, of which he estimated that fewer than half had errors. The quick reaction and the dogged pursuit of the problem by Barthomolew assisted by Nick Gimbrone had made the difference in resolving the problem. It was a relief to close this incident.

Supercomputing at Cornell

While the deployment of the array processors had taken scientific computing to a new level at Cornell, Nobel laureate Wilson continued to promote his vision of massively parallel supercomputing. When the National Science Foundation made the decision to

³⁵ “Update on Floating Point Arithmetic Errors on CornellC,” The Bulletin, Cornell Computer Services, No. 481, May 23, 1984. (First report in Bulletin 480, a week earlier)

support five national supercomputing centers, Cornell was poised to make a proposal to have such a center on its campus.

The Theory Center Is Formed

In 1984 under Wilson's leadership, the university approved the formation of a new research center, the Center for Theory and Simulation in Science and Engineering, which became known as the Theory Center.³⁷ The goals of the center were to provide scientific and programming support and computing for scientists, engineers, and scholars whose research required the development and application of complex, large-scale computational strategies. In addition to providing very powerful computers for scientists, the other main thrust of the Theory Center was to explore the strengths and uses of highly parallel supercomputers. The Theory Center was expected to have a budget of \$18 million by 1987 and a support staff of about 100.

In early 1985 the Theory Center was named as one of the five national supercomputer centers^t to be supported by NSF as a cooperative venture between Cornell, the state of New York, IBM, Floating Point Systems, and NSF. The NSF grant was for \$21.9 million over three years. Wilson, Ravi Sudan, and VP King were the principal investigators under the NSF grant. Wilson and Sudan became the director and associate director of the center when it first started. William L. Schrader was appointed the first executive director.

Included as part of the Theory Center was the Program of Computer Graphics that was to use the facilities of the Theory Center to improve graphics capabilities as well as to use those capabilities for displaying research results. Don Greenberg continued to be the director of this program.

The supercomputing program at Cornell got off to a quick start with the assistance of CCS staff and facilities. Using space renovated for computing equipment at Langmuir, the CCS systems programming and operations staffs quickly set up the Theory Center's Production Supercomputer Facility (PSF) so that it was up and running and accepting new users in May 1985.

The Production Supercomputer Facility (PSF)

The Cornell supercomputer center was operational before all the other national centers. The initial system was the CCS IBM 3081, which had been upgraded from a model D to K in December 1984. The upgrade increased memory and speed by about one-

third. This system, and the three attached FPS-164 array processors, was supplemented with additional disks and tapes and other peripheral equipment.

Planning was under way for further improvements in the computational power of the PSF and for creating space to accommodate all this new equipment and staff. In addition to all these ongoing activities by CCS staff, many of the first staff of the PSF were former staff from the Academic Computing division of CCS: Pete Seigel, Donna Bergmark, Ben Schwarz, Carol Hecht, Mark Scannapieco, and Judy Warren, all of whom had the applications software experience and who could provide the consulting services to the new larger user base. The Advanced Computing Facility, with the mission to consider new and highly parallel systems, was headed by Alison Brown, also previously with CCS.

The installation of this first supercomputer brought with it some controversy about the operating system to be used. Since 1974, when CCS first installed the VM operating system, the systems staff at CCS had become experts at extending the system and making it more useful. Even after IBM declared VM an official program product and staffed the development and support function accordingly, there were shortcomings that Cornell overcame with local modifications such as the Cowles Scheduler, SCMS, VMBatch, and the extensions for Adabas to use CMS for interactive online business transactions. As a result, Cornell systems programming staff became world-noted experts on the system and its internal operations. Given this expertise it was not surprising that this group preferred VM and was not a strong proponent of IBM's mainline MVS operating system. This bias aside, running high-quality (very stable and high performing) VM and MVS production systems would require more staff than running VM alone.

The VM-MVS issue rose to major proportions because IBM clearly favored the MVS whereas Cornell favored VM. The issue was resolved in favor of continuing with VM, as this would permit a quick start. Cornell researchers and support staff would not require any training because they had been using VM, and documentation and training could be provided for new users with a small number of support staff, using existing materials prepared by CCS. It was also argued that new users were likely to come from other higher-education institutions where VM was often the system of choice for IBM mainframes that supported academic computing. The MVS-VM issue continued to be discussed with each new computer installation until UNIX became the focus of attention. So while the use of VM was not a totally popular decision, in the end it was having scientists use the system that really mattered.

³⁷ "Supercomputers Will Tackle Simulation Problems," *Cornell Chronicle*, March 29, 1984.

Cornell National Supercomputer Facility (CNSF) Computers in CCC

The fast pace set by quick startup of the Production Supercomputer Facility (PSF) in May 1985 continued for the next two years as new, more powerful computers were acquired from IBM almost every 12 months. The need to excel and the availability of NSF funds in part drove this acceleration, which also reflected IBM's desire to position itself as a provider of supercomputing equipment. In October 1985, with the just-in-time completion of the computer room in the CCC annex, a four-processor IBM 3084QX system was installed. In addition, four new FPS-264 systems were installed, bringing the complement of FPS boxes to seven, counting the two existing FPS-164 systems and one FPS-Max system, moved along with all other PSF equipment from Langmuir to CCC. This configuration of the PSF was estimated to raise the computing power offered researchers to 200 megaflops (million floating-point operations per second).

In the summer of 1986, Lawrence A. (Larry) Lee, previously with the NSF supercomputer program, was hired as the director of the PSF. Shortly after the name was changed to the Cornell National Supercomputer Facility (CNSF), more in keeping with its mission as one of the five NSF-supported national supercomputer centers.

In October 1986, the 3084QX was replaced with a 3090-400VF system from IBM. This new system was a four-processor system and was also the first IBM system that had built-in vector facilities, much like the FPS systems, and so could execute scalar and vector arithmetic. The 3090-400VF at the CNSF was the first production machine installed outside of IBM. With software advancements, the four processors of this new system could be used in parallel, providing the first parallel supercomputer operations. When combined with the continuing FPS boxes, this new complex was rated at 600 megaflops, a rating by which the national supercomputer centers compared their power.

CISER

Although the Theory Center dominated research computing during the 1980s, other initiatives also improved access and services for researchers on the campus. In 1986, CISER and CCS cooperated to acquire an IBM 4381 to replace Cornella, previously an IBM 4341. CISER was able to obtain a grant of \$150,000 from NSF to improve computing access and services for social scientists. In effect, CISER purchased a 50 percent share of this system plus some private disk storage and shared this system with students and other casual users. The arrangement worked out well for both parties, although there were times when

student use overwhelmed the system, inconveniencing CISER users, who typically ran large batch jobs.

To simplify the access and thus extend the use of the data archives and computation resources of CISER across the campus, CISER created a user-friendly interface. To get around the powerful but complex CMS commands, a team from CISER and CCS developed a set of "full screen panels that will aid users in some of the basic (and not so basic) CMS and related software commands." The following statement from the 1986–87 CISER Annual Report nicely summarizes the accomplishments:

The current version of these panels includes the ability to alter virtual machine configuration; view existing CISER Execs; view information, create new, and alter the size of SMC minidisks; mount tapes in CMS; submit jobs to CMS VMBATCH; define temporary disks; perform file management; backup disk files to tape; and submit jobs to MVS.

The two principal staff responsible for this new interface designed to help social scientists exploit the power of computers were Tom Boggess, computing manager at CISER, and Dan Bartholomew from CCS Academic Computing. The annual report continues: "The computing group also began the development of a system that will allow users to access files from the data archive in a radically new fashion." A codebook would be used to select, store, and name datasets for later processing on the CISER computer or on users' personal computers. Bartholomew estimates that in the end there were 30 to 40 commands (panels) available to CISER users for all such functions.³⁸ Armed with this new way to access the computer and create data of interest, CISER was able to extend the use of the computer to a larger audience, increase the number of users, and promote the use of computers by social scientists at Cornell.

Information Services

The use of computers in the 1980s moved quickly to preparation and presentation of information. One of the key new developments was the use of computers to organize and create text documents (word processing) and to present stored document information to users at their CRTs using the locally developed CUINFO. The Cornell library continued to increase its deployment of computers to improve services to clients.

³⁸ Dan Bartholomew, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, July 3, 2001.

Word Processing Systems for Microcomputers— WordPerfect

When the microcomputer revolution took hold at Cornell, the options for word processing software quickly multiplied. For Macs there was MacWrite, which came bundled (free) with the operating system, later followed by WriteNow, which gained more prominence when Apple unbundled MacWrite. These applications and others, such as WordPerfect, shared the word processing software market for Macs at Cornell. For the IBM and IBM-compatibles, given the wide range of connections of faculty, staff, and students to the outside world, there were many applications on campus, including the popular WordPerfect product. Naturally, when problems were encountered, many users of these applications came to CCS for help. Often, no help could be given.

To bring some order to this chaos in late 1983 the Office Automation Group (OAG) undertook to evaluate the different systems to see which could satisfy a large number of needs and become the “officially supported” application for IBM and compatibles at Cornell. OAG considered the following: Easywriter II, Final Word, Multimate, NBI, Volkswriter, WordPerfect, Word Plus PC, Word Star, Word Vision, and XYWRITE. After due evaluation they chose WordPerfect, which was then being written by a small company in Utah.³⁹ Dr. Alan Ashton, one of the prime developers of WordPerfect and earlier a professor at the University of Utah, even visited the campus in 1985 and gave several seminars about the system and some of its advanced features. Naturally he complimented Cornell for selecting WordPerfect but also complimented CCS for trying to narrow the number of supported software packages as new ones were coming on the market every day, each touting some advantage or another. It helped that WordPerfect was available to Cornell users at a discounted price of \$198 compared with the \$495 commercial price.

WordPerfect was the main supported product during the decade until new features and interchangeability between systems became important, and Cornell support moved to Microsoft Word, which was available on Macs as well as on PCs. The WordPerfect Company was sold and resold several times, and its software was packaged with different products. Eventually product support and availability became spotty, and use of WordPerfect diminished.

Project Milestone—The Electronic Age Comes to Day Hall

In 1984 another innovative project was started: Bringing the Electronic Age to the Cornell University Executive Staff. The OAG took the leadership for

the project, and NetComm took the responsibility for wiring the offices on the third floor of Day Hall and establishing network connections. The idea was to place the machine of choice on the desks of the executives and their support staff to provide a variety of services. Direct-connect terminals could provide e-mail (still mainframe based) or access to administrative systems and information. Microcomputers could function in terminal mode for mainframe access or e-mail, or in local mode for using spreadsheets or word processing. Stand-alone or network-connected word processors were another choice. The OAG did extensive interviews with all the staff to make sure the options were well understood, and where necessary, did further training as in the case of electronic mail. The plan was to install the equipment on May 1 and, as an incentive, to allow for a four-month training period before charges began. Somewhere along the line the project name was abbreviated to “Project Milestone” as achieving the goal would indeed be a major milestone for computing at this time. Before it ended, the project was jokingly referred to as Project Millstone, given some of the challenges and difficult personalities that were encountered.

To facilitate executives’ access to online information a new capability was developed called EXECINFO. This was a menu-driven system that provided key data about Cornell and Cornell’s position among competitors. This was a joint development between Paige Ireland from Institutional Planning and Analysis and Steve Worona from CCS, the originator of CUINFO. EXECINFO was made part of the specially constructed Project Milestone menu to simplify access to the services. By year’s end most of the participants purchased the equipment that had been installed for them, as they were generally pleased to continue with the services. Most desktop installations were IBM PC-XTs, a few were CRT terminals, and there was a wide variety of printers depending on the quality of print desired. Once executives and their support staff became experienced they needed little more in the way of support from CCS. The CCS staff who participated felt that they had indeed achieved a milestone and went off to other more pressing projects.

CUINFO Grows, “Dear Uncle Ezra” Started

By 1985 CUINFO service had been extended to the DEC 2060 computer and was being used 130 times per day. In 1986 an important new item was added: the Alcohol IQ Network, which allowed people to test

³⁹ Cecilia Cowles, “IBM PC Word Processing Package Available to Cornellians at a Discount,” *The Bulletin*, Cornell Computer Services, November 9, 1983.

their understanding of this “drug.” Within six months this entry had been used over 4,000 times.

One of the more interesting offshoots of CUINFO was Dear Uncle Ezra, patterned after the Ann Landers daily column, to deal with students’ personal problems. The idea was conceived by Jerry Feist, assistant dean of students, and implemented by Steve Worona (the developer of CUINFO) and named in honor of the university’s founder. It was to be a way for students to share their personal feelings and concerns and ask specific questions electronically (and anonymously) as a supplement to counseling services. Dear Uncle Ezra was the first such service offered on an American campus when it was made available in the summer of 1986.⁴⁰ Students could submit any question or concern to Uncle Ezra using the public CUINFO kiosks. Both the question and the reply would be posted to Uncle Ezra. Students and the general public could log on to Uncle Ezra to read the latest questions and answers. Some typical comments were:

This week was my birthday and none of my friends remembered. I’m depressed.

Where is the bridge that everyone jumps off? I am considering it myself.

I am failing everything. I need serious help.

The identity of Uncle Ezra was not disclosed; confidentiality was maintained for both parties.

Desktop Publishing

Desktop publishing came into practice with the combination of the Macintosh computer, the LaserWriter printer from Apple, and the Postscript page description language from Adobe. Once other vendors like Xerox saw the power of this technology they came out with their own systems including computers, printers, and software to create a very active market. Given the large number of publications that were produced at Cornell, desktop publishing quickly became an important topic. In 1986, CCS and Apple held a full-day workshop, with representatives from vendors selling and supporting related technologies to provide education and information to the campus.⁴¹ CCS and Xerox held a similar event for the same purposes.

As with past technologies, the range of services quickly mushroomed. Although initially the work was done on a Mac or a PC and sent to a laser printer to produce “camera-ready” copy, soon it was possible to create a diskette that could be read on a higher-speed, higher-quality printer to create a document. Before the decade was over, the camera-ready copy could be sent over the Cornell network to the print shops on campus that would print the desired document. Desktop

publishing was one of those “killer applications” that made the microcomputer so successful.

One of the outcomes of this technology change was that CUTHESIS was discontinued in the fall of 1987, having served a useful purpose at the beginning of the decade but now obsolete in the face of the many other options.

Cornell Library Goes Online

Increasingly, the Cornell libraries understood they had to “automate” their information delivery services. The computer system created back in the late 1970s was one that dealt with the backroom functions—ordering books, tracking their receipt, payments, etc., and producing cards for the catalog. The card catalog itself was still to be the main way to find reference materials. The library continued to look for systems that could satisfy its needs, primarily those that would be sufficiently robust to handle the large number of books, over 5 million then in the holdings. Ryburn Ross, assistant university librarian, headed the team of library and CCS staff specialists evaluating the offerings. The new system had to be capable of supporting library acquisition, circulation, and serials subscription information as well as providing online public access for the Cornell scholarly community.

This effort was accelerated when, in 1984, the J. N. Pew Jr. Charitable Trust gave the Cornell libraries a grant of \$1.5 million to install a computer-based integrated information system.⁴² The first system considered was NOTIS, the Northwestern University Library Total Integrated System, and demonstrations of the system’s capabilities were scheduled for the campus. Other systems considered in the selection were LIAS (Library Information and Access System) offered by Penn State and Honeywell; BLIS (Biblio-Techniques Library and Information System) offered by Biblio-Techniques, Inc.; and the system offered by OCLC (the Online Computer Library Center). In 1985, BLIS was selected and a joint project was started with the libraries and CCS to implement the system, a key part of the project being the creation of a VM/CMS interface to BLIS. It was expected that the project would take five years to implement fully across the 16 libraries on the Ithaca campus and the library at the Geneva, N.Y., Experiment Station.

The BLIS system was installed in 1986 and testing had just started when the company folded, leaving

⁴⁰ Steve Worona, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, September 1999.

⁴¹ “Desktop Publishing Forum April 24,” *Cornell Chronicle*, April 17, 1986.

⁴² “Libraries Receive \$1.5 Million from Pew Trust for New System,” *Cornell Chronicle*, March 1, 1984.

Cornell and others in a quandary. According to Lynne Personius, then working for both CCS and the library system, the Cornell library went through a round of discussions with other BLIS customers to decide on the future of the system. In the end most participants formed a consortium to manage and maintain the system. Cornell chose not to join the consortium, even though the BLIS system had other advantages such as using Adabas, and instead reopened the selection process.⁴³

Another committee was formed; members included VP King, Russ Vaught, Jan Olsen from Mann Library, Jane Hammond from the Law Library, and Lynne Personius. After going through another round of vendor presentations, visits, and so forth, the committee selected the NOTIS system. After contract terms were settled work started on implementing the system that went live with the cataloging module in April 1, 1988. There was some concern about the April 1 date, but the gremlins did not prevail and the system worked well; other modules were brought into production after appropriate testing.

Mid-Decade Transitions

During the years 1984 to 1987 several significant changes occurred in the computer landscape and computing services at Cornell, including the location for the central computers and organization. There were some significant changes in staff that had an influence on computing at the university, and the dim financial outlook put a damper on activities.

The End of the Punch Card Era

In May 1986 the last vestiges of card technology—an IBM Card Reader and Card Punch—were removed from the Langmuir computer room. Punch cards had been used on campus for about 60 years by this time, although only intensively used with computers since the 1950s. By the mid-1980s, most of the punch card preparation and handling equipment had been removed from campus, and so removing card equipment from the computing site was the formal recognition that the punch card era had ended at Cornell. It was now the era of direct entry of information using CRTs or typewriter-like terminals or microcomputers networked to the computer.

The End of Mainframe Plotting Services

With the increasing prominence of microcomputers and the availability of graphics software, another

service came to an end: plotting of numerical results to display graphical representations of numerical data. From the installation of the Control Data 1604 in 1962, plotters, using different types of pens and colors and capable of different sizes of plot output, were directly connected to the mainframes. In the 1970s, plotters were remotely installed on campus to give users more control over their plotting needs. By the mid-1980s, with the availability of micros and bit-mapped screens, it was simpler to transfer information generated on mainframes or other computers to micros to display the data in WYSIWYG (What You See Is What You Get) form and then decide how to present the display.

DEC20 Removed—Vax Systems Installed

In mid-1986, the DEC20 was removed from service and sold to Digital Equipment Corporation as part payment for a new VAX 8500 system. One-third of the 8500 resources were committed to support the computing program in the business school. Larry Fresinski, systems manager, remarked that new technologies and decreasing software support for the DEC20 made this an opportune time to make the change. Since the VAX used the VMS (Virtual Memory [operating] System), this required a large adjustment for those who were used to the TOPS-20 system on the DEC20. However, it was expected that the switch would be reasonably easy, and training sessions were scheduled to facilitate the transition. Soon, the VAX 8500 became another standard offering to the campus for time-sharing central computing services.

New Deans and Executive Staff

In the mid-1980s, several significant changes were taking place in the university administration and in the economic climate. Provost Kennedy retired in June 1984 and Robert Barker, then vice president for research and advanced studies, succeeded him. Curtis W. Tarr was appointed dean of the Graduate School of Management; Tom Everhart resigned as dean of engineering to become chancellor at the University of Illinois-Urbana, and in 1985 William B. Street was appointed dean. James W. Spencer continued as vice provost. Harold (Hal) Craft, who had been director of telecommunications since 1982, was appointed acting VP for facilities, replacing Robert Matyas. In 1986, Geoffrey V. Chester was appointed dean of the College of Arts and Sciences, replacing Alain Seznec who shortly afterward became director of the university libraries.

In 1984 the S. C. Johnson family (Johnson Wax) gave a record \$20 million to the university to endow the Graduate School of Management, which then

⁴³Lynne Personius, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, April 9, 2001.

changed its name to the Johnson Graduate School of Management. A gift of this size to a graduate school of business was the first, not only for Cornell, but also for any such school.

Among other executive changes in 1987 were the appointments of James E. Morley Jr. as senior vice president, Larry I. Palmer as vice provost for academic programs, Malden Nesheim as vice provost for budgeting, and David S. Yeh as the first assistant vice provost for academic programs. Morley replaced William G. Herbster, who had come to Cornell in 1976 as the first senior VP. Morley had come to Cornell earlier as vice president of finance and treasurer, replacing Robert Horn who retired. Palmer had been in Day Hall for a number of years and in part replaced William D. Gurowitz who had served for many years as the vice president for campus affairs (student services). These appointments brought a new set of individuals into the increasingly important role of computing on campus, and each came to play a significant future role.

Dim Financial Outlook

In contrast to the expansiveness earlier in the decade, by January 1985 Provost Barker was urging caution about the continuation of income and expense increases. He noted the small margin the university had for making tradeoffs in expenditures. It was a time for assessing priorities and competing needs: salaries, facilities, programs, and other commitments. In February 1986 Barker announced the need for the university to cut expenditures by 5 or 6 percent over the next three to four years: "We must adopt a different lifestyle, living on a diet won't work." The tuition increase of 1986–87 was announced as 9.5 percent, down somewhat from the much higher increases of over 13 percent early in the decade.⁴⁴

CCC Computer Room Opens

In December 1986, the wall isolating the computer room space in the annex was removed to create one large computer room with a total of 11,000 square feet, enough space for both the CNSF production supercomputer and CCS computers. Most of the computing equipment from Langmuir was moved to CCC between Christmas Day and New Year's Day, when Cornell-based computing activity was at a minimum.

Just before this, CCS had ordered an IBM 3090-200 at a cost of \$4.06 million to replace the 3081K computer that had been in service since 1983. The new 3090-200 was scheduled for delivery in late December to avoid moving the 3081 system to CCC. It took a

year of planning and the combined effort of over 100 people from CCS and other Cornell staff and a variety of vendors to execute this move. Over 100 items of large equipment (refrigerator-sized units), one of which was the IBM 4381 computer, hundreds of boxes of supplies, and some 16,000 magnetic tape reels and cartridges from Langmuir were moved. Thousands of new network and cable connections in CCC had to be completed.

Besides this relocation, a scaled-back 3081K system continued to operate in Langmuir over the winter holiday break so that some critical year-end administrative functions could be executed before December 31. The 3081K system was returned to IBM in January while peripheral units such as disk drives were moved to CCC. The entire move went quite smoothly. The only glitch was a snowstorm that slightly delayed the delivery of the new 3090-200 computer from Poughkeepsie, N.Y., but it was available to some users on January 2 and was fully operational by January 5, 1987.

CCS Staff Move into CCC

The remodeling of the other floors of CCC was completed in the spring of 1987. Shortly thereafter most of the CCS staff were moved into CCC from their scattered locations across the campus and at Langmuir. NetComm stayed in Caldwell. The CCC building layout had Administrative Computing staff on the fourth floor, Computing Systems and the Business Office on the third floor, and Academic Computing on the second floor. The first floor housed the production staff (Operations, Production Control, and Data Entry), the Software Lending Library, and walk-in client contact functions such as building reception and obtaining computer accounts.

Plans originally called for about half the first floor for a CCS training facility and a large terminal and microcomputer laboratory accessible 24 hours a day. Despite strong efforts to follow this plan, the next



Figure 10. Computing and Communications Center (new addition to left)

⁴⁴ "Provost Says University Must Cut Its Expenditures by 5–6% in Next 3 Years," *Cornell Chronicle*, February 20, 1986.

supercomputer upgrade (the second IBM-600E) required more space than was available in the basement computer room and took over the space allocated to the training and terminal/micro facilities. Since this space was built with raised flooring, the supercomputing equipment in question, mostly associated magnetic tape and disk drives, could be installed without major room modifications other than additional air conditioning.

Office Equipment Center Merges with CCS Microcomputer Sales

In 1987 a new unit, Microcomputer and Office Systems (MOS), was formed in the Academic Computing division of CCS by combining the microcomputer sales operation of CCS with the Office Equipment Center (OEC). OEC had a long and successful history starting as the Typewriter Division in the 1950s under the sponsorship of Wallace B. Rogers.⁴⁵ The premise for the Typewriter Division had been simple and cost-effective: buy typewriters at highly discounted educational prices and lease them to departments at or below the cost of buying typewriters, writing off the standard educational pricing in three years. The arrangement was even more attractive because the Typewriter Division installed and maintained the equipment without additional cost! The key to success was the low initial cost of the typewriter itself and the relatively high resale value on the open market at the end of the lease. The big advantage to Cornell departments was being able continually to improve the productivity of typists and the quality of print as typewriter manufacturers improved their technology. Increasingly, the machines were mainly IBM typewriters as other vendors were leaving the market.

The success with the typewriter program was extended to other office equipment: copiers, then word processors, and then microcomputers. By the mid-1980s most administrative offices were down to a single typewriter for manually filling out forms, envelopes, etc., work that could not be replaced easily by electronic systems. As the future of OEC pointed more and more to selling and servicing microcomputers, it was obvious that OEC and the CCS microcomputer sales and service operations were redundant and competing for the same market. After much analysis, report writing, discussion in committees and other groups, and “to-ing and fro-ing,” the university administration decided that CCS should carry on this activity and that OEC would be phased out and merged into the CCS operation.

A new name was selected: Microcomputer and Office Systems (MOS). MOS was made part of the Academic Computing division of CCS. There followed the inevitable merger pains, but under the

capable leadership of Nancy Flynn as acting director and Bob Mindel as operations manager, the unit continued its record of service to the campus, pending the recruitment of a new director. Flynn previously had been marketing manager of PC sales in NetComm and Mindel had been director of OEC. There was also some stress as MOS was now operating out of more locations. It continued with operations at Langmuir and Caldwell but had added 110 Maple Avenue, the previous hub of OEC operations. Adding yet another complication, support functions such as consulting and the Software Lending library had moved to CCC by this time.

Ken King Leaves Cornell

Before the staff could fully settle into their new CCC offices, another major event occurred in mid-1987—VP King announced that he was leaving Cornell to assume the presidency of Educom⁴⁵ in Princeton, N.J. Educom had been struggling for several years to recruit a new chief executive, and King was an ideal choice. He left behind a grand legacy at Cornell. As the first vice provost (vice president) of computing at Cornell, he presided over an exceptionally busy and fertile period of innovation in technology at Cornell. King made major improvements and he raised computing at Cornell to a premier level among institutions of higher education. By any measure he met the objective he set out when he first came:

The administration realizes that computing is something Cornell has to excel at. This has the support of the president and the provost. If it didn't, I wouldn't be here. The climate exists for major improvement.

King also met another prediction he made when he first came. In May 1986 there were 674 interactive terminals and microcomputers in student facilities, exceeding the 650 that he was expecting would be available in five years.

Provost Barker appointed Norm Scott as acting VP after King left. At the time Scott was director of research for the College of Agriculture and Life Sciences and director of the Cornell Agricultural Experiment Station. The senior staff of CCS, namely Galloway, Vaught, and Rudan, worked with Scott on the large number of continuing issues and activities that had to be addressed. Departing from the in-house recruiting efforts of the past, Provost Barker hired

⁴⁵ Educom was the trade name for the Interuniversity Communications Council formed in 1964 and dedicated to the idea that digital computers offered an incredible opportunity for sharing among institutions of higher education.

Michael Luskin, an executive recruiter with Fischer Associates in Philadelphia, to assist in the search for a new vice president for computing and information systems.

Formation of User Groups CAUCUS and SUPER!

During VP King's tenure, computing staff continued to found user groups, a tradition that started back in the 1950s. Two user groups were started in the 1980s: CAUCUS and SUPER!

The first was CAUCUS (Cooperating Adabas Users with College and University Systems) led by King and Russ Vaught from Cornell and Steve Watson from the University of Washington at Pullman. Software AG, the provider of Adabas, was very supportive given its interest in increasing its share of the higher education market. The first conference was held in Pullman, Washington, in 1984. The second, chaired by Dave Koehler, was held at Cornell in October 1985, with over 20 attendees.

The second user group was SUPER! (Supercomputing by University People for Education and Research), organized by Larry Lee and Betsy Schermerhorn from the Theory Center with the support of IBM. Again, IBM was looking to promote interest in its emerging supercomputer equipment and to highlight some of the achievements being made at Cornell. The inaugural SUPER! meeting was held at Cornell on March 28–30, 1988, with over 160 attendees. Members of the organizing committee from Cornell were Lee, Schermerhorn, and Rudan. Other steering committee members were Carolyn Autrey-Hunley from the University of Michigan, John Connolly from the University of Kentucky, Geraldine MacDonald from Binghamton University, Nora Sabelli from the University of Illinois at Chicago, Michael Williams from Virginia Tech, and Joseph Yeaton from the University of California at Berkeley. Cornell continued to play a strong role in SUPER! for the remainder of the decade and into the 1990s. During that time, various Theory Center or CIT staff served on the program committee, chaired sessions, and presented papers.

1987 to 1989 at Cornell

The final years of the 1980s were highlighted by an increasing focus on creating new network services at Cornell, following the successful installation of the telecommunications infrastructure and telephone services. Nonetheless, Cornell Computer Services continued its strong support of the supercomputing program and computing services to the campus. The hiring of a new vice president for computing at Cornell brought considerable changes to the organiza-

tion, the policies, and the focus, which became information technologies (IT) instead of computing. The completion of the networking study created a springboard for action on networking and network services at Cornell. That was slightly deferred by the Internet worm incident that put Cornell in an international spotlight for the disruption of national and international network services.

CCS and Campus Networking Activities

Networking activity continued to expand with the addition of new technologies and services, but increasingly it became evident that Cornell needed bold action to create a more contemporary and longer-lasting network to serve the campus.

Status in 1987

In a 1987 study commissioned by VP King and Pat Nelson, director of telecommunications, James A. Muir from Cornell Computer Services reported on a number of contemporary alternatives for data communications on campus.⁴⁶ The following section from that report provides an excellent summary of computing and network services being provided by CCS at the time:

The last few years have seen equally great changes in the technology used by the Department of Computer Services. Where once most computing was done on terminals connected to two IBM mainframe computers, plus other terminals connected to minicomputers from VAX, PRIME, etc., computing is now highly distributed with about 5,000 Macintosh and 2,000 IBM (or compatible) microcomputers, several VAX minicomputers, and additional mainframe computers including those of the Cornell National Supercomputer Facility. It should be mentioned that IBM's SNA communications software is not used at Cornell. Many microcomputers are already interconnected among themselves via local area nets (LANs), which usually contain a number of microcomputers plus a file server and some kind of printer or printer server. In particular, there are ethernet networks in various campus buildings. Many of these networks are interconnected among themselves and to the mainframe computers to form a campus area network via a series

⁴⁶ James A. Muir, "Alternatives for Data Communications on the Cornell Campus." Cornell University, Departments of Computer Services and Telecommunications, July 14, 1987.

of gateways, as illustrated conceptually in Figure 2. None of these networks use the S/85 directly, although some of them use spare capacity in its wiring plant. Other terminals, microcomputers, or entire networks are interconnected via two other physical media—the campus broadband cable and the Proteon token-ring backbone campus network. A clear goal of these networks is to provide widespread access to computing resources on campus and elsewhere at low cost, consistent with high speed, and ease of management and use.

The campus area backbone network or Internet uses the TCP/IP protocol for data transmission and for addressing. For example, all Cornell addresses started with 128.84; the Applied Math Ethernet (IP) addresses start with 128.84.237, and a particular workstation on that ethernet might be 128.84.237.14. The backbone network is in turn directly connected to an increasing number of regional and national networks, including ARPANET, CSNET, NYSERNet, NSFNet, and others that use TCP/IP protocols. Indirect access is provided to the BITNET network, which uses the RSCS

store-and-forward protocol, from the CornellC computer.

For higher-speed data or video transmission in the future, the new campus wiring plant also provided two 1/2-inch CATV-type coaxial cables and (usually) two empty 1-inch subducts from the Main Distribution Frames in the telephone modules to each campus building. The relation of some of these fiber optic cables to the telecommunications switch modules is shown on a campus map in Figure 3. [not copied in this document] In addition, a fiber “backbone” was installed using AT&T 62.5 micron fibers paralleling the telephone system fibers with additional extensions to key buildings for data communications.

For reasons of security, control and maintenance, the data communications switchrooms are usually connected to but physically separate from the telecommunications switchrooms.

Instead of “Figure 2” referenced above, we include as Figure 11, which follows, a more comprehensive depiction of the Cornell networks in 1987, even though the diagram was prepared a year earlier in 1986.

CORNELL UNIVERSITY COMPUTER NETWORK

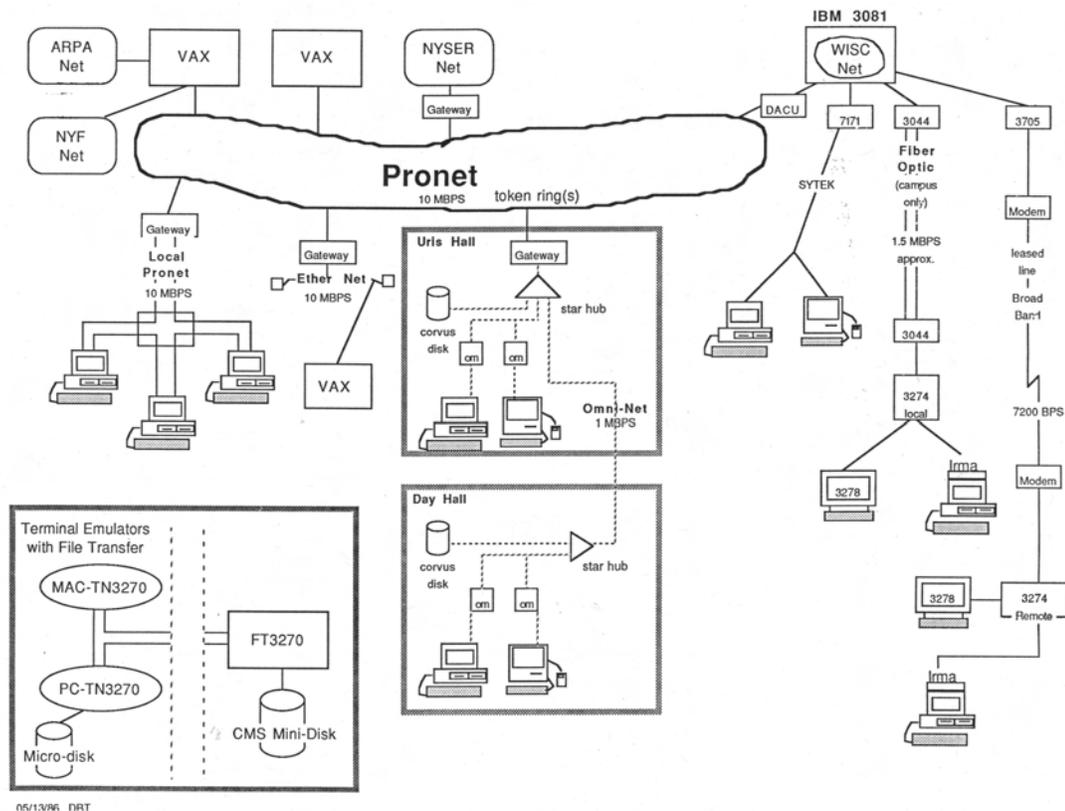


Figure 11. The campus network in 1986–1987

AT-Gateways for LAN Connections; Network Operations Center Started

Increasingly the low-cost distribution referenced in the “Gateway” in Uris Hall in Figure 11 was one of the AT-Gateways. During 1987 these devices were offered as a standard service/product offering with a pricing scheme that recovered costs but also encouraged users to connect to the better-performing campus backbone and move away from using Sytek.

In 1988 there were approximately 100 AT-Gateways installed on campus. With the increasing installation of LAN networked workstations, the AT-Gateways became the preferred alternative to get to campus desktop computers on the Internet. Accordingly, while Sytek connections were still offered, their number was declining. AT-Gateways were also increasingly being used to connect workstations on department Ethernets since Ethernet was one of the supported technologies.

Given this growing load, in mid-1988 the Network Operations Center (NOC) was organized to monitor networks during the weekday working hours, to accept calls for service, and to act as a liaison to technicians in the field. The NOC set up a cooperative arrangement with the NISC in the Theory Center for NSFNet and other external network connections.

As with most new technologies in their early deployment, the AT-Gateways had their problems. Sometimes it was hardware, but most often the failures were a result of software. To assist in the operations, SNMP (Simple Network Management Protocol) agents were deployed in both the Theory Center NISC and the CIT NOC to manage the Gateways remotely and to detect problems in advance.

The Call for Pincnet

At this time DEC machines could be networked together only by using proprietary DECnet networking technologies. In 1987, the Pronet backbone ring was not able to move DECnet packets over the network. The Peoples Independent Network Conspiracy was pushing to create a separate network—Pincnet—using the mature DECnet technologies to link their computers.⁴⁷ A feasibility test was conducted using a portion of the campus broadband network and modems and Ether-repeaters linked to a variety of computers and input devices. When this test proved successful, Pincnet promoters pressed to make this a more permanent arrangement. The issue continued to ferment for the following year, during which it became clearer that the longer-term direction lay in using the TCP/IP protocols rather than proprietary DECnet protocols, so no Pincnet was put in place.

NSFNet, NYSERNet, and GateD

During the 1987–88 period, the Theory Center and CCS continued to play increasingly important roles in developing statewide and nationwide networks, NYSERNet, NSFNet, and the development and deployment of GateD.

In addition to NYSERNet there was a growing number of similar networks connecting to NSFNet, all together forming an internet of large national research and education institutions. There was JVNCCNet at Princeton, SDSCnet in San Diego, PSCAAnet, and NCSANet, out of the other supercomputer centers, with other regional networks such as BARRNet in the Bay Area, Merit Network covering the state of Michigan, and SURANet covering higher-education institutions in the southeastern states.⁴⁸

By the summer of 1987 it was expected that 15 of the leading research institutions in New York State would be connected to NYSERNet. The Theory Center’s continuing role as the operating agent for NSFNet called for continued support and expansion of the use of GateD.

In early 1987, eight other sites had installed GateD and plans were under way to install it at all the other NSFNet sites. To support and improve the GateD software two important events occurred that year. Martyne Hallgren from the Theory Center organized a consortium of parties interested in GateD to increase the number of participating institutions and use the associated membership fees to improve the software. In return for the fee, members received the GateD software. Hallgren became the director of the GateD Consortium. In addition, NSF awarded the Theory Center a grant of \$88,422 later that year in support of GateD software development (personal communications).

Up until mid-1988, the Theory Center continued to be responsible for NSFNet operations and during this time served as the gateway between NYSERNet and NSFNet and ARPANET. In July 1988 NSFNet was upgraded from 56 Kbps to 1.544 Mbit. Shortly thereafter, management of the NSFNet backbone was transferred to the Merit Computer Network at Ann Arbor, Michigan.

Networking Study

It was during 1987–88 that acting VP Scott initiated a formal study to develop a networking strategy for linking computers at Cornell and with other off-campus

⁴⁷ Gene Zeigler, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, November 2001.

⁴⁸ “The National Networks; A Status Report,” Center for Theory and Simulation in Science and Engineering, Cornell University, *Forefronts*, Vol. 2, No. 12, April 1987.

computers. Despite the earlier-stated objective to prepare a Cornell network plan, the swirl of other activity had prevented the development of a plan for an integrated campus network. While the piecemeal networks available on campus were working, it was clear that a more comprehensive approach was needed.

Acting VP Scott appointed Alison Brown, who by this time was associate director for networking in the Theory Center, to head the Networking Task Force.⁴⁹ The group included members from different campus constituencies with Dick Cogger and John Rudan acting as staff support. Brown stated the situation as follows:

Our plans for the committee are to develop an array of solutions that are not too restrictive, but that allow interoperability of campus networks. This means, for example, that users should be able to exchange electronic mail and easily access data on other networks.

Soon after the announcement and appointment of the task force, Brown resigned her position at Cornell to become the associate director of the Ohio Supercomputer Center.^v Acting VP Scott took over as chairman and completed the assignment, the results of which are discussed later in this chapter.

Business Systems

In September 1987, Russ Vaught, director of administrative computing, issued a strategic plan for administrative information systems and shared it broadly on the campus. It had a five-year horizon and was expected to evolve as technology and institutional needs changed. It repeated the statement, "It [the plan] implicitly accepts the principle that both central administrative offices and operating units are simultaneously consumers and providers of information," continuing the reinforcement that data in systems were a university resource. A number of high-level goals were stated and the following more specific goals and objectives were listed:

- Completion of and enhancement to online integrated systems that serve the primary central administrative functions.
- Adopt where necessary development technology which provides for an integrated administrative office environment in a secure and cost-effective manner.
- Development of Decision Support and Institutional Data Summary Systems which provide vertical integration of information.
- Support of college and departmental offices' information needs.

- Development of systems that enhance faculty productivity in areas such as course enrollment and advising.

Beyond this, the document was an excellent, comprehensive compilation of recent history, technology, recommended practices, shared responsibilities for business systems, and future directions for administrative computing at Cornell and for the Administrative Computing department in particular. For example, it discussed the "information pyramid," where data flow from base transactions systems to institutional data summary systems to decision support systems. It also emphasized the importance of these transitions in system design and integration. While this report was issued when Norm Scott was acting vice president, much of the work took place during VP King's tenure at Cornell.

Improvements continued to be made in providing continuing services, such as printing, but in addition a new innovative system was developed to distribute information in the central accounting system to operating units.

Table of Central Administrative Information Systems

The 1987 Strategic Plan contained a "Table of Central Systems" that presented the status of administrative information systems in a unique way.⁵⁰ It used icons to clarify the kind of impact each system had. The information has been slightly altered to fit on a single page (Table 2). (Footnote 20 has been removed, since it has no bearing on the information presented, which remains the same as the original report.)

CUDA (Cornell University Distributed Accounting) System

In May 1987 a major improvement was made in Business Systems with the introduction of CUDA (Cornell University Distributed Accounting) to the campus. This was the first attempt to implement distributed administrative computing at Cornell, and in a way it was an early client-server implementation taking advantage of the increasing availability of desktop computers among administrative staff. The objective was to provide departments with a software tool to better manage their finances while maintaining compatibility with the information in the central accounting system.

CUDA was a joint project of the Controller's Office, Administrative Programming staff, and several key

⁴⁹ "Task Force Seeks to Link Computers On Campus and Off," *Cornell Chronicle*, March 10, 1988.

⁵⁰ Russ Vaught, "Cornell University, Administrative Information Systems Strategic Plan," September 1987.

ADMINISTRATIVE INFORMATION SYSTEMS STRATEGIC PLAN - DRAFT

*	System has strategic importance for Cornell's Mission.
\$	System has financial importance to Cornell.
👤👤	System has impact on faculty, students, employees, or external agencies.
🏠	System has information on Cornell's physical facilities.
📝	System has information on Cornell's programs.
📁	System under development.

- Developed since 1981 and in at least partial use by offices.

Academic Personnel *\$👤👤/	Microfilm Management \$🏠
Accounting *\$👤👤/ 📁	Preventive Maintenance \$🏠
Admissions *\$👤👤/	Public Affairs *\$👤👤/ 📁
Bursar / Cornellcard \$👤👤 📁	Residence Life \$👤👤🏠
Capital Projects *\$🏠/	Sponsored Programs *👤👤/
Cornell Annuity 👤👤	Stud Emp/Cornell Trad *\$👤👤
Customer Service \$🏠	Student Information 👤👤
Facilities / Space Inventory *\$🏠/	Telephone System \$🏠
Graduate Admissions *👤👤/ 📁	Traffic \$👤👤🏠
Grad Financial Aid *\$👤👤 📁	Undergrad Financial Aid *\$👤👤
Grad Student Information 👤👤 📁	Utilities Data System \$🏠

- Purchased system to be operational in next year

Library On-line Catalog *👤👤 📁

- Developed prior to 1981 and may need replacement. (📁 signifies some work towards replacement is underway at present.)

Affirmative Action *\$👤👤	Payroll / Personnel *\$👤👤 📁
Capital Equipment ²⁰ \$🏠/	Responsibility Center Analysis \$
Equal Opportunity 👤👤	Stores Enterprises \$🏠
General Stores \$📁	Student Loan \$👤👤 📁
Job Cost \$ 📁	

- Systems for which development has been requested.

Budget System *\$👤👤🏠/	Purchasing \$🏠/
Career Center *👤👤	Research Resources *\$👤👤🏠/

Table 2. Administrative Information Systems Status

administrative departments. JR Schulden and Clint Sidle played key roles in the design and implementation, assisted at different times by Dean Eckstrom, Steve Barrett, Kevin Hufford, and Gail Landowski from the Administrative Programming unit.⁵¹

The system design called for information from the Central Accounting database, which was updated daily, to be downloaded to a desktop system at which the staff could use pre-programmed screens to display data, or, if they desired, to write their own programs to analyze and display their information. First developed for the IBM-compatible PC and then for the Macintosh, the system proved very popular. The first release of the system was to 20 users. The 1988–89 CIT Annual Report by Dave Koehler makes the following commentary:

Brought the Cornell University Distributed Accounting (CUDA) System into full production. Approximately 60 departments use CUDA to produce customized income and expense statements, track local financial commitments, and plan, track, and produce grant reports.

Following Cornell tradition, a separate charge was made to use the CUDA system to support the incremental costs of the staff needed to keep the system up to date, to make improvements, and most important, to provide onsite training and troubleshooting when problems arose.

High-Speed Laser Printing Systems

In 1987 a new Xerox 9700 laser printer was installed to replace one of the 8700 printers installed earlier. By this time there were two Xerox 8700 laser printers, the first installed in 1984 and the other in 1986, to accommodate the increased print load that was approaching 1.5 million pages per month and trending toward their peak capacity. The 9700, compatible with the 8700, was capable of printing 120 pages per minute (vs. 70 for the 8700), had the capacity to store more pages in its internal print queue, and had a longer duty cycle between maintenance periods. In addition, owing to favorable terms from Xerox, costs were kept essentially the same. Credit for the success of the use of these high-speed lasers is due to the diligent work of Peter Baker and Peggy Roberts from central operations who provided the programming and technical support.

⁵¹ Louise (JR) Schulden, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, June 2000.

Research Computing

A major focus for Cornell Computing Services was providing computer support services to the Theory Center during the 1987 to 1989 period, particularly to assist in the move toward parallel computing, the use of multiple processors on a single problem. In addition, CCS provided support for the design of a new computer room in the center's new building. That effort got wrapped up with the design of the building itself and the travails that were encountered.

The Push to Parallel Computing at CNSF

In July 1987, an IBM 3090-600E, IBM's most powerful computer, costing a total of \$20 million, replaced the 3090-400 VF. The 600E, a six-processor system with vector facilities, was rated at 696 megaflops, giving the total configuration now with only five FPS boxes (two having been removed) a total rating of 886 megaflops.

In October 1987 a second 600E was installed in CCC following a two-year \$19.3 million grant from the NSF to the Theory Center. Even though some other equipment was relocated to the first floor this second system was “shoe-horned” into the CCC basement to satisfy the requirement that the two processors be within 40 feet of each other for operation as a single 12-processor system. This squeeze was achieved by clever positioning of the boxes and using minimum equipment clearances, which at times required the careful sequencing of the opening of equipment doors for access and maintenance! There was some concern that the air-conditioning capacity of the raised floor would cause heat dissipation problems, but besides a few hot spots this was not the case.

The credit for carrying out these less than yearly complex installations of major computer systems with so few problems is due to Richard Alexander, who designed the configuration and the interconnection of systems, and to Dave Pulleyn and his senior operations staff of Gene Caraccilo, Bill Biata, Ben Brown, and others. This rate of change, which was much greater than the previous rate of installing a system every two or four years, represented the pace at which computing was changing and evolving all over CCS and the campus.

Continuing the thrust to investigate “coarse-grained” parallel systems, that is, systems with 12 processors, it was planned that the two 600Es would operate in tandem by extensions to the operating system and development and improvement of parallel Fortran. After many attempts, this approach did not prove fruitful and the project was dropped.

This upgrade of the second 600E was the last in CCC. In November 1989, the Theory Center received another grant from the NSF, \$79 million over the next

six years, and was planning new equipment from IBM and SSI^W (Supercomputer System Inc.) for the years 1991–95. Those plans, however, became intertwined with the expectation of a new Theory Center building for which planning started to take place in 1986 with CCS staff playing a key role. However, SSI failed to build any computers. As a result some of the provisions for the SSI systems, such as additional load-bearing sections of computer room floor, were never used.

High-Performance Computing—A National Concern

The actions of the Theory Center were in concert with the report on high-performance computing issued by the Office of the President (USA) in 1987.⁵² Four recommendations were stated in the summary of recommendations pertaining to high-performance computers, software technology and algorithms, networking, and basic research and human resources. In each case it recommended long-range plans, strong support, and leadership by the federal government. While the report acknowledged the strong worldwide leadership of the United States in these important areas, concern was expressed that other nations were taking much more aggressive action that would diminish that leadership. The report defined “high-performance computing” as “the full range of supercomputing activities including existing supercomputer systems, special purpose and experimental systems, and the new generation of large-scale parallel architectures.” The report introduced the term “grand challenges,” fundamental problems in science or engineering whose solutions would be enabled by the application of future high-performance computing resources. Many examples of these grand challenges were given: fluid dynamics for the design of hypersonic aircraft, weather forecasting over short and long terms, electronic structure calculations for designing new materials, computations for speech recognition, and others. The term “grand challenges” found its way into Theory Center discussions and plans.

Mass Storage Systems

An important issue of the supercomputing initiative was “mass storage.” It was expected that supercomputers would generate enormous quantities of information, much of which was going to be shared with others or used in future computations. As such, the data needed to be stored with long-term retention in mind and yet be easily retrievable. A more effective solution was needed than simply buying more disk stor-

age when that on hand was full. It was prohibitively expensive to store terabytes (10^9 , or trillions of bytes) of information on the then-available disk technology. What was needed was a hierarchical storage system that migrated items to increasingly larger capacity and less expensive media while retaining sufficient control to protect the information and retrieve it in a reasonable amount of time. Such a system would also provide backup protection so that data stored in the system were not inadvertently lost.

The central computing organization at Cornell, OCS, had engaged in mass storage joint studies with IBM back in the 1970s, and so staff in the Systems Programming group revived the collaboration. Also by this time the IEEE had organized the Technical Committee on Mass Storage and had created an emerging standard, the IEEE Mass Storage Reference Model, for discussion and evaluation purposes, although there was no production implementation. Paul Zarnowski and Andy Hanushevsky played the initial leadership roles, and Hanushevsky assumed the longer-term role as project leader for the Cornell effort.

The first option considered for the CNSF was a working system at NCAR, the National Center for Atmospheric Research. Shortly after the first contacts were made in 1987, NCAR assigned the rights to its system to Mesa Archival Systems, a private corporation set up to develop and market its product, the Archival Data Library System. While Mesa and Cornell agreed that it would be mutually beneficial to develop such a system to run in the VM environment, almost all of 1988 was spent in trying to agree on an acceptable joint study contract, although work had started without a final contract being signed. By the time serious work was under way in 1989, Mesa ran into financial difficulties, and while work proceeded there was no exchange of materials between Cornell and Mesa, given Mesa’s condition.

Despite these difficulties in bringing a working system into production, mass storage systems (MSS) continued to be a topic of interest even though it was still more theoretical than practical. By the end of the decade, Cornell was also interested in using an MSS to store digitized library holdings that would also require massive amounts of data storage with the same cost and access problems as research generated data. In 1989, CIT sponsored a Mass Storage Issues and Directions meeting that was organized by Andy Hanushevsky and held in Ithaca. The meeting was attended by representatives from NCAR, from the other national supercomputing centers, and from other research organizations cognizant of the need for such data storage, along with members of the IEEE Technical Committee. The objective was to bring about a common understanding of the issues and the

⁵² “A Research and Development Strategy for High-Performance Computing” issued by the Executive Office of the President (USA), Office of Science and Technology Policy, November 20, 1987.

activities at each of the institutions. Some home-grown systems at other institutions were considered for Cornell, but there was no implementation of an MSS at Cornell during the 1980s.

The Theory Center Building

All along the Theory Center had plans for its own building. By 1985–86 it had outgrown its space in the Caldwell basement plus the space in a complex of large trailers situated behind Carpenter Hall on the Engineering Quadrangle. These trailers provided additional office space, meeting rooms, and training facilities. In April 1986, New York State announced a \$10 million contribution for a building to house Cornell's national supercomputing center; half of this amount was an outright grant while the other half was an interest-free loan. Plans at that time called for a 100,000 gross square foot (gsf) building located on the north end of the Engineering Quad on Campus Road across from Sage Hall. Discussions at the time noted that the Engineering College was the center of "gravity" of supercomputing at Cornell, so the Theory Center building should rightfully be close to the college. That location was soon changed when in mid-1987 the next design developed by Gwathmey Siegel and Associates called for an eight-story, 272,000 gsf building consisting of two rectangular connected units to be built at an estimated cost of \$29 million. This design would give the top four floors to the Theory Center and the lower four to the College of Engineering. The top two floors would be used mainly for computing and networking equipment. It was to

be located on what was then the Grumman Hall parking lot, south of Upson Hall but encroaching into the Cascadilla Gorge.⁵³

There were immediate objections to the height of the building and its encroachment into the tree line of the gorge. These objections by the City of Ithaca, environmentalists, and residents living across the gorge led to many reviews and meetings and design changes during the latter half of 1987. The proposed eight-story design would come within 25 to 45 feet of the gorge, but construction clearance would add another 10 to 15 feet to this encroachment and require 135 trees to be cut down. Regarding the height, representatives from Gwathmey Siegel commented that the eight-story design was taller than normal owing to "computer cabling requirements," which added 24 inches for the raised floor on each of the top two stories.^x To address the environmental concerns about the gorge tree line, a proposal was made to move the building closer to Grumman Hall and away from the gorge, reducing the number of trees to be cut to 45. A further refinement was proposed that would require no trees to be cut and would situate the building at least 20 feet from the tree line by erecting it over Grumman Hall, in effect encapsulating Grumman in the new building.

⁵³ "Theory Center Building Site Sparks Controversy," *Cornell Chronicle*, August 27, 1987.

⁵⁴ Taken from Center for Theory and Simulation in Science and Engineering, Cornell University, *Forefronts*, Vol. 8, No. 4, December 1988.



Figure 12. GE robot breaks ground for new Theory Center building, on command by the center's founding director, Ken Wilson (at right of control panel).⁵⁴

In January a new innovative design was proposed that satisfied all concerns but whose estimated cost had now risen to \$30 million. The new banana-shaped building, following the curve of Campus Road, had a block-shaped segment at the north end and a glass-enclosed staircase at each end. It was to be 30 to 40 feet from the tree line. Further, it would be only seven stories tall, rising at most 134 feet above ground instead of 160 feet. The size would be reduced to 197,000 gsf, much less than originally proposed. All parties accepted this design. In November 1988 ground was broken for the building using a GE robot shovel to dig the first spade. Construction started soon after and was well under way by the end of 1989 for occupancy to take place in 1990.

The Movement from Computing to Information Technology

Norm Scott's tenure as acting VP for computing came to an end in 1988 with the appointment of M. Stuart Lynn as vice president. Following his coming to Cornell VP Lynn started to reshape the organization and its direction to more contemporary themes and services.

M. Stuart Lynn Appointed Vice President for Computing

In March 1988 the university appointed M. Stuart Lynn as the new vice president for computing at Cornell. Lynn had extensive experience in computing in higher education serving as director of Rice University's Institute for Computer Services and Applications from 1971 to 1977 and director of the Office of Computing Affairs at the University of California at Berkeley from 1977 to 1983. During his time at Berkeley, Lynn was also professor of electrical

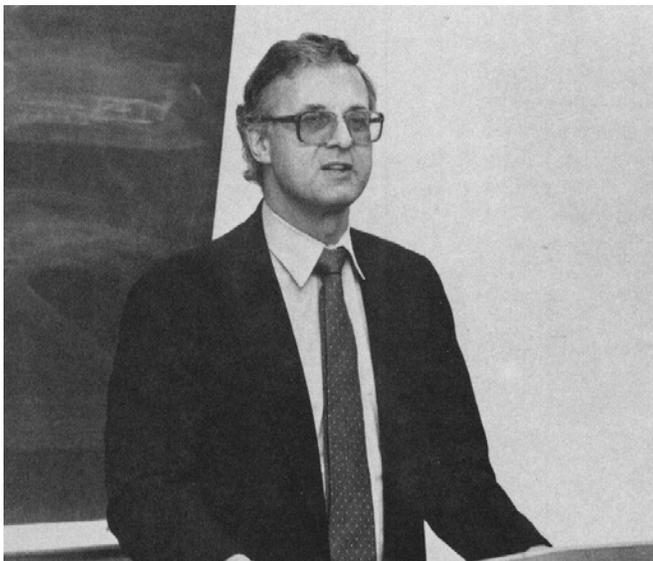


Figure 13. M. Stuart Lynn, vice president of information technologies

engineering and computer science. Before coming to Cornell, he had served as president and board chairman of Capital Technologies Corp. (USA), a financial services corporation.^Y

Cornell Computing Services to Cornell Information Technologies

VP Lynn quickly set about changing the basic orientation from computing to information technology (IT). He argued that it was time to shift the focus from the *tools* of computing to the *uses* of those tools. Besides, the tools consisted of more than the computer; they included networks, databases, file servers, software of all varieties, and multimedia workstations. His point was that concentrating on the tools alone distracted from understanding the needs of the community in putting the tools to work.

In setting this new tone on campus he separated his office, responsible for overall campus IT, from the operating and service units. His office became OIT (the Office of Information Technology), and the operating units were grouped under the name Cornell Information Technologies or CIT for short. CIT became the official name on November 1, 1988. After allowing for construction time, VP Lynn's office in Day Hall was moved into a suite on the third floor of Day Hall, closer to his vice presidential peers.

Mission of Information Technologies

After meeting with a large number of users and consulting with CIT staff, VP Lynn developed the premise that the mission of CIT should be built on three pillars: leadership, advice, and service. He developed the following mission statement.

The Mission of the Office of Information Technologies is to provide leadership, advice, and service to the university community in the academic and administrative uses of information technologies in support of the university's mission.

At its broadest level, information technologies is concerned with the electronic capture, verifications, transmission, storage, retrieval, transformation, summarization, and presentation of digital data and information; and with the interface between such digital electronic media and other information and communications media. The primary tools of information technologies are computers, data storage media, and electronic networks.

The overall mission of the Office of Information Technologies is accomplished through:

Leadership in facilitating the use and, as appropriate, development and implementation of the underlying *technologies* themselves in support of the university's needs;

Leadership in providing or nurturing the *services* necessary to enable the university community to apply those technologies to its needs;

Leadership in supporting the *uses* of such technologies and services in those areas where central support is most effective in facilitating such applications across the spectrum of university requirements;

Leadership in coordinating the definition and support of appropriate *policies, plans, standards, and control mechanisms* to provide an effective architectural framework for innovation and for the development, implementation and use of information technologies throughout the university.

These four interlocking leadership roles demand: a need for a collaborative approach driven by the requirements of the university community; a need for the highest possible standards of achievement

in each area of selected activity; a need to maintain competence and to exercise imagination in order to meet those standards; a need to exercise leadership and service with respect and consideration for others; and yet a need to focus activities and to balance limited resources carefully so as to fulfill these other needs within a diverse, complex and rapidly changing technological environment.

This mission statement is taken from an early internal-CIT version as that statement was in development. A later version with some modifications was published in the *CIT Newsletter* for public dissemination.⁵⁵

New CIT Organization

In February 1989, after a process of what VP Lynn called "Perestroika—CIT Style" (taking off from the political restructuring going on in the Soviet Union under Mikhail Gorbachev), a new organizational structure was announced for CIT. It is probably best to show the organization chart in its simplest form.⁵⁶

As is evident, this chart closely reflects the mission statement by mapping the four leadership elements of the mission—technologies, standards, applications,

⁵⁵ M. Stuart Lynn, "The Mission," *Cornell Information Technologies Newsletter*, Vol. 1, No. 2, December 1988.

⁵⁶ "A New Organizational Structure of Cornell Information Technologies," *Cornell Information Technologies Newsletter*, Vol. 2, No. 2, February 1989.

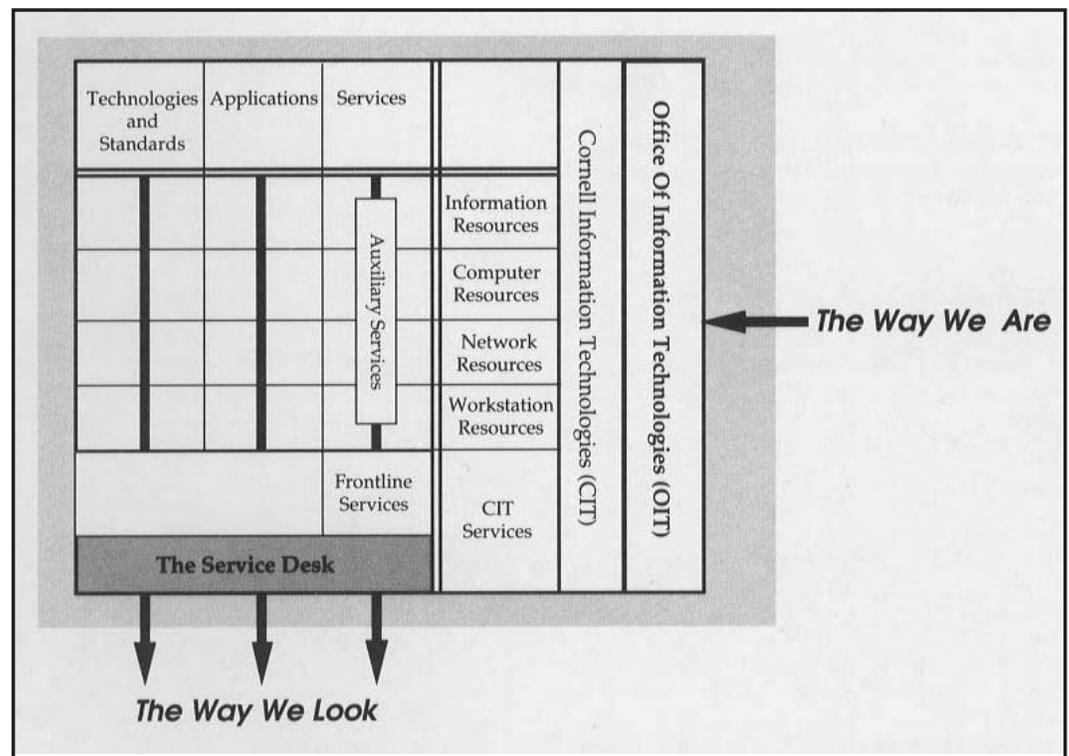


Figure 14. CIT organization model

and services—in each operating division of the organization, while reserving a special focus for the CIT services division. The services division—including the Service Desk and Frontline Services—was positioned as the principal entry point for customers as they approached CIT.

Each operating division was given primary responsibility for a segment of university users that fit best with the other responsibilities for that division. For example, Computer Resources was assigned the responsibility for research and analysis applications in part to reflect the computers it operated, the software offerings to research users, and its involvement with the Theory Center. Similarly, Information Resources dealt with a set of users reflecting the different administrative system groupings.

CIT Internal Advisory Committees: Technologies, Applications, Services

The key to cross-organizational cooperation and advancement on technical and service issues was entrusted to three infrastructure advisory committees—Technologies and Standards (later just Technologies), Applications, and Services—formed to advise CIT senior management staff. The initial chairs for these committees were Bob Cowles, Mark Mara, and Ruth Sabean, respectively. Members of these committees came from the appropriate groups in the divisions. The committees' work was directed in support of CIT's mission and shared the same basic philosophy: "to provide leadership in establishing an effective and cohesive framework for the applications of information technologies," stressing critical examination of services and products and cooperation across CIT and the campus. Further they were all to examine tactical (one- to two-year) and strategic (three- to five-year) plans that included goals, objectives, controls, and estimates of resources needed for implementation.

The Services Committee moved quickly into action because there was an obvious need to examine current service offerings in the light of changing requests for support, emerging technologies, and new applications in the midst of organizational change and the need to sort out funding priorities. Product reviews across the organization started with presentations by the relevant assistant directors in the different divisions.

The Technology Committee was charged with producing a technology "framework," a statement of direction and not a plan. The committee's initial report, "Technologies Framework—Computing at Cornell 1993–95," formed the basis for discussions and review inside CIT before it was to be released to the campus. Given the rapid changes in technology, the intention was to update this document annually.

The Applications Committee spent the rest of the decade organizing itself and coming to grips with the widespread applications crossing so many different uses and user groups.

Each committee also addressed the assignment of responsibility for different technologies, applications, and services within CIT and what cooperative efforts would need to be undertaken.

CIT Organization Chart

Table 3 lists the initial staffing of the new organization, with staff who joined later in the decade in [square brackets].

The core of the organization was made up of former management staff. Within a short while most of the vacancies, given that there were almost twice as many assistant director positions, were filled with new recruits from inside CIT or Cornell or from other agencies. The most significant changes were as follows:

- H. David Lambert was appointed director of network resources in March 1989 after a national search. Lambert had previously been assistant director for network services at Indiana University and had previously held positions at Indiana in user services and technical services.

- Dave Koehler was promoted to director of information resources, succeeding Russ Vaught who resigned in 1989 to accept the position of director of academic computing at Penn State University.

- Laverne Thomas was appointed director of the Office of Management and Budget. Thomas came from Binghamton, N.Y., where she had been budget administrator for Broome County.

- Carolyn Lambert was appointed assistant director for resource support services.

- Tom Boggess from CISER was appointed assistant director for systems programming services.

- Barbara Skoblick was appointed assistant director for research and analysis systems.

- Gordon Galloway, assistant for external relations, resigned and took a position as professor of chemistry at Michigan State University.

New Publications: *CIT News*, *CIT Briefly* Replaces the *Bulletin*

The strong program in place for training, publications, and information dissemination continued. A new publication, the *Cornell Information Technologies Newsletter*, was started in November 1988, to reach the campus not only on IT policy issues but also about specific technology or application issues and general information. By 1989 this publication became *CIT News* and was announced as part of a package of new publications.

Table 3. Cornell Information Technologies (CIT) Organization Chart, 1988

CIT Services	Ruth Sabean, director
• Consumer Services	Nancy Flynn
• Support Services	Ruth Sabean (acting) [Carolyn Lambert]
• User Services	Donna Tatro
Workstation Resources	Larry Fresinski, director
• Technologies and Standards	Tom Young
• Instructional Systems	Larry Fresinski (acting) [Carrie Regenstein]
• Personal Systems	Margie Wilson (acting) [Carrie Regenstein]
• Workstation Services	Jim Manning (acting) [Gordon Suggs]
• Workstation Systems Services	Tom Young
Network Resources	John Rudan, director (acting) [Dave Lambert]
• Technologies and Standards	Dick Cogger
• Communication Systems	Steve Worona
• Network Services	Alan Personius
• Network Systems Services	Dick Cogger
Computer Resources	John Rudan, director
• Technologies and Standards	Bob Cowles
• Research and Analysis Systems	Dan Bartholomew (acting) [Barbara Skoblick]
• Computer Operations Services	Dave Pulley
• Production Services	Jim Doolittle
• System Programming Services	Bob Cowles (acting) [Tom Boggess]
Information Resources	Russ Vaught, director [Dave Koehler]
• Technologies and Standards	Tom Dimock
• Academic and Human Resource IS	Libby Gruppuso
• Business and Financial IS	Dave Koehler [JR Schulden]
• Office Support Systems	Mark Mara
• Scholarly Information Systems	Lynne Personius
Office of Information Technologies (OIT)	
• Planning/Coordination	Agelia Velleman
• External Relations	Gordon Galloway
• Publication Services	Cecilia Cowles
• Management and Budget	George Cameron (acting) [Laverne Thomas]
• Security Officer	(to be named)

CIT Briefly replaced the weekly *CIT Bulletin* (yellow sheet) in November 1989. Number 748 was the last bulletin in the series that had been published continuously since January 1975, changing names from the OCS to the CCS to the *CIT Bulletin* along the way. The objective was unchanging: to bring timely and concise information about current events to the attention of the campus community.

A year later another publication was started—*Inside CIT*—an experiment to gather and disseminate information to CIT staff in a timely manner.

CIT Logo

In 1987–88 a CIT logo was developed to create visual identity for CIT.



Training and Workshops

After a faculty survey conducted in the summer of 1989 identified directions for future training by the late fall of 1989, a robust schedule of seminars, workshops, and training sessions was organized for the fall semester, the winter break of 1990, and the following spring.

In the fall of 1989 the training schedule was heavily weighted toward microcomputer training, about equally divided between Macs and IBM, which had recently announced their new PS/2 systems. Topics covered the operating systems for the different platforms but also applications programs such as word processors, spreadsheets, database management systems, desktop publishing, and emerging applications such as Hypercard and Hypertalk on the Macintosh systems. While the mainframes continued to be discussed, the presentations related to them shifted to the use of MAIL systems and BITNET networking technology. Fortran had finally faded from CIT training offerings, given the increased emphasis on applications and the availability of Fortran training for the supercomputers at the Theory Center. Statistical programs continued on the CIT training calendar.

Help Desk

The CIT Help Desk was organized in August 1989 as the physical and electronic (virtual) implementation of the service desk of the organization chart. It was located on the first floor of CCC. Almost all business transactions with CIT—such as computer accounts, special concerns, access to publications, and problem reporting—came through the Help Desk. It was to be the “one-stop shop” from the casual interaction to scheduled appointments for more extensive consultations. Resource Services staff was supplemented with experienced students to provide more hours of coverage for the convenience of users.

Online consulting service was introduced that same year to allow any user with access to one of the mainframe systems to use electronic mail to send a problem description to the Help Desk (SHDX@cornellA) and receive a reply via e-mail. A prompting feature, CONSULT, was offered on all the CIT mainframe systems where, once a user was logged on to the system, typing CONSULT would initiate a narrative to assist a user to send a problem to the Help Desk.

Microcomputers/Personal Computers Update

MOS Becomes Part of the Services Division

When MOS (Microcomputers and Office Systems) was first formed in 1987, it was part of the Academic Computing division of CCS. When CCS evolved into CIT, MOS became part of the Services Division. To cut down on the number of separate operating locations, to take advantage of the space at 110 Maple Avenue, and to release needed space in Caldwell Hall, the MOS sales office was moved to 110 Maple Avenue in 1988. At this time it was considered as one of the frontline components in the Services Division.

The following year, 1989, the Distribution (equipment warehouse and pickup point) and Repair Services units of MOS became part of Workstation Resources and moved from Langmuir to a nearby new facility at 33 Thornwood Drive. CIT had made leasing arrangements with a developer to build a designed facility for its needs at 33 Thornwood, in the area near the airport being developed by Cornell Real Estate.

Product Fairs Evolve into Educational Distribution Events

In 1988 the first educational distribution event took place at the beginning of the fall semester to facilitate the distribution of large numbers of purchased computers to new or returning students. Previously, product fairs had been held in different public locations, typically the west lounge of the Statler Hotel or Willard Straight Hall, at different times of the year and involved only a single vendor at a time—IBM, Xerox, HP, Apple, or others. The 1988 fair, held in Willard Straight Hall with IBM and at Lynah Rink with Apple, was the first timed for the start of the semester. This experiment was so successful that it became an annual event.

The second event, in 1989, was held in Lynah Rink in a “fair” type environment with balloons and decorations and with booths for vendors that included Apple, AST, Epson, Hewlett-Packard, IBM, Kodak, NeXT, and Toshiba. MOS sales staff was on hand to take orders for equipment that would be delivered to incoming students that day or in a day or two. Faculty and staff and everyone on the campus considering the purchase of a computer or printer or other peripheral were invited to attend and in this way collect information quickly and conveniently. CIT and vendor staff were on hand for those with questions. CIT staff not directly involved with microcomputer sales and support were asked to volunteer in all kinds of capacities. When the fair closed, all workers, including vendors, participated in a grand celebration party. Nancy Flynn,



Figure 15. CIT staff taking orders for personal computers in Lynah Rink

Peggy Fluman, Joan Manheim, and others deserve credit for their efforts in organizing these distribution events.

NeXT Workstations

The NeXT workstations were introduced on campus in January 1989. The NeXT “scholar’s workstation” was selected by the Department of Computer Science to support instruction (CS 314, Introduction to Computer Organization and Architecture; CS 432/33, Databases; and CS 412, Compiler Writing), and 30 workstations were installed in Upson Hall. Cornell was one of the first institutions to select this system and to help test some of the software. The NeXT system was distinguished for its all-black color and one-foot cube CPU, something typical of Steve Jobs. More important and more interesting were several technical innovations, most notably the removable read/write magneto-optical disks, the new screen management based on Display Postscript, and the use of the Mach Operating System, compatible with Unix BSD 4.3. Display Postscript from Adobe was the first truly “WYSIWYG” system and became dominant for laser printing. The NeXT operating system later became the core for Mac OS X when Apple bought NeXT in the 1990s.⁵⁷

Steve Jobs, founder of Apple Computer and founder of NeXT, came to the campus in February 1989 to demonstrate his vision of the workstation that could be used for education innovation in the 1990s and beyond. Jobs’s presentation was quite a spectacle, held

in Bailey Hall, which was filled to capacity not only to see the NeXT system but to see and hear this pioneer of desktop computing.

Cornell Faculty Continue to Win Awards for Software

In 1989 two Cornell projects won EDUCOM/NCRIPTAL Higher Education Software Awards.⁵⁸ NCRIPTAL, the National Center for Research to Improve Post Secondary Teaching and Learning, was cooperating with EDUCOM in its quest for better tools for classroom teaching.

- The Best Engineering Software Award was given to J. Robert Cooke, E. Ted Sobel (both of Cornell’s Department of Agricultural and Biological Engineering), and D. C. Davis (Washington State University) for MacPoisson: Instructional Finite Element Analysis for Solving Poisson’s Equation with the Macintosh.

- The Best Biology Software Award was given to Marcia L. Cordts, Ronald Beloin, and Jane Gibson (Department of Microbiology, the Ecosystems Research Center, and the Biochemistry Section of the Department of Molecular and Cell Biology, respectively) for their Tutorial in Recombinant DNS Technology.

⁵⁷ Gordon Galloway, “The ‘NeXT’ Computers Come to Cornell,” *Cornell Information Technologies Newsletter*, Vol. 1, No. 2, December 1988.

⁵⁸ “CIT Congratulates Cornell’s EDUCOM/NCRIPTAL Software Award Winners,” *CIT News*, Vol. 2, No. 6, September 1989.

MacPoisson was available as a PC version, PC-Poisson, whereas the Tutorial on Recombinant DNA was only available for Macs since it relied on the HyperCard software, unique to Macs.

Microcomputers/Personal Computers on Campus

Toward the end of the 1980s it was too difficult to keep track of all the projects and faculty innovations because the number of active projects became very large. While Project Ezra had concluded by this time, the Mac•Ed Center was still accepting proposals for projects to develop “courseware” for Macintosh systems. More important, faculty members were continuing to develop such software, introducing it into their courses and discussing it at various local and national forums.

It was commonplace at this time to have a personal computer in each faculty office and to find extensive use of such computers in classrooms and in homework assignments. The growth pattern was exponential, going from an estimated 500 systems in 1980 to roughly 10,000 systems by 1989, a 20-fold increase in 10 years.

MUGWUMP, the local Macintosh users group, continued to exist and had a membership list of over 200 members. Meetings were held as often as four times a month, including a general monthly meeting on campus and other meetings in the evenings devoted to special interests, such as desktop publishing.

Networking and Telecommunications

The Networking Task Force report was issued in 1989 and formed the basis for moving ahead with future improvements in network services from those offered at the time, depicted in Figure 16.

Networking Task Force Report

In March 1989 the Networking Task Force under the direction of Norm Scott issued its report on the future of networking at Cornell.⁵⁹ They envisioned the network at Cornell to be ubiquitous, transparent, extensible, and adaptable, secure and capable of integrating existing and yet-to-be developed technologies. They proposed the following goals:

- Increase scholarly productivity and collaboration through improved access to and integration of information sources, processing, and technologies.
- Enhance student learning through the integration of information technologies with every aspect of the learning environment and by using the network for new educational environments.
- Enhance the campus library system as the central resource for scholarly information including electronically encoded information; provide access at the scholar’s workstation to electronic bibliographic, full text, numeric, image and all types of data files, available locally and across the world.
- Enhance administrative systems by providing for secure, efficient, and effective management of the flow of administrative data between the central administrative offices and operating units of the university.
- Facilitate external cooperation by providing access to and operation of an electronic information system to enhance and broaden the interaction between the university, the

⁵⁹ “Cornell University Networking Task Force Report,” March 1989.

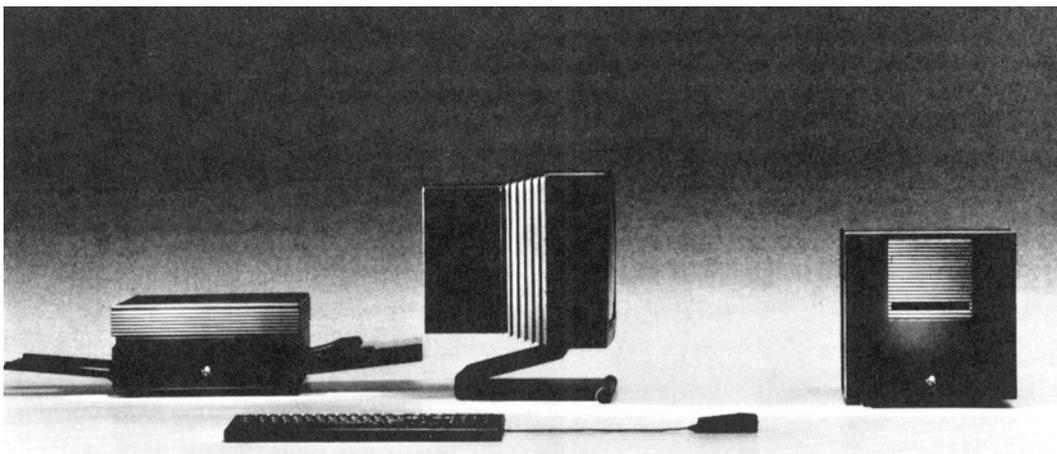


Figure 16. NeXT workstation—printer, keyboard, monitor, CPU

local community, New York State audiences, and industry in support of research, education, and extension processes.

- Establish a system of governance by creating a structure, establishing a process, and defining a policy to ensure equitable access to networking resources for a broad community of users.

The report went on to describe the then-current situation and steps that needed to be taken to achieve these goals.

The report, the subsequent appointment of H. David Lambert to the position of director of CIT Network Resources, and the transfer of the Department of Telecommunications to the division of Network

Resources (August 1989) provided the springboard for action to implement new networking services in the future. The Cornell network offerings at the end of the decade are depicted in Figure 17.⁶⁰

Other Communication Services

Various communication service improvements were made before the close of the decade. Fax machines became an increasing presence for transmission of information both on and off the campus. While most machines operated using a telephone line connection,

⁶⁰ "Connecting to Campus Networks Through AppleTalk Boosts Your Resources," *Cornell Information Technologies Newsletter*, Vol. 2, No. 3, March 1989.

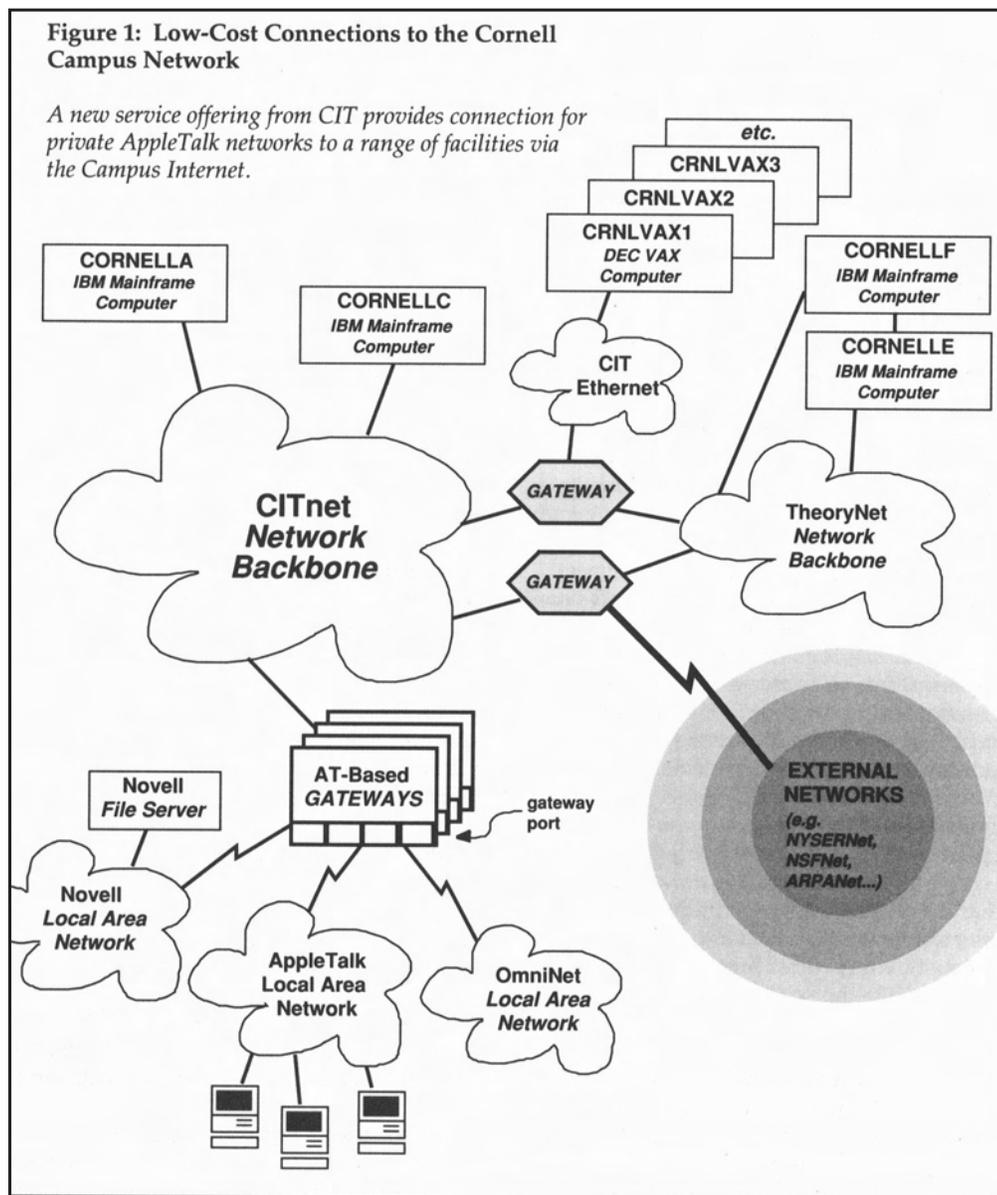


Figure 17. Cornell networks in 1989

it was also possible to send or receive a fax from a personal computer by having the proper hardware features and software.

“Voice-mail” was introduced on the campus in October 1989 with the implementation of AUDIX (Audio Information Exchange) on the Cornell telephone system.

Information Services

Information services continued to evolve, and during these last years of the decade a new mainframe-based e-mail system was installed. CUINFO continued to expand its offering and availability at public kiosks on campus, and different applications were available for sharing information over networks. Dear Uncle Ezra continued to be popular, and a compendium of previous exchanges was published as a book. The library automated its operations with the installation of a new computer-based library management system and moved to provide an information system for scholars to find and retrieve information from different media and sources.

Electronic Mail—New System for IBM Computers

Throughout the 1980s, mainframe-based e-mail continued to be the main provider of e-mail services, whether on the IBM or DEC systems at CCS or on other systems at Cornell. In early 1988, VP Lynn appointed a committee to consider options for a new electronic mail system for the IBM mainframes at CIT. The committee, chaired by Dave Koehler, evaluated Mail2 and a new Mail3 (both developed locally by Worona) and “Rice-mail,” and recommended conversion to Rice-mail. Rice-mail was a shareware system developed at Rice University and installed at many other universities, especially those in the BITNET community. Shortly afterward Rice-mail was installed and labeled as the CMS Electronic Mail System. Although some controversy surrounded this change, it should not obscure the significant contribution made by Worona, who brought the first e-mail system to Cornell, well in advance of other institutions. The fact that Rice-mail was in use at other universities, was being supported and improved collectively, and provided new features for saving mail, however, outweighed the disadvantages of continuing on an independent path and some loss in the “user-friendliness” of commands. To provide for a transition period the old system was continued for three months.

By this time, there was an increasing number of mail systems from different vendors running on different computers. Some of the most popular were available on Macintosh computers and Mac LANs. There was Intermail from Interactive Network Technologies,

Inc., Inbox from THINK technologies, and QuickMail from CE software. QuickMail proved to be the most popular and long lasting. When first started, QuickMail and the other mail systems ran only on Macs on single AppleTalk LANs. But because these independent LANs, within or between buildings, were linked with gateways using installed telephone plant wire to create a campus internet of LANs, use of these systems grew. The mail system evaluation committee started to address the question of how mail could be routed across all the networks. There was some discussion of creating a central mail server, but no action was taken.

E-mail etiquette became an increasing concern, especially the avoidance of too much abbreviation and cryptic phrases, a carryover perhaps from mainframe message practices. Humor crept in with the increasing use of emoticons such as :-) to represent a smiling face or other expressions :- (.

CUINFO

By 1989 there were 12 public CUINFO kiosks located in libraries, in Gannett Clinic, in Willard Straight Hall, in the Dean of Students office, and other high-traffic locations. Also by this time the content had grown to include some 35 items, including Cornell Dining’s co-op menus, the weekly list of job opportunities at Cornell, the current weather report, Career Center news, General Stores catalog, student and staff directories, and the Campus Code of Conduct.⁶¹ During academic year 1988–89 CUINFO was accessed a total of 330,000 times, an increase of 10 percent over the previous year. Also by this time the system had been exported to Penn State, the University of Rochester, and Carleton University in Canada, among others.

Dear Uncle Ezra

Dear Uncle Ezra was popular from the time it started in 1986. By 1988–89 it was one of CUINFO’s most popular items, answering a wide variety of questions and concerns. In a typical month some 175 students posed questions and perhaps ten times as many read the exchanges. In September 1989, Jerry Feist and Steve Worona, the co-developers, published *The Best of Uncle Ezra*, a selection of items that appeared in Dear Uncle Ezra between September 1986 and July 1987. Copies were on sale at several outlets on campus.

A new experimental feature, “Conversations with Mr. Chips,” was introduced as part of CUINFO to communicate anonymously with the Office of the

⁶¹ “CUINFO-Information at Your Fingertips,” *Cornell Chronicle*, January 19, 1989.

Vice President for Academic Programs. That same month over 60 screens worth of questions and answers appeared on a variety of issues.

Bulletin Boards and Electronic Conferencing

One of the manifestations of networked computers, whether on campus or across the world, was the development of other forms of communication for the exchange and sharing of information. In 1989, the *CIT Newsletter* carried articles about USENET and BITNET LISTSERV to bring these applications to the attention of users. (Vol. 2, No. 1, January 1989 and Vol. 2, No. 6, September 1989, respectively.) USENET's "newsgroup" metaphor allowed readers to explore topics of interest at any time, catching up on recent "posts" from around the world. The LISTSERV mailing list software on BITNET similarly offered discussions on many varied topics, but only "subscribers" would get each message.

Regarding USENET, the *CIT Newsletter* wrote, "The USENET News System is the largest electronic bulletin board in the world. Over 8,000 computer systems exchange thousands of messages, called articles, posted every day by subscribers to the news systems. Articles are organized into categories, called newsgroups, covering hundreds of topics from computer systems to movie reviews." The article went on to say that USENET was one of the earliest bulletin boards, created in 1979 by students at Duke University to connect Unix systems at Duke and the University of North Carolina using the UUCP networking protocol. Now USENET traffic crosses many different networks including BITNET, the TCP/IP Internet, and various DECnet networks. To get access to the discussions, interested parties were advised that "several departments participate in USENET, including Computer Science, the Theory Center, and Information Technologies. CIT offers public access on their VMS systems, CRNLVAX5, to all members of the Cornell community."

Regarding LISTSERV mailing lists on BITNET, the *CIT News* (note the change in name) wrote, "Spanning the globe is a multitude of people—invisible to you and to me and even to each other—who have formed their own universe of electronic conferences. At all hours of the day and night, they (users) are busily discussing their ideas and opinions about an unimaginable breadth of topics. . .from religious beliefs to organizational behavior, from intercultural relationships to the latest in copy-proof paper." "The best way to learn about the variety of conferences available is to look at the file LISTSERV GROUPS on the BITNET disk." All one had to do was to type "OBTAIN BITNET" on any of the Cornell IBM mainframes to get access to the BITNET disk and go on from there.

These services supplemented, and then finally supplanted, the local "bulletin board systems" (BBSs) that sprang up to allow local discussions, and even rudimentary networked discussions, exchanging messages via telephone connections.

Library System Goes Online

The initiative to automate the library system operations culminated in November 1988 with the installation of the NOTIS system, a joint project of the Cornell University Library and CIT. The most notable feature was the online catalog that could be searched from over 100 dedicated terminals throughout the libraries and another 150 terminals reserved for staff use. The catalog was a union catalog showing every holding of the system, regardless of which of the 16 campus libraries housed the item. By December 1989, the catalog was accessible from home and office computers.

While the online catalog was powerful alone, future plans also called for automating circulation, providing remote access to the catalog, keyword searching, and downloading entries into a personal database. A major impediment to enabling these processes was that few individual library records were in machine-readable form. It was necessary to bar code all the holdings to automate circulation. Following this, the library needed to develop a patron database for identifying individual borrowers. These were large and expensive projects.⁶²

Libraries Provide More Information Resources

Beyond the circulation of materials, the libraries were moving aggressively into being providers of information using electronic sources.⁶³ Different initiatives at the time included database search services using commercial services such as DIALOG or COMPASS search service performed by librarians or by the individual. In all cases, there was a charge for using these commercial services. CD-ROM databases were another option to get information, and they were starting to replace the microfilm/microfiche distribution of information. Nine of the Cornell libraries had CD-ROM workstations at which individuals could insert the relevant disk such as ERIC, Medline, or GPO (Government Printing Office) both on SilverPlatter and Drugs Database. Specialized databases were available in the different colleges, schools, and departments and were a growing activity. It was noted that the total cost of 38 CD-ROMS acquired by Cornell exceeded

⁶² "Computer-Literate Users Ease Library Automation," *Cornell Chronicle*, May 12, 1988.

⁶³ "Access to More Information Sources," *Cornell Information Technologies Newsletter*, Vol. 2, No. 4, April/May 1989.

\$90,000! Peter Martin, former dean of the Law School, was creating the first “treatise and law library” research tool that would allow full-text search and retrieval of legal materials. Given the amount of material, the proposal was to use CD-ROM technology for storage. Mann Library and CIT were working on a joint project, the “Scholarly Information System,” to make large citations of scientific journal literature available on campus. The CISER Data Archive was yet another source of information.

Business Systems

A new administrative systems and data policy was formulated and adopted in the last years of the 1980s. As part of this policy a new executive oversight group for business systems and new departmental data administrators were created in defined functional areas. New high-speed laser printers introduced color printing for printing special forms and color highlights.

Administrative Systems and Data Policy

The previously issued “Administrative Information Systems Strategic Plan” had barely been digested on campus when a new policy document—“Administrative Systems and Data Policy,” prepared by Russ Vaught and VP Lynn—was issued in 1988 and adopted in 1989.⁶⁴ The essence of this policy is contained in the following summary:

This policy provides a coherent framework for the management of administrative information systems and data within the university. Particular emphasis is placed on those systems and data subject to central purview (as defined by this policy). The policy provides a taxonomy for the classification of systems and data. The policy recognizes central data as a university asset which may be produced and used by multiple units; central, college, or departmental. The policy establishes organizational responsibility for different facets of activity; and also provides a framework for effective planning.

The policy covers information about different views of data, different classifications of systems and data, and ownership responsibilities of systems in addition to many other aspects of the functioning of this important activity. Several notable items provide perspective on the evolution of business systems.

- The policy formally established the principle that data are a university asset that is to be shared across the university.

- The policy defined the role of the responsible data administrator—RDA, for each functional area responsible for creating data definitions that are consistent and useful across the university. This role would typically be at the vice presidential level.
- The policy defined the role of the responsible system administrator—RSA, with overall responsibility, coordination, and accountability for the development, maintenance and operation of the assigned central system
- The policy established a high-level Administrative Systems Steering Committee, ASSC, to be appointed by the provost and sr. vice president, who would also be members. Other members would be the vice presidents for information technology, planning and budgeting, and finance and treasurer. The director of information resources, the university auditor, and the university budget officer would be nonvoting members. While acting as advisory committee to the provost and sr. vice president, the major responsibility of the committee was for the execution of the policy.

The Office of Information Technology (OIT) was charged with the responsibility for planning for systems and data that were subject to central university purview with two levels of planning. The strategic plan had a three- to five-year planning horizon, to be updated every year on a rolling basis. Tactical plans were to be prepared yearly and tied to the budget cycle. As with other such policy formulations, this policy served as the framework for the development of systems until the decade ended.

High-Speed Laser Printing

The upgrading of high-speed laser printer technology continued. In 1989, the remaining Xerox 8700 was replaced with the first Xerox 4090 printer to provide highlight color printing (red, blue, or green) while maintaining total print capacity at a reduced cost and with a smaller footprint. In that year, 50 new special report print routines were developed, adding to the 177 already in production. The end of preprinted special forms was fast approaching.

Status of Mainframe Computers at CIT

As the 1980s were ending, the load on the mainframe computers at CIT continued to increase, causing overload conditions and poor service times. These conditions were overcome mostly by upgrading the operat-

⁶⁴ Administrative Systems and Data Policy,” Office of Information Technologies, May 1, 1989.

ing system software. The use of these computers for statistical computing continued to be a strong service offering.

Mainframe Computers Struggle to Keep Up with the Load

Toward the end of the decade the mainframe computers, CornellA, CornellC, and VAX5 were beginning to seem like “dinosaurs” of computing. The client-server (C-S) model—a desktop personal workstation linked by a network to any number of servers—was getting increased attention as the new model for future computing. Nonetheless, as access to these older “workhorses” grew rapidly with the proliferation of network-connected terminals and workstations, and as these systems were the main source of large-scale administrative and research applications and of electronic mail, mainframe capacity was pushed to the limit.

As in the past, the first approach for CornellA and CornellC was to upgrade the operating systems. In 1988–89, VM/XA (Virtual Machine/Extended Architecture) was installed on these IBM systems for improved throughput and with added features such as larger memory spaces for virtual machines. With these changes some features were lost, such as line-mode access (line-at-a-time transfers between the computer and the remote terminal), but by 1989 this was trail-

ing-edge technology as full-page transfers were the preferred method. In addition, for these IBM systems the change in operating system forced a change in disk storage, so the IBM 3370 disks were upgraded to 3380 technology, thus providing more storage at lower cost and better performance from increased transfer rates.

On the VAX 5 computer, the VMS operating system was upgraded to version 5, which significantly improved system performance by changing the way multiprogramming was done on these multi-processor systems.

Statistical Computing

Statistical programs for both simple and complicated analyses of data continued to be a strong offering of CIT during the decade. Their use on the mainframe systems probably reached its highest point by the end of the decade both for being the sole source of many such programs and for the rich variety of options available. In the next decade, desktop systems became the preferred choice for many of the statistical analyses because of their increasing storage and processing power, ease of data entry, and the ability to adapt analyses quickly, as well as the availability of software. The richness of the offerings is best displayed by the summary in Table 4, taken from Vol. 2, No. 7, the December 1989 issue of *CIT News*.

Statistical Software on CIT Mainframe Computer Systems

CMS ON CORNELLA	CMS ON CORNELLC	MVS ON CORNELLC	VMS ON CRNLVAX5
SAS	SAS	SAS	SAS
SAS/STAT	SAS/STAT	SAS/STAT	SAS/STAT
SAS/GRAPH	SAS/GRAPH	SAS/GRAPH	SAS/GRAPH
SAS/IML	SAS/IML	SAS/IML	SAS/IML
SAS/ETS	SAS/ETS	SAS/ETS	SAS/ETS
SUGI	SUGI	SUGI	
MPROBIT			
MLOGIT			
SPSS-X	SPSS-X	SPSS-X	SPSS-X
		LISREL	
	BMDP	BMDP	
EQS			
IMSL	IMSL		IMSL
			GLIM
MINITAB	MINITAB		MINITAB
MATLAB	MATLAB		MATLAB
			TSP

The table above shows the availability of mainframe statistical software under each of the major operating systems (CMS, MVS, VMS) on CIT's large computers (the IBM mainframes CORNELLA and CORNELLC, as well as the VAX super-minicomputer, CRNLVAX5).

Table 4. Statistical Software Available on CIT Mainframes

The Cornell Worm Incident

The decade came to a close with Cornell University and the Department of Computer Science gaining more national prominence than they wanted when in November 1988 a graduate student in Computer Science reportedly released a “virus” that infiltrated thousands of computers nationwide and disrupted network services as a result. The student was identified as Robert T. Morris, a first-year graduate student, who shortly after the incident took a leave of absence from the university. Using e-mail and exploiting some known loopholes in the operating systems, his rogue program attacked VAX 5 systems running Berkeley Unix as well as SUN systems running OS3 or OS4. Over 100 computers at Cornell were affected and thousands globally were affected as the program replicated. The rapid replication clogged the NSFNet because of network traffic being generated as the “virus” spread. However, since the CNSF supercomputers were not using the UNIX operating system they were not affected.²

Provost Barker asked VP Lynn to head an investigative commission to review this incident. Early on it was determined that no one at Cornell in a position of responsibility authorized or had knowledge of Morris’s project. The investigative commission with representatives from the Department of Computer Science, Cornell Legal Counsel, and other departments was to gather material on Morris’s alleged involvement and violation of Cornell computer policies and practices. Since the FBI was responsible for investigating any violations of federal policies VP Lynn was also working with them. A report on the computer virus⁶⁴ was issued in early 1989. The report was comprehensive in scope and detailed in its evidence of the sequence of events that took place, supported by many appendices. Appendix 1 in particular, “A Tour of the Worm” by Donn Sealey of the Department of Computer Science at the University of Utah, provides an almost minute by minute chronology of how the event unfolded.

Provost Barker delayed the release of the report until the responsible U.S. attorney decided what legal action would be taken. In August 1989, Morris was indicted on a federal felony charge by a federal grand jury and released on his own recognizance. The felony count under the 1986 Computer Fraud and Abuse Act carried a maximum penalty of five years in prison, a \$250,000 fine, and a provision for restitution to those affected by the criminal act. The case was not settled until early in 1990.

Following the “virus” incident the University Computing Board issued the following statement of principles:

The use of computers and network systems in no way exempts you from the normal requirements of ethical behavior in the Cornell University community. Use of a computer and network system that is shared by many users imposes certain additional obligations. In particular, data, software, and computer capacity have value and must be treated accordingly.

Legitimate use of a computer and network system does not extend to whatever you are capable of doing with it. Although some rules are built into the system itself, these restrictions cannot limit completely what you can do and can see. In any event, you are responsible for your actions whether or not rules are built in, and whether or not you can circumvent them.

The following specific principles apply to all users of Cornell computers and network systems:

- Respect the privacy of other users’ information, even when that information is not securely protected.
- Respect the ownership of proprietary software. For example, do not make unauthorized copies of such software for your own use, even when that software is not protected against copying.
- Respect the finite capacity of the system and limit your own use so as not to interfere unreasonably with the activity of other users.
- Respect the procedures established to manage the use of the system.

Those who cannot accept these standards of behavior will be denied use of Cornell computers and network systems. Violators may be subject to penalties. . . under state and federal laws.⁶⁵

While in most ways this closed the “Morris” incident, its aftereffects continued well into the next decade, both internationally and at Cornell.

⁶⁴ “The COMPUTER WORM, A Report to the Provost of Cornell University on an Investigation Conducted by The Commission of Preliminary Enquiry,” February 1989.

⁶⁵ “Principles for proper use of computer networks,” Cornell Chronicle, January 18, 1990.

Other Campus Computing Activities

In this section we continue with the story of the Dairy Records Processing Lab as recorded by Lyle Wadell. Wadell's story ends in this decade although the lab continued operations into the 1990s.⁶⁶ We also highlight the computer-related stories taken from the *Cornell Chronicle* during the 1980 to 1989 period.

Dairy Records Processing Lab

As our load on the system 370/138 grew it started to impact the efficiency of our data entry operations to such an extent that employee productivity was dropping. To take care of this we installed an IBM 4331 computer and IBM 3370 disk drives in early 1981 to handle the data entry operation with capabilities to switch data files to the system 370/138. The system 370/138 was finally replaced by an IBM 4341 in May 1981 and in the fall of 1981 the IBM 4331 was returned.

In the fall of 1984 we started purchasing (as opposed to leasing) our equipment for the first time in our history. This started with the installation of a new IBM 4381 computer, additional IBM 3375 disk drives, and a new 3705 teleprocessing controller. This was done as a financial move to try and control escalating hardware costs. At the same time in 1984 we replaced all of our cathode ray tubes by purchasing IBM 3180 CRT's under a state contract arrangement resulting in significant savings.

At the same time we made these changes to attempt to partially control costs, the cow numbers dropped from 651,000 in 1984 to 625,000 in 1985. Thus far in 1985 we are seeing some of the cows come back with a projected volume of 644,000 by January 1, 1986, being handled by a staff of approximately 44 full-time equivalents.

CADIF (Computer-Aided Design and Instructional Facility)

A part of a 1981 grant from the J. N. Pew Jr. Charitable Trust gave the Engineering College \$1 million to start a computer-aided design facility. Dean Everhart was quoted as saying, "The facility is going to be extremely important for engineering education at Cornell. The progressive industries are already using computer-aided design. . . . We want our students to have the opportunity to gain as much experience as

possible in what is already becoming a standard design tool for engineers."⁶⁷ Shortly after, CADIF was opened in Hollister Hall occupying the entire south corridor. The facility essentially duplicated the software and hardware developed by the Program for Computer Graphics and ran on VAX 780 and 750 computers that connected to 30 terminals and six high-end Evans and Sutherland vector graphics displays. The expectation was that advances made by Computer Graphics as part of their research program could be migrated to this educational environment.

CISER (Cornell Institute for Social and Economic Research)

Capitalizing on an earlier effort to form a research initiative for faculty and researchers interested in social and economic data, and statistical and other analysis tools, a new cooperative faculty-led effort resulted in 1982 in the formation of CISER. The major supporters were the Colleges of Arts and Sciences (A&S) and Agriculture and Life Sciences and the School of Industrial and Labor Relations. Robert McGinnis, professor of sociology in A&S, was CISER's first director. After only one year in operation, CISER had acquired over 200 computerized data sets of information in demographics, economics, and social statistics. In 1986, with the help of a \$150,000 grant from NSF, CISER acquired a half share in an IBM 4381 in cooperation with the Provost's Office and CCS and was able to greatly expand computer use by its members. In 1989 there were 354 CISER members out of the 400 estimated social scientists on campus.

Department of Computer Science

In 1981 the Department of Computer Science received an NSF grant of \$2.5 million for a new computer system and support staff. It was one of five such NSF grants to upgrade research computer activities and, as a result, the department needed a proper computer room. The university trustees approved \$160,000 for creating a computer room on the fifth floor of Upson Hall. This increased space lasted until mid-1985 when there was an urgent need for more space and the trustees approved building a two-story addition on the east wing of Upson Hall to enlarge the fifth floor for the rapidly growing department. In 1986, the department was awarded another NSF grant

⁶⁶ Lyle Wadell, "History of the North East Dairy Records Processing Laboratory," November 1985, part of the Cornell Computing Oral History along with commentary by H. W. Carter and J.-D.-Burke.

⁶⁷ "\$3 Million Pew Gift to Aid Computing," *Cornell Chronicle*, February 5, 1981. (\$1 million for Computer-Aided Design Instructional Facility)

of \$3.6 million to upgrade facilities. In May 1987 the addition was completed and the department acquired 22,000 square feet of new space. For the first time in 22 years, the department had its own designed space and facilities.

The number of courses offered in the department grew from 58 in 1980–81 to 66 in 1989–90 although the number of students was less, decreasing from 3,555 to 2,886 over this same period. Because CS 100 was changed from 3 to 4 credit hours, the number of credit hours over this same period remained essentially the same being 10041 in 1980–81 and 10709 in 1989–90.⁶⁸

Endnotes

^a Doug Van Houweling and I should get partial credit for bringing Ken King to Cornell. Since none of the external candidates, including Fred Harris from the Visiting Committee, was significantly better than either of us and since neither of us was acceptable to Cornell, a stalemate was created. Bob Cooke, to his credit, and with the support of others in the administration and the Board of Trustees, was able to have the position elevated from director to vice provost to be able to pursue and recruit a different class of individual for this leadership position.

^b The director of computing at Tulane was Bob Woodruff, who earlier had been a candidate for the leadership position at Cornell.

^c Wassails had quite a formal structure. They could not begin until the Wassail whistle was blown. For the first few years, Ben Schwarz, systems programmer, was the official whistle blower and he took his responsibility seriously, making sure he was in King's office promptly at 5:00. The drink selection was "programmed" to be either vodka and orange juice or scotch and water with whatever amount of ice the person chose. Different people took on the task of making sure the supply of ice and liquids was adequate. The main purpose for the filing cabinet in King's office was to store the supply of hard stuff. Other drinks were brought by attendees who also occasionally brought their favorite munchies. It was said that the main reason King resisted having his office on the third floor of Day Hall was so that he could host these Wassails!

The most famous of these Wassails that I recall was the one where King brought back the first 128K Mac to campus. He had been to an Apple University Consortium meeting and had been given one of these systems to show off on the campus. Without much instruction or reading of the manual (if there was one,) several of the staff started to use the system and for some reason it hung up and the diskette containing the operating system could not be ejected to reboot the system. We were stuck. King recalled that Steve Jobs, Apple president, gave the attendees his personal telephone number and told them to call him directly if they had any problems. That's what King did. On dialing he reached Jobs, who advised him to insert a straightened paper clip into the hole above the diskette and push it hard to eject the disk. To great jubilation, someone did that, the floppy ejected, and the system was rebooted so that others could play around. One cannot imagine the amazement in the group (most of whom had been brought up on IBM's restricted vision for computing systems) in response to Apple's design of this unique feature. That paper clip ejection of a stuck diskette on the Mac probably changed a few minds about Apple that day.

^d Van Houweling was at Carnegie-Mellon until 1984 when he was appointed vice provost for information technology at the University of Michigan and director of the Merit network that served all the state of Michigan education institutions. He held a number of high-level leadership positions in information technology at the University of Michigan and also in statewide initiatives. In 1997 he was appointed chief executive officer and president, University Corporation for Advanced Internet Development (UCAID). UCAID is the consortium of universities and institutions that cooperated to develop the then-new Internet technology and that will provide orders-of-magnitude increases in speed and capabilities for the new and future Internet.

^e Grimison joined IBM and was a key member of the staff, developing, promoting, and improving IBM's supercomputing offerings.

^f Gale was director of computing at Nebraska for 11 years, during which he also served a year, 1990–91, as director of NSFNet. Following his term at Nebraska, he was director of computing at George Washington University for three years, after which he took the position of director of OARnet, the network division of the Ohio Supercomputer Center (OSC).

⁶⁸ Department of Computer Science, Annual Reports.

^g Blackmun enjoyed a long and successful career as director of Computing and Network Services at the University of North Carolina at Charlotte. In 1997 he was appointed to the position of associate vice president of information technology and CIO (chief information officer) of the North Carolina Community College system.

^h My appointment to this senior management position was sprung on me without notice. Lori Somerville (King's administrative assistant) was working with King on a draft announcement of this reorganization in preparation for his using it in discussions with the involved staff, including me. In an apparent miscommunication between them on the meaning of the word "OK," Lori sent the announcement via e-mail to all of CCS. Telephones started to ring, congratulations and complaints started flowing across the organization. I was personally surprised by congratulations from a close associate, after which I hurriedly went off to read the e-mail message. It took a quite a bit of deft footwork by King to calm down those who were affected, surprised, or disappointed. I was definitely surprised but also pleased by this show of support.

ⁱ One of the unheralded services that OCS, CCS, and CIT provided was the scanning of mark sense forms as an alternative means of data entry into the computer. Mark sense forms recorded information in prescribed positions on preprinted forms using circles or rectangles that were darkened with a #2 pencil. These markings could be sensed (read) by the scanner, recorded appropriately to indicate their position, and sent to the computer for storage or analysis. The largest application was exam scoring for true/false or multiple choice questions. At the end of every semester, a large workload had to be processed in a short period, and this method worked well. Once established, other applications that could take advantage of this technology quickly materialized, so that a routine central service was established. Gene Holleran was in charge of this service for many years.

^j At its peak in 1991–92 BITNET connected some 1,400 organizations in 49 countries for the noncommercial exchange of information in support of research and education. It ended operations at the end of December 1996, when the successor organization, CREN, the Corporation for Research and Educational Networking, believed that the Internet was the longer-term means to achieving needed network services.

^k Engaging in microcomputer sales brought with it a host of problems, including legal and town-gown relationships. Vendors proposed new type contracts that had to be reviewed by the Cornell legal staff. The knotty issue was how to reinforce the restrictions on sales and resales in a responsible manner. Locally, the Chamber of Commerce and individual stores complained about how Cornell was undercutting their markets with the highly discounted prices that vendors were offering Cornell.

^l One of the first deliveries of microcomputer equipment by the new sales organization was a shipment of DEC Rainbows in late December, the last day the university was open before the winter holiday break. To get the systems to the purchasers in time for them to use them over the break, Alan Personius and his wife, Lynne, worked late that Friday, personally delivering the systems out of the CCS van. It was well after working hours, but those individuals who were expecting systems waited until the delivery was made!

^m Alan Personius relates a story about how the state-of-the-art security system designed by CCS staff member Bruce Johnson led to his surprise encounter with Cornell security (police) staff. Johnson had hooked up the system without advising anyone. On a weekend visit to the facility with his family, Personius was met by the campus police cars pulled up to the door and officers with their hands on their gun holsters.

ⁿ The Mann Library publication notes that 16 Apple II computers were placed in Agricultural Engineering in 1981 to support course use in that department but that there was little movement after that to place more such facilities in the College of Agriculture and Life Sciences.

^o Chris Pelkie, an avid Macintosh user and supporter, relates the following story to support his contention of being one of the early Macintosh computer owners in Ithaca. "I had played with the Lisa at DCS a few times and was psyched to have such a 'non-green-screen' experience, but couldn't abide the \$10K price tag. Then the Mac came along, appearing to be a smaller Lisa, and \$3K was in my grasp. I bought the eighth Mac to arrive in Ithaca [in 1984]. I know that because I preordered at Computerland after learning that Cornell wouldn't get their first ones until late summer. Ben Hermann at C-Land told me I placed the last order they accepted, as they were allotted only eight in the first wave after the Superbowl announcement."

^p Documents at the time refer to Pat Paul and later Pat Searle, one and the same Pat Nelson!

^q It was during this hectic time after 1982 when new projects were coming at a much faster rate than old projects could be appropriately ignored or almost completed, when I came up with the term "treadmills of opportunity." It was quite fashionable and at times provided comic relief to use this term when the next such project was being discussed.

^r One of the complications concerning the final design and layout of the interior of CCC was the need to placate the New York State Historical Preservation Society. In 1985, the university believed it had obtained the proper clearances to tear down Stone Hall on the southwest corner of the Ag Quad, at the intersection of Tower Road and Garden Avenue. Within hours after a wrecking ball had knocked down a top corner of the structure, local preservationists had obtained a court order forcing the university to leave Stone Hall in its partially dismantled state until it could go through further reviews and appeals to obtain another clearance. Accordingly, extreme care was taken so that any modifications done to CCC were done with the explicit permission of the preservationist group.

There were a number of sticky points. An easy one was to leave all the large and impressive windows facing the Ag Quad as they were. They could be upgraded with new glass, but no changes could be made to the number and shape, etc. Another was that the building roof line could not be altered, so that air-handling equipment on the roof had to be carefully designed and installed to blend in. The most difficult one was the balcony that surrounded the inner portion of the third floor and opened onto the second floor facing the windows. The proposed plan left the balcony intact, in fact made it larger, but had large glass see-through panels above the balcony rail height wall to provide sound isolation between the floors and permit daylight to enter. When the inspection team from the historical group visited the site during construction, I still remember their "chief" standing on the second floor, by those magnificent windows, and on pointing to one of the interior corners of the third-floor ceiling, saying we could do anything we wanted as long as his view to that corner was not obstructed from that point! Out went the plan for the see-through panels along the periphery of the balcony and in came the long-lasting animosity of staff who felt their working space was being sacrificed to bureaucrats. Finding a good use for that balcony space continued to be an issue until the Legal Office was moved to the third floor of CCC and used the balcony for storing legal books.

^x Schrader left Cornell and became full-time president of NYSERNet in 1986. In 1987 he took the position of director of the Syracuse University Parallel Architecture Center, later the Northeast Parallel Architecture Center, or NPAC. In 1989, he founded PSINet, Inc. (Performance System International), where he served as chairman and CEO. PSI was one of the first Internet service providers (ISPs) and was located in Reston, Virginia.

^t At the time the national supercomputer centers were started, three of the five centers—University of Illinois, San Diego, and the Pittsburgh consortium—had the usual Cray systems, considered to be the real supercomputers of the time. In contrast, the Princeton consortium was to use a new, just-designed system from ETA, a subsidiary of Control Data, while Cornell aligned itself with IBM. The general feeling at the time was that Cornell and Princeton would not last long. That was true in the case of Princeton, for ETA/Control Data never delivered their system and reneged on their commitment and so the Princeton center was disbanded. It is ironic that by the mid-1990s, Cornell was in many ways the leading supercomputer center, not only providing first-class production services but leading technical innovation with highly parallel supercomputers made up of hundreds of small, cheap, and fast RISC systems (RS-6000s) running parallel Fortran and creating a whole new paradigm for such systems. In a way, Cornell verified Wilson's predictions for future supercomputers. Unfortunately, the national Cornell Supercomputer Center was not funded by NSF in 1995 and was phased out in 1997.

^u Wallace (Wally) B. Rogers spent over 40 years at Cornell, starting in the purchasing department and retiring in 1987 as director of business operations. Well-known in higher education for his innovations in university business practices, he was once referred to by the Chronicle of Higher Education as the Guru of Institutional Travel. He was widely recognized for starting a number of "first" practices in purchasing and university operations.

^v Alison Brown's leaving Cornell precipitated somewhat of a crisis for Cornell and the Theory Center. Her husband, Ken Wilson, Nobel laureate and director of the Theory Center, announced his intention to leave Cornell and accompany his wife to Ohio.

^w The SSI system designed by Steve Chen, the top computer designer at Cray, never was built and delivered to Cornell. Since the SSI system itself was being designed as a special super-fast system while the Theory Center building was under design, consideration of this system added over half a million dollars to the building costs. The SSI CPU was expected to be extremely heavy, and so the area where it was to be located in the machine room had to have additional bracing to support the load. Further, since the hardware components were too large to be lifted to the seventh-floor machine room by elevator, provisions were made to have a skylight built into the roof through which the system components could be lifted by crane and placed in the computer room.

^x The remodeling of Comstock Hall to CCC was constrained by the height of the floors determined when Comstock was built early in the 1900s. Fortunately, that floor to ceiling height was over 10 feet on some floors and as much as 12 feet on others. The basement, however, which was committed to computer room space for security purposes, being mostly underground and out of the way, was close to the 10-foot height. After allowing for air ducts in the new suspended ceiling, the cavity under the raised floor came out to be 12 inches high, just meeting IBM specifications. There was lots of discussion about the possible future problems this height would cause, as the chilled water pipes under the floor were about 8 inches in diameter with insulation. When raised an inch from the floor and layered with 1-inch-diameter "bus and tag" cables connecting all the computer equipment, in places there was at best an inch or two of free space for airflow. But there also was no easy way to make the cavity taller, and so the 12-inch design prevailed. The good news in the end was that there were never serious problems from the design. But despite creative equipment layouts, there were definite hot spots when the room was overfilled with CNSF hardware. The even better news was that when the Theory Center computer rooms were designed, the raised floor cavity was over 20 inches tall to avoid congesting the space. Somewhat ironically, soon after the Theory Center room was completed, IBM introduced fiber optic interconnecting cables, one-quarter of an

inch in diameter, to replace the much thicker bus and tag copper wire cables. Another design change eliminated the need for direct water cooling of hot component; the new parallel systems consisted of hundreds of smaller computers for which chilled air was sufficient.

^y As with many senior-level appointments at Cornell, existing staff members reporting directly to the vacant position were typically accorded an interview to give their own assessments to the search committee and responsible executive. The interview of Lynn stands out in this regard. Galloway, Vaught, Cardman, and I were given the last afternoon time slot just before Lynn's departure from Ithaca. We all drove to the former Sheraton Inn and were going to talk over a drink in the bar. That conversation went quite well for a few minutes until the sound system started blaring. The bartender turned the volume down as best he could, but we still had to speak loudly to communicate with each other. After tolerating this situation for a while, we left and drove Lynn to the airport feeling we had accomplished as much as we could. I should also mention that I had quite a time finding the Cornell Chronicle announcement of Lynn's coming to Cornell. While it made the front page, in contrast to King's small inner-page announcement in 1980, it was lost under the headline "Two named to key executive positions," with a picture of Inge Reichenbach, the other appointee! Notwithstanding this low-key announcement, Lynn was to have a significant impact on computing at Cornell.

^z The day after the Morris worm was released and was clogging the national networks with the traffic it generated as it spread, there was an important IBM visit regarding the supercomputer program. I recall some major event in Rockefeller A, and many of the CIT and Theory Center staff were concerned about the success of a demonstration using the system. Just before the event got under way, a worried Bruce Johnson came to the area looking concerned and noting the unexplained network slowdown in the national networks and the possible need to disconnect the supercomputer from the networks. However, to play it safe, Scott Brim, then director of networking at the Theory Center, had already decided to temporarily disconnect the supercomputer from the network for about 30 minutes. Since the demo only used the campus network connections, it went off without a hitch. The next morning we knew all about the Morris incident from extensive reports in the different media.

1990 to 1999

The Decade of the Network

Contents

1990 to 1999—Industry Overview	175
Hardware	175
CPU—Chips, Packaging, and Capabilities	175
Microcomputer Types	175
Personal Data Assistants (PDAs)	175
RISC Systems and Mainframes	176
Magnetic Disks and Tapes	176
Software	177
The World Wide Web and Network Browsers	177
Internet Relay Chat	178
Microcomputer Operating System Software	178
Applications Packages	179
Programming Languages	179
Services	180
Client-Server Computing	180
Networking Technologies	181
The Internet	181
Network Security—Kerberos	182
Telecommunications Act of 1996, Communications Decency Act	183
Vendors	183
Legal Actions	184
Technology Terms, Acronyms, and Buzz Words	184
1990 to 1992 at Cornell	185
Tight Financial Situation on Campus	185
New Executive Staff and Deans	185
CIT Priorities	186
Help Desk Update	186
General Computing Accounts Replace (“Free”) Distribution Accounts with Computer Resource Units (CRUs)	186
CornellC Mainframe Upgraded to IBM 3090-200J	186
Growth in CIT Servers—The Server Farm	186
Training and Publications—CITNews, “InCITes”, and Inside CIT	187
SCRIPT Support Phased Out	187
Microcomputer Sales, Back-to-School Fairs, Selecting a PC Clone	187
New Facilities Management Service; Software Contracts	188
New Information Technologies Committees: CUBIT and ADSPAC	188
Formation of 2001 and Instructional Computing Committees	189
Internal CIT Committee Status—Technologies, Applications, Services	189
Financial Restructuring	189
Need for Additional Building Space	189
CIT Reorganization; Organization Chart	190

<i>Cornell Computing Directors</i>	191
<i>Network Pilot Project in Residence Halls</i>	191
<i>Cornell Worm Case Settled</i>	191
<i>Another Cornell-Initiated Network Incident</i>	191
<i>PSINet Replaces Telenet Network Services</i>	192
A New Network for Cornell.	192
<i>New Backbone Network</i>	192
<i>Campus Building Networks—UTP Ethernet</i>	193
<i>Network Management Center</i>	193
<i>Upgrading External Links</i>	193
<i>Integrating the Electronic Communication Environment</i>	193
New Network Services	194
<i>Project Mandarin and Bear Access</i>	194
<i>Mandarin Consortium</i>	195
<i>EZ-PRINT/Local</i>	196
<i>Terminal Emulators</i>	196
Information Services.	196
<i>Electronic Mail and Directory Services</i>	196
<i>CUINFO—Gopher Client</i>	197
<i>Digital Library, Historical Preservation Project</i>	198
<i>Mann Library Digital Library Projects</i>	199
<i>Cornell Electronic Publishing Working Group (CEPWG)</i>	199
<i>Chaos Corner</i>	199
<i>Instructional Computing</i>	200
<i>Instructional Resource Center and Course Statistics</i>	200
<i>The “Getting Started” Program</i>	201
<i>Learning Technologies Program</i>	201
<i>Classroom Software (Courseware) Awards</i>	201
<i>Intellectual Property Issues</i>	202
Research Computing	202
<i>CIT Joins the “Smart Node” Program at the Theory Center</i>	202
<i>Supercomputing at the Theory Center</i>	202
<i>IBM 3090-600s Upgraded, Array Processors Removed</i>	203
<i>Theory Center Building Opens; Computers Moved from CCC</i>	203
<i>“Garden” of Supercomputers at the Theory Center</i>	203
<i>Mass Storage Systems</i>	203
Business Systems	204
<i>Administrative Data and Systems Policy, Strategic Directions</i>	204
<i>New Systems: CHRISP, Purchasing (APPS), Inventory Management (IMS)</i>	205
<i>Office Support Systems Projects</i>	205
<i>High-Speed Laser Printing</i>	206
<i>CADE Data Entry System Upgraded</i>	206
<i>Exemplar Users Group</i>	206
Campus Planning Initiatives	206
<i>Strategic Planning Initiative and College Priority Planning</i>	206
<i>Quality Improvement Process (QIP)</i>	206
<i>The 2001 Committee Report</i>	207
Migrating to Client Server Computing, EZ-LINK	208
<i>Transition to the Desktop</i>	209
<i>Access to the Network, NetID Established</i>	209
<i>Getting Organized for EZ-LINK: START and SPIRIT Committees</i>	209
<i>Project STOP—Removing CornellC and Software Upgrades</i>	209

1993 to 1996 at Cornell	210
<i>Building EZ-LINK</i>	210
<i>GCA Migration Project</i>	211
Microcomputer Innovations and Issues	213
<i>The Big Red Machine</i>	213
<i>Workstation Ownership Studies</i>	213
<i>Software Piracy</i>	214
<i>Back-to-School Fair</i>	214
Research	214
<i>IBM SP1 Installed at the Theory Center</i>	214
<i>Mass Storage—EPOCH Server</i>	214
<i>New CISER Computers—Gaea and Hera; CornellA Removed</i>	215
<i>Supercomputing—IBM SP2 Installed; KSR1 Removed</i>	215
<i>First URL (Universal Resource Location) at Cornell</i>	215
<i>Theory Center Celebrates 10 Years; TC Building Named Rhodes Hall</i>	215
<i>CISER Takes Responsibility for Statistical Systems</i>	216
Information Services	216
<i>The Making of America Project</i>	216
<i>Online Digital Library</i>	216
<i>CUPID Project and Award</i>	217
<i>Internet Relay Chat</i>	217
CIT Activities, Issues, and Transitions	217
<i>Information Technology Review—Merten Committee</i>	217
<i>FABIT—Faculty Advisory Board on Information Technology</i>	218
<i>The Need for More Building Space</i>	219
<i>CIT Staff Departures, Information Resources Split into Two Divisions</i>	219
<i>VP Lynn Leaves Cornell, Lambert Appointed Acting VP</i>	219
<i>New Deans, New President, and Other Executive Staff</i>	219
<i>Strategic Planning Outcomes; Information Technology Recommendations</i>	220
<i>QIP Fades Away</i>	221
Networking and Network Services	221
<i>ResNet—Networking Connections in Cornell Residence Halls</i>	221
<i>Travelers of the Electronic Highway (TEH)</i>	222
<i>IthacaNet First Conference</i>	222
<i>Bear Access, Version Control, SideCar</i>	222
<i>New EZ-LINK Services: EZ-Publish, EZ-Backup</i>	224
<i>Distance Learning</i>	224
<i>GateD</i>	225
<i>Misuse of Networks: Network Use Policy</i>	225
Business Systems	225
<i>CourseEnroll</i>	226
<i>Employee Essentials</i>	226
<i>Project STOP Stops</i>	226
<i>Information Access Project (IAP)</i>	226
<i>Data Warehousing Activity</i>	227
<i>CIT Mainframe “Refreshment” Project—New CornellC</i>	227
Information Services	228
<i>CU-SeeMe</i>	228
<i>CUINFO—Gopher Phased Out</i>	230
<i>CIT Training and Publications—CIT Website</i>	230
<i>Chaos Corner Comes to an End</i>	230

Digital Library	230
Tulip Project—EPOCH Mass Storage Phased Out.	230
Instructional Computing	231
FABIT Report and Recommendations	231
Student Laboratory Upgrades.	231
FABIT-Supported Teaching Space Upgrades	231
CIT Activities.	231
CIT Services and Staff Move to Maple Avenue	232
Lambert Reorganizes CIT	232
Cornell Policy and Law Institute Formed	232
CIT Training Becomes Technology Training Services.	233
Year 2000 (Y2K) Millennium Issues	233
CIT Receives Recognition and Awards	234
CAUSE Award for Excellence in Campus Networking	234
Bill Gates Recognizes Cornell Achievements in Networking	234
Cornell Graduates Start The Globe.com CourseInfo/Blackboard	235
Lambert Recognized as Innovative Leader	235
Networking and Network Services.	235
Network Services	236
Celebrating 10 years of Cornell Telecommunications/Telephone Installation	236
“Cells in Frames,” Asynchronous Transfer Mode (ATM), and Workstation Telephony	236
Viruses	237
The Computer Decency Act	237
Electronic Mail.	237
IthacaNet	237
Business Systems	237
Steps Leading to Project 2000	238
Project 2000—Creating a Best-Managed University Announced	239
Project 2000—Best of Breed vs One-Vendor System Selection and Plan C	240
PeopleSoft Selected as Vendor for P2K systems	241
Project 2000 Activities and Teams	241
CornellC MVS and VM Software Upgrades	241
1997 to 1998 at Cornell University and CIT.	242
Theory Center Loses NSF Support, Reinvents Itself	242
Another CornellC Mainframe Upgrade	242
CIT Electronic Publishing Phased Out; Use of CUPID Ends	242
CIT Transfers Microcomputer Sales and Support to Campus Store	243
CIT Data Entry Operations End	243
CIT Servers and Network Management Center Moved to Rhodes Hall.	243
VP Lambert Leaves Cornell.	244
Tom Dyckman Appointed Acting VP.	244
CIT Staff Transitions; Cogger, Regenstein, and Siegel Leave Cornell	245
CSARS Use Ends	245
Business Systems	245
Year 2000 Committee Appointed	245
Project 2000 Organizing Steps in 1997	246
Selecting a New Database Management System—Informix	246
COLTS—Cornell Online Time System	247
Development of Production Architecture.	247
Survey of Administrative Processes and Campus Readiness for Change	248
Project 2000 in 1998	249

P2K—New PeopleSoft Human Resource/Payroll System Installed	249
CUDA Fades Away	250
Networking and Network Services.	250
Cornell Network Upgrade	251
Net-Print—New Network Printing Service	251
Bear Access Upgraded	251
EZ-Remote and ResNet	251
Cornell Award Grant to Join NSF High-Speed Network—vBNS	251
Cornell Office of Distance Learning Formed	252
IthacaNet—Van Houweling Keynote Speaker	252
CU-SeeMe Fades Away	253
Instruction.	253
More FABIT-Funded Teaching Classroom Upgrades.	253
Computer-Based Training Introduced.	253
1999 Highlights	253
Polley McClure Appointed VP for Information Technologies	253
Formation of the Faculty of Computing and Information Science	253
CIT Staff Transitions and Turnover	254
Fred Rogers Leaves Cornell	254
CornellC Mainframe Upgraded to Model 9672-R25	255
Local Cable TV Network Services	255
IthacaNet—Schrader Keynote Speaker.	255
Network Security	255
Software Piracy	255
More Network Upgrades—On and Off Campus	256
EZ-Remote	256
E-mail Activity	256
Central Calendaring Service	256
New CIT Publications/Computing at Cornell Web Site	256
Computer Science Course Statistics	256
FABIT-Supported Teaching Classroom and Laboratory Upgrades	256
Bear Access with Runway	257
EZ-Backup Growth	258
Net-Print Activity	258
Information Technology and the Libraries; HeinOnline	258
Theory Center Switches to Dell Equipment	258
New DataMart, Old Data Warehouses, Mark IV Fades Away	258
The Millennium Arrives	259

Figures

Figure 1. An early Bear Access screen	195
Figure 2. VP Lynn meeting with the Cornell Electronic Publishing Working Group.	199
Figure 3. Bear Access for different workstation operating systems	211
Figure 4. Bear Access launch pad showing services available	223
Figure 5. SideCar	223
Figure 6. Growth in CornellC Adabas commands, 1984 to 1995.	227
Figure 7. CU-SeeMe demonstration by Tim Dorcey	229
Figure 8. Project 2000 logo.	239
Figure 9. Cornell data networks after 1999 upgrades.	257
Figure 10. The Network Operations Center at the Y2K rollover	259
Figure 11. VP Polley McClure and CIT staff at Y2K rollover.	259
Figure 12. Action in the NOC during Y2K rollover	259
Figure 13. CIT logo, 2003.	260

Tables

Table 1. Cornell Computer Services Organization Chart, 1992 to 1993	190
Table 2. Hosts and Services for Terminal Emulators Supported by CIT in 1990.	197
Table 3. Number of Courses and Enrollment Supported by CIT, 1990 to 1992.	200
Table 4. Number of GCA Accounts by User Category in 1991	211
Table 5. Number of NetIDs by User Category in 1994	212
Table 6. Use of Network Services by Category of Users in 1994	212
Table 7. Trends in Workstation Ownership, 1989 to 1993	213
Table 8. Applications Used by Undergraduate Students, 1992 to 1993	213
Table 9. Cornell Computer Services Organization Chart, 1996 to 1997	233
Table 10. Project 2000, Application Access Types	248
Table 11. Technology Resources in Administrative Areas on Campus, 1997	249
Table 12. Department of Computer Science Course Statistics during the 1990s.	256

1990 to 1999—Industry Overview

This industry overview describes the major hardware, software, and service innovations in the 1990–99 decade. In hardware the theme was smaller, faster, with less energy consumption. The outstanding developments during this decade were the Internet browsers and the spread of the Internet all around the world and almost to every installed personal computer. The installation of new network technologies improved network services as well as the speeds of transmission, enabling applications such as real-time video and the downloading of large files such as movies to become possible.

Hardware

The number of computers worldwide continued to increase rapidly as uses multiplied and prices declined; the items were sold as commodities in many outlets, with entry-level models priced at less than \$500. By the end of 1999, it was estimated that there were over 400 million personal computers worldwide, with over 140 million in the United States alone. That same year, it was estimated that personal computers were being shipped at the rate of 43 to 49 million per year in the United States alone—five times higher than the 9 million in 1990.¹

CPU—Chips, Packaging, and Capabilities

Moore's Law, which predicted computing power of a silicon chip to double every two years, continued to be in force. Intel, Motorola, Cyrix, and Advanced Micro Devices (AMD) were the principal vendors of these CPU chips, and they typically leap-frogged each other's chip offerings to gain competitive advantage. We will use the offerings from Intel to demonstrate the advances made over the decade in the design of these chips.

	Intel 386, 1989	Intel Pentium III, 1999
Rating	20 Mhz	500–750 MHz
Number of Transistors	855,000	28,000,000
Sales Price	\$150 in 1,000 units	\$776 for "large" orders

Memory chips also followed Moore's law so that main memory on personal computers in 1999 was 64–128 MB and expandable, compared with 1–2 MB at the beginning of the decade. CMOS chips that required less power were increasingly used to reduce the power draw of chips in all types of systems, especially in laptops or portable devices, and to extend the operating time when using battery power.

Microcomputer Types

Microcomputers themselves continued to evolve in three planes—the high-end servers, the desktop workstations, and the portable (or laptop) computers.

Servers took advantage of the increasing speed and features of CPU chips from Intel and other vendors to challenge mainframes for speed and functionality. As will be discussed later, the new RISC-based processors started to play a larger role, particularly in the scientific area where Unix software was more popular.

Vendors tried different approaches to attract buyers to their desktop systems. Apple made a bold move at the end of the decade, introducing their iMac computers, first in "Bondi Blue" then later in five bright colors; a translucent cover encapsulated the CPU and associated components, with a larger case over the monitor. This broke the uniform use of beige for all such hardware, and vendors started to use different colors for encasing components. One of the innovative ideas that attempted to piggyback on the greater availability of network connections in the home (direct or dial-up) was the attempt to develop and produce the "network computer." This would be a special microcomputer but "dumbed down" so that it would be simple to use for the mainstay network services such as e-mail and surfing the Internet. Further, it would be connected to the home TV, eliminating the need for a separate, expensive monitor, so it could be offered at a price well below \$500. The major vendors agreed to a common standard, although each promoted its own hardware implementation. The product failed to create any excitement or market because the targeted buyers were grandparents, who proved to be quite facile with point and click personal computers and who also appreciated other applications such as word processors and spreadsheets on these systems.

To increase the usefulness of laptop computers so they could be used in and out of a person's office or home, vendors developed "docks" for their laptops. Initially, laptop configurations were limited, omitting features such as floppy drives or CDs to conserve battery life. A laptop inserted into a desktop dock could run on electric power and serve as a network-connected desktop computer with a large screen, a full-size keyboard, and access to all the network facilities and equipment. Toward the end of the decade, laptops had essentially the full functionality of desktops.

Personal Data Assistants (PDAs)

Continuing the trend to smaller, more portable computers, a new development in the 1990s was the

¹ Information principally from Ken Polsson, 2003 Chronology of Personal Computers, www.islandnet.com/~kpolsson/comphist/

hand-held computing/information technology device, dubbed the personal data assistant, or PDA. The first notable and heavily promoted device was the Apple Newton, introduced in 1993. Early models were somewhat bulky (6 inches long, 4 inches wide, and 1 inch thick), expensive, and full of bugs. The handwriting recognition software, a key to exploiting the commands and features of the device, did not work well. Although later models fixed most of the problems, the Newton did not take off as expected.

The Palm Pilot PDA, first marketed in 1994, was widely popular. The Pilot was about the size of a deck of cards, making it not only handheld but convenient to carry in pockets or purses, with buttons and a stylus used to initiate commands. The initial and popular applications were address book files (for names, telephone numbers, etc.) and electronic schedules. Soon the device could be connected to a workstation and server for uploading or downloading information. After the end of the decade they were made wireless/network capable and could receive and send e-mail and carry out other Internet functions

New vendors such as Handspring (started by the original Palm Computing founders) entered the market, some using the Palm operating system but others using the Microsoft CE system or some other choice. Before the decade was out, cellular (cell or mobile) telephones started competing with the handheld devices, offering e-mail and Internet browser services and the future potential of including a digital camera for taking and transmitting images.

RISC Systems and Mainframes

IBM and others such as HP started to exploit the RISC (reduced instruction set computer) architecture to build servers and computers for scientific computation and then for mainframes. IBM is an example of how these systems developed and were deployed. Introduced in 1990, the IBM product built on RISC technology, the RS6000 (RS for RISC System), became very popular for scientific computations. They relied on a number of IBM innovations, such as a new microchannel architecture and advanced CMOS chip technology. Depending on the model, these computers were rated from 27 to 41 MIPS, or 7 to 13 MFLOPS.²

Once the RS6000 machines proved themselves, IBM took advantage of their capabilities by grouping them together to form their first parallel computing system—the SP1 (scalable power parallel system 1). Unveiled in 1993, the SP1 was touted as the first microprocessor-based supercomputer. The first one could have up to eight processors, or eight nodes. The next year this was surpassed by the next generation system, the SP2, built with faster RS6000s and offering up to 64 processors. First deliveries included Cornell.

This was the beginning of a new breed of supercomputers called massively parallel systems.

Cray Research was among others to adopt this same technology, creating massively parallel supercomputers using clusters of very fast RISC machines linked together. Combined with new operating systems and applications software, most notably Fortran, which effected parallel operations, these systems were put to work on the “grand challenge” problems³ for large-scale computing. One great advantage was that all such systems could be cooled with conditioned air rather than by piping chilled water directly to hot components; as a result energy savings were considerable. The use of CMOS chip technology, which was noted for low power consumption, made this possible.

IBM continued to dominate the mainframe market; in tune with the focus on client-server computing, the machines were referred to as “enterprise” servers—acting as a large-scale file server that was used by the entire organization, the enterprise. With the success of their SP1 systems, IBM started to assemble mainframes using multiple RISC processors instead of specially designed subsystems

Magnetic Disks and Tapes

Improvements in data storage technology continued on several fronts, most notably with increasing recording densities and, in the case of magnetic tape, with the use of different materials for the recording media. Different packing of the disks improved the capacity and data availability, and transfer speed increased in line with improvements in channel speeds.

In the early 1990s, it was common for microcomputer systems to have internal 20 MByte hard drives. By 1999, a 10 GByte or 20 Gbyte hard drive was quite common, increasing storage by a factor of 1,000 at reduced cost.

Different means of external storage (using removable media) also became available for personal computers during the 1990s. CD-ROMs became more popular in this decade, especially for the distribution of software and data, before the Internet became useful for these purposes by the end of the decade. However, even with the Internet, in 1999 roughly 60 percent of retail software was still distributed on CDs. While the capacity of CDs remained the same, the transfer rate improved by a factor of 24 during the decade.

Continuing this development of auxiliary storage, in 1995 Iomega introduced their ZIP drive and ZIP

² IBM news release, “High-Performance RISC Workstation Family Announced by IBM,” February 1990.

³ Executive Office of the President (USA), Office of Science and Technology, “A Research and Development Strategy for High-Performance Computing,” November 20, 1987.

disks, removable 3.5-inch diameter floppy disks that held as much as 100 MB of storage and could connect to different types of data transfer ports. By the end of the decade, ZIP drives and disks capable of storing 250 Mbytes were being built into many desktop as well laptop computers. Later, Iomega also introduced the JAZ drive, removable disks capable of holding as much as 1 gigabyte with transfer rates as high as 10 Mbytes/second.⁴

For fixed-head disk storage, new RAID disk technology was developed to increase storage capacity and improve data availability and data transfer speeds. RAID (originally, Redundant Array of Inexpensive Disks, but later changed to Redundant Array of Independent Disks) operated by using arrays of disks and various data-striping methodologies. Data-striping meant that incoming data was distributed across a number of disk drives (an array) and so it could be retrieved much faster than from a single drive. In addition, some designs were able to provide data redundancy by recording additional information so that the data could be reconstructed if a single drive failed. Different levels of RAID technology offered tradeoffs between transfer speeds and redundancy. This technology came to dominate the choice of disk storage for large systems during the 1990s, but with increasing success, it was made available on larger desktop and server systems using different data transfer interconnect technologies. Capacity, depending on the needs, ranged from the 10s or 100s of gigabytes, and by stringing units together, to terabytes.

Magnetic tape technology for large-scale systems and applications continued to evolve during the decade, increasing the amount of information that could be stored on a physical tape in a smaller package. Again, using IBM as the example, starting in the mid 1980s, the 3480 tape cartridge was able to store 200 Mbytes in a 5.5-inch square cartridge and transfer the data at 3 Mbytes/second, compared with the 3420-11 tape reel, which stored 180 Mbytes on a 10.5 reel of tape with a transfer rate of 1.25 Mbytes/second.⁵ Successive upgrades to IBM 3490 models increased the storage to 800 Mbytes per cartridge, then to 2.4 Gbytes per cartridge. In the mid-1990s IBM introduced the Tape Library Data Server that automated tape handling by using robotics technology to find cartridges and make them available for use without operator intervention.

Software

A defining statement for software during the 1990s can be summed up in a humorous “law” made by Lincoln Spector, who as a counterpoint to Moore’s Law for hardware, composed Spector’s Law:⁶ “the time it takes your favorite application to complete a given task doubles with each new revision.” That was a way

of saying that regardless of the software involved, as it grew in size and complexity it took away some of the gains in hardware performance.

The World Wide Web and Network Browsers

The killer developments of the decade were the World Wide Web and the associated Internet browsers. Tim Berners-Lee created the World Wide Web (WWW) in 1991, which had the ability to link documents and create “web” pages. He defined the addressing scheme, the Universal Resource Locator (URL), which gave each page its unique address. He built the WWW based on earlier work he had done in the 1980s defining a hypertext system for linking documents to each other, for example, by words in the title of a document. This system later came to be known as HTML (HyperText Markup Language), the language by which content developers created and linked web pages in a site. In its initial development, text commands were used to get around the information sources because Berners-Lee’s initial work predated point-and-click graphical interfaces.⁷

The first well-known graphical network browser was Mosaic, developed at the National Center for Supercomputing Applications (NCSA) at the University of Illinois by Marc Andreessen, then a student employee at the NCSA. He and a group of other students developed Mosaic in 1993 as a visual way to navigate the Internet. When it was introduced, the introduction read: “NCSA Mosaic provides a consistent and easy-to-use hypermedia-based interface into a wide variety of information sources,” which included the concept of storing and retrieving images. To capitalize on the interest in Mosaic,⁸ Andreessen formed a partnership with Jim Clark, a Silicon Valley entrepreneur, and they created Netscape Communications to produce a commercial version named Netscape Navigator. In a departure from standard software distribution, in 1994 Netscape Navigator could be downloaded free on the Internet without going through commercial channels.

At about this same time, Microsoft made the strategic decision to orient its resources to developing products for the Internet. In 1994 Microsoft introduced Internet Explorer (IE) as their browser. In a controver-

⁴ www.iomega.com

⁵ R. Bradshaw, C. Schroeder, “Fifty years of IBM innovation with information storage on magnetic tape,” *IBM Journal of Research and Development*, vol. 47, no. 4, July 2003.

⁶ “Plugged into a New Millenium,” *Inforworld*, October 26, 1998.

⁷ “Time 100 Scientists and Thinkers—Tim Berners-Lee,” www.time.com/time/time100/scientist/profile/bernerslee.html

⁸ “10 years ago, who knew what his code could do?” *USA Today*, March 10, 2003.

sial move, IE was bundled with the Windows operating system, a source of irritation for users and application software vendors and a source of litigation with computer and operating system software vendors.

The dramatic impact of Microsoft's entry into the browser market is demonstrated by the following market share information:⁹

	Netscape Navigator	Microsoft Internet Explorer	Other
1994	54.6 percent	29.5 percent	15.9
1999	23.0 percent	75.0 percent	2.0

This situation led to lawsuits, discussed in more detail in the section "Vendors" that follows.

Internet Relay Chat

Internet Relay Chat, or IRC, became the popular interactive use on the Internet during the 1990s. It was described as the network equivalent of CB radio but with an extended range enabling people all over the world to participate in real-time conversations. IRC was developed by Jarkko "Wiz" Oikarinen in 1988 at the University of Oulu in Finland.¹⁰ It is built on the client-server architecture where the client on a user's system connects to an IRC server, for which the client has been enabled. The servers are all interconnected and pass information over the network. Once connected to a server, and after establishing a "nickname" or moniker, a person can join different chat channels to participate in discussions of one's choosing or liking. Typically, the person setting up a channel topic then moderates that channel as the channel operator, although private conversations are allowed. IRC gained international fame during the Gulf War of 1991 and the coup against Boris Yeltsin in 1993, when IRC users from those areas were able to send reports about the events to others around the world.

Microcomputer Operating System Software

Microcomputer operating system software from Microsoft continued to take a greater share of the market. In 1990 Microsoft introduced Windows 3.0, and 2.75 million copies were shipped before the year was out. In 1991 Microsoft announced Windows NT (New Technology) for high-end desktop and server-type platforms. In 1995, there followed Windows 95, an upgrade/replacement of the Windows 3 systems. According to Polsson, when Windows 95 was released in mid-1995, 300,000 copies were sold the first day, and within a month over 7,000,000 copies were sold.

A sense of how Microsoft continued to dominate the operating system offerings during this decade is demonstrated by the following market share information for desktop operating systems (from Polsson):

- In 1991 Microsoft DOS is estimated to have 75 percent market share.
- In 1998 Microsoft is estimated to have an 87 percent market share.

(Windows 95 at 63 percent, Windows NT at 20 percent, DOS at 4 percent, Macintosh at 6 percent, and others at 7 percent)

Since IBM itself had long since lost its dominance in the microcomputer industry, the term "Wintel" started to be used to denote IBM or IBM-compatible systems running the Windows operating system and built using Intel chips.

IBM tried other operating system alternatives during the decade. In 1990, it entered into a cooperative agreement with Microsoft to develop OS/2 for its new line of PS/2 workstations, but this agreement ended within a year and IBM took over responsibility for OS/2. Despite heavy promotion, OS/2 never became a solid competitor to the Windows systems. In the above market share comparison, OS/2 had so faded that it did not even deserve separate mention and was included with "others."

In an attempt to reduce Microsoft's dominance of the operating system market, nine vendors, including IBM, DEC, and Hewlett-Packard, formed the Open Software Foundation.¹¹ The plan was to promote "open computing" by developing a common operating system and interfaces based on developments in Unix and the X Window System. OSF/1 was announced as the industry's first open operating system in 1990. OSF subsequently defined the Distributed Computing Environment (DCE), a vendor-neutral set of distributed technologies that would be accepted as an industry standard. In time the organization became known as OSF DCE, whose pledge became the promotion of its technology in three key areas of computing: security, the World Wide Web, and distributed objects.

The only new serious competitor to Microsoft Windows was Linux, developed by Linus Torvalds from Finland in 1991.¹² Linux was an operating system variant of Unix that Torvalds distributed freely around the world as a freeware offering for IBM-compatible workstations. Linux captured the imagination and hearts of many "techie" looking for an alternative to Windows. By the end of 1998, shipments of Linux operating system software reached almost 3 million as vendors such as Red Hat made it commercially avail-

⁹ Ken Polsson, 2003 Chronology of Personal Computers, www.islandnet.com/~kpolsson/comphist/

¹⁰ <http://daniel.haxx.se/irchistory.html>, and www.mirc.com/irc.html

¹¹ "Open Software Foundation," <http://wombat.doc.ic.ac.uk/foldoc>

¹² <http://ragib.hypermart.net/linux>

able and provided technical support. Even IBM started to promote the use of Linux for its workstations and in 1999, for example, announced the distribution and technical support of Linux on its systems. Toward the end of the decade, many vendors released their applications packages, for example, WordPerfect, to work under Linux.

In 1999 Apple, continuing on its own path, introduced Mac OS X (OS ten) as the new server-only operating system. OS X was based on a version of Unix first used with NeXT computers, the company that Apple bought after Steve Jobs returned to Apple. A year later, in 2000, Apple introduced Mac OS X for personal workstations and kept the same spirit as the original Mac OS.

Applications Packages

Early in the 1990s, Lotus 1-2-3 was the most popular spreadsheet program for IBM and compatible computers. Microsoft introduced Excel as the competitor spreadsheet for Windows systems. Excel grew out of Multiplan, first developed in the 1980s for the Macintosh as a GUI-powered spreadsheet. At this time, according to Polsson, Lotus had about 50 percent of the market while Excel and Quattro from Borland each had about 14 percent market share, with a host of others making up the rest of the market. This pattern changed considerably by the end of the decade.

The most popular word processing package at the beginning of the decade was WordPerfect, which held about 70 percent of the market according to Polsson, with Microsoft Word and other products making up the rest of the market.

In the interests of providing competitive advantage, the major software vendors started packaging their applications software into “suites” that bundled word processing software, spreadsheet software, and other commonly used applications such as database systems and presentation preparation software into “office” bundles. Not unexpectedly, Microsoft came to dominate this market with its offerings, and by late in the decade, in 1997, the three major offerings and market share were: Microsoft Office—79 percent; Lotus Smartsuite—13 percent; WordPerfect Office—5 percent.

Programming Languages

Application programming languages changed considerably during the 1990s, although the old standby, BASIC, first developed by Kemeny and Kurtz at Dartmouth in the 1960s, continued to be available. In 1990, according to Polsson, Microsoft shipped the 50 millionth copy of BASIC.

One language that played a prominent role in supporting the growth of the Internet was Java, developed by a group at Sun in 1991.¹³ Java was an entirely new, processor-independent language that was designed to move media content across networks of heterogeneous devices. Processor independence was achieved by using a very old technology—interpreters from the 1950s—for which a common language could be run on different processors by having an interpreter on each processor translate the source code to its unique object code. By the use of Java “applets” (a small, re-usable part of a larger Java written application), Java offered the capability to move “behavior” of the application along with the content. HTML alone could not do that. This whole new area became known as object-oriented programming. The portability of Java applets and the later incorporation of Java into Netscape Navigator made it possible to create WWW content at one site and be able to execute these on different computers at other sites. Sun developed a Java standard that most vendors followed, except Microsoft, which resulted in yet another case of failure to agree on a common standard across platforms.

Two other programming languages for developing network applications came to be popular in this decade for different reasons: Pascal and “C.” Niklaus Wirth defined the Pascal language in 1972 for teaching computer programming. Wirth’s principal objectives for the language were that it (a) be efficient to implement and run and (b) allow for the development of well-structured and well-organized programs. With the increasing popularity of the Wintel systems and the availability of the popular Turbo Pascal from Borland as the decade started, Pascal was the dominant programming language for instruction at educational institutions.

Later, the C programming language and its extension, C++ (with capability for object-oriented programming), gained increasing use and popularity and pretty much displaced Pascal for teaching purposes. C was developed in the early 1970s at Bell Labs as the systems implementation language for Unix, then itself in its early stages of development. C gained in popularity, riding along with the increasing popularity of Unix as the operating system of choice for most of the high-end server platforms (and clustered super-computing platforms). Most vendors of these platforms had pretty much standardized on Unix, or the company variant, as the operating system of choice, and C came along with these different platforms. This was a harmonious association. Both ANSI (American National Standards Institute) and ISO (International

¹³ <http://java.sun.com/features>

Organization for Standardization) agreed on a standard C programming language in the 1980s, and that definition continues to hold.¹⁴

Fortran continued to be the most popular language for programming scientific applications, and it evolved with the changing technology. In 1990, Fortran 90 was defined as the new standard language with new and useful features. In 1992, High-Performance Fortran, HPF, was developed to take advantage of massively parallel computer architectures. While HPF was not an official standard, it became very popular on the new supercomputer and mainframe architectures. A subsequent version, Fortran 95, was defined, which was a revision of the Fortran 90 standard. With the involvement of ISO and ANSI, the useful life of Fortran is likely to continue as the language evolves to keep up with new developments in hardware and software technology.¹⁵

Oracle Corporation extended its dominant position in providing relational database management systems and tools by developing a wide range of tools for large-scale business applications. It was one of the leaders in developing enterprise resource planning (ERP) solutions and enabling Internet-capable solutions over its entire product line—database, business applications, and decision support tools. It positioned itself as the world's leading supplier of software for information management and the world's second-largest independent software company.¹⁶

It was a testament to software's increasing capabilities when, in 1996, IBM's "Deep Blue" computer held its own in a chess match with Garry Kasparov, world chess champion. In the first-ever traditional chess match between man and computer, Deep Blue won the first game, although Kasparov ultimately won the match four games to two. In the match that was played in March 1997, the new Deep Blue was a special combination of hardware and software using an IBM SP2 system "capable of examining 200 million moves per second—or 50 billion positions—in the three minutes allocated for a single move in the game." In that match, Deep Blue won 3.5 to 2.5 games. The series continued during the decade, with other Deep Blues and with systems from other vendors, and increasingly the computers became the champions.¹⁷

Services

As the 1990s began, delivering computing services continued on two different paths. On one path, online services between the user at a terminal and the host mainframe followed the time-sharing model that had been developed in the 1970s. This was often referred to as the host-dumb terminal model, as all the computing was done on the mainframe and the terminal

was used only for data entry and the display of results. In this decade, however, with advances in networking technology and the presence of a personal computer workstation on the user's desk, when time-sharing services were needed, the dumb terminal was emulated on the desktop computer.

The second path that was followed was at the level of the local area network, the LAN, where a server was used for common and large-scale functions such as file sharing. In this model, the server stored databases and on request delivered all or a portion to the user's desktop computer where the computing was done. This model did suffer from some deficiencies in that the user's workstation had to be sufficiently powerful to execute the work in a reasonable time, and the network had to be of sufficient capacity to transfer files from the server to the workstation in a reasonable time.

Client-Server Computing

The new model that emerged was called the client-server model, which consisted of developing applications with a shared responsibility; the processing demands were divided between the server and the client according to which was best suited for which task. One tried to balance the computing power of the server and the user workstation and the bandwidth of the network. In this client-server architecture, only the client component was resident on the user's workstation, reducing the demand on that system, and the server component was at the server end. With designed intelligence built into the applications, this cooperative processing could take advantage of the power of the computers and the connecting network. This mode was also called "distributed computing." In the early days, clients performed much of the processing work and were referred to as "fat" (robust) clients. In time the balance shifted somewhat, and clients became known as "thin" clients, with more of the associated processing being done by the server.

A good example of the evolution to the client-server environment is electronic mail. In the host-dumb terminal model, the e-mail system was resident on the host, and the terminal simply sent and read mail; all the files were resident on the host and all processing was done by the host. In the client-server model, the

¹³ <http://java.sun.com/features>

¹⁴ Dennis M. Ritchie, "The Development of the C Language," <http://cm.bell-labs.com/cm/cs/who/dmr/chist.html>.

¹⁵ "A brief history of Fortran," www-h.eng.cam.ac.uk

¹⁶ www.oracle.com

¹⁷ IBM press release, "May 1, 1997: Ancient game, modern masters," *IBM Research, Deep Blue Overview*.

server acted as the post office, sending, receiving, and forwarding mail, and the client did the composing and retention of mail.

Networking Technologies

The development of new network technologies continued apace with other technology developments. Ethernet, a technology that was almost 20 years old, was improved so that it could be used with plain copper wire. Up until the 1990s, Ethernet could only be reliably sent over limited distances using coaxial cable. In the 1990s developments permitted the use of copper wires, referred to as UTP Ethernet—Unshielded Twisted Pair Ethernet—that provided 10 megabits/second transmission. As the need for more network capacity grew to keep up with the amount of information being transmitted and with new applications such as video transmission, new network technologies were developed.

FDDI—Fiber Distributed Data Interface—operated at 100 megabits/second using a token-passing technology over a pair of fiber optic rings. It became one way of providing higher-capacity networks or serving as a backbone to a lower-speed Ethernet.

ATM—Asynchronous Transfer Mode—was another technology that operated at 100 megabits/second. It had the capability of being used for both local and wide area networking and was capable of carrying voice and video traffic as well.

Ethernet technology itself continued to evolve. By the end of the decade, devices were available that operated at 100 megabits/second using optical fiber, and discussions were under way to provide gigabit Ethernet services. At the same time, Ethernet technology was available as switched Ethernet, which could provide 10 to 100 megabits to the desktop instead of sharing this capacity with all the other devices on the local area network. Essentially, network technology was keeping up with new developments in computer technology and new applications.

Modems increased in speed during the decade, typically starting at 9,600 baud but going up to 56 kilobits/second at the high end. Wireless networking technology was in the development stages during the decade and wasn't commercially available until after the year 2000.

The Internet

In 1990 the Internet was not the Internet we think of today in the year 2004. It consisted of the NSFNet backbone and an increasing number of subnets, such as NYSERNet, with connections to similar research or educational networks around the world. The predominant protocol in networking was TCP/IP, which

formed the basis of all the networks. The predominant clientele up until 1994 were educational and research institutions, with a small but increasing number of commercial clients.

As the international, national, and local networks were expanding and changing, the evolution during the decade resulted in a three-tiered collection of network service providers. Tier 1 providers were those companies, for example, PSINet and MCI Worldcom, that operated the very high speed fiber optic backbones that tied the Internet together. Tier 2 providers were the large national or regional network carriers, such as NYSERNet, that connected to the Tier 1 backbone. Tier 3 providers, Lightlink in Ithaca, for example, connected to a Tier 1 or 2 provider and were local service providers with a limited reach, typically referred to as ISPs, or Internet Service Providers. Typically, a company or individual dealt with the service provider that offered the service desired, including line speed, number of hops or different network connections, availability, and support at an affordable cost.

Although the term “Internet” is used with an assumed understanding of what is meant, the following definition is worth repeating. In 1995 the Federal Networking Council unanimously passed a resolution defining the term Internet; the definition had been developed in consultation with members of the Internet and intellectual property rights communities.¹⁸

The Federal Networking Council (FNC) agrees that the following language reflects our definition of the term “Internet.”

“Internet” refers to the global information system that (i) is logically linked together by a globally unique address space based on the Internet Protocol (IP) or its subsequent extensions/follow-ons; (ii) is able to support communications using the Transmission Control/Internet Protocol (TCP/IP) suite or its subsequent extensions/follow-ons and/or other IP-compatible protocols; and (iii) provides, uses, or makes accessible, either publicly or privately, high-level services layered on the communications and related infrastructure defined herein.

The following is a brief outline of the key developments of the Internet as well as some highlights over the period from 1980 to 2000:¹⁹

¹⁸ Internet Society (ISOC), “All about the Internet: History of the Internet,” www.isoc.org/internet/history/brief.shtml

¹⁹ “Hobbes Internet Timeline v7.0,” www.zakon.org.robort/internet/timeline/

- ARPANET was operating around the globe with many national and international connections, using the TCP/IP protocols.
- In 1984 the Domain Name System (DNS) was introduced. DNS provided every computer on the Internet with a unique address, its IP address, and a corresponding mnemonic name for universal resolvability. In 1985 the Information Sciences Institute at USC is given responsibility for DNS management.
- In 1986 NSFNET created the national backbone connecting the five supercomputing centers and other regional networks. The backbone speed was 56 Kbps. A year later, NSF signed a cooperative agreement with Merit Network, Inc., to manage the NSFNET backbone.
- In 1988 CERT, the Computer Emergency Response Team, was formed by DARPA following the Morris worm incident to prevent and respond to such occurrences in the future. CERT was located at Carnegie-Mellon University.
- In 1990 ARPANET ceased to exist.
- In 1993 InterNIC was created by NSF to provide specific Internet services such as directory and database management services, registration services, and information services.
- In 1995 NSFNET reverted to a research network, vBNS, the very high speed backbone linking the national supercomputing centers. Later this became known as Internet2. The main U.S. backbone traffic was routed through interconnected network providers. That same year, registration of domain names was no longer free, and a \$50 annual fee was introduced.
- In 1996 Internet telephony caught the attention of U.S. telecommunications companies, who asked Congress to ban the technology. In this same year, MCI upgraded the Internet backbone from 155 Mbps to 622 Mbps.
- In 1997 ARIN, the American Registry for Internet Numbers, was established to handle the administration and registration of IP numbers previously done by InterNIC.
- In 1998 the U.S. Department of Commerce entered into an agreement with ICANN, the Internet Corporation for Assigned Numbers, to establish a process for transitioning DNS from the U.S. government to industry.

By the end of the decade, the new Internet was growing extremely rapidly. According to information from CISCO, the number of Internet hosts grew from 80,000 in 1989 to 29,670,000 in 1998, while the number of web sites grew from 100,000 in 1996

to 1,834,710 in 1998. This phenomenal growth of national and worldwide networking gave rise to the term “dot-com bubble” when overexpansion in network facilities and companies outpaced the market at the end of the decade.

Network Security—Kerberos

Security became increasingly important with the spread of the Internet. There was always a concern about security because access to services and networks relied largely on passwords. With the development of network “sniffers” that watched network transmissions, it became relatively easy to capture such passwords, and of equal importance, to steal important information such as Social Security numbers and credit card numbers. This was the beginning of electronic identity theft. To counter the deliberate or accidental discovery of private or personal information, network transmissions were often encrypted. This deterred theft and discovery but did not stop it, for the encryption was often “breakable” with modest effort.

In 1991 Phil Zimmerman wrote and freely distributed an encryption program, PGP or Pretty Good Privacy, to protect files and electronic mail. PGP uses the public key method in which one key is public and the other private. The public key is given to those from whom one wants to receive messages; the private key is then used to decrypt the received message. Zimmerman built on work done at MIT using federal grant funds, and for a time was sued by the federal government for distributing the software in the public domain and making it available to non-U.S. sources. That suit was settled and PGP became widely used.

Project Athena, a project at MIT to develop a campuswide networked environment to integrate computing technology into the curriculum, faced the same problem of protecting records and transmissions.²⁰ To that end MIT developed Kerberos, a network-based user authentication protocol named for the three-headed dog that guarded the entrance to Hades in Greek mythology. Kerberos was designed to provide “strong” authentication for client-server applications by using secret-key cryptography. It is an especially secure method of authentication because the user’s password is never sent over the network, not even encrypted. In simple terms, by having a part of Kerberos in the client computer and another in the server computer and by encrypting transmissions between the two, Kerberos granted “tickets” with a timed expiration for access to services. In effect, Kerberos authenticated that “you are who you say you are” and made sure that in the event of an unau-

²⁰ <http://web.mit.edu/kerberos>

thorized use of a ticket, the expiration time of the ticket would deny access at some later time. Services that used Kerberos for authentication were said to be Kerberized, offering this extra level of protection to transactions.

Telecommunications Act of 1996, Communications Decency Act

In 1996 Congress passed the Communications Decency Act as part of the Telecommunications Act of 1996,²¹ the first major overhaul of telecommunications legislation since 1934, and it was expected to have far-reaching effects on the lives of Americans at work and in the home. Much of the act dealt with deregulating telephone services, where it was expected that increased competition in local and long distance markets would reduce charges. The Communications Decency Act focused on the activities of the Internet.

This act proposed that the current law, which stated that it was a crime to use a telephone in interstate commerce to make “any comment, request, suggestion, or proposal which is obscene, lewd, lascivious, filthy or indecent” be extended to include telecommunication devices. A second clause extended the telephone harassment component to telecommunication devices as well. The complication that resulted was that an individual or a service provider could be liable under the law for transmitting or otherwise making available the message, not just for creating the message. Thus, providers of online services had to exercise care if they knowingly transmitted an offensive message. This was going to have to be resolved by future legal and other actions.

Vendors

The 1990 decade had no less activity for vendors than the 1980 decade, with firms vying for dominance in the market, bringing out new products or new releases of old products, and using the courts to adjudicate claims of copyright infringement. As in the past, new firms also were on the scene and old firms faded or were bought out by competitors.

The major new player in the manufacture of personal computers was Dell Computer Corporation. Dell was started in 1984 by Michael Dell, who sold computers out of his dorm room at the University of Texas. The company followed a different model of selling that included being able to make and ship custom-ordered systems directly to customers in a matter of days.²² By early in the 1990s Dell had become a competitor not to be ignored, and by the end of the decade it had a market share of 25 percent of the personal computers sold in the United States, surpassing others who were dominant at the beginning of the decade,

notably Compaq, Hewlett Packard, and IBM.

During this decade Compaq acquired several old-time vendors such as Tandem Computer (\$3.5 billion) and Digital Equipment Corporation (DEC, for \$9.6 billion), as these companies could not keep up with the changing market yet had intellectual properties and name brand recognition that were considered of value.

Even newer companies faded away. Lotus Development went out of business when IBM bought it in 1995. NeXT followed when Apple bought it out in 1997. Commodore, one of the early microcomputer vendors from the 1970s, simply went out of business.

Apple Computer struggled to survive during the decade. Despite continuing to produce innovative and interesting products, such as Quicktime multimedia software and the early Newton PDA, Apple computers continued to lose market share. Apple’s share of the U.S. market fell from a high of 12 percent in 1990 to an estimated 3.5 percent by the end of the decade. There are various explanations for this decline, not the least of which was considerable management turnover and staff defections during this period. To improve conditions, Apple tried both licensing its operating system to run on other platforms and entering into agreements to have other vendors make Macintosh-compatible hardware. It entered into a number of product alliances with IBM on hardware (the PowerPC) and software (Taligent) and with others, such as HP, as well. It also was involved in a number of lawsuits with other vendors, most notably with Microsoft on a number of issues. It also tried different hardware and software solutions in order to be able to run the Windows OS on its computers. Shortly after Apple bought out NeXT Corporation, it named Steve Jobs as president again. Jobs undertook a bold program to revive the Apple product offerings, the first resulting in the “Bondi blue”—colored iMac (other colors came later), which was followed by iBooks and PowerBooks and other cost-competitive offerings. This strengthened the company but did not restore its former dominance.²³

Just after the decade ended, the Cray computer name faded away. Cray Research, the company that continued Seymour Cray’s original work to develop supercomputers, merged with Silicon Graphics (SGI), a maker of popular systems for computer graphics in 1996. It was expected that SGI systems would form the basis for Cray’s massively parallel systems. In 1999 SGI created a separate Cray Research business

²¹ “FCC—Telecommunications Act of 1996,” www.fcc.gov/telecomm.html

²² www1.us.dell.com

²³ Ken Polsson, 2003 Chronology of Personal Computers, www.islandnet.com/~kpolsson/comphist/

unit to focus on the unique requirements of high-end supercomputing, and in 2000 this unit was sold to the Tera Computer Company which undertook to market supercomputers with their own branding.²⁴

One of the powerful and influential network vendors that rose to prominence during the decade was CISCO Corporation.²⁵ CISCO became the leading vendor of hardware, software, and service offerings to IP-based networking. Claiming that its name is synonymous with the Internet, it developed network service solutions for business, education, government, and home networking. Founded in 1984 by several Stanford University computer scientists, it enjoyed sales of \$69 million in 1990, rising to \$12,154 million in 1999.

Legal Actions

Somewhat the same roiling activity prevailed for software vendors as for hardware vendors, although in this case there was much more legal action over copyright issues. For example, Lotus Development Corporation sued Borland International, claiming that their Quattro spreadsheet had copied the “look and feel” of the Lotus 1-2-3 presentation. In the end Borland was forced to remove the contested code, but within a year it sold the Quattro software to Novell Corporation, which later in the decade bought the WordPerfect package to create the WordPerfect Office Suite.²⁶

Much of the controversy between vendors during the decade was about web browsers, and it was played out between Netscape Communications and Microsoft and the U.S. Justice Department. These were the so-called browser wars. As noted earlier, while Netscape was first to release its Netscape Navigator, Microsoft soon followed with Internet Explorer embedded in the Windows operating system, making it difficult for consumers to use the Netscape software. It didn't help that when Steve Jobs was resuscitating Apple in 1997, he made an agreement with Microsoft that made Internet Explorer the default browser for all new shipments of Apple computers. Sun also initiated legal action against Microsoft for shipping Internet Explorer 4.0 with a nonstandard implementation of the Java programming language.

In 1997 the U.S. Justice Department asked a federal court to hold Microsoft in contempt for forcing PC makers to distribute Internet Explorer as a condition of selling Windows 95. Late that year Judge Thomas Jackson issued a preliminary injunction against Microsoft, requiring Microsoft to allow Windows 95 licensees the option of not including Internet Explorer. Following this ruling the Justice Department and 20 attorneys general filed antitrust lawsuits against Microsoft for this practice of bundling the software. Before the year was out, Microsoft applied for and

got a stay of this court order and issued an appeal. In midyear 1998, an appeals court ruled that Microsoft did not violate its consent decree by bundling Internet Explorer with the Windows operating system and overruled Judge Jackson's ruling. By the time this was all over in late 1998, America Online (AOL) had bought Netscape Communications, which passed into history after its brief but technologically significant software development. Microsoft in the end ameliorated its practice of bundling the two software components together to avoid further legal actions.

Technology Terms, Acronyms, and Buzz Words

POP (Post Office Protocol)—the protocol used in the client-server model, allowing e-mail to be sent and retrieved between the client and the server.

SLIP (Serial Line Interconnect Protocol)—the protocol that allows TCP/IP transmissions over dial-up serial telephone lines. It typically was used in the home to get direct Internet access using a simple modem connected to the home telephone line.

TCP/IP (given the importance of this protocol, we repeat the discussion from the 1970s decade)—TCP, the Transmission Control Protocol, was first proposed by Vint Cerf, who wrote the paper “A Protocol for Packet Network Interconnection” with Bob Kahn.²⁷ Later Cerf, with the participation of Yogel Dalal and Carl Sunshine, formally described TCP in 1974. The objective was to define a standard information packet transmission protocol so that it was possible to build a network of cooperating computers. It was designed to be flexible enough to handle the physical differences in host computers, routers, and networks and still allow such physically different entities to transmit data between themselves despite differences in packet sizes. Later in the decade TCP split into TCP/IP (Internet Protocol), where IP was responsible for routing packets of information and TCP was responsible for creating packets, error control, and the retransmission and reassembly of packets. A whole new technology and industry was started to develop the software and hardware needed to produce fast and inexpensive gateways for routing information over networks stretching around the world. By the end of the decade TCP/IP was reasonably mature and had become the standard military protocol for networking.

²⁴ www.cray.com/company/history.html

²⁵ [http://newsroom.cisco.com/dlls/company overview.html](http://newsroom.cisco.com/dlls/company%20overview.html)

²⁶ Ken Polsson, 2003 Chronology of Personal Computers, www.islandnet.com/~kpolsson/comphist/

²⁷ *IEEE Transactions on Communications Technology*, vol. COM-22, no. 5, May 1974.

1990 to 1992 at Cornell

Several bold initiatives to bring campuswide leadership to computing and to improve computing services on campus were started early in the decade. The most important was the installation of a contemporary campuswide network as part of a joint study with IBM. The new network provided a consistent technology across the campus and enabled the development of new applications, in particular, Cornell's Bear Access, one of the earliest network navigators that included a robust suite of applications for access to services. New services were developed for the three campus constituencies: research, instruction, and administration, and continuing services were strengthened and improved. CIT reorganized its staff resources and formed internal committees and task groups to start implementing services based on client-server technology.

New information technology advisory committees—CUBIT (the Cornell University Board for Information Technology) and ADSPAC (the Administrative Data and Systems Policy Advisory Committee)—were formed to advise the provost and senior vice president on IT issues. The provost formed two campuswide committees, one to study what information technology would look like in the year 2001 and the other to consider instructional computing support and services. The 1992 report from the 2001 Committee established a vision of the influence of information technologies on Cornell for the 1990s. CIT also participated in several campuswide planning and service quality improvement initiatives.

Tight Financial Situation on Campus

As the 1990s started, the predictions Provost Barker was making in early 1989 about “hard times” from New York State budget cuts seemed to be coming true. Tight budgets continued across the university but most particularly in the statutory colleges, which relied heavily on funding from the state. In 1990 an 8 percent tuition increase was approved for the endowed colleges to raise faculty salaries by 10 percent and financial aid by 18 percent. There was serious discussion of various options: layoffs in the statutory colleges, increasing the student-faculty ratio by 10 percent, cutting faculty and staff positions by 4 percent, lengthening the financial planning horizon, and taking a longer-term look at cost reductions.²⁸ In 1992 when President Rhodes gave his State of the University address, he declared that the next five years would be difficult for Cornell. In particular, he mentioned that \$3 million in administrative costs were being cut from the 1992–93 budget and that an additional \$1.2 million in cuts were being sought for next year. There was the clear sense that the financial problems would linger.

In addition to cutting expenses, in 1990 the university launched a very aggressive campaign to raise \$1.25 billion over the next five years. The most ambitious campaign ever launched by any university, this campaign demonstrated the commitment to underpin the future with additional funds. Reports in the early years indicated excellent progress on meeting the goals of the campus, and there was optimism that the campaign would meet and surpass its goal, which it certainly did by raising \$1.507 billion when it ended in 1995.

New Executive Staff and Deans

While President Rhodes started his 12th year as president in 1990, changes in the executive and academic staff continued. In early 1990 Mal Nesheim, who had been assistant provost for budget and planning, was appointed provost; Robert Barker was appointed senior provost, with primary responsibility for planning. Don M. Randel, who with Philip E. Lewis had been named associate dean of the College of Arts and Sciences in 1989, was named dean in 1991. Alan G. Merten, a specialist in computer and information systems, was appointed dean of the Johnson Graduate School of Management in late 1989.

Frederick A. Rogers from Carnegie-Mellon was appointed chief financial officer and vice president for finance and treasurer as of July 1, 1990. Rogers, who had been director of planning and administrative systems before becoming vice president for business affairs at Carnegie-Mellon, was expected to play a key role in improving planning efforts to better manage Cornell's resources.²⁹ Controller John S. Ostrom retired in 1991 and was replaced by Yoke San Reynolds, and the Division of Campus Life was phased out and became part of the responsibility of the Division of Academic Programs, headed by vice president Larry Palmer.

Norman R. Scott, now vice president for research and advanced studies, and Malvin H. Kalos, director of the Theory Center, formed the duo guiding the future development of supercomputing at Cornell. Both had been appointed to their positions in 1989. Last, Tim P. Mount was appointed director of CISER, the Cornell Institute for Social and Economic Research, succeeding Robert McGinnis, founding director, who had served since 1982. In various ways all these individuals would play key roles in influencing the developments in computing/information technology at Cornell in the 1990s.

²⁸ “Barker urges 10 percent increase in student to faculty ratio,” *Cornell Chronicle*, May 3, 1990.

²⁹ “New treasurer, vice president joins CU July 1,” *Cornell Chronicle*, February 1, 1990.

CIT Priorities

As the decade started, CIT was continuing to meet the challenges of its mission statement, enunciated by VP Lynn when he first came to Cornell in 1988. To repeat, these challenges were:

- leadership in facilitating the use, and as appropriate, development and implementation of the underlying *technologies* themselves in support of the university's needs;
- leadership in providing or nurturing the *services* necessary to enable the university community to apply those technologies to its needs;
- leadership in supporting the *applications* of such technologies and services in those areas where central support is most effective in facilitating such applications across the spectrum of university requirements; and
- leadership in coordinating the definition and support of appropriate *policies, plans, standards, and controls* to provide an effective architectural framework for innovation and for the development, implementation, and use of information technologies throughout the university.

Help Desk Update

The Help Desk celebrated its first anniversary in the summer of 1990, having been established in 1989 as the one-stop shop for CIT front-line services. When started it was squeezed into a corner on the first floor of the CCC building. The Help Desk took on much of the previous one-on-one consulting load, but increasingly the consult feature and e-mail and the telephone were being used to answer front-line questions previously done face-to-face with users. In the first 10 months there were 13,000 contacts with the community, of which 90 percent were resolved immediately.³⁰ A year later, there were more than 28,800 inquiries, more than double those of the previous year, and more than 4,000 requests for more extensive consulting.

General Computing Accounts Replace Free Distribution Accounts with Computer Resource Units (CRUs)

The previous distribution accounts, now known as General Computer Accounts, or GCAs, with some exceptions were provided at no cost to all faculty, students, and staff for use of services on CIT host computers. Students were pre-assigned user IDs in their registration materials and had to come to the Help Desk to activate their ID. Students in the Hotel School and the Johnson Graduate School of Management were excluded, as they had their own computing support.

A subtle but important change was made in resource accounting. GCAs were allocated CRUs—Computer Resource Units—instead of dollars even though the calculation was the same; only the frame of reference changed, reflecting service considerations rather than fund considerations. Accounts were started with 50 CRUs for students and 100 CRUs for faculty or staff, which was expected to be sufficient for a week's work, and refreshed to that amount weekly.³¹ It was expected that these accounts would be used by students and faculty to access CUINFO, electronic mail, the library, and other host-based services, i.e., time sharing and batch processing in support of instruction on CornellA and CrnlVax5, and by staff for similar services on CornellC.

CornellC Mainframe Upgraded to IBM 3090-200J

In the summer of 1990, CornellC, now used primarily for business systems processing, was upgraded to a model 3090-200J to gain about 50 percent more capacity over the 3090-200 installed in 1986. The cost for the upgrade was \$3.1 million. The driving factor was the increasing use of the mainframe to run old and new online administrative applications during the peak load of daytime work hours. Transaction response times had reached an unacceptable level by this time, and there were difficulties in meeting overnight batch processing commitments. Once the upgrade was completed, all e-mail use for CIT staff was moved from CornellA to CornellC to relieve overloading on CornellA. CornellA also was overloaded as a result of increased student use of GCAs for classwork and e-mail services and continued growth in the CISER load, so this shift of CIT staff e-mail use helped balance the load on the two machines.

Growth in CIT Servers—The Server Farm

The number of server computers in the CCC computer room continued to increase. This pattern started in the late 1980s and the early 1990s as new applications were run on servers instead of mainframes. Smaller computers were added one at a time, usually in response to CIT's proposing a new service or evaluating mini- or microcomputer-based products, or in response to requests by departments to house their computer in a secure and supervised computer room. Accordingly, operations problems escalated. A rough count at the start of the fall semester in 1991 identified as many as 30 such computers scattered in

³⁰ Office of Information Technologies, Resource Services, annual report 1989-90.

³¹ "General Computer Accounts Available to Staff, Faculty, and Students," *CITnews*, July-August 1990, vol. 3, no. 3.

the machine room. Anticipating continued growth, the computer operations staff developed a special area with racks to house a larger number of servers. This area became known as the server farm. Further, given the wide range of CPUs and operating systems, the computer room operators, well versed in overseeing mainframe operations, had to be trained on these systems before they could assume responsibility for the usual operator functions such as backup, restart, etc. Once this was completed for CIT, accommodating servers became one of the facilities management agreements that CIT offered to the campus. By the end of 1992 there were 12 official servers in the “farm” (after removing from the count the machines that were not servers as then defined.)

Training and Publication—*CITNews*, “InCITes,” and *Inside CIT*

Training continued to be a key activity, with CIT providing short courses to introduce basic software tools to the faculty and staff. In a change from the past, training was now a cost-recovery activity, and departments were expected to pay for training their staff. A small number of students were allowed to attend at no cost.

The July 1990 summer training schedule was devoted to topics such as Introduction to the Macintosh; Introduction to Microsoft Word or Word for the Power User; Desktop Publishing Using Pagemaker; Hypercard; Network Administration Using Appleshare; Getting Around in PC Windows; Microsoft Word for Windows; and Using FileMaker II. Only one course titled Introduction to the IBM Mainframe was offered, and no Fortran course was offered. The priority was teaching about the desktop computer tools.

Two new general information publication outlets were started: *CITNews*^a and “InCITes.” *CITNews* was first published in mid-1990 as a successor to the previously published *CIT Newsletter*. It was a bimonthly news magazine featuring policy and technology issues as well as news about current IT-related events occurring in CIT, on the campus at large, and in higher education in general

“InCITes” was introduced in mid 1991 as a regular column in the “Networking” section of the *Cornell Chronicle* to provide technology information to employees on the campus. This was the latest such attempt to meet the need for up-to-date, user-friendly information about technology using the regularly published and widely distributed Cornell newspaper. The first “InCITes” column was written by Nancy Flynn and Sharon Marcus.

Inside CIT was started in 1990 to keep CIT staff informed and up-to-date on all that was going on in

OIT (the vice president’s office) and CIT and was intended to be strictly internal and informal so issues could be openly presented. As the first issue states: “The Internal Information Dissemination Work Group (IDWG) of the Services Committee decided to produce an internal CIT newsletter as an experiment to see if information could be gathered and disseminated on a fairly timely basis without too much effort on the part of any one staff member.” Cecilia Cowles volunteered to coordinate the experiment, which turned out to be more long-term, and the publication continued for a number of years.^b

SCRIPT Support Phased Out

In mid-1991 CIT phased out support of Waterloo SCRIPT on CornellA and CornellC. SCRIPT had been the primary application package for text processing on IBM mainframe computers since the mid-1970s and had been eclipsed by the newer word processing systems on desktop workstations, which were more versatile and more generally used. The application no longer warranted support at its diminished level of use.

Microcomputer Sales, Back-to-School Fairs, Selecting a PC Clone

The sale, installation, deployment, and support of microcomputers continued to grow every year, not only for faculty and staff, but also for students. OIT annual reports from the 1991–92 period give sales figures of 3,500 to 4,000 new workstations a year, with the repair of about 6,000 systems a year.

Back-to-school fairs continued to be held once a year at the beginning of every fall semester, usually in late August. Incoming freshman students were encouraged to pre-order equipment so that it would be set aside for them to pick up when they first came to campus, typically when their parents were still around. The fairs were held in Lynah Rink, where sales staff and consultants were available as well as representatives from vendors and different Cornell units that were directly involved, such as networking staff, or indirectly involved, such as the Safety Division, which advised on protection and prevention measures to avoid theft.

In 1990 it was estimated that 6,000 individuals entered the area within the ice rink boards, where all of the displays, consultants, and sales desks were stationed on the concrete surface. This estimate was obtained with a good old-fashioned mechanical counter, a “clicker” used for counting items by depressing a button.³²

³² *Inside CIT*, September 7, 1990.

In late 1990, responding to requests from campus users to consider offering a PC-clone instead of just IBM systems, Workstation Resources formed a task force to consider the issue and make a recommendation. The task force, composed of CIT staff across different divisions and representatives from departments across campus, became known as SIAM, the Small Intel Architecture Machine task force. In May 1991, after evaluating 12 separate vendors including Dell and DEC, SIAM announced Everex Systems as its selection. The company was reputed for producing quality hardware, offering a broad range of different machines, and having two product lines—entry level and high performance. Once arrangements were completed, MOS announced the availability and pricing for these systems.

New Facilities Management Service; Software Contracts

A new Facilities Management service was started in CIT Services to provide fee-based help to those units that wished to have their workstations and local area networks designed, installed, and maintained by CIT. This service relieved those units of developing their own in-house expertise. Graham Hall, the first manager of the service, negotiated with both the client and the operating unit in CIT that was to provide the service.

CIT also assumed a leadership role in university-wide implementation of software contracts and licenses. By negotiating licenses for the entire university, it was possible to achieve significant cost savings. Depending on the involved software, the cost, and the prospective users, CIT provided such software at no cost or low cost to the participating parties.

New Information Technologies Committees: CUBIT and ADSPAC

In 1990 two new information technology committees were organized: CUBIT (the Cornell University Board for Information Technology) and ADSPAC (the Administrative Data and Systems Policy Advisory Committee).³³

CUBIT, which replaced UCB (the University Computing Board), was charged with providing advice to the provost and the senior vice president on policy and strategic directions concerning the use, development, operations, support, and financing of all aspects of information technologies within the university. The responsibilities of CUBIT were

- to have general oversight review and advisory responsibility for the adoption and execution of IT policy;

- to review the proposed plans of OIT to see that they integrated with plans of other campus units;
- to advise on the priorities of the needs and directions for campus information technologies.

Members were to be the provost, the senior VP, the VP for IT, the VP for finance, the VP for planning, the VP for research and advanced studies, the university librarian, and a dean to be selected by the provost. The first meeting of CUBIT, held on October 9, 1990, was attended by Mal Nesheim, Jay Morley, Stuart Lynn, Fred Rogers, John Weisenfeld, Norm Scott, Alain Seznec, and Alan Merten.

The CUBIT charter called for the creation of a University Advisory Committee for Information Technology (UCIT), whose chair was to serve on CUBIT. Conversely, the chairpersons of ADSPAC and CUBIT would be members of UCIT. I found no documentation about UCIT or evidence that it was ever formed.

ADSPAC was responsible for coordinating the strategic planning and implementation of information data and systems consistent with the Administrative Data and Systems Policy first issued in May 23, 1989, and since updated. ADSPAC replaced the Administrative Systems Steering Committee and was to advise the provost and senior vice president and CUBIT in the context of CUBIT's overall responsibilities.

ADSPAC was to be chaired by the treasurer and vice president for finance, and members were to be the "responsible data administrators" (RDAs), responsible for coordinating overall policy and planning for a given functional area. These RDAs were from Business Affairs, Decision Support, External Relations, Facilities, Finance, Human Resources, Materials Research and Advanced Studies, and Student Services as well as the VP for IT. The CIT director of information resources (IR), the university auditor, and the university budget officers were to be nonvoting members.

The responsibilities of ADSPAC were

- to coordinate the development of policies and plans governing each functional area;
- to coordinate the development of a strategic plan integrating all of the functional areas;
- to track implementation of such policies and plans;
- to recommend to CUBIT actions that needed to be taken to facilitate such policies and plans, including actions concerning resource requirements;
- to facilitate a transition from a "systems view" of administrative systems to a "data view";

³³ M. Stuart Lynn, "Information Technologies at Cornell University: Advisory Structure," October 1990.

- to coordinate processes to support the definitions of central data to ensure completeness and integrity; and
- to coordinate with other campus leaders dealing with issues of policies, plans, priorities, and resources.

As a result of these actions, the high-level oversight groups for information technology and information systems were now in place and ready to play their role in the coming decade.

Formation of 2001 and Instructional Computing Committees

In early 1991 VP Lynn discussed how CIT services would be affected by future developments in technology. One of the key underpinnings, for which planning was already under way, was the deployment of a campuswide data communications network. To take advantage of this new development and to address the future needs of IT, Provost Nesheim chartered two campuswide committees: the 2001 Committee and the Committee on Instructional Computing. The first, chaired by Professor Simon Levin, was to establish a realistic vision of the influence of information technologies on Cornell over the next decade; the second, chaired by Professor Keith Dennis, was to study the near-term campus requirements, plans, and policies concerning all aspects of instructional computing.^c In announcing these two studies, VP Lynn stated that if the 1980s had been the decade of personal computing, the 1990s would be the decade of collaborative computing, and that CIT's role would be to support this shift through the 1990s and beyond.

Internal CIT Committee Status—Technologies, Applications, Services

The three parallel committees—Technologies, Applications, and Services—continued their activities to satisfy the CIT mission statements given at the beginning of the decade. The Technologies Committee, under the leadership of Bob Cowles, continued to update the Technologies Framework. In September 1990 version 0.8, covering the period 1990 to 1995, was issued and circulated on campus for comment. Later another update was issued (version 2.0 draft) for the period 1995 to 1997, with reasonable conjectures about the technology that would be available in this period to guide the planning process in the interim period. The Applications Committee changed leadership in 1990 when Steve Worona became chairman. While continuing to consider its mission, it did not come to closure with the wide-ranging applications that were rapidly changing with new innovations. The Services Committee continued to focus on

evaluating old and new proposed services and recommending changes to improve services to users and keep current with changing technology.

Financial Restructuring

While all the activity described so far was taking place during the first several years of the 1990s, CIT itself was undergoing change in its way of conducting business. Spilling over from the 1980s when CIT was first organized was the issue of restructuring CIT's internal accounting procedures to better represent the organization then put in place. Once this was done, a financial restructuring project was undertaken to revise the charges and charging philosophy so as to be in tune with the organization, the services, the technology, the times, and university policies.

The issue addressed was how to create appropriate cost transfer mechanisms internal and external to CIT, including both direct and indirect costs. Again, the issue was raised of whether rates, given their complexity, were the best mechanism for transferring costs, or whether in some cases the direct transfer of costs would be more appropriate. In the end, CIT was treated as a pseudo-enterprise, and a new blend of charges was put in place, including direct billing (using rates), indirect billing (by cost transfers), and centrally absorbed costs. Generally speaking, CIT services were to be recharged/billed to departments, while applications were covered by the university budget. Besides trying to define, design, and implement these changes, CIT also was facing budget cutbacks of 5 percent a year—close to \$1 million total over the years 1990 to 1992.

Need for Additional Building Space

With the merger of Telecommunications and CIT and the expanded networking program—and the growth in staff creating and supporting new services while continuing the old services—the space crunch that occurred in 1990 and 1991 was inevitable. When Theory Center staff vacated the basement of Caldwell and moved to their temporary trailers behind Carpenter Hall, the CIT Network Management Center and support staff moved into that space. The Information Resources Division located on the fourth floor of CCC had filled that space to capacity and moved the Finance and Business group of 11 staff to rented space at 33 Thornwood Drive, in the same building where CIT Sales and Repair was located.

In 1992, after the Theory Center relocated its equipment to Rhodes Hall, CIT Equipment and Operations support staff moved to the basement computer room, and the first floor of CCC was assigned to user support functions. Improved and larger space was assigned to

the Help Desk, to the Instructional Resource Center, and for other training activities. All space planning and assignments were coordinated by the CIT Space Committee, which also started to concern itself with future space needs and the distribution of CIT staff to a growing number of buildings, mostly off campus.

CIT Reorganization; Organization Chart

In 1992 there was a restructuring of both the CIT Workstation Resources and the CIT Services divisions to regroup several of the service providers and to refocus on the calls for service being placed on the organization. Workstation Resources, under the continued direction of Larry Fresinski, was renamed the Division of Sales and Support and assumed responsibility for all CIT's enterprise activities. Microcomputers and Office Systems (MOS), responsible for selling and servicing personal computers, workstations, and servers, was

transferred from CIT Services to this new organization, along with the Facilities Management service. They took as their slogan "making technology work for you." In this new arrangement Sales and Support offered increased support for desktop and associated technologies.

Instructional Support Services was moved from the former Workstation Resources division to CIT Services, under the continued direction of Carrie Regenstein and with a renewed focus on improving the use of technology in support of instruction. Carol Lambert, General Support Services; Donna Tatro, Departmental Support Services; Nancy Flynn, Outreach; and Cecilia Cowles, Service Integration, rounded out the rest of the management staff and units in CIT Services.

There also was reorganization in Network Resources, including the assumption of support services for new

Table 1. Cornell Computer Services Organization Chart, 1992 to 1993

CIT Services	Ruth Sabeau, Director
• Outreach	Nancy Flynn
• General Support Services	Carolyn Lambert
• Departmental Support Services	Donna Tatro
• Instructional Support Services	Carrie Regenstein
• Service Integration	Cecilia Cowles
CIT Sales and Support	Larry Fresinski, Director
• CIT Sales	Jeanne Carey
• Technologies and Standards	Tom Young
• Workstation Services	Gordon Suggs
• Workstation Systems Services	Tom Young
Network Resources	Dave Lambert, Director
• Advanced Technologies	Dick Cogger
• Networks and Telecommunications Services	Pat Nelson
• Network Systems Services	Mark Oros
• Collaboration Systems	Barbara Skoblick
Computer Resources	John Rudan, Director
• Technologies and Standards	Bob Cowles
• Research and Analysis Systems	Dan Bartholomew (acting)
• Computer Operations Services	Dave Pulleyn
• Production Services	Jim Doolittle
• System Programming Services	Tom Boggess
Information Resources	Dave Koehler
• Technologies and Standards	Tom Dimock
• Academic and Human Resource IS	Libby Gruppuso
• Business and Financial IS	J. R. Schulden
• Office Support Systems	Mark Mara
• Scholarly Information Systems	Lynne Personius
Office of Information Technologies (OIT)	
• Assistant for Planning/Coordination	Agelia Velleman
• Management and Budget	Laverne Thomas
• Special Assignments	Steve Worona

technologies and applications. Unit leaders were Pat Nelson, Networks and Telecommunications Services; Mark Oros, Network Systems Services; Dick Cogger, Advanced Technologies; and Barbara Skoblick, Collaboration Systems. Network and Telecommunication Services formalized the merger of the previously independent Networks and Telecommunications groups. Collaboration Systems formalized the support of electronic mail and audio and video teleconferencing into one group. Steve Worona, who previously headed this group, joined OIT as a special consultant. Lynne Personius, head of Scholarly Information Systems in the Information Resources Division, now supported CUINFO, assuming Worona's previous responsibility. Table 1 reflects the organization chart for CIT after all these changes.

Cornell Computing Directors

In the early 1990s, CIT formalized its recognition of the group known as the Cornell Computing Directors (CCDs). The group, formed in the 1980s, was composed of representatives from laboratories, departments, and research centers that had their own IT expertise, largely built around DEC systems. The group was brought into greater prominence in the late 1980s when they lobbied CIT to create PINCNet, a DECNet ring that would connect all DEC computers across campus. Following that, the CCDs were increasingly used as a sounding board and advisory group for CIT proposals. With formal recognition of the group, a mailing list of members was created and monthly meetings were held with CIT Services staff and other invited CIT staff members. The CCDs played a major role in the evaluation and selection of UTP Ethernet, and CIT continued to call upon them as needed for advice as new services and technologies were considered for evaluation and deployment.

Current information about the CCDs can be found on their web site: www.ccd.cornell.edu.

Network Pilot Project in Residence Halls

While some students were very anxious to have high-speed network connections in their rooms, there was some hesitancy to move ahead with a full project to provide network connections in all dormitory rooms, given the costs involved and the likely subscription rate. To address this issue, in early 1991 a joint committee of CIT staff and Residence Life staff was formed to conduct a pilot project giving students access to the campus network in their rooms.³⁴ The pilot project, using UTP Ethernet, was conducted during the 1991–92 academic year and involved 328 students in Dickson and Donlon Halls. Its success led to a formal project to place network connections in all dormi-

tory rooms later in the decade. Giving students access to computer terminals in dormitory complexes was studied in the 1970s, was partially implemented in the 1980s, and now it seemed likely that every dorm room would be wired to the campus network.

Cornell Worm Case Settled

The first month of the 1990 decade brought closure to the "Morris incident" of 1988. Robert T. Morris, at that time a graduate student in computer science, was convicted of a felony charge for releasing a computer program that overwhelmed computers nationwide in November 1988. The internal Cornell investigation³⁵ led by VP Lynn found Morris responsible for creating and spreading the worm and for acting "with reckless disregard for those probable consequences." While acknowledging that Morris probably did not intend for the worm to destroy data or files, the report stated that he probably did intend for it to spread widely, although not replicate uncontrollably as it did. On January 3, 1990, Morris was convicted of a felony charge for releasing his computer program,³⁶ and in May 1990, Morris was sentenced by Judge Howard G. Munson to three years of probation and ordered to pay a \$10,000 fine and perform 400 hours of community service. While Morris planned an appeal, that sentencing effectively brought the Morris^d incident to an end for Cornell.

Another Cornell-Initiated Network Incident

In 1992 there was another network virus incident at Cornell. In February two students were arrested and charged with computer tampering and allegedly launching a computer virus embedded in three games. Both students were employed by CIT, and one of them worked in a campus Apple Macintosh facility from which it was believed the virus had been launched. The national Internet incident tracking service CERT (Computer Emergency Response Team) at Carnegie-Mellon had quickly alerted Cornell that a Macintosh computer virus embedded in three computer games had been launched through a Cornell computer. The students were arrested and charged with a Class A misdemeanor and held in jail until their bond was posted.³⁷

³⁴ "Report from the Pilot Project Committee Phase 1 Report Model Design Statement," May 21, 1991.

³⁵ "The COMPUTER WORM: A Report to the Provost of Cornell University on an Investigation Conducted by the Commission of Preliminary Enquiry," February 1989.

³⁶ "Morris convicted for releasing 'worm' into computer network," *Cornell Chronicle*, January 18, 1990.

³⁷ "Students charged with releasing computer virus," *Cornell Chronicle*, February 27, 1992.

In May, a Tompkins County Grand Jury indicted David Blumenthal and Mark Pilgrim for computer tampering in the first degree, a class E felony, and in June another student, Randall Swanson, was indicted for participating in this endeavor. Closure came quickly this time; in October Judge Betty Friedlander sentenced each of the participants to pay \$2,087 in restitution to Cornell and \$366 to Medtronic Cardio-Care, where employees claimed to have suffered damages. In addition, Blumenthal and Pilgrim were ordered to perform 520 hours of community service; Swanson was ordered to perform 450 hours of service, but he declined to accept the order, which he was permitted to do under the terms of the violation. VP Lynn's comment at the conclusion was, "This was no experiment; it was a tale of misused talent, deception, and self-gratification at the expense of other people's time."³⁸ Again, the university experienced an upheaval and once again emphasized the need for responsible computer behavior, fully understanding that despite all attempts and vigilance, such incidents were only the beginning of the network revolution and future similar incidents.

PSINet Replaces Telenet Network Services

CIT terminated Telenet Public Data Network services for off-campus access to CornellA and CornellC computers in February 1990.³⁹ Telenet was installed in 1974 as part of the initiative to increase the number of off-campus users, but newer and less expensive alternatives were now available. The combination of lower costs and the use of Internet protocols made PSINet a better choice at this time. PSINet connected to Cornell through NYSERNet. Users would have to connect with a PSINet site, a chargeable transaction, but after that the use of NYSERNet would be at no charge. PSINet was available in most of the major cities and by the end of the year was expected to have 60 sites in the country.

A New Network for Cornell

In April 1991 VP Lynn announced that Cornell would undertake one of the most extensive university projects in the country to install a campuswide high-speed optical fiber network.⁴⁰ Although as noted in other sections, work had already started on the project. The initial deployment was to be a 100 megabit-per-second network that would be fast enough to transmit the text of an encyclopedia in 10 seconds. This was to be a joint project between Cornell and IBM who would share in the total cost of \$5.9 million, of which IBM was contributing \$2.8 million in equipment and support. Work was already under way to install an estimated 28.9 miles of fiber optic cable connecting 105

campus buildings. H. David Lambert, newly appointed director of network resources, was the project manager for Cornell. The basic plan was to use the fiber optic network as the backbone of the campus network, which would interconnect individual lower-speed networks throughout the campus and provide the interface to national and international networks.⁴¹

Building the new campus network was divided into a number of discrete steps:

- building the backbone
- building the networks in campus buildings
- improving operational and support services
- upgrading external links
- integrating the electronic communications environment

This plan recognized the transitions needed from the current situation to achieve the new integrated solutions. The years 1990 and 1991 were busy with a large number of "construction" projects to build the network, elements of which were operational in late 1991, and all were pretty much completed by the end of 1992.

New Backbone Network

The proposed solution to building the backbone was to implement the emerging high-speed standard, Fiber Digital Data Interface (FDDI) protocol, operating at 100 Mb/s. Optical fiber was to be extended to 105 major academic buildings on the campus, and interconnections were to be made between CITNet, TheoryNet, and the Engineering Ethernet, the three current backbones on the campus. IBM was to supply the building-level routing hardware (6611 routers) in exchange for Cornell's developing key software for routing IP (Internet protocol), AppleTalk, and IPX (Novell) protocols to interconnect LANs using these different technologies. The Network Technologies group under Dick Cogger was to play a key part in this project. Since the mid-1980s this group had been involved with developing the software programs for the AT-Gateways so that both Macintosh- and PC-based LANs (AppleTalk and Omninet) could be connected to the then-running Pronet (TCP/IP) backbone and other related network technologies such as

³⁸ "Judge calls computer virus a 'malicious' act," *Cornell Chronicle*, October 15, 1992.

³⁹ "PSINet Makes Cornell Debut," *CITnews*, November-December, 1990.

⁴⁰ "CU to install campuswide fiber-optic data network," *Cornell Chronicle*, April 4, 1991.

⁴¹ H. David Lambert, "Building the Communicating Campus," *CITNews*, November/December 1990.

terminal emulators. In other words, they were an experienced and capable group for this task.

Campus Building Networks—UTP Ethernet

Constructing the building networks meant selecting a technology that took advantage of the copper wire installed for the voice telephone network (which was itself ubiquitous on campus) and that could support multiple workstation technologies, provide a smooth migration path, and be reliable and reasonably priced. A technical committee of campus users and CIT staff appointed by Lambert to study the offerings selected the UTP (Unshielded Twisted Pair) Ethernet and selected David Systems as the supplier. Several large prototype installations in the Cornell Theory Center, the College of Human Ecology, and the Division of Nutritional Sciences using UTP Ethernet technology were successful enough to support this choice. It was envisioned that shared Ethernet technology to the desktop would replace and improve the services offered by the hodge-podge of current technologies that included Sytek, Omninet, and 3270 connections, and support the current LAN technologies. Not only would this new service be faster (shared 10 Mbit/second transfer rates to the desktop instead of the 1 Mbit/second AT-Gateways), but a key feature of this new connectivity was a low flat-rate monthly fee of \$11 per workstation port and \$30 per LAN. There was an installation charge of \$70 per port and \$300 per LAN to effectively enable the “B” side of each telephone wall outlet for these communications. (When telephone wall outlets were installed in 1984–85, there were two lines per outlet that could be used for a telephone [typically the A side] or telecommunications.) These attractive network prices were made possible by keeping the rates for telephone services at their current level; the full debt of the telephone system installation was paid off in 7 years instead of the planned 17 years, effectively subsidizing the new network.

Network Management Center

Improving operational and support service included the establishment of a single point of contact for network management. In July 1990, the Network Resources division of CIT created a Network Management Center (NMC) in Caldwell Hall and assumed responsibility for managing both CITNet and TheoryNet, Cornell’s main central campus Pronet backbones. Old-timers still referred to the NMC as the NOC, the Network Operations Center started in the 1980s, and the two terms were used interchangeably. New network monitoring tools and equipment, taking advantage of the NSFNet-developed SNMP (Simple Network Management Protocol),^e were installed.

They allowed more extensive testing and diagnosis of Ethernet networks, from the Network Management Center to the individual workstation.

Improving installation services was another key component, and improvements were expected from the integration of the staff of Telecommunication Services into the CIT Network Resources division and the implementation of a single management/supervisory structure for both services.

A key component of improving future services was the development of campus networking standards: if all users followed the same standards, management and operations would be simplified, and CIT could focus its support efforts on achieving better service.

Upgrading External Links

Upgrading external links and community access required Cornell to maintain its leadership and participation in national and international organizations. The National Science Foundation had allocated \$600 million over the next six years to expand and upgrade NSFNet as part of the High-Performance Computing Act then working its way through Congress. The plan was for NSFNet to eventually evolve to the National Research and Educational Network (NREN). Cornell had provided leadership in establishing NSFNet and was expecting to be doing the same for NREN.

Cornell had been chosen as one of a small number of institutions to participate in the upcoming upgrade of the national data backbone from 1.5 to 45 megabits/second. In December 1990, Cornell in cooperation with SPRINT was awarded an NSF contract for providing and managing all international NSFNet connections. SPRINT was to provide the circuits while network operations services were to come from Cornell. The award of this contract fit well with CIT’s plans to provide round-the-clock staffing in the NMC/NOC, improving responsiveness to problems for all customers internal to Cornell and external as well. At about this same time, the Theory Center decided to discontinue its networking operations. This activity, along with staff and equipment, became part of the NMC/NOC.

Integrating the Electronic Communication Environment

The elements for creating the integrated electronic communications environment included

- “humanizing” electronic mail, the most basic and oldest of all network applications, by using more workstation-based systems that would provide the look and feel of workstations.
- enhancing public information services such as CUINFO, Cornell’s pioneering work; new offerings

and extensions to take advantage of more powerful workstations were one aspect, and extending offerings to off-campus outreach activities was another.

- enhancing voice and data network services; possibilities going beyond the newly installed AUDIX voicemail services included integrating workstation and telephone services, providing cellular telephony and paging services, and adapting other coming innovations.

- providing network services to the dormitories and improving services to homes; in late 1990 modem services were being upgraded to provide 9,600 baud dial-up access while another project was under way to provide direct access to the Cornell backbone using IP and AppleTalk protocols over serial phone lines using SLIP (Serial Line Internet Protocol).

Taken together, these network projects would make real the visions of the Networking Task Force report: to develop the “communicating campus” and provide the integrated, seamless environment where one could connect from Cornell to anywhere at any time.

New Network Services

Coincident with the installation of new networking technology at Cornell, new network-based applications were being introduced to the campus, almost all of them building on the client-server architecture gaining popularity.

Project Mandarin and Bear Access

In 1990 CIT and Apple Computer embarked on Project Mandarin, a creative endeavor to merge complex technology with information resources and an innovative workstation interface to enable administrators to run billion-dollar universities more efficiently and effectively.⁴² In loose terms, the project was to create an “administrator’s workbench” similar to other efforts going on at the time to create specialized interfaces and tools—for example, the efforts in the Theory Center to create the “scientist’s workbench.” The name Mandarin was explained in the Project Mandarin general information document: “The dictionary defines mandarin as a high civil servant thought to exercise large undefined powers without publicity or political control. Mandarins were highly respected and highly educated figures in Chinese history”—that is to say, a bureaucrat having high influence, like a university administrator. As the project was initially envisioned, it would provide an administrator with a rich set of tools for using internal and external data in the decision-making process. The project was underwritten and supported by Apple Corporation, with Dick Szymanski, local Apple sales representative, playing a lead role. The plan called for

Mandarin to be ported to IBM PC architecture and to Unix systems so that it would be useful in the diverse environment at Cornell and other similar institutions.

As the project was taking shape, the Office Support Systems unit in Information Resources (IR) was looking for a winning application of client-server computing, the technology underlying the project infrastructure. The first application chosen displayed course grades to students. A survey of the students who stood in long lines at the entrance to the registrar’s office determined that the majority of them were waiting for a clerk to verify their identity and display and print their grade reports. With the support of David Yeh, assistant VP for student and academic services, and Gloria Howell, the registrar, a team from Office Support Systems wrote a Macintosh program that would display grade reports when a student’s ID was entered. The team members were Peter Bosanko, Doug Hornig, and Angela Mennitto from the Office of the University Registrar. A Macintosh computer was put on the counter at the entrance to the registrar’s office in March 1991, and students were invited to use the system. According to Mark Mara, director of Office Support Systems, this service was so popular that in the first week over 6,000 different students had used that single computer.⁴³ A clear winner and an “awesome idea,” as one student described it.

This successful project went by various names, one of them being Bear Access, until that name became more appropriate for grouping a number of services and applications. To get student involvement, the Office of the University Registrar and the Office Support Systems group, with the encouragement of Apple, decided to hold a contest to create a name for the service, by then being referred to as “Student Info.” As an incentive, Apple offered a Macintosh computer for the winning entry. Ken Chung, a senior, came up with “Just the Facts,” the winning name, and was awarded a brand new Macintosh Classic computer.⁴⁴

Once Just the Facts proved to be such a popular application, there was a two-fold thrust to extend the services. One was to extend the services with new applications, and the other was to build a robust underlying technology. VP Lynn was interested in creating a whole set of services that could be packaged together with an easily accessible common interface. Services such as e-mail and the library catalog, for example, could be included on this collection

⁴² “Project Mandarin to improve administrative computing,” *Cornell Chronicle*, May 10, 1990.

⁴³ Mark Mara, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, July 1999.

⁴⁴ “Student Info System Renamed Just the Facts,” *CITnews*, summer 1992.

that was now referred to as Bear Access. This was a continuation of his desire to eliminate technical terms and complicated access to services. VP Lynn likes to relate the story of how when he first arrived at Cornell, access to the library catalog had to be obtained by using the command "DIAL LIBRARY" once a user was in the VM operating environment on the CornellC host. He recalls questioning Bob Cowles, then head of Systems Programming, why the word "dial" was needed. Cowles insisted this was part of the VM command structure, but within several weeks was proudly able to tell Lynn that a solution had been developed to eliminate "dial" and henceforth only "library" needed to be typed. While this development may seem trite in hindsight, it was but one example of the tone being set for simplification and improved access to technology as the workstation GUI was transforming the human-machine interaction.

Mara and his team were interested in creating a better underlying Bear Access technology that could be used to access all the new services. They proceeded to build a set of reusable tools, a tool kit, that could lace together applications on the client workstation with technology at the server end. The innovative application that followed Just the Facts was written in 1993 for the Cornell Campus Store by the Office Support Systems staff. Using "Campus Store," a student could submit his or her ID and the program would check the courses the student was officially registered to attend and provide a list of books for those courses, giving them a printable shopping list of books to buy. The Campus Store application also provided general information about the store. Other applications using this same technology also were developed for Public Affairs and Sponsored Programs.

VP Lynn, working with Tom Young in Workstation Technologies, developed a HyperCard-based presentation tool to aggregate the Bear Access services on a single screen tableau; a user could then point and double click on a function to bring it into play. So, for example, "Library" would appear as an icon on the Bear Access screen; double clicking on it would execute the link to the Cornell Library System on CornellC. From there on, the Library application would execute as designed. In this case, Bear Access was automating the previously typed command line "dial library," which had been simplified to "library," which now was even further simplified as a graphic clickable link. When the technologists and support staff saw this new approach, well before the web was introduced, they saw the potential for improving and simplifying access to services using the power of the Macintosh's graphical interface.

Efforts on Bear Access continued along these two different paths as VP Lynn and the management staff

and technologists in CIT and OIT debated the merits of the two application and data access approaches: continue with HyperCard or switch to the Mandarin technology. The argument often likened Mandarin technology to a powerstrip into which one could easily plug different applications, not easily done with HyperCard. It was finally resolved prior to the EZ-LINK announcement in late 1992. Mandarin technology was favored, as it provided a more general and long-term solution that was extendable to other platforms and applications. From that year onward, Bear Access became the flagship service for the campus and served as a vehicle to introduce all newcomers to the world of Cornell network computing.

Although there was sometimes confusion when the terms Bear Access and Mandarin were used interchangeably, Bear Access was the implementation of Mandarin technology at Cornell; other institutions that used Mandarin technology, such as Penn State, called their network navigator by other names. What is clear is that Bear Access was ahead of its time and well in advance of network browsers and the web, positioning Cornell as a leader in adapting client-server technologies to users' needs.



Project Mandarin

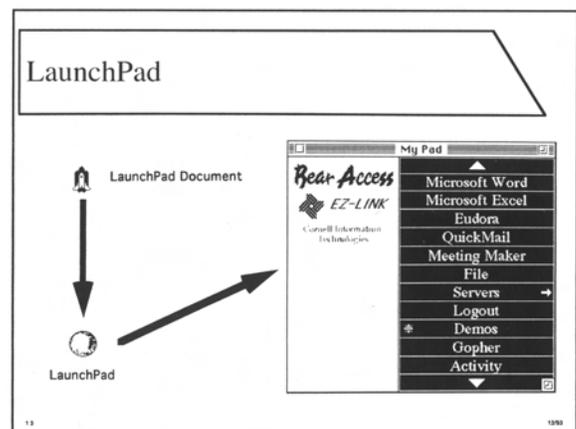


Figure 1. An early Bear Access screen

Mandarin Consortium

Project Mandarin then took a different direction from the one stated at its inception. Instead of developing an "administrator's workbench," it created a robust and flexible infrastructure for the development of client-server applications on the Macintosh platform. Most notably, CUSSP (Cornell University Stateless Server Protocol) was developed by Tom Dimock, Peter Bosanko, and Andy Wyatt to provide a link between the client and server platforms. This work led to the

development of the complete Mandarin tool kit that would be used to build different kinds of applications. Recognizing the potential for this general program, VP Lynn decided to see if other institutions would be interested enough in the technology to underwrite its future extensions and maintenance. This led to the formation of the Mandarin Consortium, which started with MIT, Stanford, and Cornell as its first three members; it reached a peak of 26 member institutions before disbanding later in the decade.^f Cornell and all those institutions were thus able to introduce new client-server applications on their campuses at a low cost, well ahead of the mainstream development of the technology. IBM decided to support Cornell and consortium efforts to port Mandarin technology to their platforms so those campus users of IBM systems would not be disadvantaged.

EZ-PRINT/Local

In 1992 a new network service was introduced, EZ-Print/Local, which enabled information generated on CornellC to be printed at LAN printers rather than on the mainframe-connected printers. This service was of particular interest to administrative units, because they could get reports printed faster and at less cost. The service operated over all the Cornell LANs and on different equipment and was implemented by a combined group from Computer Resources (CR) System Programming and Information Resources (IR) Technical and Applications Programming. The key staff included Mark Bodenstein, John Voigt, Mike Hojnowski, Sheila Patterson, Mark Sincock, Ted Pless from CR, and Doug Honig and David Wakoff from IR. The service proved quite popular. There were plans to create a companion service, EZ-Print/Public, as a pay-per-page printing service in CIT public labs and perhaps in dormitories. That came much later.

Terminal Emulators

The work that had gone on during the 1980s at Cornell, as well as at other commercial and noncommercial enterprises, had produced a rich set of terminal emulators for use with differing workstations and host computers. Table 2, taken from the November-December 1990 issue of *CITNews*, page 30, best captures the status of these emulators at the time.

Information Services

New information services from both commercial and nonprofit sources were evolving along with the technologies for delivering those services. New services included new electronic mail, Gopher-based CUINFO, and digital libraries. A pilot project for the

digital preservation of decaying library books, a joint study with Xerox, proved successful and was continued on a larger scale.

Electronic Mail and Directory Services

At the beginning of the 1990 decade, e-mail was still predominantly mainframe based using "Ricemail," now referred to as Mail90 on the IBM hosts and VMSMail on the DEC host. BITNET was predominantly used to send mail around the campus and the world that BITNET reached. Internet use based on the TCP/IP protocols was beginning to spread. Most Cornell computers had a domain name so that mail sent to an address in the form "vax5.cit.cornell.edu" reached the person with a computer account on Vax5. The use of e-mail increased rather dramatically with the widespread distribution of general computing accounts (GCAs) to the Cornell community.

As the number of workstations and local area networks increased, Quickmail became increasingly popular for those users with Macintosh systems. By late 1990, it was estimated that over 1,000 staff and faculty were using Quickmail, and CIT was moving quickly to make this a fully supported product in 1991. After stress testing in late 1990, GatorMail-Q software was installed to enable transfer of mail between Quickmail and other mail systems, including Ricemail on the IBM hosts and Unix-based mail systems. Improvements continued to be made, so that in 1991 it was possible to send and receive Quickmail between Cornell's medical college in New York City and the Ithaca campus when the medical college installed a product that linked Quickmail to the Internet. By 1992 QMRelay, the gateway system linking CITNet to the Internet, was accommodating 3,000 Quickmail users on campus and was uniquely operating at Cornell on two separate Mac SE30s, one for incoming and the other for outgoing mail. This was further extended by a link to TheoryNet, so that mail could flow between users connected to this campus network.

To improve connectivity between these disparate systems, and to achieve one of CIT's goals of seamless services, CIT undertook the creation of an electronic directory. This service would enable finding e-mail addresses while simplifying the presentation of the addresses, hiding the complicated routing structure needed to move information over networks. During 1992 and 1993, as part of the migration to client-server computing, CIT introduced a new e-mail service known as POP (Post Office Protocol) mail, a service with the potential to serve the very large number of users expected to be active over the campus, something that mainframe-based and LAN e-mail services could not accommodate. Eudora was the Mac client,

Hosts and Services for Terminal Emulators Supported by CIT										
Type of Connection	Type of Workstation	Application	HOST							
			IBM Mainframes			VMS Systems			UNIX Systems	
			connect	file transfer	graphics	connect	file transfer	graphics	connect	file transfer
Sytek (including modems)	Macintosh	Comet	Y	Y	N	Y	N	N	Y	N
		C19 (1)	Y	Y	N	N	N/A	N/A	Y	N
		VersaTerm	Y	Y	Y	Y	Y	Y	Y	Y
		MacKermit	Y	Y	N	Y	Y	N	Y	Y
		MacTerminal	Y	N	N	Y	N	N	Y	N
	IBM	C19	Y	Y	N	Y	N	N	Y	N
		MS-Kermit	Y	Y	Y	Y	Y	Y	Y	Y
Internet (TCP/IP)	Macintosh (2)	Comet	Y	Y	N	Y	Y (3)	N	Y	Y (3)
		Cornell TN/TN3270 (1)	Y	Y	N	N	N/A	N/A	Y	N
		NCSA Telnet	N (4)	N/A	N/A	Y	Y	Y	Y	Y
		Brown TN3270	Y	Y	Y (5)	N	N/A	N/A	N	N/A
	IBM	Cornell TN/TN3270 (6)	Y	Y	N	Y	Y (7)	N	Y	Y (7)
		NCSA Telnet (8)	N	N/A	N/A	Y	Y	Y	Y	Y
		Clarkson TN/TN3270 (9)	Y	Y	Y	Y	Y	Y	Y	Y

Notes:
 Except as otherwise noted, graphics terminal emulation is a Tektronics 4014. Except for the Cornell software's use of FT3270, file transfer is via the Kermit protocols over a Sytek/modem connection and via the FTP protocols over the campus internet.

(1) This software will be superseded by the production version of Comet.
 (2) The Macintosh Telnet/TN3270 packages (except for the old Cornell TN and TN3270 programs) all support both LocalTalk and Ethernet connections, and all can use Apple's MacTCP drivers. Only the Cornell-written packages support Omninet connections.
 (3) Via the TFTP file transfer protocol.
 (4) A commercial version of NCSA Telnet for the Mac, TCP Connect, does support TN3270 connections to IBM hosts.
 (5) IBM 3179g graphics terminal emulation.
 (6) The Cornell package supports connections via Omninet, AppleTalk, Novell IPX, and Ethernet boards with packet driver interfaces.
 (7) Cornell TN/TN3270 supports FTP only when used with a packet driver network interface.
 (8) NCSA Telnet for the PC supports most popular Ethernet interface cards. A test version is available that supports connection via AppleTalk.
 (9) The Clarkson emulator supports Ethernet connections, including packet driver interfaces.

Table 2. Hosts and Services for Terminal Emulators Supported by CIT in 1990

and NuPop the PC client, that communicated with POP servers acting as post offices to send, receive, and store mail.

One of the features provided by most of the mail systems, whose use was supported and assisted by CIT, was the formation of e-mail lists. While users could create their own mail list on their computer, centrally accessible lists had to be registered with CIT, with someone taking responsibility for keeping the list up-to-date. These lists got off to a slow start with respect to widespread use for distributing official information across the campus because the number of users was relatively small, but there was brisk use of central lists to circulate information that was important to selected groups of users.

CUINFO—Gopher Client

Information provided through CUINFO continued to move apace. In 1990, a popular service within CUINFO was ALERT (Alcohol and Drug Related Education, Research, and Training), which had been used over 18,000 times since being introduced three years earlier. This popularity reflected the growing concern for increased consumption and abuse of these substances and the university's efforts to temper demand. By the fall of 1991, one of CUINFO's most frequently referenced items was the Cornell staff directory, with up to 10,000 accesses per month. There was also a directory of BITNET sites. Additionally, it was now possible to forward display screens to one's e-mail address for later printing.

As part of the movement toward the client-server model for computing services, in 1993 CUINFO was

moved from the CIT mainframes and was only available via a Gopher server on the Cornell network. That also was part of the move toward using Gopher and WAIS (Wide Area Information System) technology to link different information sources in advance of the World Wide Web.

Digital Library, Historical Preservation Project

Another major thread of the first years of the 1990 decade was to develop an “electronic library” as well as the ability to deliver electronically stored documents across the campus network. To that end, in mid-1990 Cornell, Xerox Corp., and the Commission on Preservation and Access (established to foster and support collaboration among libraries and allied organizations to ensure the preservation of published and documentary records) undertook a collaborative pilot project that would involve scanning 1,000 Cornell Library volumes into a digital image storage system.⁴⁵ The objective of the project was to test an advanced technology for recording deteriorating books as digital images and producing, on demand, multiple high-quality copies. Worldwide, about one-third of the books on the shelves of research libraries were deteriorating because of the chemical processes used to make paper since the mid-19th century. The commission provided partial funding, Xerox provided advanced technology, and the Cornell Library and CIT provided staff and facilities for the project. A book stored as a series of scanned digitized images of each page could be distributed over the network or reproduced as needed; however, there was no way to search an index or the text. It was expected that the scanned document could be converted to an alphanumeric text-searchable format at some later date.

A special scanner was installed in Olin Library, where the books were carefully torn apart and each page scanned to accurately digitize the image of the page. The images were then transmitted over the network to CCC, where they were stored on a file server. To print all or part of a book, the relevant images were sent to the special Xerox printing system installed in CCC. These images were stored at the printer and then printed as required.

Another project was started to create a mass storage system capable of storing large numbers of books in the form of book pages. Because efforts by the Mass Storage group in CIT had not yet produced an operational system, an EPOCH Infinite Storage File Server was installed. This unit used a combination of magnetic and read/write optical technology to store up to 100 Gbytes of data in a fairly small footprint.

When the new Xerox printing system was initially installed it was not yet a released product, so a secure space had to be found for this “stealth printer,” as it

was called. A section of the lockable printer room on the first floor of CCC was cordoned off with a hospital-like white curtain that hung from the ceiling to the floor. Even if a person entered the room, the special printer was hidden from view, and the curtain was rolled back only in special circumstances. Later on this system was announced as the Xerox Docutech Publishing System.

Cornell was one of the first, if not the first, to install this powerful printing system. Because it printed at 1,250 or 2,500 dpi (dots/inch), in contrast to normal laser printing at 300 or 600 dpi, the print produced was so sharp it was difficult to tell it had been printed on a laser printer. When the Cornell Library produced a book using the Docutech printer on special acid-free paper, except for the newness of the cover, it was hard to tell the original from the reprint. Because the books chosen were well out of print, there was no danger of copyright infringement by making such reprints. However, it would be a while before images of the complete book could be transmitted around the campus because the electronic image files were so large in volume.

During this period one of the Cornell Library’s priorities was the opening of the new Carl A. Kroch Library for special collections, an underground facility located between Goldwin Smith and Stimson Halls and completed in the fall of 1992. However, librarians and technologists from CIT were inching closer to the electronic library where “light-emitting diodes will replace paper and ink as the principal medium for words.” Alain Seznec, university librarian, said that “the basic problem of the Library is how to maintain the traditional uses of the library and find resources for the new ones.” Seznec noted, for example, that the Cornell Library still carried 55,000 journals that had to be bought, checked in, cataloged, and stored.⁴⁶ Besides cost, space was becoming an increasing problem, and the library was anticipating increasing its off-site storage facilities. While electronic technology would not replace some of the rare books given a place of honor in the Kroch Library, digital electronics was viewed as a partial replacement for most day-to-day publications.

According to Lynne Personius, assistant director of Scholarly Information Systems at CIT and director of library technologies in the Cornell Library, Cornell was well behind its peers when it installed its online catalog in 1988. Now, however, in addition to the services being provided by Mann Library, Cornell was offering the LEXIS database for legal cases in the

⁴⁵ Office of Information Technologies, Cornell University, annual report 1989–90.

⁴⁶ “Ushering in the electronic era at the library,” *Cornell Chronicle*, November 19, 1992.

Law Library, the Dow Jones News Retrieval System in the library of the Johnson Graduate School of Management, and other notable electronic services.

The movement toward computing and network technology was gaining momentum and continued throughout the decade. The *CIT News* issue of summer 1992 was dedicated to the significant changes taking place in libraries, library technology, and library and information services and provides a good summary of the state of these issues at the time.

Mann Library Digital Library Projects

Jan Olsen, head of the Mann Library, was particularly active in promoting new services. In early 1990 Mann Library was the principal investigator in the CORE study, the Chemistry On-line Retrieval Experiment, which would provide 150 Cornell faculty and students with seven years' worth of 20 journals of the American Chemical Society at their workstations. It was estimated that the database would amount to some 600,000 pages and be stored on a variety of media. The project was expected to set an example for full-text journal delivery systems in other scientific disciplines.⁴⁷ By mid-1992 the project had reached the state where tests were under way to evaluate different approaches to delivering full text on computer screens and to evaluate how this presentation would supplement the printed journals.

At about this time Mann Library also opened an "electronic gateway" as a simple means for computer users to search numerous databases from their offices or homes. Accessible databases ranged from the Department of Agriculture crop estimates to the giant Research Library Information Network, which held 32 million entries from library catalogs of the nations' research universities. Such gateways were viewed as an emerging technology to provide users a central point of access for searching databases inside and outside the university and to provide an intellectual road map to help users navigate among resources.⁴⁸

Cornell Electronic Publishing Working Group (CEPWG)

In 1992 VP Lynn was the prime motivator to form the Cornell Electronic Publishing Working Group (CEPWG), a group that brought together all the printing and publishing organizations on campus, including the libraries, into a cooperative venture to promote the increased use of electronic publishing. By the fall of 1992, electronic publishing was available from CIT (in CCC), the Media Services print shop in Martha Van Rensselaer Hall, the Cornell Campus Store, and the University Print Shop. CIT, Media Services, and the Print Shop had Xerox Docutech printers, while the Print Shop also had a competi-

tive system from Kodak. The summer 1992 issue of *CITNews* has an excellent summary of electronic publishing on campus at this time.



Figure 2. VP Lynn meeting with the Cornell Electronic Publishing Working Group

Chaos Corner

Responding to the somewhat chaotic conditions of rapidly changing technology and new developments in the early 1990s, in June 1991 Bob Cowles created an electronic newsletter he called *Chaos Corner*, his attempt to raise consciousness of the Internet as a valuable resource. In his words this newsletter consisted of "random and (in my opinion) interesting things I have come across recently that may be interesting to some of you. Since I spend a lot of my time (when not in meetings) reading about and banging on new things, I think it is important to pass along what I may have discovered." He signed off with the moniker "Dr. Chaos (I have a master's degree) rdc@cornell.cit.cornell.edu." The first issue⁴⁹ covered file transfer protocol (FTP), network news, mailing lists, a number of technology updates on Windows3 applications, Unix security, the use of Postscript, new CD-ROM drives available, and other topics. By 1993, 600 subscribers around the world received *Chaos Corner*. Each year

⁴⁷ "Jan Olsen: An innovative leader in information technology," *Cornell Chronicle*, January 1991.

⁴⁸ "Mann Library opens electronic gateway," *Cornell Chronicle*, April 23, 1992.

⁴⁹ "A New Feature—Chaos Corner, A Collection of Technology Tidbits," *Inside CIT*, June 28, 1991.

the columns were printed and packaged together as an annual *Chaos Corner* publication with an index for easily locating items of interest. As the World Wide Web came more into prominence, Cowles simply made URLs known so that interested parties could get more information by going to that site.

Instructional Computing

Responding to some of the concerns about instructional use of technology, CIT attempted to make improvements, one of which was establishing an instructional resource center to assist faculty in this endeavor. A program was started to introduce incoming students to information technology and the resources at Cornell. CIT partnered with Apple Corporation to develop and teach technology courses. Cornell faculty continued to garner awards for their courseware development, even though the issue of ownership of such software was still unresolved.

Instructional Resource Center and Course Statistics

The Instructional Resource Center (IRC) was established in the summer of 1990 in Room 124 in the Computing and Communications Center (CCC) to assist faculty in learning more about instructional technologies that they could bring into the classroom.

Carrie Regenstein was the first director of the center, and Charlotte Kiefer was the first instruction support coordinator. Information collected by Kiefer⁵⁰ (see Table 3) gives a picture of how much support CIT was giving instructional computing in the early 1990s and the mix of technologies that were being used by the courses.

Although the data are not complete, they do show growth in the number of courses using information technology over this three-year period, even though the number of students declined. In addition, the information in the notes supports the conclusion that over this brief period there was a distinct shift to using personal computers and workstations in support of instructional computing and away from the use of time-sharing on mainframes. This information is consistent with efforts by CIT to move to more contemporary workstation-based technology.

In the same period, statistics from the Department of Computer Science show that the department taught 74 courses in 1991–92, with an enrollment of 3,431 students; 3,506 students were enrolled the following year, 1992–93. The number of courses taught during this time was considerably higher than the range of 45 to 58 courses taught over the 1980 decade; enrollments, ranging from 2,800 to 4,200 students, did not increase by the same degree.

Table 3. Number of Courses and Enrollment Supported by CIT, 1990 to 1992

<i>Semester</i>	<i>Number of Courses</i>	<i>Number of Students Enrolled</i>	<i>Notes</i>
1990 fall	54	4,284	
1991 spring	60	4,000	See Note 1
1991 fall	80	6,200	
1992 spring	80+	5,000	See Note 2
1992 fall	85+	5,200+	See Note 3

Note 1. Twenty courses used CornellA, 8 used Vax5, 4 used Crux1, 1 used AIX on a Theory Center machine, 11 used NeXT machines, 8 used Macs, 11 used MacIIci's in Upson, and 3 used IBM PCs.

Note 2. Eighty percent of the students used workstations, and 20 percent used time-sharing systems. One mathematics course, one electrical engineering course, and seven computer science courses used the SUN lab in Upson. The Mac was the most popular platform, used by more than 60 percent of students.

Note 3. Sixty percent of students used the Mac, 14 percent used CornellA, and 10 percent used Bear Access; Crux and IBM PS/2s were used by the rest.

⁵⁰ *Inside CIT*, December 1991.

The “Getting Started” Program

To cope with the increasing number of students, mostly freshmen but others as well, coming to campus in the fall and wanting to use their own machines or instructional labs, CIT embarked on the “Getting Started” program. This got under way in 1990 with the publication of *Introduction to CIT*, which listed services, locations, etc., and a series of “Getting Started” documents that provided succinct but sufficient information for a user to get started using Cornell’s information technology services.

In the fall of 1990, CIT offered several “Getting Started” courses—one- to two-hour lectures or workshops on different applications of likely interest to new users. That year over 1,000 undergraduate students participated in this program, with 600 students doing hands-on training and another 400 attending the lectures.

Learning Technologies Program

With the cooperation and support of Apple Computer, in 1991 Instructional Support Services developed the new Learning Technologies Program (LTP)—a series of workshops designed to facilitate the adoption and development of learning technologies in the higher education curriculum.⁵¹ Hands-on sessions for participants were very much part of the program. The program’s underlying premise was that technology must be presented in the context of the faculty’s instructional goals. The core curriculum was based on presenting non-discipline-specific tools, such as presentation software, Hypercard, and computer projection, combined with peer-to-peer workshops known as the Exemplar Teaching Series, in which faculty who had successfully integrated those tools into the curriculum would share their experiences and insights. The Learning Technology Program used the resources of the Instructional Resource Center in CCC, where consultants were available to help faculty review instructional products (hardware and software) in an unhurried environment.

Apple and CIT continued to cooperate on other training programs of interest to both Cornell and the higher education community. CIT became one of four sites Apple designated as an “Apple Developer University,” whereby CIT would train application developers at other higher education sites as well as conduct “train the trainer” programs to increase expertise outside of central computing organizations. Both endeavors became segments of the Learning Technology Program when an agreement with Apple was concluded in 1990–91. The first LTP session was a pilot workshop held in the fall of 1991 with 20 attendees, roughly half from Cornell and half from

other institutions, including Ithaca College, Syracuse University, and Hamilton College. The program was continued for a number of years, with sessions hosted at Cornell and other institutions around the country.

Classroom Software (Courseware) Awards

Cornell continued to win awards for the development of courseware software that could be used in classrooms. In 1990 EDUCOM recognized four Cornell faculty members as winners of the EDUCOM/NCRIPTAL Higher Education Software Awards Competition. NCRIPTAL, the National Center for Research to Improve Postsecondary Teaching and Learning, was cooperating with EDUCOM in their quest for better tools for classroom teaching.⁵² The awards went to Beverly West and Doug Alfors for Analyzer* 3.0 and to Peter L. Jackson and John A. Muckstadt for the Manufacturing System Development Game.

In 1991 EDUCOM recognized eight faculty members by acknowledging the excellence of their software in the Joe Wyatt Challenge, issued by Joe B. Wyatt, the chancellor of Vanderbilt University. Wyatt challenged faculty to develop innovative uses for computers in the classroom. The presentations, also known under the title “101 Success Stories,” were made at the 1991 EDUCOM Conference, where Cornell was the most strongly represented institution.⁵³ The award winners were

- Richard W. Conway for XCELL+ Factory Modeling System
- J. Robert Cooke for Instructional Finite Element Analysis
- D. Peter Loucks for Interactive River System Simulation
- Anita Racine for Integration of Computer-Aided Design into an Apparel Design Curriculum
- Lee W. Schruben for Simulation Graphical Modeling and Analysis (SIGMA)
- Donald F. Sola for InterLex Lab
- Michael L. Thonney for Cornell University Beef Cow Herd Simulation Program (CUBEEF)
- Paul Velleman for Data Desk

These awards attest to this continuing activity of developing instructional courseware that got off to a significant start in the mid 1980s with Project Ezra

⁵¹ “The Learning Technologies Program,” *CITnews*, summer 1992.

⁵² “Cornell Shines Again in EDUCOM/NCRIPTAL Awards,” *CITnews*, February-March, 1991.

⁵³ “Instructional Insights: Cornell Faculty Recognized in EDUCOM’s Joe Wyatt Challenge,” *CITnews*, fall 1991.

(supported by IBM) and the Mac*Ed Center (supported by Apple) and grants to faculty by IBM, Apple, and other vendors.

Intellectual Property Issues

The issue of ownership rights to software and other course materials being developed in support of classroom instruction continued to simmer throughout the 1980s and into 1990. After eight years in the making, in June 1990⁵⁴ the Cornell trustees approved a policy that would give Cornell the copyright claim to new computer software developed by faculty or other employees whose work made “substantial” use of university resources. This policy had been first proposed in 1988 and had been reworked several times by faculty representatives and Cornell legal staff. The Faculty Council of Representatives rejected the policy despite its being approved by their Research Policies Committee prior to the trustees vote. The key words “substantial use,” which earlier had been “extraordinary use,” were not considered clear enough to define the policy unambiguously. Even though individual faculty members who had developed the software were permitted to use the work for noncommercial purposes, faculty argued that this struck at the heart of ownership of intellectual property rights and created disincentives for the creation of publishable software. Given the lack of public information about this policy for the remainder of the decade, it appears that Cornell enforced this policy with benign neglect.

Research Computing

Supercomputers at the Theory Center increasingly became the focus of research computing at Cornell. CIT continued to provide operational support services to the Theory Center through computer upgrades, in planning for the new computer room in the center’s new building, and in planning and executing the move of the supercomputers from CCC into their new room. The Theory Center’s offerings continued to change with new computer installations. Efforts to produce a full-function and operational Mass Storage System continued.

CIT Joins the “Smart Node” Program at the Theory Center

In 1990 CIT officially became a “Smart Node” by joining this outreach program of the Theory Center. The main objective of the program was to strengthen ties between Cornell researchers and the center; but more importantly, it provided several benefits, a key one being a “free” test account to evaluate the supercomputing services supported by specially trained CIT

consulting staff. At its peak, over 60 percent of the clients at the Theory Center came from Cornell. The center took strong steps to attract more users nationally and outside of Cornell so that it would truly be considered a national supercomputing resource.

Supercomputing at the Theory Center

In the early 1990s, CIT and the Theory Center continued their cooperation on several fronts. CIT provided the computer support services for the center’s computers, which were housed in CCC as the decade began; as mentioned earlier, CIT had taken over responsibility for the former TheoryNet network as well as for the national and international network connections.

Mal Kalos continued as director of the Theory Center, and in May 1990 two new directors were announced: Peter Siegel was named director of the Cornell National Supercomputer Facility (CNSF), and Jay Blaire was named executive director of the center. Siegel, who had most recently been deputy director of the CNSF and who had held a number of management positions in CIT and the Theory Center, replaced Lawrence A. Lee, who had left Cornell to become executive director of the newly established North Carolina Supercomputing Center.

Despite a high level of activity, productive research outcomes, a nationally recognized outreach program, and other initiatives, in 1991 the center suffered a setback when NSF reduced its level of funding, forcing the layoff of 15 staff members. The announcement noted that in the previous year more than 2,000 scientists and engineers from 150 research institutions working on over 500 projects had used the supercomputers at the center.⁵⁵

During this time the Theory Center continued the tradition of assisting Cornell researchers in winning awards for their work. In 1990 Steven Pope, professor of mechanical and aerospace engineering, won first prize in the first IBM Supercomputing Competition. Several other Cornell researchers were awarded honorable mentions in the competition that year. At the same time, the Theory Center and NSF sponsored the SuperQuest competition, which invited teams of three to four students and their teachers/coaches from high schools throughout the country to submit proposals requiring supercomputing for their solutions.

⁵⁴ “Trustees OK policy governing computer-software copyrights,” *Cornell Chronicle*, August 9, 1990.

⁵⁵ “Theory Center lays off 15 staff members,” *Cornell Chronicle*, March 26, 1991.

IBM 3090-600s Upgraded, Array Processors Removed

The supercomputers at the Theory Center continued to evolve and be upgraded. In 1990 the two 3090-600 computers in the CCC computer room were upgraded to model Js, providing an increase in computing cycles. The project to link these two machines to create a 12-way parallel processing computer system continued.

The era of the Floating Point Systems Array Processors came to an end this same year when they were all removed. They had served their function admirably since the first array processor was installed in 1978, but by this time their functionality was built into the 3090 processors, making the units redundant.

Theory Center Building Opens; Computers Moved from CCC

On June 4, 1991, the Theory Center celebrated the opening of its new \$35.2 million building, the Engineering and Theory Center Building (later renamed Rhodes Hall) and held the symposium “Supercomputing: The Next 50 Years,” which was attended by high-ranking scientists and dignitaries from campus, corporations, the NSF, and other research organizations.⁵⁶ During this same year Congress passed the High-Performance Computing Act to provide funding to advance the state of supercomputing and high-speed networking in support of science and research initiatives.

After the completion of the Engineering and Theory Center Building and the Theory Center’s computer room, there was a delay of about a year before the computers could be moved from CCC, due to lack of funds; NSF had cut the Theory Center’s budget and the university was not willing to provide the funds. A year later in 1992, the equipment was moved and the center took over responsibility for their computers and operations and systems support previously performed by CIT. After the move, the Theory Center and CIT continued to cooperate by providing mutual backup sites to each other and for offsite storage of critical materials.

“Garden” of Supercomputers at the Theory Center

By late 1992 the center was promoting their “garden” of supercomputers. This collection consisted of an IBM ES 9000, a cluster of 32 IBM RISC 6000 machines (prototype of the HPSSL—Highly Parallel Supercomputing Systems Laboratory) to be delivered in 1993, and a KSR1 scalable parallel computer from Kendall Square Research. When first installed in 1991, the KSR1 was a 32-processor unit, which was upgraded to a 64-processor unit consisting of two rings of 32 processors.

These were the first steps toward achieving the vision promoted by Ken Wilson, founding director, that the future of supercomputing was in clusters of smaller machines working in parallel instead of the big, superfast sophisticated processors such as the Cray systems.

Mass Storage Systems

The initiative to produce a mass storage system (MSS) for use by the Theory Center, the Cornell Library, and other units at Cornell continued on two fronts—one, to define the formal reference model, and two, to produce a working system. Andy Hanushevsky from CIT became an active member of the IEEE Technical Committee on Mass Storage Systems and Technology, writing and presenting position papers. Cornell representatives from both CIT and the Theory Center attended the symposiums organized by this group for the exchange of information. As the other national supercomputer centers had the same interests as Cornell’s, along with the need to share data between their sites, various position papers were prepared and submitted to the NSF as a way of keeping them informed and to build support. In 1991 Cornell hosted a meeting of the IEEE Storage System Standards Working Group to further discuss and refine the reference model. In 1992 Hanushevsky completed a grant proposal entitled “Access Control for Distributed Storage Systems,” which was funded by NASA-Lewis. It was expected that this design would be transferred to vendors in exchange for free use of their mass storage system.

Getting a mass storage system into production proved to be a hard task, beset by all kinds of cross-currents and differences between hardware platforms and operating systems of the servers and host/client platforms. Further, there were differences in how data were to be stored and accessed on the different physical tapes and disks and in how to deal with a distributed environment across distances. It was a complicated mess.

To add to the complications, various operational systems running in limited environments were touted as candidates for being “hardened” into commercial products, and companies were being organized to take on this work. So, for example, HADES (Heidelberg Automatic Datamanagement and Editing System), a system developed in Germany, was of interest because it ran on IBM hardware and could run under the VM or MVS operating systems. Those features had a certain appeal, as did high-capacity storage offerings using

⁵⁶ “Supercomputer program to dedicate Engineering/Theory Center Building,” *Cornell Chronicle*, May 23, 1991.

IBM hardware and software. A test version of HADES was installed at the CNSF for evaluation purposes.

Of increasing interest were offerings that ran on Unix platforms, because converting to the use of the Unix operating system was the future direction of the CNSF and all the other supercomputer centers. Another early candidate was the DataTree, the distributed file management system from General Atomics. This was an enhanced version of the widely used Common File System developed in 1979 at the Los Alamos National Laboratory. While it ran on IBM hardware, it had interfaces to most of the other systems and technologies.

Cornell's choice was finally UniTree, a Unix-based hierarchical file and storage management system for networked multivendor computing environments. It had been developed at Lawrence Livermore National Laboratory and had been in operation there since 1988. It was a complete implementation of the IEEE Mass Storage Reference Model (an emerging standard). The mid-1990 release was to provide basic client /server functionality via NFS or FTP access methods. A joint study agreement was reached with DISCOS, Distributed Computing Solutions of San Diego (the same division of General Atomics that also offered DataTree). Under the leadership of Hanushevsky, a system with Cornell modifications was put into production in 1991. By the end of the year, approximately 150 users had stored more than 62 Gbytes of data and use was growing at the rate 5 Gbytes per month.

A competing approach to creating a data archive and retrieval system was the installation of IFS (Institutional File System), which was being developed at the University of Michigan with the support of IBM and others. IFS was built on the Andrew File System, work done at Carnegie-Mellon to create a file system that could support tens of thousands of workstations. The plan was to use IFS as the front end interface to other systems, such as UniTree, which handled the work of actually storing and managing the data on a server, while IFS would be the client interface that called for the service. IFS ran on IBM computers under the then three main IBM operating systems of VM, MVS, and AIX. The system was installed at the CNSF at about the same time as UniTree.

Business Systems

Advances continued to be made in developing new business systems and updating existing ones and their relevant support services. A new data and systems policy was announced. The policy emphasized that data was a university resource to be shared and that the strategic direction was to newer technologies and distributed systems. A users group was formed to share

applications developments from institutions using Adabas as the database management system.

Administrative Data and Systems Policy, Strategic Directions

In early 1990 the Information Resources (IR) division of CIT responsible for supporting Cornell's administrative use of IT was highlighted in *CIT News*. David W. Koehler, director, emphasized that the new administrative data and systems policy placed emphasis on data rather than systems as the more important university resource.⁵⁷ The direction was clearly toward distributed processing systems as opposed to the older mainframe-based central systems; all new development was taking this direction.

The following strategic directions were stated for Information Resources:

- Provide an electronic authorization system to reduce the amount of paperwork and time required to process a document.
- Provide a financial model that will enable departments to use the new system features.
- Provide an adequate security system to allow storage, retrieval, and access to appropriate individuals.
- Educate the campus in the use of distributed computing capabilities, rights, and responsibilities.
- Work with university administrators to define appropriate data archives for providing decision support data at all levels of the university.
- Continue to encourage the development of professional staff as "information specialists" and to provide training in the areas of analysis, communication, and consulting.
- Develop Project Mandarin, a desktop environment for the administrator to facilitate interactions with administrative data and information in central and distributed databases and to provide convenient and intuitive access to public information, communication systems, and scholarly information.

⁵⁷ "CIT's Information Resources: Supporting Administrative Computing—from the Office...to the Machine Shop," *CITnews*, March-April 1990.

- Provide leadership in the effort to find technological solutions for preserving scholarly information, such as university library data and other archival databases.

This was a large agenda but continued to provide guidance for the remainder of the decade.

New Systems: CHRISP, Purchasing (APPS), Inventory Management (IMS)

All projects of the early 1990s were guided by the above strategic directions. For example, CHRISP, the Cornell Human Relations Information System/Payroll project, introduced the concept of electronic forms to the campus and was promoted as moving to a “paperless” campus. CHRISP eliminated a number of paper forms; many human resource and payroll transactions could now be “securely” initiated by staff in their own offices rather than handled by paper forms sent through the mail, with data entry done at the central office. CHRISP was prototyped in the spring of 1991 and operational later that year.

Two new systems were under active development in 1990: APPS (the Automated Procurement and Payment System) and IMS (the Inventory Management System). The team working on these projects included Jim Storelli, the project leader, and Ann Santiago and Beth Bement. IMS went into use in 1990, but only by General Stores and Maintenance and Service Operations, with plans to eventually have others use the system to manage their inventories. APPS would provide online requisitioning from a workstation, electronic approvals, procurement support, online receiving, and invoice payment. APPS was made available to the campus in 1991. A direct link was planned between APPS and IMS so that goods ordered and received in APPS would be directly entered into IMS.

For student records, the Admissions System was enhanced to record more data on prospective applicants and track all communication with them. Also, the Student Employment System automated an employee appointment form. A new Loan Management System was purchased from Information Associates in Rochester, N.Y., to improve the tracking, billing, and servicing of each Cornell student’s loan portfolio. Along with installing this new system, improvements were planned for improving the student billing statement.

The Public Affairs System was enhanced to add more information on gifts donated to Cornell. A new Budget System was designed to automate the budgeting and forecasting of all university funds. It was expected that this new system would alleviate time bottlenecks while providing tools to improve the cre-

ation of budgets and allow better understanding of the process. The implementation took place in phases starting in 1990, with complete integration achieved by 1992.

The work on these administrative systems was being done in the Adabas environment. (Cornell was recognized as the first, in 1981, and now largest site to use Adabas under VM.) One measure of the growth of computing in support of business systems at Cornell is that the number of Adabas commands grew from 10,000 commands per month in 1984 to over 120,000 commands per month in 1990. By 1990 there were 300 simultaneous users accessing Adabas on CornellC at peak times.

Office Support Systems Projects

Office Support Systems, under the direction of Mark Mara, took an increasingly aggressive approach to developing tools and techniques for administrative offices to exploit the power of workstation resources, working in concert with the host-based business systems. In a way, this group complemented another initiative in this same section—to build tools to expedite and simplify the development of local departmental information systems. Frederick (Rick) E. Jones was the first manager of this new initiative. The first staff members were Tom Weyer, Tom Lane, and Keith Kubarek. Andy Wyatt and Tony Damiani soon joined them.

Some of the early projects (in the Career Center and the College of Engineering) were the first implementations of database systems on workstations that integrated information with the central systems. As Jones noted in a personal communication: “The customer base eventually spanned the gamut from academic departments like the physics department in A&S, to large administrative divisions such as Facilities and Business Operations (F&BO), to, in one case, other higher-ed institutions (Hamilton College).” The Hamilton College project is worth mentioning, for it was the last remaining use of SPIRES on the CornellC computer. A professor at Hamilton College, Jones noted, had developed “a massive, one million plus, and extremely complex record of all aspects of political activity in the Soviet Union.” The challenge was to convert that to another technology on a desktop machine. That was accomplished, although as Jones noted: “The translation between the two types of storage architectures proved to be far more complicated than originally assumed, but it was eventually completed successfully.” After that, CIT was able to remove SPIRES from CornellC.

High-Speed Laser Printing

The Xerox laser printers supporting the production of reports and other output from all the central business systems continued to change and take advantage of new technologies and lower prices. In 1990 another 4090 printer and a 4050 printer (a slower-speed 4090) replaced the 9700 printer, bringing the complement to three printers (two 4090s and one 4050) and offering 50 percent more print capacity. The highlight-color feature of the 4090s enabled “Cornell” to be printed in red and other type in black. This feature enhanced the “special form” capability of printed output by adding color to mailings, for example, to alumni or prospective students.

CADE Data Entry System Upgraded

In 1991 a new Unisys 6000 data entry system was installed for creating input to business systems. This new system replaced the 11-year-old CADE system, offering the same capabilities as well as maintenance cost savings sufficient to recapture the additional purchase costs within two years. It was expected that this system would last until data entry was phased out in 1994, with more and more data entry being done online by staff working directly with the newer systems.

Exemplar Users Group

In keeping with the long-standing Cornell tradition of starting user groups and other cooperative ventures, in 1992 the Information Resources division of CIT and an equivalent division from Penn State announced the formation of the Exemplar Consortium.⁵⁸ The two principal organizers were Dave Koehler from Cornell and Ken Blythe from Penn State. The purpose of Exemplar was to provide educational institutions nationwide with the best of administrative computing applications. The consortium, which charged a membership fee, selected administrative software that represented the best practices among peer institutions and made that software available to all consortium members.

To create an incentive for members to make their software available to the consortium, a member whose software was selected was paid a fee, an amount that depended on the number of members who wanted to use the software. This was a novel twist on the perennial problem of sharing business software between institutions, which had been attempted since the early days when data processing programming consisted of wiring boards for punch card processing machines. Cornell was the first contributor to Exemplar,^g submitting its Inventory Management System (IMS) for possible use by others. Penn State installed IMS, claiming

that it helped them avoid \$1 million of development costs and generated \$300,000 per year in savings.

Campus Planning Initiatives

Based on planning that had started earlier in the decade, in 1992 Cornell committed itself to three interlocking initiatives:

- the Strategic Planning Initiative, to provide guidance to university decision making;
- the College Priority Planning, to have each college and major administrative unit determine internal priorities; and
- the Quality Improvement Process (QIP), to improve the quality of services across the campus.

Also during this period the 2001 Committee issued its report.

Strategic Planning Initiative and College Priority Planning

President Rhodes and Provost Nesheim were responsible for the Strategic Planning Initiative and the College Priority Planning initiative.

VP John Weisenfeld carried out the Strategic Planning Initiative with support from the staff in Institutional Planning and Research (IPR) and with advice and recommendations sought from all other constituencies. It was to be a comprehensive examination of the internal and external factors most likely to influence Cornell’s future in the coming decade.

The Strategic Planning Initiative was to provide the framework for the second initiative, College Priority Planning. Each college and administrative unit was invited to examine its own programs, priorities, and resources and build on efforts already under way across the campus.

Quality Improvement Process—QIP

In 1992 senior VP Morley and other academic and administrative staff spent the early part of the year defining QIP and carrying out some pilot programs. They were assisted in this effort by Qualtec Quality Services, a consulting subsidiary of Florida Light and Power, which had had outstanding success with implementing TQM, the Total Quality Management program. TQM was in vogue in the early 1990s as U.S.-based corporations strove to match the Japanese in making goods of superior quality. The success of TQM in manufacturing was extended to services and all other company activities, including dealings with customers.

⁵⁸ “CU, Penn State form computing consortium,” *Cornell Chronicle*, October 15, 1992.

QIP was defined in a publication as “Cornell’s version of Total Quality Management (which is being implemented in industry and in a host of universities). QIP is a process that engages teams of Cornell people in solving problems cooperatively and improving the quality of services, education, and scholarship the university provides. It is a process rather than a program because it’s designed to become part of the way all of us manage our work—administration, support services, teaching, or research—now and in the future.”⁵⁹ The QIP mission was “to advance learning, scholarship, and outreach at Cornell by improving the quality of our individual and collective efforts.”

After some refinements, the basic principles of QIP were defined as follows: to respect people; to understand and satisfy those whom we serve, our “customers”; to value scholarship and learning; to exercise stewardship; to act on facts; and to improve continuously. So that QIP would be pervasive throughout Cornell, teams were formed at the different levels: a university-wide Quality Council to lead teams in colleges or departments as appropriate, functional teams, task teams, and various cross-functional teams in the operating units. A considerable amount of training was done to support the outlook and methodologies of TQM in the Cornell context. QIP was to create a whole new culture for doing business at Cornell, with a consistent methodology understood and applied to all issues. Early estimates were that over 1,300 staff members, roughly 15 percent of the university staff, would be trained at some level of expertise in QIP, and this would be reflected in the continuity of the process and the future improvements.⁶⁰

CIT started three QIP teams in the second wave of setting up 22 teams across the campus. These three teams were under the direction of the CIT Lead Team that consisted of CIT directors or a designate and Cecilia Cowles, assistant director for service integration. One team considered the use of electronic systems to improve the handling of the consulting load at the Service Help Desk. One team was from CIT Sales and Support and tried to develop success measures on their conduct of business activity. The third team was set up to examine the complaint that “it takes too long to produce production procedures”—that is, there was an extended delay in documenting a completed business application prepared by staff in Information Resources to be run in production by staff in Computer Resources.⁶¹ Following standard QIP procedures for examining this theme, data were collected on the time required for this action. The data showed that the time was acceptable, so no further action was necessary. There are no records on what happened with the other two studies. In simple terms, QIP did not make a material impact on CIT.

The three planning initiatives created an optimistic environment on campus, in contrast to the dismal fiscal outlook for the prevailing budget problems. Better than that, the initiatives were creating a new way of going about everyday business.

The 2001 Committee Report

In late 1992, two years after its formation, the 2001 Committee issued its report.⁶² It believed that CIT’s “Vision for the Nineties: At Any Time, from Any Place—Collaboration through Technology” pointed in the right direction and needed to be supported if the university was to continue to fulfill its mission. The committee made these specific recommendations:

- There should be a transition in the support structure to a distributed model that places responsibility for and control of the deployment and use of information technologies as close as possible to those who benefit the most. This recommendation is a fundamental one and supports all of the succeeding recommendations.
- The promise of information technology is to increase effectiveness and service by exploiting the power of information networks in breaking down barriers of form, time, and space. Such gains multiply exponentially when everyone has access to these networks, and time- and cost-saving applications can be developed under the assumption of such universal access. Cornell must work toward providing such universal access as soon as practicable so every student, faculty, and staff member will have high-speed connection to a broadly distributed campus information network. Over the decade this network must come to integrate voice, video, and data into a unified broadband network.
- The university must facilitate the use of information technologies to the point where any faculty member can make use of such technologies in teaching, and every student will experience the benefits in most of their courses. To this end, the committee

⁵⁹ “QIP: Put your ideas to work. At work,” Office of Communications Strategies and Office of Publication Services, March 1993.

⁶⁰ Cornell University QIP Design Team, “QIP Implementation Plan for Teams,” December 17, 1992.

⁶¹ “QIP,” *Inside CIT*, vol. 4, no. 2.

⁶² 2001 Committee, “Cornell in 10 Years: The Influence of Technology,” November 9, 1992.

recommends the establishment of programatically oriented support teams in colleges and department to assist faculty directly and of a central design and evaluation group to provide coordinated support to these teams. In addition, the campus infrastructure needs to be upgraded—in particular, classrooms, residence hall rooms, and other facilities—to ensure network connectivity and an adequate technological environment.

- By the year 2001, the electronic library must become an important adjunct to the paper library, consuming a major portion of library resources. To further this objective, the Cornell Library should engage the community in the development of a proactive strategic plan and take leadership, in partnership with others both at Cornell and elsewhere, to further the progress of the electronic library; to engage in the development of model projects; and to exploit wherever possible the use of electronic technologies over traditional technologies whenever this can be shown to be both cost effective and to meet or exceed the needs of scholars.

- Institutions like Cornell depend in part on how well we disseminate information. By the year 2001, much of this information will be created, stored, and published electronically. Working with others, Cornell should exercise leadership in facilitating the changes that will occur, building on both the inter- and intra-campus partnerships (such as an expanded version of the Cornell Electronic Publishing Working Group). Cornell should embrace the use of electronic technologies wherever these are cost effective and lead to improved service. This will require careful examination of the organizational relationships among units that traditionally have had overlapping roles.

- A key focus for the university is to provide easy, online student, faculty, and staff access to all information and data required to support student and other administrative services. Students and others should not be hampered by the continuation of manual, burdensome, or bureaucratic processes. Such easy access from across the network should be directed at reducing or eliminating paperwork, duplication of data entry, manual handling of routine processes, and

waiting in lines. Systems must be built that integrate data from points-of-entry all the way to ultimate points-of-use, driven by the needs of individual users.

Although it is clear that CIT's actions and directions for the first several years of the decade were well under way toward achieving the recommendations of the 2001 Committee, and were in fact confirmed by the report, nonetheless, these recommendations served as a useful guide for the rest of the decade.

Migrating to Client-Server Computing, EZ-LINK

In late 1991, VP Lynn formulated "A Vision for the Nineties: At Any Time, from Any Place—Collaboration through Technology," which was to be the guiding principle for CIT. To implement this vision, he defined the building of EZ-LINK, "the Cornell Electronic Information Network, helping people get to what they need, where, when, and how they need it." The catch phrase was "annihilating both time and space," taken from a statement made by Ezra Cornell, the university's founder, when in 1884 he commented to his wife on the wonders of the recently invented telegraph: "Is not space annihilated?" He meant that distance no longer hindered communication between people. Now of course, time was also annihilated, because people would be able to get what they needed, where, when, and how they needed it.⁶³

The goals of EZ-LINK, building the information network, were to change the campus by

- revolutionizing service to students;
- enriching the learning process;
- extending the scholar's reach;
- speeding the transfer of information and technology; and
- streamlining the management of our resources.

To realize this vision, CIT committed itself to lead the community to make productive use of information technologies, to partner with others, and to render exceptional service—all in order to put information to work for the whole community. The task that lay ahead was to build EZ-LINK.

Building EZ-LINK, whose basic theme was to help people connect with access tools, information sources, support, and service, was based on three main areas of change:⁶⁴

⁶³ "A Vision for the Nineties: At Any Time, from Any Place—Collaboration through Technology," *CITnews*, fall 1991.

⁶⁴ "EZ-LINK—Making Technology Work for You," *Cornell Chronicle*, August 26, 1993.

- transition to the desktop;
- access to the network; and
- making it easier with Bear Access.

These components had been under development during the early years of the decade; the EZ-LINK theme put them all together under an umbrella structure, which not only packaged the then-current services but also provided a platform for the development of new and future services. Some of these earlier activities were based on two key threads of emphasis:

- to simplify computing access, that is, to eliminate much of the “techie” terms needed to get computers to do the intended work, and
- to develop the “electronic library” and the ability to deliver electronically stored documents across the campus network.

Transition to the Desktop

Transition to the desktop reflected the emerging client-server technology. Stated simply, mainframe-based services would be replaced by specialized servers capable of working in harmony with the desktop workstations to make tasks faster and cheaper and more intuitive. The arcane command structure of the mainframes was giving way to the graphical user interface (GUI) of the personal workstation. Large and powerful mainframes would become servers for large collections of data, such as the library and administrative systems, while services such as e-mail would be moved to the new environment.

Access to the Network, NetID Established

Access to the network via client-server architecture was based on a person’s identifying himself or herself to the network to obtain services. This was done by means of a network ID.⁶⁵ After much deliberation, the scheme chosen was to use the person’s initials in the official student or personnel databases maintained by the university, followed by a sequence number for that particular set of initials. So, for example, jwr1@cornell.edu became the network ID for John W. Rudan, where the 1 represented the first person to use this set of initials. By this time also, cornell.edu was the domain name for Cornell University registered with the domain registry agency and thus known around the world. In effect, the network ID for most applications replaced the former computer ID. A password was still required for network access, and this was made more secure with the addition of Kerberos authentication technology to authenticate the user—that is, the user is who he/she says he/she is.⁶⁵

Getting Organized for EZ-LINK: START and SPIRIT Committees

To create CIT working groups that were more project-oriented, in 1991–92 two new cross-organizational committees/teams were put in place: START and SPIRIT. START (Strategic Technology and Architecture Resource Team) was to consider how to bring technology changes into practice; Mark Mara took on the leadership role for this group. SPIRIT (Service and Product Integration Resource for Information Technologies) was to focus on getting services and products into production status. Cecilia Cowles took on the leadership of this group. The general thrust was that one of these teams would act as an umbrella and coordinating team for other smaller teams, which would actually carry out the projects themselves. The formation of both of these teams came at a time when one of the university’s campus-wide themes was improving quality of services through QIP, although neither was in direct response to that broader initiative.

During the period 1992–93, the projects under the two umbrella groups were:

START: Bear Access, Kerberos, and SLIP

SPIRIT: EZ-Account, EZ-Remote, Bear Access, EZ-Print, migration to client-server computing, Big Red machines, front-line survey, fall events planning

Almost all of these were successfully deployed, as some were already in production status but needed more work and fine tuning, while others, such as EZ-Account were never implemented.

Project STOP—Removing CornellC and Software Upgrades

One of the important projects that was started in 1992–93 was STOP, as in “stop the mainframe, I want to get off.” At the time CornellC use was escalating sharply from use generated by Bear Access and other services. Every access to Just the Facts or the Cornell Library system—all from Bear Access—along with all the use by the hundreds of staff in administrative offices using central business systems, went to CornellC. As related earlier, new Adabas applications were coming every year, and if all these trends continued, another new expensive mainframe upgrade would be needed before 1995. That became the objective of STOP: to find alternative server-type platforms, to move Adabas applications and the Cornell Library system, the two key components of the load, to those platforms, and to use this exercise to move more aggressively to distributed computing. At the time, the

⁶⁵ “The Big Red ID: Are You Who You Say You Are? Technology in Transition,” *CITnews*, winter 1993.

assumption was that a new Unix-based Cornell Library system would become available in early 1995, which fitted well with the plan. January 1995 was also the month when the IBM 3090-200J CornellC mainframe would be released, and so a switch at this time would make funds available for the alternative systems. January 1995 became the target date.

Based on the above objectives, the STOP project was divided into three subprojects. The first was to build a distributed computing infrastructure so that selected Adabas applications could be transferred from CornellC to other less expensive platforms. The second was to extend the life of CornellC by improving system and application performance for those applications that remained. The third was to prepare the client-server infrastructure to support the re-engineering of applications, possibly using the Mandarin technology to a larger extent. Three teams were formed to address each of these items.

Within a year significant progress was made on improving CornellC performance by a combination of factors. A new version of MVS/ESA was installed in 1992, followed by a similar upgrade for VM when VM/ESA was installed in 1993, and the changes kept performance ahead of the growth. Various other changes—such as converting 100,000 lines of PL/I code to NATURAL code, removing a significant number of IBM system-dependent codes, changing how the Adabas environment was accommodated in the VM systems, and others—brought some immediate relief. However, the relief was considered to be short-term, as the use of Adabas was continuing to increase at the rate of 15 to 20 percent per year. Work continued on the infrastructure projects, as this required installing new hardware, operating systems, and applications support programs in order to evaluate the performance of Adabas in these new environments.

1993 to 1996 at Cornell

The availability of the World Wide Web and network browsers started to change the way instruction, research, and business activities took place, enhancing the building of EZ-LINK, the Cornell Electronic Information Network. There was a successful conversion to client-server computing for basic campus services, such as electronic mail, and instructional computing. Network connections were placed in the dormitories; new network services were made available; and greater interest in regional networking was explored with several area conferences. Several new network services were introduced, the most noteworthy of which was CU-SeeMe, which used inexpensive cameras to transmit video images between workstations connected to the Internet.

New administrative applications using Bear Access were developed. These put more strain on mainframe computing resources; the conversion of Adabas to alternate platforms did not prove feasible, and so CornellC was upgraded to obtain additional capacity. Initial steps were taken to implement data warehouses to improve access to central information databases by offices across the campus. Information services continued to evolve, with changes in technology and the success of digital library projects. The Theory Center installed the first massively parallel supercomputer from IBM.

The Strategic Planning Initiative was completed and issued its report. Because information technology itself was not part of this planning initiative, the provost decided to conduct an information technology review by a faculty committee. Recommendations from the committee's report formed the basis for future CIT projects and priorities and led to the formation of FABIT, the Faculty Advisory Board on Information Technology. CIT published a series of reports on the ownership of workstations by undergraduate students between the years 1989 and 1993.

Changes continued in the executive staff of the university; the most important to CIT was the resignation of VP Lynn and the appointment of Dave Lambert in an acting role, followed by his permanent appointment as VP for information technology. Cornell and CIT were nationally recognized for their achievements in information technology services and innovations.

Planning took place on several major initiatives. One was the increasing concern about the effects of the rollover to the year 2000 (the Y2K problem), when systems that had imbedded the use of the two-digit year, that is, 1999 shortened to 99, would possibly fail with undetermined outcomes. The other initiative was Project 2000 (or P2K), which was going to be a total makeover of business systems at the university.

The lack of space for CIT staff and operations became a more critical issue and was addressed by giving up space in CCC in exchange for more space at Maple Avenue, where many CIT staff and services relocated.

Building EZ-LINK

Putting EZ-LINK into practice became the number one priority of CIT once the components were in place. To repeat, the elements of EZ-LINK were to move services to the desktop, provide access to the network, and make it easier to connect to information sources, support, and service using Bear Access.

By 1993 the new networks using FDDI for the backbone technology and UTP Ethernet to the desktop were pretty much in place across the campus. Bear Access was now over a year old, and new applications

were constantly being added to expand and improve services to users. Microcomputer workstations running on LANs connected to the campus network were the primary means of obtaining network-connected information and computing services. Network connections on campus were becoming faster with the introduction of the new campus backbone and with individual workstations being connected with shared 10 megabit/second transfer rates of local area UTP Ethernet-based connections. At this time the CIT public computer labs and two of the residence halls, Donlon and Dickson, had these direct high-speed connections. Others who needed modem access to the network were encouraged to use the new EZ-Remote network-connection services. EZ-Remote/low was the slower-speed dial-up service available at no cost (2,400 baud), while EZ-Remote itself was a fee-based higher-speed service available through subscription (9,600 to 14,400 baud).

GCA Migration Project

The GCA (General Computing Account) Migration Project, which started out as a grassroots, self-organized effort to put EZ-LINK into practice and to deal with service issues as noted above, was over time officially sanctioned as the project to implement this major technical and cultural change across the university. The members of the project team were Cecilia Cowles, Barbara Skoblick, Nancy Flynn, and Donna

Tatro.⁶⁶ Their efforts started in 1991 and continued through 1992 and 1993, when the team became part of the SPIRIT initiative. They cleverly used a combination of incentives and restrictions, “carrots and sticks,” in their terms, to move the campus from mainframe-based services to client-server-based services. A sense of the magnitude of the task can be gained from considering the information in Table 4. Even though the number of users represented only 20 percent of the total population of potential users, it was a formidable task to change the technology environment of 6,000 users.

Table 4. Number of GCA Accounts by User Category in 1991

User Category	Number of GCA Accounts	Percent of Category
Undergraduate students	3,000	25%
Graduate students	1,500	25%
Staff	750	10%
Faculty	750	50%

The EZ-LINK theme of “making it easier with Bear Access” involved exploiting the suite of services available with this first simplified network access tool, which predated the World Wide Web by several years.

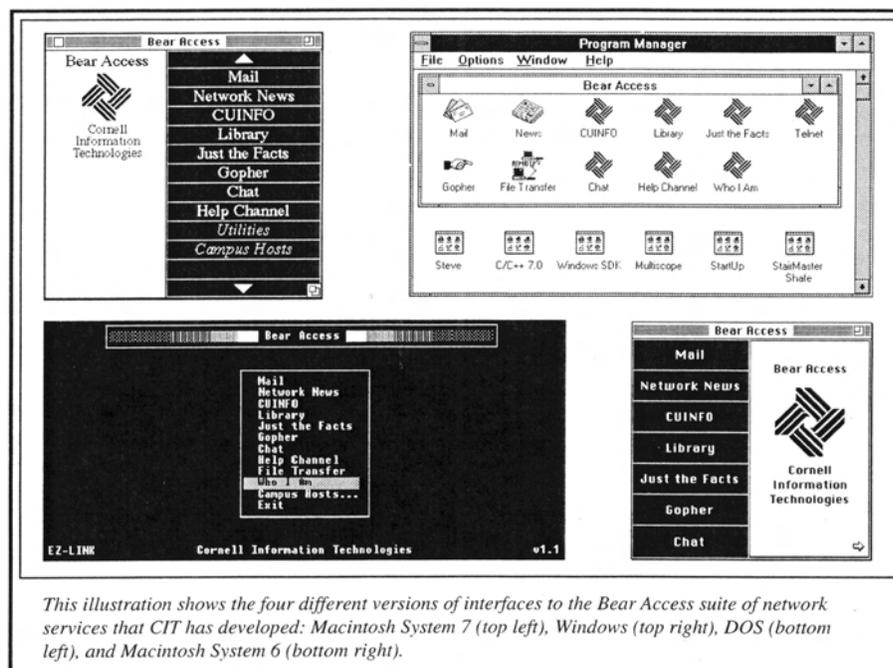


Figure 3. Bear Access for different workstation operating systems

⁶⁶ Cecilia Cowles, Barbara Skoblick, Nancy Flynn, “Power to the People: Cornell Migrates from the Mainframe to the Desktop,” 1994.

The term EZ-Access started to be used in reference to Bear Access to emphasize that its graphical “point and click” way of obtaining network services was much easier and more intuitive than the old command way of getting at services. When announced in 1992, the services included electronic mail, the online library catalog, and access to other information resources such as CUINFO, Network News, the use of FTP, and several homegrown services. Who Am I connected anyone to the Cornell network ID directory to look up e-mail addresses and other pertinent information, and Just the Facts displayed a student’s current course schedule, grades, and bursar and CornellCard accounts, and for undergraduates, their financial aid information.

The GCA Migration Project was nearly completed by 1993–94. As mentioned above, the motivational strategy consisted of “carrots and sticks” to implement the transition. The “carrots” strategy, used in 1992, was based on offering incentives to current users for voluntarily changing to the new services. (New students simply started with these new client-server services.) For those users with workstations, client software was made available in the Bear Access desktop environment for previously mainframe-based services such as e-mail, library access, and CUINFO. Corresponding servers were installed to support these services. For those without an adequate desktop environment, UNIX-based timesharing accounts were provided to facilitate access to those same services during this interim period.

In 1993 the strategy shifted to “sticks”: the mainframe computers were removed and disconnected, thus ending the service. On November 1, 1993, the DEC Vax5 was removed from service, and the IBM Cornella was no longer available for general use.

The statistics in Table 4 from 1994 attest to the magnitude of the accomplishment in moving the campus to the new client-server-based services. The numbers represent more than a four-fold growth in users, compared with the 6,000 holders of GCA accounts three years earlier, and approach the limit of potential users. The GCA migration was an outstanding success and resulted in a significant transformation and cultural change in delivery of services, which brought with it a whole new set of challenges to CIT.

Table 5. Number of NetIDs by User Category in 1994

<i>User Category</i>	<i>Number of NetIDs</i>	<i>Percent of Category</i>
Students	18,300	97%
Staff	5,870	75%
Faculty/Academics	2,008	88%
Total	26,178	

Several other fascinating statistics, shown in Table 6, illustrate the degree to which different categories of users were using the new EZ-Access and e-mail services after the GCA migration.

Table 6. Use of Network Services by Category of Users in 1994

<i>Service</i>	<i>Category</i>	<i>Percent Using Service</i>
Bear Access	Undergraduate students	70%
	Graduate students	64%
Electronic Mail	All students	60%
	Faculty/academics	56%
	Staff	42%

One of the outcomes of the GCA migration was that use of the central POP e-mail service soon replaced mainframe e-mail and several LAN-based services, primarily Quickmail, as the system of choice on campus. That was a positive event for campus e-mail users, CIT, and many network administrators. Eliminating Quickmail allowed CIT to remove the routing server that distributed e-mail between different systems on the campus and out to the world and to eliminate the support for that hardware and software. Network administrators were pleased with not having to support a package on their LAN and LAN server, with all the attendant problems. Finally, the growing use of e-mail demanded a system that could support a prospective load of 1,000,000 messages a day. E-mail users got a more robust and reliable e-mail service without the bottlenecks and instabilities of the Quickmail routers. Large computer operations, such as the Department of Computer Science, continued to run independent UNIX-based e-mail services on their systems, but the great majority of campus users began to rely exclusively on the CIT e-mail system. In the first two months it was offered, about 10,000 campus users vol-

untarily switched to the POP e-mail service (personal communication from Barbara Skoblick, e-mail administrator at the time). The POP service also paved the way for the migration to PC-based desktop computers. Quickmail was designed as a Mac-based application and did not work well on PCs. However, several POP clients were available for both Macs and PCs. CIT site-licensed Eudora, which is still the primary e-mail client in use almost 14 years later.

Microcomputer Innovations and Issues

Sales and deployment of microcomputers continued at a high level; some 3,500 to 4,000 systems were sold each year to university departments as well as to students, faculty, and staff for their personal use. To keep up with the times, new approaches were taken to simplify the introduction of new customers, mainly students, to this new technology.

CIT Sales continued to expand its offerings as products from vendors changed. It began offering high-performance workstations from Sun Microsystems, while continuing to offer systems from Apple, IBM, Hewlett-Packard, and others. In 1992 Everex Systems, chosen as the vendor for IBM PC-compatible systems in 1990, declared bankruptcy, although they expected to come out of their Chapter 11 filing and continue operations. Regardless of those expectations, that condition put a damper on sales and they soon fizzled. Sales continued to be done from 110 Maple Avenue, while the inventory and technical shop was now at the new Thornwood Drive site in Cornell Research Park near the Tompkins County Airport.

The Big Red Machine

In 1993 a new innovation was introduced: the Big Red Machines, a project out of the SPIRIT initiative. When first introduced these were prepackaged Macintoshes that were preloaded with software that worked in the Cornell environment. The software was mainly Bear Access, including such applications as electronic mail, web browsers, course information, and access to the library catalog.

Workstation Ownership Studies

In 1993 Agelia Dumas, who had been conducting annual studies on the ownership of personal computers by undergraduates at Cornell, published a summary report for the period 1989 to 1993.⁶⁷ In 1989 the study, in the form of a telephone survey, was conducted for freshmen. In subsequent years the studies were extended to include the upper classes. Information taken from the study, not all of which was collected for every year, shows the key trends and conditions, presented in Table 7.

Table 7. Trends in Workstation Ownership, 1989 to 1993

<i>Student Classification</i>	<i>Percent in 1989, 1990</i>	<i>Percent in 1993</i>
Freshmen (1989)	39%	59%
Sophomores (1990)	43%	56%
Juniors (1990)	43%	53%
Seniors (1990)	48%	56%
<i>Types of Systems Owned by Freshmen</i>		
<i>Types of Systems</i>	<i>Percent in 1990</i>	<i>Percent in 1993</i>
IBM compatibles	17%	36%
Macs	59%	48%
IBM brand	19%	14%
<i>Freshmen Who Brought Computers to Campus</i>		<i>Percent in 1993</i>
IBM compatibles		42%
Macs		33%
IBM brand		16%

Table 8. Applications Used by Undergraduate Students, 1992 to 1993

<i>Application Used (Weekly Basis)</i>	<i>Percent in 1992</i>	<i>Percent in 1993</i>
Word processing	94%	
Analytical tasks	50%	
Programming	31%	
Electronic mail	29%	43%
CUINFO	73%	75%
Online library catalog (infrequent)	78%	83%

Several other items are worth noting:

- Of the students surveyed in 1992 who reported NOT owning a workstation, 83 percent reported using workstations in public labs, 81 percent reported using a friend's, and 68 percent reported cost as the reason for not making a purchase.

⁶⁷ Agelia Dumas, "Workstation (Personal Computer) Ownership among Undergraduates at Cornell University, Report on the Ongoing Longitudinal Study Sponsored by the Office of Information Technologies, Conducted 1989-93," 1993 (there are similarly titled reports—a summary and one with details).

- In 1993, 66 percent of undergraduates in Arts and Sciences reported owning workstations, compared with 58 percent in Engineering, 48 percent in Agriculture and Life Sciences, and 37 percent in Human Ecology.

Taken together, it is clear that individual ownership and use of workstations was increasing in the first several years of the decade and that a major use of these workstations was for word processing and increasingly for electronic mail. The personal computer had indeed met the prediction that Vice Provost Ken King made 10 years earlier—that students soon would be bringing computers to campus much as they then were bringing with them calculators and stereos.

Software Piracy

Software piracy continued to be discussed, and awareness of its existence and penalties was promoted of campus. Misuse of software was receiving increased attention thanks to the Software Publishers Association, nicknamed the “software police.” This industry group had filed more than 100 lawsuits since 1989 and forced the collection of penalties as high as \$350,000. Federal armed marshals bearing court orders had surprised more than one corporate official. The association was responsible for a \$130,000 settlement at the University of Oregon for copyright infringement. The cost of software for courses was a favorite excuse of students and others for the illegal copying of software.⁶⁸

With the removal of special protections in the software and the increasing use of “shrink-wrap” licensing (the act of removing the shrink-wrap cover was an acceptance of the contract terms), copying became quite easy. At the university such practices were a violation of the code of academic integrity and the judicial code and were not to be tolerated. VP Lynn acknowledged that enforcement was a difficult task but nonetheless the Cornell community needed to be responsible and acknowledge that copying was a violation of the developer’s intellectual property rights, which had to be respected. A program was developed to bring this issue to the attention of the campus. At the same time, vigorous efforts were undertaken to lower the cost barriers through lower-priced student versions of software, special contracts with vendors, and software sharing agreements. Vigilance continued to be practiced to make sure Cornell was not subject to a sweep by the software police, although it was clear that full compliance could never be achieved in practice.

Back-to-School Fair

The eighth annual Back to School Fair was held in Lynah Rink in late August 1995, offering one-stop

shopping, particularly for incoming students. It was open to all on campus, with participation from the industry’s leading vendors: Apple, IBM, Lexmark, Hewlett Packard, US Robotics, and Microsoft. The cost-effective Big Red Machines continued to be offered, particularly to incoming freshman students, as a convenient and cost-effective way to start their use of IT at Cornell. In 1995 close to 11,000 individuals came to the fair.⁶⁹

Research

CIT continued to cooperate with both CISER and the Theory Center to enhance and promote the use of their facilities by researchers on campus. CIT tracked the continuing changes in the Theory Center’s computer services because CIT continued to be part of their Smart Node program.

IBM SP1 Installed at the Theory Center

In April 1993 the Theory Center was the first customer to install IBM’s new Scalable Powerparallel (SP1) system, which consisted of 64 (RS-6000) RISC-based processors. The system was purchased in part with a grant of \$12.3 million from the state of New York as part of its program to advance high-performance computing in the state.

In late 1993 the KSR1 system was a 128-processor unit made up of two 64-processor units working in parallel and was rated at 5.12 Gflops. By this time it was supporting over 50 research groups working in part to meet the “grand challenges”—fundamental problems in science or engineering whose solutions would be enabled by the application of future high-performance computing resources, as posed in the High-Performance Computing national initiative.⁷⁰

Mass Storage—EPOCH Server

CIT focused its mass storage initiatives on support of the EPOCH server; other projects had become the responsibility of the Theory Center with the transfer of operational responsibilities and systems support for the supercomputers to the center. Starting in 1994, the Theory Center continued using the Unitree system, although IBM was now supplying the Unitree system based on the National Storage Labs (NSL) version of the software. At about this time, the Theory Center was involved in a proposal with the national supercomputer centers to create a single image of com-

⁶⁸ “CU issues warnings on misuse of software,” *Cornell Chronicle*, September 24, 1992.

⁶⁹ “CIT’s annual computer fair offers one-stop services,” *Cornell Chronicle*, August 24, 1995.

⁷⁰ Cornell Theory Center, *Forefronts*, vol.8, no. 3, winter 1993.

puting and other resources (a metacenter of distributed heterogeneous computing) so that a user could use any resource from any location. Not part of those deliberations, CIT was primarily concerned with supporting the EPOCH box for the Cornell Library project to store scanned books and other documents, as well as for CISER.

By 1993 the CornellA IBM 4381 shared between CISER and CIT was overloaded, with customers from both groups receiving inadequate service. To provide some additional planning time for CISER to seek alternative computers to meet its needs, CISER's extensive data archive was placed on the EPOCH storage system. Not only would this facilitate the migration to new computing facilities, but it would enable CISER's users to think more in terms of data files and less in terms of the particular tape or disk where the file was stored. Toward the end of this period, CISER was using 15 Gbytes for online storage of its data archives; by the end of 1996, CISER was storing 150 Gbytes of information for its users on the EPOCH system.

New CISER Computers—Gaea and Hera; CornellA Removed

In anticipation of the removal of CornellA, in April 1993 CISER installed an IBM RS6000 system that would be their future server in the new computing client-server model.⁷¹ According to Mariann Carpenter, manager of the CISER Computing Group, they chose the name Gaea—the “first deity, the origin of all things and the personification of the earth”—for the first of several anticipated servers that would be similarly named. Because Gaea would operate with AIX, IBM's Unix-based operating system, and it would require both experienced technical resources and time to convert all the users, CISER entered into a formal facilities management agreement with CIT to support their new computing environment. Gaea was installed in the CIT computer room in CCC. CISER users converted their data, programs, and operating style to this new environment during the transition year, 1993–94. In January 1994, another CISER server, Hera, was installed. In this two-server configuration, the use of Gaea was intended for access to the data archive, for the manipulation of data prior to analysis, and for small-scale analyses; Hera was to be the cpu server for large-scale analyses.

The end of mainframe computing for research and instructional computing, and 13 years of shared use of CornellA between CISER and CCS/CIT, came on June 30, 1994, when the IBM 4381 was taken out of service and sold.

Supercomputing—IBM SP2 Installed; KSR1 Removed

In early 1994 the Theory Center installed a new IBM Scalable POWERparallel (SP2) system computer to replace the SP1 installed a year earlier. The SP2's POWER* RISC processors would have twice the performance of the SP1 processors, eight times more memory, and four times greater internal bandwidth. While the initial installation consisted of 64 processors (nodes), it was expected that all 512 nodes would be installed by the end of the year. When fully configured, the SP2 was capable of performing 136 billion calculations per second, and the 512 nodes could all be working on a problem in tandem. Peak performance was rated at 136 gigaflops, with each node rated at 266 megaflops per node.

Kendall Square Research announced in September 1994 that they were discontinuing the manufacture and sales of the KSR1 computer, and the KSR1 was removed from service on March 31, 1995. Users were given enough lead time to move their work to the SP2 or other systems.⁷²

First URL (Universal Resource Location) at Cornell

With the announcement of the installation of the SP2, a new innovation was also announced and that was World Wide Web access to information about the SP2 and research projects using the National Center for Supercomputing Applications Mosaic¹ browser. This was one of the first and early URLs at Cornell, <http://ibm.tc.cornell.edu>, which allowed information about the Theory Center and its systems to be available to a wide variety of workstations across the country and the world.

Theory Center Celebrates 10 Years; TC Building Named Rhodes Hall

In May 1995, the Theory Center celebrated its 10th anniversary with a scientific symposium and an open house. At this time it was noted that more than 5,000 researchers in diverse fields had used the computing resources at the center. The SP2 had the full complement of 512 processors that provided 1/8 teraflops of computing capability.⁷³

In June 1995, the Engineering and Theory Center building was renamed Rhodes Hall in honor of President Rhodes, who was retiring at the end of June.

To illustrate the dramatic increases in supercomputing power that had taken place in 10 years, the new

⁷¹ Cornell Institute for Social and Economic Research, annual reports, 1992–93 and 1993–94.

⁷² Cornell Theory Center, “KSR1 to be removed from service, March 31,” *Hottip 94HT27*, November 28, 1994.

⁷³ “Hundreds help Theory Center celebrate anniversary,” *Cornell Chronicle*, May, 25, 1995.

capability of the SP2 was an increase of 625 times the estimated 200 megaflops of computing power available when the IBM 3084 QX system and the seven supporting FPS array processors were installed in 1985. The center announced a plan to triple its supercomputing capability in 1997 using Power2 Super Chip (P2SC) nodes to replace the IBM RS/6000 processor nodes.

CISER Takes Responsibility for Statistical Systems

CISER continued to play an increasingly greater role in providing campuswide statistical services as the decade progressed. In 1995 CIT dissolved the Statistical Computing Support Group, and the site license for SAS became the responsibility of CISER rather than CIT. Except for CIT's providing facilities management for its two server systems in the CCC computer room, CISER's only direct link to CIT was in the continued use of the EPOCH storage server to contain CISER's ever-increasing use of online storage for data files.

Information Services

The success of the 1991–93 prototype digital library project to scan books to create and store digital images of pages led to some further new projects to extend the availability of library materials in digital form.

The Making of America Project

In late 1993 the success of the prototype project resulted in "The Making of America" project. With support from the Culpeper Foundation, the project was to digitize material from Cornell's collections "that document the development of America's infrastructure—transportation, communications, and the built environment—between 1860 and 1960." The project leaders planned to enlist the support of other institutions to digitize 100,000 volumes. The project was a joint venture between the Cornell Library and CIT, working in collaboration with Xerox Corporation, which was providing the technology. The effort also had the support of the Commission on Preservation and Access and of Sun Microsystems.⁷⁴

Online Digital Library

The project to make the digital images of the scanned books available for online viewing continued to move ahead, according to William R. Turner, the lead technical person, who provides this description (personal communication, 2003).

The scanned images were stored as 600 dpi TIFF images in a Xerox proprietary document format called XDOD; the Xerox-sup-

plied software could transmit an XDOD document in its entirety to the Docutech printer. Programmers in the library defined a Cornell Digital Library document architecture that supported identifying key parts of the document, such as the table of contents, chapters, and indices, and supported jumping to the desired place in the document and paging through it either forward or backward. Conversion software was written to export the documents from the XDOD architecture to the Cornell Digital Library architecture and to store 50 dpi scaled images of each page for online viewing, in addition to the 600 dpi images to be used for printing. A client-server architecture and protocol was defined to support both the viewing of images online and requesting printed copies, and these were published as Internet RFCs. Finally, the client and server software was developed. Sun Microsystems contributed a robust server with large disk capacity to store the images so that online performance would be better than was available with the EPOCH magneto-optical jukebox, and 800 volumes were converted from the Xerox XDOD format to the Cornell Digital Library document format. Unfortunately, with the departure of VP M. Stuart Lynn, CIT's interest in the project shifted, and work on the Cornell Digital Library software was stopped. Instead, an effort was made to obtain software developed by the University of Michigan for a different project, and when this effort failed, the focus was shifted to another protocol developed at Cornell, Dienst.

The movement toward using computing and network technology was gaining momentum and was sure to continue throughout the decade. In 1993, for example, Mann Library was the recipient of the American Library Association/Meckler Library of the Future Award⁷⁵ for "exemplary work in the development of an overall program that demonstrates the capabilities and applications of information technology." In this same period, Mann Library had initiated a project to computerize core agricultural literature as a worldwide reference library and became the Internet site for USDA statistics of more than 140 agricultural data

⁷⁴ "Digital library on campus to record 'Making of America,'" *Cornell Chronicle*, December 16, 1993.

⁷⁵ "Mann named Library of the Future," *Cornell Chronicle*, May 13, 1993.

sets, which were made available at no cost and were open to anyone with a network connection.

CUPID Project and Award

About the time that the digital library project first got under way, a separate but related project on printing technology was started. It took the name CUPID—for Consortium for University Printing and Information Distribution—with VP Lynn playing a key organizing role. CUPID was a cooperative effort of 10 major research universities interested in changing the model by which all books and periodicals, reprints of articles, and other such materials were distributed. The idea was to move “bits not books” and capitalize on the availability of the Internet. So, for example, a reprint of an article could be stored at a cooperating site and when a copy was requested, CUPID would move the electronic version of that article to the requestor’s site, and there a copy would be printed using equipment such as the Docutech. Books and journals could be printed as bound, finished documents using technology similar to the Docutech. Steve Worona led the effort to define the CUPID protocols, and the prototype software to carry out these functions was written at Cornell (Sailesh Gurnani) and Harvard (David Greenlie) and demonstrated at various conferences.

In 1994 CUPID won the Innovation of the Year Award from XPLOR International, a nonprofit association of the electronic printing industry—the first time the award was given to an academic organization.⁷⁶

Internet Relay Chat

One of the popular offerings on Bear Access was Internet Relay Chat (also known as Chat or IRC), which was discussed in the winter 1993 issue of *CITnews* with the title “Internet Relay Chat: A Chance to Join the Global Village.” That particular issue was dedicated to the theme “Bear Access Delivers the World to Your Workstation.” The article mentioned the wide range of possibilities for conversations (“a universal electronic cocktail party”) to which Cornell users could connect. There was a special Help Channel that CIT staff monitored to provide assistance to any member of the Cornell community. Opportunities for use included onscreen-ongoing news, electronic class meetings, electronic town meetings, and others left to the imagination of the campus community. CIT operated one of the several IRC servers on campus to support the service.

CIT Activities, Issues, and Transitions

Despite all the planning activity going on during the early years of the decade, and the understanding that

information technology was relevant in all areas of the university, no particular activity was directed to IT even though it commanded the considerable attention of various planning committees. At the same time, the release of the 2001 Committee report recommended some attention be directed to the development of IT on the campus.

Information Technology Review—Merten Committee

In April 1993 Provost Nesheim appointed a committee to review CIT and the ways that university faculty and staff used and managed information technologies—particularly computers and communications networks—throughout the campus. The chairman of the committee was Alan Merten, dean of the Johnson Graduate School of Management, who held a doctorate in computer science and who had guided corporations through similar studies. Accordingly, this committee became known as the Merten Committee.⁷⁷ Other members were John Hopcroft, associate dean of the College of Engineering and a professor of computer science, John Wiesenfeld, vice president for planning, and Charles McClintock, associate dean of the College of Human Ecology.

VP Lynn welcomed the review, noting that IT now permeated all campus activities, that what had been something of a convenience was now a necessity, and that deployment of these resources was increasingly important. The committee invited consumers of information to make their comments and concerns known to the group and considered their job as evaluative and not visionary. They asked academic and administrative units to explain why information technology is critical to their mission, to assess their needs for the next five years, and to suggest changes in the mission of CIT and other organizations in light of those needs.

In preparation for the Merten Committee review, VP Lynn and the senior management group prepared a report displaying statistics and other summary information to show the extent to which units on campus took advantage of CIT services in support of their academic and administrative activities.⁷⁷ It was an excellent snapshot of the degree to which information technology had penetrated the campus and showed that all departments and colleges depended on CIT for infrastructure, voice and data networks, and major administrative systems. CIT was the provider of first resort for many departments and colleges and of last or intermediate resort for those departments and colleges that were more self-sustaining. Many statistics were

⁷⁶ “Computer Program named ‘Innovation of the Year,’” *Cornell Chronicle*, October 13, 1994.

⁷⁷ M. Stuart Lynn, “A Report on Campuswide Use of CIT Support and Services,” May 1993.

presented to display the use of CIT facilities and services by college. Here are a few examples to show both the breadth of services and their use.

- Use of the Help Desk grew from 14,587 consultations in 1989–90 to 38,804 consultations in 1992–93. In 1992–93 there was a total of 49,000 Help Desk interactions, of which 43 percent came from university staff.
- The number of attendees at workshops and “Getting Started” seminars increased from 1,786 to 2,228 between 1989–90 and 1992–93.
- The number of students and courses using CIT facilities increased from 4,400 students enrolled in 60 courses in the fall semester of 1990 to 5,400 students in 97 courses in the spring semester of 1993.
- In 1991–92, the most recent full year studied, CIT Sales sold over \$9.4 million of computers and peripheral equipment and office equipment on campus.
- In the first nine months of 1992–93, there were 1,765 requests for equipment repair services, at a total cost of \$207,053 to the campus.
- In 1991–92 there were over one million invocations of e-mail on the CIT mainframes, of which 86 percent was from academic units and 14 percent from other units. (This volume was only a portion of the e-mail traffic, as by this time there was LAN-based and early client-server e-mail.)
- In 1991–92 there were 119,000 invocations of CUINFO on the CIT mainframes, of which 84 percent came from academic units and 16 percent from other units
- The Cornell Library’s NOTIS system handled an average of 260,000 transactions on a daily basis, with 1.5 million input/output operations and 2.9 million database calls.
- In April 1993 there were 487,000 searches on the online catalog and 130,000 charges and renewals
- For administrative systems there were 7,569 users campuswide, 39 percent using academic and resource systems and 61 percent using business and financial systems.

The Merten Committee issued its report⁷⁸ in December 1993, and the summary conclusion was that “Cornell is rich in technological infrastructure but less effective, in comparison, in meeting the needs of its community for everyday applications of information technology.”

The team found an institution with apparent contradictions: advanced technology second to none and a showcase electronic library existing side-by-side with a number of faculty members in “pockets of poverty” who were still not connected to the network. They also noted administrative and student systems that fell

significantly short of expectations. They advocated a new shared responsibility between CIT and Cornell units, noting that such units needed to take responsibility for their IT applications, in effect seeking a new balance between CIT and the units.

VP Lynn concurred wholeheartedly on the need for CIT and Cornell’s units to strike a new relationship, one in which the mission of CIT was set jointly across the university. Consulting with users and being driven more by demand than by technology were key parts of the recommended changes. For example, despite the technological and practical value of Bear Access, some on campus felt that CIT accomplished less than it could because it did not consult sufficiently with the user community to ascertain needs. Acknowledging that CIT was the target of both widespread praise and criticism, the report concluded that, on balance, campus customers held a negative view of CIT. “Some of this negative view derives from ‘sins of commission,’ but more appears to come from ‘sins of omission.’ This is definitely the case with respect to administrative systems.”

The review team also noted that tracking expenditures on information technology across the university was very difficult and no one knew the total scope. The largest identified component was the CIT budget of \$24 million in fiscal 1992–93, including \$11.5 million in direct university funding, an amount that actually declined in constant dollars during the decade.

FABIT—Faculty Advisory Board on Information Technology

In late 1993 VP Lynn formed a new Faculty Advisory Board on Information Technology (FABIT), which still exists as of this writing. Forming this board was not a response to the Merten report—it had already been in the planning stages—but the timing seemed to be in sync with the times. The 13-member board with wide campus representation was broadly chartered with advising VP Lynn on key policy issues, requirements, directions, priorities, and concerns about IT.⁷⁹

One of the key issues FABIT considered was the relationship between CIT and the university’s academic units in dealing with the phenomenal growth of IT applications across the campus. Other proposed topics included the following:

- What should be the directions for instructional computing, how should it be supported, and how should centrally managed facilities articulate with college facilities?

⁷⁸ Alan G. Merten, et al., “Report on the Provost’s Information Technology Review Team, Cornell University,” September 2, 1993.

⁷⁹ “Faculty board to advise on technology,” *Cornell Chronicle*, December 9, 1993.

- To what extent should Cornell accelerate the diffusion of learning technologies into the curriculum, and who should play what role?
- What are the faculty's perceptions of the need for access to campus and worldwide information resources and for collaborative tools, and how should these be deployed?
- How can Cornell facilitate completion of the transition from centralized mainframe computing to distributed computing, support faculty through the transition, and finance deployment of resources?

FABIT was able to influence IT policies on campus, in particular the upgrading of instructional facilities, as will be seen from reports later in the decade.

The Need for More Building Space

Insufficient space (both quantitatively and qualitatively) for academic and administrative support services continued to be a problem for CIT during the 1992–94 years. As noted earlier, a number of moves were made in CCC after the CNSF supercomputers were relocated to Rhodes Hall, and some Information Resources programming staff were moved out to Thornwood Drive, near the airport. These changes brought some relief, but the pressure continued. There was pressure for CIT to vacate the space in Caldwell Hall and to release other space in CCC. Pressure also came from the need to vacate Sage Hall, which was to be fully renovated for the Johnson Graduate School of Management, and to vacate smaller units in Day Hall so that units moved out of Sage might be placed there for continued access by students. The emerging view was that the space in CCC was prime central campus space, needed more for academic or academic support functions than for service functions.

The CIT space committee continued to meet and consider different alternatives and plans, such as moving all or part of its operations off campus to the Humphreys Service Building or to the former Cornell Laundry building at 120 Maple Avenue. CIT even looked at the old Ithaca Gun site on Gunshop Hill. The final decision was for CIT to give up the third and fourth floors of CCC and all of their space in Caldwell and to move that staff, plus the IR staff at Thornwood Drive, into new, renovated quarters at 120 Maple Avenue. Planning on the Maple Avenue project began in early 1994.

CIT Staff Departures, Information Resources Split into Two Divisions

The years 1993 and 1994 saw significant staff turnover in CIT. Dave Pulleyn,^k longtime assistant director for operations, took early retirement in early 1993 and

was replaced by Jim Doolittle, who took on that role along with his responsibilities for production control and data entry.

In early 1994, Dave Koehler,^l director of Information Resources, left Cornell for a position as director of information systems at Stanford University. John Rudan was appointed acting director of IR, pending recruitment of another director. At about this time the IR division was split apart, forming a smaller IR unit under the temporary direction of Rudan and a Distributed Technology (DT) division under the direction of Mark Mara.^m In effect, the core of IR staff members who had been dedicated to new technology (Bear Access and Project Mandarin) and in serving a broader campuswide clientele were separated from the staff serving administrative customers and supporting current business systems.

VP Lynn Leaves Cornell, Lambert Appointed Acting VP

In April 1994 VP Lynn announced his resignation to return to California.ⁿ He left in July 1994 after serving just over seven interesting, innovative, and exciting years in the position. He deserves significant credit for moving Cornell to new client-server technologies and applications and for striving to enlarge, improve, and simplify access to IT resources. His leadership led to the installation of a contemporary network. He almost single-handedly moved the Cornell Library into the new world of digital libraries, opening new vistas for handling the acquisition and display of books and scientific journals. In summary, his efforts significantly increased use of information technology on campus, enabling Cornell to continue its leadership role at the forefront of information technology use in higher education in the country. When he left he was quoted as saying, "These have been wonderful years. Cornell is a terrific institution, and the staff at CIT is absolutely fabulous." H. David (Dave) Lambert was appointed acting VP pending recruitment of a new vice president for information technologies.

New Deans, New President, and Other Executive Staff

Throughout 1994 there were a number of significant changes in the executive staff and college deans. Early that year, John E. Hopcroft was appointed dean of the College of Engineering, while in the previous month Alan G. Merten had been re-appointed to another five-year term as dean of the Johnson Graduate School of Management. In March of that same year, President Rhodes announced that he would step down as president by July 1995, initiating a search for a replacement. In June Susan H. Murphy, dean of admissions and financial aid, was appointed vice president for student and academic services, replacing Larry Palmer,

who returned to a full-time teaching position in the Law School. At about this same time Provost Nesheim announced his intention to step down in another year. In September Norm Scott was appointed to another five-year term as vice president for research.

In December Hunter Rawlings was appointed president of the university, succeeding Frank Rhodes, who had been president since 1977. Shortly afterward John R. Weisenfeld, vice president of academic programs and planning, was replaced in an acting capacity by Ronald G. Ehrenberg, professor in the School of Industrial and Labor Relations. By mid-1995 Ehrenberg and Lambert were confirmed as vice presidents on a permanent basis. Fred Rogers was appointed senior VP and chief financial officer, replacing Jay Morley, who left Cornell to become head of NACUBO. Don Randel was appointed provost, succeeding Mal Nesheim, who returned to his faculty position. While all these individuals came to play an important role in setting the priority and growth parameters for information technology to the end of the decade, it was the troika of Weisenfeld/Ehrenberg, Lambert, and Rogers that played the key role in the mid-decade period.

Strategic Planning Outcomes; Information Technology Recommendations

The outcome of university-wide strategic planning is contained in the reports “Consultation Draft Plan” and “Task Force Reports” issued in 1994. A revised working statement,⁸⁰ “Toward 2000: Mission and Values,” was issued for Cornell, proposing a focus on three areas:

- We must be first-rate in education at every level, through the strength of our academic offerings and the enrichment that both undergraduate, graduate, and professional students gain from learning in a research environment. Enrichment opportunities include participation in the creation of knowledge and exposure to intellectual and social learning environments outside the classroom.
- We must be first-rate in scholarship, in research, and in creative endeavors by recognizing that the need to know and to understand is one of humankind’s most basic characteristics and that a strong program of fundamental research and scholarship is our best response to the needs and goals of society.
- We must be first-rate in knowledge transfer. Cornell’s intellectual and instructional

activities must not stop at the boundaries of the campus or on the pages of academic journals. Outreach is one of our fundamental responsibilities, linking the fruits of our research to agricultural productivity and profitability, industrial innovations, and the needs of the family, the classroom, and society at large, in keeping with our land-grant role. Additionally, our academic programs must be informed continually through interactions with the outside world.

The draft report then went on with proposed strategies and recommendations to meet the mission core values and goals.

In the “Ways and Means” section of part III (“Objectives and Proposed Strategies”) of the draft report, the objective for dealing with communication and information technologies is stated as follows:

We will develop our communication and information technologies to promote more effective learning, to extend interactions with and beyond Cornell, and to enhance the quality and effectiveness of academic programs and support services.

Several task forces considered this objective and developed these final recommendations:

- Streamline administrative structures with a view to placing decision-making functions closer to the workplace. (Strategy A.3.2, Task Force #2, The Generation and Application of Knowledge)
- Ensure that Cornell maintains an integrated voice, data, and video communication infrastructure that is second to none, linking all faculty, students, and staff within the campus and with external communities; and develop technology and funding strategies that ensure network access and support remain competitive with national and international standards of excellence. (Strategy A.3.3, Task Force #2)
- Assume leadership in the use of new electronic modes of dissemination; and make available and provide access to electronic depositories of scholarly and other information sources, locally and on a national and international basis. This task should be pursued in collaboration with other institu-

⁸⁰ Institutional Planning and Research, “Cornell University, University-wide Strategic Planning, Consultation Draft Plan, Task Force Reports,” May 1994.

tions, taking advantage of consortial opportunities to share costs and resources, and to advance standards. (Strategy A.3.4, Task Force #2)

- Develop campuswide technology and funding strategies that guarantee ubiquitous network access and support that remains competitive with evolving national and international standards of excellence. (Recommendation A.5, Task Force #2)
- Use information and communication technologies to improve information flow and enhance inter-institutional cooperation. (Strategy B.4.5, Task Force #2)
- Exploit opportunities afforded by world-wide voice, video, and data electronic networks to facilitate interchanges with scholars and practitioners abroad, to support international institutional partnerships and collaborative projects, and to import and export educational information to and from Cornell. Through these media Cornellians can benefit from the use of videoconferencing, e-mail, and other collaborative tools. (Strategy E.3.8, Task Force #2)
- We recommend the active and vigorous use of re-engineering approaches to develop new and much more effective ways of operating the university, rather than a focus on simply “downsizing” the university’s operations in response to budget reductions. (3.b, Task Force #3, Exercising Effective Stewardship)
- We recommend the implementation of standard administrative data systems and the development of greatly enhanced campuswide access to and the usability of centralized information databases and electronic services. (3.c, Task Force #3)

All these statements played into the priority given to the activities of CIT in the years that followed. In particular, when the last two statements about administrative operations are linked with the statements in the Merten Committee report about administrative and student systems falling significantly short of expectations, it should not be a surprise that these issues became of paramount future importance. Clearly IT had moved to the “front burner” of university issues, and CIT had a long list of guiding statements for its activities in the short and longer term.

QIP Fades Away

The departure of so many of the executive staff in 1994 and 1995 was a serious setback to QIP, the Quality Improvement Program, and the Strategic Planning Initiative. Prior to his departure, Morley had transferred oversight of QIP to Beth Warren, director of human resources, and the program started to lose some momentum after two years. The greatest contribution of QIP was in the training of large numbers of staff at different levels, the approaches it promoted to teamwork and problem solving, and the breakdown of “stovepipes” at all levels of the organization. Some departments kept the process going, and in some it evolved into other staff improvement training. While there was no formal end to the campuswide program, it quietly faded away.

Networking and Network Services

Network use continued to increase dramatically, forcing changes not only in the network infrastructure but in the servers and computers supporting the services. An article from the *Cornell Chronicle* in September 1996⁸¹ gave some sample statistics for the academic year 1995–96:

170 million e-mail messages handled by the e-mail servers⁸²

1,032,000 Just the Facts accesses

39,000 Employee Essentials accesses

282,000 library catalog accesses

At this point there were over 36,000 NetIDs and roughly 24,000 users of e-mail services. These statistics reflect a 25 percent growth in NetIDs and a 50 percent growth in the number of e-mail users in about two years. To accommodate the growth in e-mail use, two new larger-capacity e-mail servers replaced the five older servers during the summer of 1996. The configuration of two servers continued throughout the decade.

ResNet—Networking Connections in Cornell Residence Halls

ResNet was the name chosen for the service placing network connections in all rooms in Cornell residence halls and all associated living sites for students. In other words, where there was a student bed or tele-

⁸¹ “CU’s network services have gotten a boost from computer upgrades,” *Cornell Chronicle*, September 16, 1994.

⁸² 170 million e-mail messages stated in the original article seems out of line with tracking statistics maintained by CIT support staff, which report 64 million for the same period.

phone jack, there would be a network connection. ResNet was one of the priorities mentioned in the 2001 Committee report and later in the strategic planning reports. It had strong support from CUBIT.

Following the success of the pilot project conducted during 1992 and 1993, the ResNet project started in earnest in the fall of 1994 as a full-fledged pay-as-you-go subscription service for UTP Ethernet network services. The Residence Hall Network Planning Group of Fred DeWolfe and Peg Lacey from Residence Life and Dave Lambert and VP Lynn from CIT, along with Larry Palmer, VP for academic programs, gave strong support to the project.

In the fall of 1994, M. Scott Walters was hired to direct and move the program ahead. That fall there were 3,000 connections available, and approximately 18 percent or 540 subscriptions were taken.⁸³ While such a subscription rate may seem low, the services were limited (only e-mail and Gopher), and the student subscribers were mostly computer science majors and the more computer-literate students. Installation of connections continued during the fall semester, with a planned installation of 6,000 connections or ports. When the final project approval document was signed in the summer of 1995, the project cost was estimated to be \$1,050,000, to be financed with \$850,000 from Campus Life room rates and \$200,000 from CIT current funds.⁸⁴ A rate of \$82.50 per semester per connection (\$165 per year and \$40 for the summer) was expected to recover operating costs (personal communication, M. Scott Walters). The rate stayed the same throughout the decade.

In 1995–96 the number of subscriptions went up by a factor of four to over 2,300. CIT scrambled to support this unanticipated rapid growth, adding enough staff and resources to provide adequate services in those first few hectic weeks of the semester. By 1996–97, subscriptions had grown to half the student population after a modest increase in the number of connections when some of the smaller living units were wired for service.

Over this three-year period the number of Macintosh systems decreased from 41 percent to 23 percent of the subscriptions, while the “Wintel” systems (Windows OS with Intel chip processors) increased from 59 percent to 77 percent. Given the way the industry was moving, that was not unexpected, although it was disappointing to the Macintosh supporters.

Travelers of the Electronic Highway (TEH)

All new Cornellians were given a NetID and password upon registration. Technologically savvy students could begin using their computers immediately after subscribing to ResNet during the Back to School Fair, or using their own modems. To expedite and shorten

the IT learning experience, all new incoming students had to take a one-session mandatory course, Travelers of the Electronic Highway (TEH), which covered Bear Access basics along with network policies and recommended behavior.⁸⁵ The TEH course was organized in 1994 in response to the recommendations of the FABIT report, which called for such broad-based introductory training. The course was built on the experience of the “Getting Started” courses given earlier in the decade. Students who failed to attend the course had their NetIDs disabled after seven days and then had to make special arrangements to restore their network access.

IthacaNet First Conference

In the Ithaca area, the issue of networking continued to draw attention and interest. In 1995 a loose-knit coalition of interested parties, calling itself IthacaNet,⁸⁶ was organized to encourage the growth of computer communications in Ithaca and Tompkins County. Members of IthacaNet included representatives from educational institutions (Cornell, Ithaca College, BOCES, and the South Central Research Library Council), local government agencies (such as the City of Ithaca and Tompkins County), and private corporations (NYNEX and NYSEG and other agencies).

The first IthacaNet conference was held in April 1995 with the following conference description: “The IthacaNet Organizing Group invites you to participate in the first of what we hope will be an annual event. Networking Tompkins County—Building Infrastructure, Providing Access, and Creating Content will bring together businesses, government agencies, and other organizations in Tompkins County to chart a common course on the Information Superhighway.” The conference was held at the Ramada Inn on Triphammer Road.

Bear Access, Version Control, SideCar

In the fall of 1995, Bear Access had graduated to a new level of sophistication. With the availability of new terminology coming from the rapid development and deployment of the World Wide Web, Bear Access was now known as an “on ramp” to the Internet. Bear Access could be installed on Macintosh, Windows,

⁸³ M. Scott Walters, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, December 2000.

⁸⁴ Campus Life Residence Halls Networking Project, Project Approval Request (PAR), August 3, 1995.

⁸⁵ “New students get electronic ‘travel’ lessons from CIT,” *Cornell Chronicle*, April 25, 1996.

⁸⁶ www.ithaca.ny.us/Events

or DOS-based computers. New services were being made available by campus service providers who had access to Cornell-developed tools to make such additions easy. The Office of Student and Academic Services offered Just the Facts, Faculty Advisor, and CoursEnroll; Human Resources offered Employee Essentials; and the Cornell Library offered Library Resources. Several Cornell instructors had started to incorporate Bear Access services into their courses—for example, the distribution of assignments using web pages and the increasing use of e-mail for communication between instructors and students. VP Lambert was quoted as saying, “Bear Access is the envy of every university in the country,” because Cornell had led the way in the transition from mainframe-based computing to client-server technology with the development of Project Mandarin technologies and their deployment as Bear Access.⁸⁷

In 1995 one of the new features introduced as part of Bear Access was SideCar. SideCar was an application that extended Kerberos protection to online services that did not have Kerberos built in, providing electronic verification of the individual seeking to use the service.⁸⁸ SideCar was developed by Project Mandarin, and the selection of the name is attributed to Tom Weyer of the Mandarin team (who also provided Figure 5).⁸⁹ The first versions of SideCar were developed by Andy Hanushevsky for the Mac and by Kevin Leonard for the PC. SideCar was extended to variants of Unix by CIT staff when CIT took responsibility from Project Mandarin and has continued to be in use until the present time. As with Kerberos, SideCar could be downloaded from Bear Access.

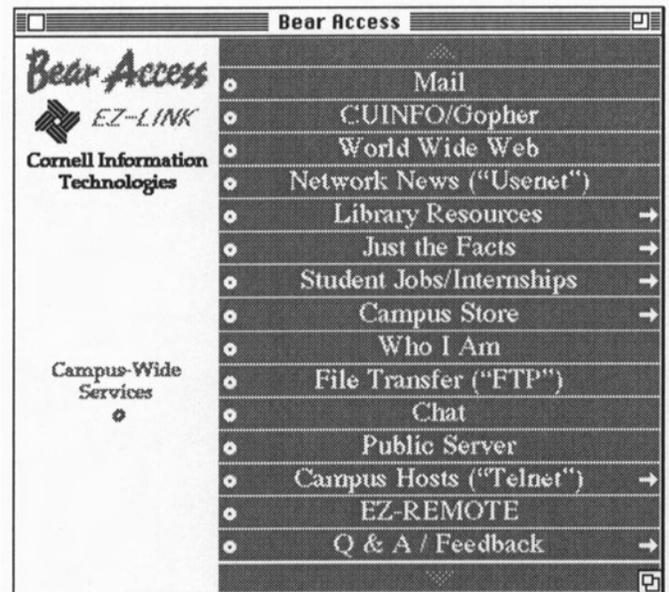


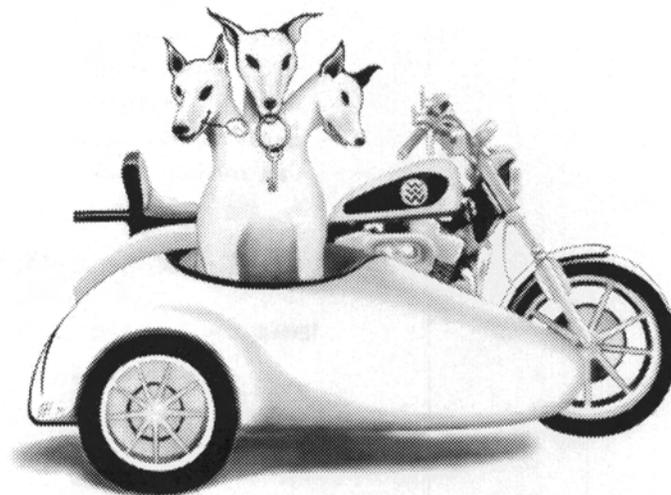
Figure 4. Bear Access launch pad showing services available

One of the other important features of Bear Access was its ability to update itself and the applications it contained. A click on a service button automatically checked the version of the software on the workstation. If the version was older than the version on the

⁸⁷ “Bear Access software ‘graduates’ to a new level of sophistication,” *Cornell Chronicle*, October 19, 1995.

⁸⁸ “About Kerberos and SideCar,” www.cit.cornell.edu/kerberos/about.html.

⁸⁹ Project Mandarin, Inc., and Cornell Information Technologies Publication and Information Group, “Project Mandarin SideCar,” April 1996.



When used with Kerberos software, SideCar effectively restricts access to data stored on a server. In Greek mythology, Kerberos was the three-headed dog who guarded the gates of Hell. In this illustration, created by Tom Weyer of Cornell University, Kerberos and SideCar are depicted together since SideCar is often used in the authentication process.

Figure 5. SideCar

central server, the user was reminded to download and update the service and thus stay current with changes and avoid obsolete software. The notification also described the impact the download would have on storage. This “version control” ability was another key innovation developed at Cornell. During the period from June to December 1995, almost 1.2 million software updates were made through Bear Access, compared with fewer than half that many—402,796—in the comparable period in 1994. The alternative of distributing upgrades on diskettes or CDs would have been expensive, not to mention probably less successful in terms of actual upgrades made, given the differential in convenience.

New EZ-LINK Services: EZ-Publish, EZ-Backup

Two innovative network services were started in 1993–95, taking advantage of the fact that the network reached all across the campus. In keeping with the EZ-LINK nomenclature, they were called EZ-Publish and EZ-Backup.

EZ-Publish capitalized on the availability of the Xerox Docutech printer that was first installed as part of the Cornell Library book preservation project. The Docutech was capable of producing plain, bound, or stapled documents, with insertions of different color and weight of paper so that the final product was a high-quality document, close to that which could be made using print shop technology.⁹⁰ CIT itself published some of the last editions of *CIT News* using this technology. EZ-Publish enabled staff around the university to produce professionally prepared documents right from their workstations, a useful extension of desktop publishing. When announced, the service had a charge of \$.025 per image (per page in most cases) plus additional fixed charges for special services, such as binding.

EZ-Backup, as its name implies, was the service that could back up the disks from any workstation on the network whether running MacOS, Windows, or Unix. Paul Zarnowski, a member of the Systems Services group in Computer Resources, developed this service. Zarnowski was looking for an automated way to back up the servers in the computer room in CCC.⁹¹ At about this time there were almost 50 servers in the server farm, each with its own backup procedure, which had to be performed daily to ensure against loss of information. It was driving the operators to distraction having to manually perform this many independent daily backups, each with its own operating system and different database technology. In addition, others on campus were looking for a similar solution for their backup needs. In these early days of workstation computing, individual owners were responsible for backing up their own critical data by putting data on diskettes,

storing those in some safe place, and following routine manual procedures to stay current. As workstations started to be clustered into LANs, the task of backup was usually assigned to the LAN administrator, who used either the same procedure or newly available software and larger-capacity storage to back up all the workstations on a LAN using the LAN server. As more and more of the day-to-day work of the university was done on workstations, and as file retention and backup became more important to uninterrupted operations, EZ-Backup provided a university-wide solution.

At about this time a number of software products for backup came on the market, and IBM’s ADSM product—ADSTAR Distributed Storage Manager—was selected for EZ-Backup. It had the advantage of accommodating MacOS, Windows, and Unix operating systems along with a broad range of other options that fit into the Cornell environment quite nicely. A secondary but important feature of ADSM was that backups could be defined and controlled by the workstation owner, including specifying the type of archiving and the length of time a backup would be saved. A further influencing feature was that data was compressed at the workstation and backups were only incremental, that is, they only backed up those files that had changed since the last backup. This combination made it possible to transmit this information over the campus “Internet” without any noticeable performance impact. Last, EZ-Backup offered the advantage of being off-site, that is, in a different location than the primary data source. EZ-Backup started as a fee-based production service in January 1995, with an initial load of 95 computers (nodes) and storing 129 gigabytes of information.

Distance Learning

The concept of distance learning, that is, the use of computers and networks to conduct education at remote sites, started to be discussed in earnest in many parts of the university in the mid 1990s. Videoconferencing (using CU-SeeMe or satellite uplinks and downlinks) had been used for particular events and now was entering into classroom use. For example, in April 1993, CU-SeeMe software developed by CIT staff was used for an international videoconference.⁹² The Cornell Law Schools’ Legal Information Institute, the College of Human Ecology,

⁹⁰ Jim Doolittle and Sharon Marcus, “The Xerox Docutech Redefines Publishing Technology,” *CITnews*, summer 1992.

⁹¹ Paul Zarnowski, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, November, 2000.

the Veterinary College, and the School of Industrial and Labor Relations (ILR) were some of the first users and promoters of this technology. The ILR School first implemented the use of compressed two-way video instruction in 1991 for its master's program in New York City and for international education in 1992.

GateD

In September 1995 the responsibility for GateD software was transferred to the Merit Computer Network in Ann Arbor, Michigan, and the Cornell-supported GateD Consortium faded away. In effect, the Cornell GateDaemon Project became the Merit GateDaemon Project. The Merit GateD Consortium intended to foster the same kind of practices that Cornell had followed—using membership fees to augment grant and proposal funds to support the project.

During the five years preceding this transfer, the responsibility for GateD activities had shifted from the Theory Center to CIT as the Theory Center phased out its network activities. Martyne Hallgren continued to direct the GateD Consortium during this period. In addition to the support from membership fees, Scott Brim was able to secure a two-year cooperative agreement with NSF in the amount of \$315,918 to support the project. At the end of this grant, with the shift of technical staff to other Cornell projects, an evaluation of continuing the consortium found the transfer to Merit as the best choice.⁹³ The Merit GateD Consortium continued in operation for an additional five years. When the consortium disbanded, several of the Merit staff organized NextHop to continue a trend that had been occurring on the side: to produce hardened, extended, commercial routing software based on GateD.

Misuse of Networks: Network Use Policy

With increasing use of the network and associated services came increased misuse. In November 1995, Barbara Krause, Cornell judicial administrator, rendered a judgement against four freshman students at the university who had authored and distributed an e-mail listing 75 reasons women should not have freedom of speech. She judged that the authors of the "75 reasons" e-mail had not engaged in sexual harassment and did not misuse the university's computer resources.⁹⁴ She also concluded the students did not violate the Campus Code of Conduct. Her reasoning was that the students distributed the e-mail to a handful of friends whom they believed would not find the contents offensive. Apparently one of the recipients sent a copy on to others and soon it spread across the country and to Canada. Although some of the content offended and angered and distressed many people in

the Cornell community and beyond, Krause noted that by far the widest distribution came from people who were offended by its content rather than from the four authors. Nonetheless, it was another case that brought undesired attention to Cornell.

Close to Thanksgiving 1995, CIT issued warnings of an e-mail hoax. An e-mail purporting to be a confidential memo from Krause discussing the recent action involving the "75 reasons" e-mail was sent by a group calling itself OFFAL—Online Freedom Fighters Anarchist Liberation. They posted messages to various electronic bulletin boards claiming responsibility for the hoax.

Notices about electronic junk mail ("spamming" became the favored term) surfaced as an issue in 1996 as the amount of unwanted and unsolicited e-mail increased. Retaliation in the form of flooding the offending site with numerous or large messages (that is, "mail bombing" the site to disrupt their operations) was discouraged as a violation of the Responsible Use Policy.

In December 1995 CIT issued a new Responsible Use of Electronic Communications Policy. In announcing the policy, Marjorie Hodges, the policy adviser in the Office of Information Technologies, noted that Cornell's policy did deal with the issue of acceptable and nonacceptable behavior in the networked world. At the time Hodges commented that Cornell, whose policy had been in effect since 1994, was getting many calls from other institutions that were just starting to worry about the issue.

In early 1996 Pat McClary of the University Counsel's office commented that copyright infringement was widespread on the Internet, bringing attention to the issue at Cornell.⁹⁵ Her point was that the same rules that apply in the print world are applicable to other media and that the Cornell community had to take care not to violate copyright law, also a violation of the Campus Code of Conduct.

Business Systems

While deliberations about the future of Cornell business systems were taking place as part of the Strategic Planning Initiative, the Merten Review, and the deliberations of the Administrative Systems Planning and Advisory Committee (ADSPAC), new and important system changes and additions were introduced during

⁹² "International videoconference uses computer software developed by CIT," *Cornell Chronicle*, April 29, 1993.

⁹³ www.ifm.liu.se/~peter/doc/gated/node8.html

⁹⁴ "Judicial administrator renders decision in offensive e-mail case," *Cornell Chronicle*, November 30, 1995.

⁹⁵ "Copyright infringement on the Internet widespread, say CU lawyers," *Cornell Chronicle*, March 14, 1996.

this mid-decade period. Several of them were built on the Bear Access infrastructure, others just streamlined services, and others set new directions using data warehousing technologies to improve access to business data.

CoursEnroll

In the fall of 1995, online registration was introduced using CoursEnroll, a feature added to Just the Facts. Students no longer had to stand in line but could access CoursEnroll to register and enroll in courses from any computer, on or off campus, and could at any time determine their status and schedule. Improvements were made during the following year so that as early as mid-August, before their return to campus, continuing students could access their course schedules and find out if they were registered. With course selection for the coming semester completed the previous spring, the only thing left for students to do was check their schedules and add or drop courses. (Add/drop remained a manual process until 2003.) The improvements in systems had cut the spring enrollment and registration period from four days to two days and with a lot less hassle for all. The key was advance preparation with systems that supported the automation of the process.⁹⁶

In January 1997⁹⁷ David Yeh, assistant VP for student and academic services, said:

With all the streamlined processes we have on campus now—classroom scheduling, CoursEnroll, faculty advisory access to student data, SCAMP (the Student Communications and Mailing Project, which streamlined mailings of information to students), student employment—it is hard to realize that only a few years ago we were still doing all this by hand. Registration used to be a three-day ordeal and it is now a four-hour process. Eventually it should take as little as 20 minutes.

Clearly significant improvements had been accomplished. According to a personal communication from Yeh, his view was that SCAMP^o was one of the more successful QIP initiatives that continued as a practice until the end of the decade.

Employee Essentials

Less spectacular but nonetheless evident improvements in business systems were also being made on other fronts. In February 1995 Employee Essentials was first released, enabling faculty and staff to see their group life, health, and accident insurance cover-

age and some address information. The next release in December 1995 displayed more complete address information and allowed faculty and staff and student employees to update some information themselves, including home address and pertinent campus information such as work mailing address, work location, telephone and fax numbers, e-mail address, and emergency contact information. These first releases were limited to users with Macintosh computers (going against the industry norm of first releases being on Windows systems) and in a short while were available for Windows systems as well.

Project STOP Stops

By late 1994 Project STOP (to move Adabas applications from the CornellC mainframe to other computers) came to a stop itself and went in two different directions. One direction that had started earlier was the Information Access Project (IAP) to develop a distributed information access environment. This came about because increased contacts with administrative offices led to the conclusion that administrative staff needed improved access to existing data. The other direction was to continue to provide acceptable service conditions for the Cornell Library system and Adabas applications on CornellC; all efforts to offload work from CornellC had come to a dead end, and an increase in capacity was needed.

Information Access Project (IAP)

The Information Access Project was led by Bob Cowles, with the assistance of Tom Boggess and staff from the Computer Resources and Information Resources divisions. In addition to creating a distributed information access environment and the supporting infrastructure, IAP took on the objectives of educating the community about distributed computing and service expectations and promoting the re-engineering of processes using client-server technology.

Once Project STOP stopped, STOP became the acronym for avoiding the “senseless transfer of obsolete programs,” and efforts were directed toward creating a new information access environment. In building the new environment, the project team was able to build on all the work that had been done investigating the uses of server technology in support of administrative applications, and it was considering ways to exploit the Mandarin technology for developing applications. By early 1995, over 200 staff had attended a

⁹⁶ “Streamlined enrollment, registration, and orientation offer welcome start,” *Cornell Chronicle*, September 5, 1996.

⁹⁷ “CU students who once stood in line now go online,” *Cornell Chronicle*, January 23, 1997.

two-day class on introductory distributed computing concepts, and CIT gained a better understanding of their needs.

Data Warehousing Activity

One of the more successful and longer-lasting parts of the Information Access Project was the effort to develop data warehousing services and support,⁹⁸ including the hardware and software, database administration functions, and technical support. Oracle was used as the database management system and GQL from Andyne was used as the query language. The response to this service offering was immediate, and by late 1994 Randy Naegely had implemented a prototype data warehouse for the payroll system (personal communication). By early 1995, the payroll warehouse was in production, and prototype developments were under way for Human Resources/Payroll/Budget, Accounting, the Graduate School, and the Utilities Data System. All these were placed into production later in the year. Development ended at about the time the staff was assigned to Project 2000, which was going to provide the new and future business systems environment.

CIT Mainframe “Refreshment” Project—New CornellC

One of the key blocks to the removal of the CornellC mainframe from service was that a Unix-based library system was not going to be on the market and available for installation at Cornell before January 1995. In addition, moving even a portion of the Adabas applications was estimated to be a very expensive, long-term project. The ways Cornell had chosen to

exploit the power of Adabas, the way the applications were written, and the way Adabas was implemented on other alternative platforms made it difficult, if not impossible, to move to non-IBM servers.

The need to increase computing capacity at the time is best displayed in Figure 6, which shows the growth of the Adabas load. It was expected that the trend would continue given all the activity that was already in the pipeline or being discussed.

To determine what action to take, a project team of CR and IR staff, under the direction of Jim Doolittle, made a case for acquiring another more powerful IBM mainframe. To avoid using the term upgrade, the project was referred to as CIT Mainframe Refreshment, which in part included upgrading the magnetic disk storage and capitalizing on an early upgrade to new magnetic tape drives using a robotic system.

After evaluating the options, the decision was made to install an IBM 390 Parallel Enterprise Server, model 9672-R32, system to replace the 3090-200J and increase computing capacity by 30 to 40 percent. The 9672 was a more powerful mainframe in a smaller box (one-fourth the footprint) with significant savings (estimated to be on the order of \$200,000 per year) because it was air cooled rather than water cooled. As in the past, the conversion was done over the winter holiday break, in December 1995. To make sure that this system would last five years and accommodate the expected 15 to 20 percent yearly growth in Adabas commands, steps continued to be taken to remove

⁹⁸ Later, CIT would reference these single central office databases as data “marts,” as the term “warehouse” was used to refer to a compilation of data from multiple office databases.

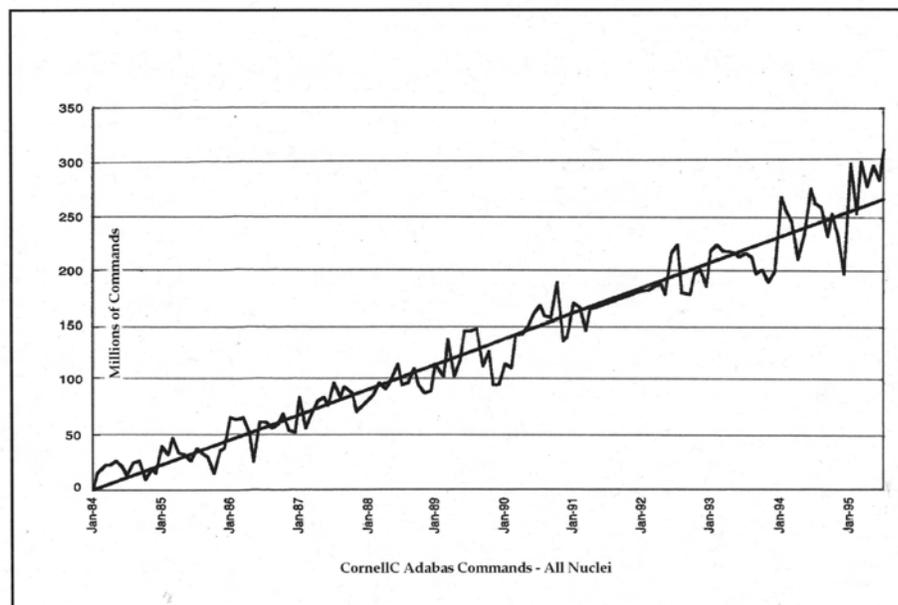


Figure 6. Growth in CornellC Adabas commands, 1984 to 1995

load from this new CornellC and to implement all possible other improvements. In 1995 one of the load components removed was BITNET. Until then, CornellC had been a main node of BITNET, routing considerable traffic around the network and using an estimated 5 percent of the computing power for this purpose.

Contributing to improved mainframe or enterprise server performance was the addition of fiber-based ESCON channels with a transfer rate of 17 Mbytes/sec—almost four times faster than the 4.5 Mbytes/sec transfer rate of the “bus and tag” copper wire cables they replaced. A further benefit of the fiber channels was the shrinkage in the diameter of the connecting cables, from over 1 inch to 0.25 inch, which improved the air flow to the hardware components because less space was taken by the cables in the airflow cavities under the raised floor. The staff, who no longer had to manhandle the heavy older cables during hardware reconfigurations, welcomed the significant reduction in weight of the fiber cables.

As part of this refreshment project, all the installed IBM 3380 magnetic disks were replaced by IBM RAMAC magnetic disks, offering improved performance and continuous availability. These new disks operated at new ESCON channel speeds and had large intermediate storage caches. The underlying RAID 5 technology was the basis of the performance and availability improvements, and the increased recording density provided the equivalent of 90 gigabytes of storage space in a box the size of a four-drawer file cabinet.

In May 1995, several months before the “refreshment” of CornellC, an IBM tape robotics system—Total Storage Enterprise Automated Tape Library 3494—was installed. The plan was to move away from the magnetic tapes/cartridges being used for storing intermediate data and items such as print files and instead use online disk storage that had shorter retrieval times for these purposes. Besides, disk storage was becoming less expensive as a result of technology developments and competitive market conditions. The 3494 would be used to migrate and store MVS datasets that were saved on disks to this less expensive and automated device. To fully exploit this automated storage, IBM’s Hierarchical Storage Management (HSM) software system was installed to automate dataset migration from active disk storage and take care of long-term data retention. When installed, the 3493 had an initial online accessible storage capacity of 600 cartridges, yielding a total raw and uncompressed storage of 480 Gbytes. Since the 3494 robotic system used the newer and higher-capacity 3490E cartridges, the unit was not compatible with the 3480 cartridges that had been in use since the late 1980s, having them-

selves been installed to replace the 3420 tape reels in use for decades by that time

After all the above hardware changes, the amount of computer room floor space used was substantially reduced. According to Jim Doolittle, the new disks, the new tape robotics system, the removal of individual tape drives, and the installation of the new IBM 390 system/model 9672-R32 mainframe freed up about 50 percent of the floor space. In 1996 this space was used to relocate the Production Control staff from the first floor of CCC and for the next relocation of the Network Management Center/NOC from Maple Avenue.

With the move of the NOC to the CCC computer room and the installation of new disk and tape robotics technology, the opportunity was there to begin merging the NOC and computer operations functions. With the use of new technology, users were able to eliminate significantly the need for on-demand tape mounts and large print jobs from weekend processing. This allowed the operations staff to take on the responsibilities of the NOC during weekends and, over time, to merge into one group managing both services.

Information Services

Information service offerings continued to change and improve with the availability of new technologies and to take advantage of exploiting the World Wide Web. Cornell was at the forefront of using networks for videoconferencing with the development of CU-SeeMe by CIT staff. CUINFO, CIT training, and information dissemination changed to WWW Netscape offerings while CHAOS Corner ended.

CU-SeeMe

In December 1991, attending a networking conference, Dick Cogger saw a rudimentary demo of live video transmitted across the country via the Internet using a Sun workstation. He wondered if the same could be done with less expensive equipment and began experiments back at Cornell on the Macintosh. In the spring of 1992, Tim Dorcey, a statistical consultant in CIT, approached Cogger, looking for an interesting project on which he could develop his programming skills, and so a partnership began that subsequently led to CU-SeeMe.⁹⁹ Dorcey first coded what was then called MacVideo as separate applications; one to send and the other to receive live black-and-white video images on the Macintosh. The enabling innovation in Tim’s work was a new compression

⁹⁹ “Internet Video: I See You, You See Me,” Brian Gallagher, *Boardwatch Magazine*, June 1994.

algorithm, coded in assembler, manipulating several pixels in parallel with register arithmetic. The code was able to compare all pixels in subsequent video frames, searching for differences greater than a threshold, giving extra weight to those near in the image, all with an average of two machine instructions per pixel (personal communication from Dick Cogger). While the images were small, and the telephone was used in parallel for audio, this first version was vivid proof the Internet could be used for more than e-mail and transferring files.

With support from Dave Lambert and VP Lynn, these early demo programs were combined into a two-way application, now given the name CU-SeeMe, which was passed around, via the Internet, among acquaintances in the university community. In late 1992 a program officer at the NSF asked if CU-SeeMe could be extended for multiparty conferencing and used to support a project of his involving primary schoolchildren. To implement the needed multiparty conferencing, Cogger came up with the reflector concept. The “reflector,” a separate program to run on a Unix server, was developed by John Lynn, and CU-SeeMe’s rudimentary protocols were extended with design critique from Scott Brim. For a multi-party conference, each participant was connected to the reflector, which copied each video stream to the other participants. Each user could show up to eight windows on screen and choose which participants to view at any given time.

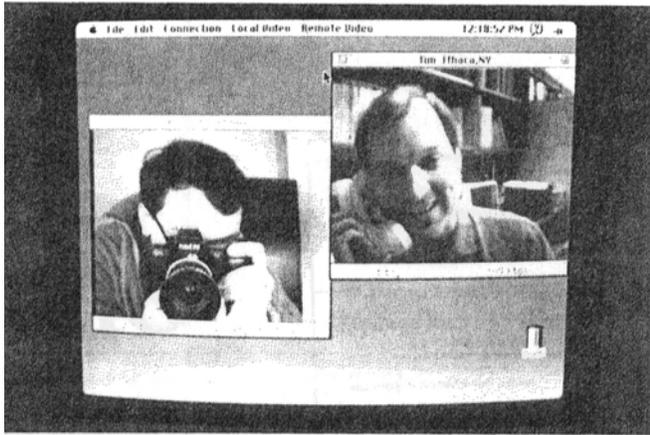


Figure 7. CU-SeeMe demonstration by Tim Dorcey

CU-SeeMe faced—and passed—its first major test, supporting an international videoconference between students in California, Tennessee, Virginia, and London, England, who were participating in the “Global Schoolhouse Project” sponsored by NSF.¹⁰⁰ Following this successful use of CU-SeeMe, Cogger and Brim were able to obtain a three-year grant from NSF to support an agenda of improvements to the program. For this work the technical team was expanded and acquired collaborators. Steve Edgar, like

Dorcey before him, coding on his own time, ported Dorcey’s compression code to write the first version for Windows. Collaborators included Steve Erde at the Cornell University Medical College (CUMC), who had an interest in medical applications, and Charlie Klein from the University of Illinois, who contributed audio code, both of whom were supported by the grant. Further refinements included improving the reflector, extending the protocols, and adding audio and color-image support, a plug-in architecture, still-image transfer, a text-chat facility, bandwidth management, and a number of other items. Eventually, besides Cogger, Dorcey, Brim, Lynn, and Edgar, the Cornell development included work from Rich Kennerly, Larry Chace, Aaron Giles, Jeff Han, Jill Charboneau, Tom Parker, Joy Veronneau, Pete Bosanko, Andy Wyatt, and others.

In 1995 CU-SeeMe was used to broadcast the entire speech of President Lee Teng-hui of Taiwan when he spoke at Cornell in August. The NSF team in Antarctica communicated to the United States via CU-SeeMe during the long months they were wintered in. A grad student defended a thesis via CU-SeeMe, and a team member joined an NSF funding evaluation remotely.¹⁰¹

In that same year, a CU-SeeMe consortium was formed with institutional and commercial partners, including CUMC, NYSERNet, Cisco Systems, and White Pine Software, to further extend improvements in CU-SeeMe. (At some point the names CuSeeMe, CU-See-Me, and CuseeMe came into use as the software evolved. We will stay with the original name.) In addition, the software was licensed to White Pine, who planned to offer an enhanced version as a commercial product and to act as a licensing agent for other commercial developers. Martyne Hallgren, executive director of the CU-SeeMe consortium, was quoted as saying: “CU-SeeMe software is to videoconferencing what Mosaic was to information retrieval on the Internet,” an indication of the success CU-SeeMe was experiencing.¹⁰²

At its peak, approximately in 1997, CU-SeeMe for both the Mac and PC-compatibles was in daily use by over a million users worldwide, was being enhanced by several programmers at Cornell and collaborators elsewhere, and was the subject of scores of stories in the media. In one story, a couple just married recounted originally meeting via CU-SeeMe. By this time, users

¹⁰⁰ “International videoconference uses computer software developed by CIT,” *Cornell Chronicle*, April 29, 1993.

¹⁰¹ Dick Cogger, *Oral Histories—Cornell Computing and Information Technology*, www.cit.cornell.edu/computer/history/, 2004.

¹⁰² “CU videoconferencing software is licensed,” *Cornell Chronicle*, June 22, 1995.

had the choice of obtaining the “freeware” version from Cornell or the commercial product from White Pine.

CUINFO—Gopher Phased Out

CUINFO continued to evolve with web technology. Gopher itself was phased out in mid-1996 when CUINFO services moved to the web. During the previous two years when CUINFO had been available using Gopher and the web, the use of Gopher had declined while the use of the web had increased substantially. Given this trend and the additional cost of maintaining two systems the decision was made to offer only the web-based service.

CIT Training and Publications—CIT Web Site

CIT training and publications and distribution of public information continued to evolve over the decade in response to changing technologies and users’ needs. In late 1995 the information about Cornell that had previously been in CUINFO now had its own web address, www.cornell.edu, and a web page, Welcome to Cornell University. With this technology, future students, and anyone else for that matter, could find out a lot of information about Cornell from a remote distance. Various other academic and administrative units on campus, the new so-called “content providers,” could link to this page and provide a comprehensive text and pictorial tour of the campus and all its different facets.

Later that year, in keeping with the “@” trend, a new monthly column, “@cornell.edu,” was started in the *Cornell Chronicle* to bring announcements of CIT activity and technologies to the *Chronicle’s* large audience. The tagline for the first column in the November 2nd issue read, “This is the first monthly column from Cornell Information Technologies. Columns will address issues of interest and concern to campus computer users.”

Chaos Corner Comes to an End

Chaos Corner, an electronic newsletter Bob Cowles started in 1991, ended publication in November 1996. The sign-off in Cowles’s words read:

Chaos Corner—Its Time Has Come. When we started *Chaos Corner* nearly 5 1/2 years ago, Dr. Chaos and I wanted to promote the Internet—to encourage people to become more aware of the information resources available out there in cyberspace. It appears that task is no longer necessary, so with great thanks to our loyal readers we declare this to be the last issue of *Chaos Corner*.

My thanks also go to Cornell University for providing the computing resources that made *Chaos Corner* possible, and a salary that helped pay the bills—it will all be missed. Dr. Chaos and I are moving on to other things. We’ve had the opportunity recently to dabble a bit more in web technology, so who knows where we may pop up again!

Digital Library

In December 1993, CIT installed an IBM SP1 computer in CCC for expected use by the Digital Library project. The SP1 was obtained as part of a Shared University Research (SUR) grant from IBM. Expectations were that the SP1 would serve as the front end to a file server for materials from the digital library, or as a fallback plan, it would serve as the platform for the new library management system being considered at the time as part of the STOP project. The system was also thought to be capable of serving as an alternative platform for processing Adabas transactions, also part of the STOP project.

As it turned out, the system was not useful for any of the above applications; its deployment as a general-purpose server did not exploit its potential. CIT explored other uses. Discussions were held with the Theory Center and with other research centers at Cornell regarding other possible uses of the SP1. In the end the SP1 was transferred to the Department of Computer Science in 1995 to be used in its research program.

TULIP Project—EPOCH Mass Storage Phased Out

From 1993 to 1996, CIT continued to expand the use of the EPOCH mass storage system in support of digital library projects. An aptly named TULIP project was being considered to scan and reproduce scientific journals online in conjunction with the Dutch-based publishing corporation Elsevier. This was to be a collaborative project involving CIT, the Cornell Library, and the Department of Computer Science. As a result of a joint study with the vendor, a second EPOCH server was installed in CCC and was to be used to store these images. In exchange for a no-cost lease of this second server, CIT agreed to write a Mac client that would allow the EPOCH server to be used as a secure backup for Macintosh computers.

Unfortunately, by the time all the arrangements had been made and the work defined, interest in the EPOCH box was waning due to other developments in data storage technology and the departure of VP Lynn. With neither Lynn’s leadership nor a clearly defined need, the EPOCH boxes were removed in 1996.

Instructional Computing

During the mid-1990s, the use of information technology in support of instruction became a major focus of FABIT, the Faculty Advisory Board on Information Technologies. Following their 1995 issuance of a report on learning technologies services, additional funds were made available for the installation of information technology support services in teaching spaces. CIT organized a unit dedicated to assist in these classroom upgrades. As part of the effort, CIT upgraded the equipment in several laboratories.

FABIT Report and Recommendations

In early 1995, FABIT issued its report “Planning for Learning Technologies Services.”¹⁰³ Many of the recommendations were not one-time actions but ongoing efforts toward upgrading and improving services. FABIT’s recommendations included

- implementation of the recommendations of the 2001 report that called for the creation of a central technical organization for the entire campus as well as support teams in each academic unit.
- establishment of special funding incentives for faculty to create innovative ways to use new technology.
- improvement of student access to computing facilities and networks.
- an upgrade of all teaching spaces to a minimum standard defined as Tier 1—an overhead projector with a liquid crystal display to project computer displays and a network connection. There would be fewer Tier 2 sites, which would have phone jacks for teleconferencing and a satellite download. Tier 3 sites, with state-of-the-art technology and individual student terminals, would be even fewer.

While the recommendations came at a time when the university was facing economic cutbacks, the sense was that these requirements needed to be laid out as a basis for developing future plans and funding sources. It was especially noted that many of the improvements would have to be a joint effort between the colleges, the university, and CIT. At the time, acting VP Lambert noted that CIT had already started to upgrade some of the student computing labs and improve access to computers and networks.

Student Laboratory Upgrades

In 1996, a year after the FABIT report was issued, CIT announced major upgrades to three instructional labs (those dedicated for use in conjunction with classes)—in Martha Van Rensselaer Hall, Upson Hall, and Sibley Hall—following the FABIT recommendation that computer labs be upgraded on a three-year cycle. Power Macintoshes and IBM Pentium 90s replaced the older and slower equipment. A trickle-down strategy was followed so that equipment from these labs replaced even lower-end equipment in the general labs (multipurpose use), and as a result the whole campus saw some benefit. At this point CIT had general labs in Noyes Community Center, Robert Purcell Community Center, Carpenter Library, Uris Library, and Mann Library.

In support of FABIT’s recommendation to upgrade teaching spaces to three different tier levels of technological capability, the provost provided several rounds of funding starting in 1996. To solicit support for projects, academic and other units had to provide matching funds, that is, real money or staff time.

FABIT-Supported Teaching Space Upgrades

In July of 1996 Provost Don Randel committed \$2 million to FABIT for teaching space upgrades to be carried out between 1996 and 1998. The FABIT Subcommittee on Technology for Teaching Spaces requested that all academic units prepare proposals for review and approval by FABIT. By this time FABIT had committed \$735,000 to the first round of proposals, which had been quickly solicited in April for implementation during the following year.

In support of this increasing activity, Academic Technology Services organized a Classroom Technologies Group, which developed classroom design guidelines for upgrading the technology in classrooms and acted as an advisory group to academic units and the subcommittee.

CIT Activities

CIT experienced a number of significant changes in the location of staff and the hiring of new senior management staff in 1995–96. With the confirmation of Dave Lambert as vice president for information technologies, the CIT organization started to change to reflect his priorities.

¹⁰³ Faculty Advisory Board on Information Technologies, Cornell University, “Planning for Learning Technologies Services,” January 9, 1995.

CIT Services and Staff Move to Maple Avenue

In December 1995 and January 1996, over 165 CIT staff moved to 120 Maple Avenue, a project for which planning had started in 1994. The Information Resources (IR) staff started preparing in early December to vacate the fourth floor of CCC. They and the IR staff at Thornwood Road moved during the winter holiday break. Both groups occupied the first floor of 120 Maple Avenue—the first time in over five years that all the IR staff were together in one place. The Distributed Technologies (DT) staff moved from CCC to the second floor of 120 Maple Avenue.

After the move, the Network Resources (NR) staff, who had been divided between Caldwell Hall and Surge 3, were all located on the ground floor of 120 Maple Ave., another consolidation of CIT staff into one place. The Network Management Center/NOC was also moved to 120 Maple Avenue. Along with some staff in 110 Maple Avenue, basically three entire divisions of CIT were now housed in 110 and 120 Maple Avenue.

Even with the completion of these moves, CIT was still split over three locations: Maple Avenue, CCC, and Thornwood Road (microcomputer sales and repair). The fourth floor of CCC was turned over to the Student and Academic Services' Learning Skills Center, Instruction Support Unit, and Technology Services.

Lambert Reorganizes CIT

After Dave Lambert was confirmed as vice president for information technologies in July 1995, he set about reorganizing OIT and CIT with his own ideas for how the organization would meet expectations. The plan had been developed during his interim time as acting vice president and was announced soon after his appointment. Basically, the OIT and CIT separation was maintained, but CIT was collapsed into three major divisions—academic computing, administrative computing, and network and computing services. This consolidation was done in part to reduce the number of directors—and as a result, the previously more distributed responsibilities—and to effect cost reductions. The goal was to reduce the overall CIT operating budget by at least \$1 million and to release those funds for innovations. In addition to the reduced number of directorships, all the clerical and support staff were reorganized into a centralized Integrated Business Service Center (IBSC) to reduce costs and streamline operations. Combining such small separate administrative support groups was the “model office” movement being promoted across the campus to create service centers serving more than one organizational unit and in this way achieve significant efficiencies in

administrative operations. Liz Colucci was appointed head of the IBSC.

In January 1996 two new directors were appointed to the three new director-level positions. Helen T. Mohrmann was named director of Administrative Systems and Distributed Technologies (ASDT). ASDT combined the previous Information Resources and Distributed Technology divisions in VP Lynn's organization. Mohrmann was to report jointly to Lambert and senior VP Rogers and serve as a member of both management teams. Her primary responsibility was deploying administrative systems and building the infrastructure of the client-server architecture for new systems.

Ann Stunden was named director of Support Services and Academic Computing. Stunden was responsible for information technology services supporting the digital library, public information services, support of desktop and network operating environments, support of faculty to bring technology to the teaching and learning environment, and improved technology support for researchers. Later Stunden's division, renamed Academic Technology Services, was organized as three groups: Academic Technology, Technology Integration and Communications, and Technology Support Services.

A series of “town meetings”—open meetings to which all CIT staff were invited—were held during the spring to discuss the new CIT organization and to deal with any questions and concerns. It took almost another year for the third director to be appointed to complete the expected level of staffing.

In April 1997 Peter M. Siegel was named director of Network and Computing Systems. Prior to this appointment, Siegel had been executive director and director of corporate partnership for the Theory Center. In his new role, Siegel was responsible for the university's voice, data, and video networks, central computer systems, and file servers.

After all the changes were made, the OIT/CIT organization looked as shown in Table 9.

Cornell Policy and Law Institute Formed

During 1996, at the urging of Worona^P and Marjorie Hodges Shaw, VP Lambert created the Cornell Computer Policy and Law (CPL) Institute, expanding on the initiative taken by VP Lynn in 1995 when he appointed Shaw as a computer policy adviser in OIT. Shaw, a graduate of the Cornell Law School, had been Cornell's judicial administrator and was familiar with many of the recent policy violations at Cornell. This familiarity had given her a useful perspective as she started formulating and documenting policies.

Hodges and Worona became co-directors of the Cornell CPL Institute and carried out extensive work-

Table 9. Cornell Computer Services Organization Chart, 1996 to 1997

Academic Technology Services	Ann Stunden
• Finance, Planning, Administration	Pat Nelson (Searles)
• Academic Technology	Carrie Regenstein
• User Technology Services	Jim Lombardi
• Technology Support Services	
Administrative Systems	Helen Mohrmann
• Finance, Planning, Administration	Agelia Dumas
• Accounts, Planning and Architecture	Rick Jones
• Development Management	Mark Mara
• Technical Services	Nancy Van Orman
• P2K Technical Project Management	JR Schulden
Network and Computing Systems	Peter Siegel
• Finance, Planning, Administration	Greg Lee
• Advanced Technologies	Dick Cogger
• Networks and Operations Support	Jim Doolittle
• Network/Computer Systems Engineering	Carl Remon
Integrated Business Service Center	Liz Colucci
Office of Information Technologies (OIT)	
• Finance, Planning, Administration	Alan Personius
• IT “Evangelist”	Steve Worona
• Computer Policy Program	Margie Hodges Shaw
• Security	Barbara Skoblick
• Human Resources/Organizational Development	Judy Hart
• Communications/Public Relations (1/4 time)	Cecilia Cowles

shops and seminars both at Cornell and elsewhere in the country after they allied themselves with the Department of Continuing Education at Cornell. The institute also set up a web site that linked to over 700 institutional computer use policies, providing a worldwide resource to guide others in creating their own legal and ethical use policies, a topic of considerable interest given the extensive worldwide reach of any single workstation.

CIT Training Becomes Technology Training Services

In 1996 the CIT training program was renamed Technology Training Services with the intent of making major changes in the courses offered. Courses would now focus on the web: navigating the web, creating web pages, using hypertext markup language (HTML) versus Adobe Acrobat (PDFs) for web publishing, setting up web servers, digitizing images, and incorporating video and sound into web pages. Workshops on using word processing, spreadsheets, database programs, and desktop publishing applications were discontinued, given that alternative training was often available in the Ithaca area or by special arrangement.

The rapidly increasing demand for this web-oriented training quickly outstripped CIT’s ability to supply it. In the fall of 1996, the training approach shifted to “train the trainer”—technology support staff around campus would be trained in the new and emerging technologies by CIT staff or other experts and then be expected to train their own local staff. CIT also embarked on creating a campuswide team of trainers to create a set of training services and materials, including online tutorials, for the entire community. Additional information was available on the Technology Training web site (<http://training.cit.cornell.edu>).

Year 2000 (Y2K) Millennium Issues

During the mid-1990s, people all over the world began to worry about the year 2000 problem, or Y2K “glitch” or “millennium bug,” which arose due to two-digit abbreviations for years (for example 58 for 1958) used on the early mainframes when main memory storage was very expensive and limited.⁹ This practice started in the days of punch cards, when only the last two digits of the year were stored in the formats dd/mm/yy or yy/mm/dd, thus conserving 2 of the 80 columns on

a card. This shortcut worked fine as long as the years increased to 1999, but when 2000 would appear as 00, sorting sequences would not work, as 00 would be interpreted as 1900. This shortcut had become a standard approach in software development.

The great question was: *What* would stop working when December 31, 1999, rolled into January 1, 2000? The problem applied to all kinds of software on large and small computers, and was also thought to be hard-wired into computer chips installed in all kinds of devices. It was predicted to be a nightmare and even gave some support to those who were predicting the end of the world. Emotions aside, CIT started taking preliminary steps to address this issue, including simply raising awareness on campus. More aggressive actions came later in the decade.

CIT Receives Recognition and Awards

The following subsections describe a number of awards CIT received in 1996 for outstanding achievements in information technology. One significant award came from CAUSE (College and University Systems Exchange) for networking activities; another for VP Lambert as a network innovator. Indirectly, Bill Gates, founder of Microsoft, credited Cornell for significantly influencing Microsoft product development directions. Lastly, several Cornell graduates were recognized for starting new software companies built around the World Wide Web.

CAUSE Award for Excellence in Campus Networking

In May 1996 Cornell received the 1995 CAUSE Award for Excellence in Campus Networking. Jane Ryland, president of CAUSE, came to the campus to present the award to President Hunter Rawlings. The award recognized CIT's exemplary campuswide network planning, management, and accessibility as well as effective use of the campuswide network to enhance teaching, learning, research, administration, and community service. Cornell was one of 17 colleges and universities that completed applications to be recognized in this way. A statement accompanying the award said, in part: "Cornell has evolved into an Information Age institution, where 95 percent of the faculty, almost 100 percent of the students, and 90 percent of the staff are connected to and use the campus network."¹⁰⁴

Bill Gates Recognizes Cornell Achievements in Networking

A somewhat different form of recognition also came in 1996, from Bill Gates, legendary CEO of Microsoft Corporation. According to the *Cornell Chronicle* of

October 3, 1996,¹⁰⁵ "Gates had pretty much left the Internet out of Microsoft's corporate strategy until Steve Sinofsky '87 sent him an e-mail with that terse announcement ('Cornell is wired')." According to the *Chronicle* story, Sinofsky was snowed in during a recruiting visit to Cornell and spent a day on campus. He visited the libraries, talked with students, spent almost the full day with Steve Worona, who demonstrated the features of Bear Access and Just the Facts, and visited a CU-SeeMe event going on that day. He was clearly impressed by the degree to which information technology was being used at Cornell, and on returning to his office he sent Gates his now famous note.

In *Business @ the Speed of Thought*, published in 1999, Gates comments on this incident.

In his e-mail report Steve marveled at how 'wired' the school was. About a third of the students had PCs, some school departments provided PCs, and kiosk PCs were available in public spaces. E-mail use by students was close to 100 percent. Many of Cornell's instructors were communicating with students online, and students were pestering their parents to get their own e-mail accounts. A wide variety of information, including much of the Cornell Library catalog, was available online. A student could view her current course schedule, her previous grades, her outstanding accounts, financial aid information, and a directory of the school community online. Many faculty members were communicating with students online and used online chat services to collaborate with each other. There was a "huge movement" to make all sorts of information available to students via the web. Steven even saw real-time videoconferencing over the Internet.

What struck Steve Sinofsky was how thoroughly this technology had become integrated into campus life "in practically no calendar time" and how students took it totally for granted. He said that for students "the online services are as ubiquitous and expected as regular telephone service" and that "this pace of change in information access is faster than for any other technology I have seen in my lifetime, including the personal computer itself. Students were even complaining that they couldn't sign up for classes online."¹⁰⁶

¹⁰⁴ "Excellent networking," *Cornell Chronicle*, June 6, 1996.

¹⁰⁵ "Computer lore; 'CU is wired' led Microsoft to reverse course," *Cornell Chronicle*, October 3, 1996.

¹⁰⁶ Bill Gates, *Business @ the Speed of Thought, Succeeding in the Digital Economy*, Warner Books.

There is probably no better testament to the transformation that had taken place at Cornell. Further, this recognition from Bill Gates was a credit to the inspired and spirited, yet difficult and at times overwhelming work, of the leadership and the staff in CIT over the past 20 years.

Cornell Graduates Start The Globe.com, CourseInfo/Blackboard

Kudos also came to Cornell through the publicity garnered by several students who started Internet companies. In 1996 Stephen Paternot and Todd Krizelman were recognized nationally for their company WebGenesis, which was operating out of offices in Collegetown. Their web site was called The Globe (www.theglobe.com) and at the time was reported to be the largest Macintosh web site, with nearly 150,000 visitors per month. It all started in 1994 in their dorm room, in pre-Netscape days, when they first experimented with setting up an information site. The site grew into an active “chat room” in which users could engage in real-time conversation and bring graphs and images into their conversation.¹⁰⁷ By 1997, it was based in New York City attempting to grow into a larger Internet company, one of the many so-called new-economy, dot-com companies. In 1998 when it went public, it held the record for the largest first-day gain of an IPO on NASDAQ, reaching as high as \$97 per share on the first day of trading. As part of the “dot-com bust” in the year 2000, the company started to falter; Krizelman and Paternot turned the company over to a more experienced hand and removed themselves from day-to-day duties. By 2001 the stock was trading at 25 cents a share and was de-listed. Various ways were sought to transform the company, but nothing seemed to work, and the company shrunk in size and activity although the name persists (information from a variety of news reports and articles).

In 1997 Daniel Cane, Class of '98, with several of his classmates and friends, founded CourseInfo. The idea behind the software was to allow instructors to provide web pages for their students by using a template rather than writing in HTML. Such a web site would give course announcements and links to pages about course content, assignments, and communication and would be easily developed with templates. The point was to keep programming simple for instructors who wanted to use “e-education” for their courses, making available course syllabi, rosters, sample exams, and links to other relevant web sites. Instructors also could include project chat rooms and discussion boards and students could link to their own personal information. The use of CourseInfo at Cornell grew quickly. In October 1996 it was estimated that as many as 25 courses were using the product. According to Caroline Hecht of

CIT Academic Technologies, in the fall of 1997 there were 125 course sites using what was then called the “Teacher’s Toolbox” (personal communication).

As with other student startups, the company that became CourseInfo^F started in a room in Cane’s Collegetown house but quickly grew to having 11 undergraduate and one alumni employee in the Student Agencies building in Collegetown.¹⁰⁸ Before the decade was over, CourseInfo was bought out by Blackboard and the offices were moved to Washington, D.C. The CourseInfo name continued to be used until version 4, and the product became known as Blackboard with version 5.

By the year 2003, there were roughly 2,500 sites using the commercial product. This is quite an accomplishment in just seven years, not to mention the large number of course instructors using the product across the country and around the world.⁵

Lambert Recognized as Innovative Leader

In October 1996, VP Lambert was recognized as one of 25 “Innovative Network Technology Drivers” by *Network Computing* magazine. The particular innovation recognized by the magazine was “cells in frames,” a method for running ATM over Ethernet networks. The article noted that this technology was “being developed by Dick Cogger and Scott Brim in CIT’s Advanced Technology Planning division.” Cells in frames would enable the Cornell data network to carry voice and eventually video at far higher speeds than was possible at the time without a heavy investment in new hardware. It was particularly satisfying to Dick Cogger that a magazine that tracked technology developments was paying attention to developments going on in universities.¹⁰⁹

Networking and Network Services

Increasing attention was being paid to improving network services to keep up with the quickly expanding use of the network for all the applications being developed or brought to the market and deployed at Cornell. The campus celebrated the 10th anniversary of the new telecommunications/telephone system, which was the forerunner of new network developments on campus. Networking continued to be a topic of considerable interest in the Ithaca area. The U.S. Congress passed the Computer Decency Act, which caused some concerns on campus about surveillance

¹⁰⁷ “A student-created company is the talk of the Web,” *Cornell Chronicle*, April 11, 1996.

¹⁰⁸ “Senior’s company helps to produce Web pages for college courses,” *Cornell Chronicle*, October 16, 1997.

¹⁰⁹ “Magazine cites CIT’s Lambert for computing innovations,” *Cornell Chronicle*, October 3, 1996.

of e-mail content. Viruses continued to be a concern as they were still being sent over networks and could cause potential harm.

Network Services

In late 1995, the EZ-Remote modem “pool” (a bank of modems reachable by dialing a single number) for off-campus access to Cornell services were upgraded. The cost-recovered HI (high speed) services went from 14.4 kbps to 28.8 kbps, while the number of modems was doubled for the no-cost LO services that continued to operate at 2,400 bps.

During 1995 there was a concerted effort to phase out the use of AppleTalk networking and convert local area networks (LANs) using this technology to Ethernet technology. EtherTalk (Ethernet for Apple LANs) had the potential to operate 140 times faster and so accommodate the increasing use of video-conferencing and of transmitting graphics as part of applications. By the end of 1995, over 180 AppleTalk LANs had been converted to EtherTalk. CIT phased out AppleTalk support in January 1996, a year later, although departments could continue using AppleTalk for their LANS for local print or file services, as long as they provided their own support.

Celebrating 10 Years of Cornell Telecommunications/ Telephone Installation

In 1996 CIT celebrated the 10th anniversary of the installation of its telephone and telecommunications systems. While the installed Cornell PBX (private branch exchange) system easily accommodated an increase from 11,500 telephones using 5,000 lines in 1986 to 16,000 telephones and lines in 1996, at no increase in rates, there were two other important benefits that came with this project. One was the switch from analog to digital transmission. The second was the installation of a new and contemporary wire plant, including optical fiber, with spare capacity for expansion.

During the period, AUDIX voicemail and features to personalize telephone operations were introduced, keeping up with contemporary offerings of PBX systems at the time. The PBX controller (an embedded proprietary computer system) had to be upgraded once, midway during these 10 years, from System 85 to Definity, to accommodate the increased use.

“Cells in Frames,” Asynchronous Transfer Mode (ATM), and Workstation Telephony

Two events came together in the 1995–96 time frame that led to several innovative network projects. One was called “cells in frames” (CIF), which attempted to exploit the availability of an emerging network tech-

nology—ATM, or asynchronous transfer mode—to increase network speeds without installing a new wire plant. The other was network telephony, which would take advantage of ATM and CIF on the Cornell networks to provide workstation-based telephone services and so avoid another telephone system upgrade.

ATM had the capability to change the way bits were moved over the network. ATM recognized the difference between transmitting an e-mail, for example, which didn’t rely on continuity of transmission, and transmitting a telephone conversation, which required the bits to arrive in a manner that presented a smooth conversation. The vendor-proposed approach to implement ATM would have required Cornell to upgrade its networks, which in turn would require new routers and a new ATM card in every desktop computer, replacing the current Ethernet card. Cogger’s proposed “cells in frames” would repackage the ATM cells inside larger data chunks called frames and thus use the existing Ethernet cards in all the workstations. Additional software and a small “\$100 box” would be needed at each workstation, but this was estimated to be a lot less expensive than buying 15,000 new ATM cards. Cogger and Scott Brim secured an NSF grant of \$710,153 to validate the cells-in-frames concept.

A short time later, Cogger and Brim secured another multiple year grant in the amount of \$1,216,155 to develop workstation telephony based on the CIF/ATM model, working in conjunction with IBM. This was a more complex project requiring the development of not only software but also special chips for the inter-connecting box. There was interest on the campus and in networking circles in this new use of networking technology. The campus plan, outlined in the *Cornell Chronicle*, was for field trials to start in 1997 and implementation later that year.¹¹⁰ To expose the plan more broadly, Cogger and Brim, both well known in the networking field, took the initiative to organize a Cells-in-Frames Alliance, which held a meeting in March 1996, attended by representatives from 35 companies.

Despite their potential, these projects did not lead to any concrete changes at Cornell or in networking technologies. ATM simply did not capture the market interest, in part because of the development of higher-speed Ethernet (Gigabit Ethernet) that increased network capabilities while maintaining the same underlying technology. The tradeoff between using a new technology with the current physical plant or staying with proven technology and upgrading the physical plant at higher cost favored the higher-cost, same-technology alternative. At Cornell the pressure to eliminate the non-Y2K-compliant IBM routers, no

¹¹⁰ “Campus plan links phones to computers,” *Cornell Chronicle*, March 14, 1996.

longer marketed by IBM, was also a motivating factor to stay with Ethernet technology in a subsequent network upgrade.

The workstation telephony grant effectively ended when the networking development staff left Cornell en masse later in the decade.

Viruses

Viruses continued to recur. In December 1996 there was a warning about the Irina virus, which was not a virus but a hoax similar to the Good Times hoax in 1994.¹¹¹ Nonetheless, these messages caused concern because simply reading the e-mail supposedly placed the virus on the workstation. Instead of simply passing along concern about possible viruses it was recommended that computer users first call the CIT Help Desk or another reliable authority. One such source was the IBM antivirus online web site, which provided both virus alerts and hype alerts. Hacking and hackers became a concern, as it was now possible for an ill-intentioned group anywhere in the world to do damage to any computer on the Internet.

The Computer Decency Act

The passing of the Computer Decency Act in 1996 raised some concerns about Cornell's response to the transmission of indecent materials over the campus networks. This act made it a crime to use a telephone or a telecommunication device in interstate commerce to make "any comment, request, suggestion, or proposal which is obscene, lewd, lascivious, filthy, or indecent." While it was Cornell policy not to ordinarily monitor private messages or files, Marjorie Hodges, CIT policy adviser on these issues, stated that such action could be taken if CIT were to become aware of behavior that violated the law or Cornell policies. Putting to rest some concerns that CIT had recently acquired some monitoring software, Alan Personius, acting director of network resources, assured the campus that such software was to measure traffic volume and not content.¹¹² A *Cornell Chronicle* article advised that "users should think carefully about what they transmit.... A long-standing rule of thumb is not to send anything across the Internet that you would not want to see attributed to you on the front page of the *New York Times*."

Electronic Mail

During the two fiscal years 1995 and 1996, electronic mail use at Cornell rose from 41 million to 69 million messages sent or received annually by the CIT-supported campuswide system. A new version of Eudora was in test mode by the end of the 1996 calendar year

to provide more features without having to change the hardware on either end. However, given the growth rate, there was increasing concern that additional servers, and perhaps additional network capacity, would be needed soon.

IthacaNet

The second annual Networking Tompkins County conference, sponsored by IthacaNet, was held in 1996.¹¹³ About 300 participants attended the conference held at Ithaca College. In the year between the first and second conferences, three private companies started offering Internet access (Lightlink, Baka, and ClarityConnect), and at least 100 businesses had home pages on the web. The conference theme was "Spreading the Net," focusing on ways to get more people online, and how to use the new capabilities. David Lytel, formerly a City of Ithaca alderman but now a staff member in the White House Office on Science and Technology, was the keynote speaker.

Dick Cogger, also a member of the Ithaca Cable Commission, demonstrated how the local television coaxial cable could provide high-speed Internet access to homes and businesses. Much issue was made of fiber optic cable as the future and sparked much interest about when it would be available into individual homes. Taking advantage of renewing the charter with Time Warner Cablevision—the new owners of the cable TV franchise—the City of Ithaca was able to lay some fiber optic cable from City Hall to outlying units and to connect the local school district to high-speed communications. Builders of new large apartment complexes in the Collegetown area also were looking to the future and installed cable/wiring capable of providing Ethernet services in the buildings, independent of Time Warner offerings.

Business Systems

The recommendations of the reports from the 2001 Committee, the Strategic Planning Initiative, and the Merten Committee led to a concurrence that the major information technology issue in 1994 was the need to upgrade the Cornell business systems. The Merten Committee report, which noted that administrative and student systems fell significantly short of expectations, brought a focus to the deliberations of ADSPAC (Administrative Systems Planning and Advisory Committee), which had been struggling with

¹¹¹ "Computer viruses: Separating what's real from what's hype," *Cornell Chronicle*, December 12, 1996.

¹¹² "University policy hasn't changed in wake of communications law," *Cornell Chronicle*, February 22, 1996.

¹¹³ www.ithaca.ny.us/Events

such issues for many years. At the same time, the success of Bear Access and all the hoopla concerning the performance and cost advantages of client-server technology, combined with early concerns about the Year 2000 problem, created the environment for taking innovative new steps.

Several other elements contributed to this desire to introduce change and upgrade systems. One was increasing attention to the role that the Internet would play in future business systems. Another element was the prominence being given in the popular and business press to reengineering processes responsible for improvements in customer satisfaction and, simultaneously, in lowering cost. A negative consideration was the unsuccessful team-based reorganization of Information Resources^t and the general dissatisfaction with current administrative systems that users had tolerated for years only because of the promise of change. Last, with Rogers as chairman of ADSPAC and Lambert as the new VP of IT—and both feeling they could make a difference—all the elements needed to make business systems the IT focus for the latter half of the decade were in place.

Steps Leading to Project 2000

In 1993–95, ADSPAC created a vision statement¹¹⁴ to set the tone for the future: “Students, faculty, staff, and others who interact with the university should clearly experience an environment in which they conclude: *Cornell has its act together.*”

The vision itself was condensed into three main points:

- Core technologies will be available and in common use.
- Information dependencies will be widely understood and accepted.
- Customers will experience “fully informed” responses and services.

Agreements reached in this process and accepted by the ADSPAC members were that Cornell would buy and not build systems, and that the first initiative would be to identify a vendor for the human resources (HR) system.

During this same time frame, this exercise became known as Project 2000^u or P2K, with the understanding that this was not just an upgrading exercise but a pervasive change that would transform all the business systems and practices on campus. The original vision was changed slightly and became: “Reengineer Cornell’s management policies and practices to focus all human, financial, and capital resources on excelling at the core missions of education, research, and public service.”

A number of guiding principles were stated for this project:

- Standardize practices; don’t customize software.
- Eliminate unnecessary work.
- Reduce authorizations, expand access.
- Distribute work, aggregate information.
- Follow the “rules of one”: data stored once at the source; data stored and defined only once; transactions approved once.

To get an outside perspective on the situation as it was unfolding, Rogers and Lambert retained Computing Contractors (J. Michael Duesing) in late 1994 to perform a situation-assessment study.¹¹⁵ Duesing was charged with conducting a strategic review of Cornell’s administrative information systems to identify opportunities and recommend how to proceed to achieve effective cost benefits in the future. In a period of several months, Duesing interviewed 60 staff members, covering both technical and administrative staff; staff from central offices, colleges, and research units; and vice presidents and college deans. He covered suppliers and supporters of business systems as well as users of those systems at different levels of the university and in different operating units. He found much skepticism about this large transformational project. These views came in part from individuals who were not involved in the development of P2K, and in part from those whose past experiences with new systems had not made their work simpler or easier. Most people interviewed neither understood nor appreciated the P2K vision and the difference it could make in their work lives.

In response, Rogers and others took the “show on the road,” talking to many offices—formal and informal groups of college and central office administrators—to explain Project 2000 and gather support. When word got out on the campus that the project was going to save \$20 million per year, a large concern arose about cost cutting and layoffs. There were several meetings in Bailey Hall, presided over by President Rawlings, to speak plainly about how the savings would be gained and redirected to academic programs.

¹¹⁴ Fred Rogers, “Campuswide Approach to Systems Planning,” circa 1996–97.

¹¹⁵ J. Michael Duesing, “Cornell University Administrative Information Systems Review Report (Version1),” January 30, 1995.

Project 2000—Creating a Best-Managed University Announced

In March 1996 President Rawlings formally announced Project 2000—Creating a Best-Managed University, a strategy for organizational change designed to make Cornell a model for effective university administration and to enable the university to target its resources on academic excellence.^v The project included replacing five core information systems: Student Information, Human Resources/Payroll, Sponsored Programs, Finance, and Alumni Affairs/Development, and reengineering administrative processes around them by rethinking the way administrative work is carried out with new technologies.¹¹⁶



Figure 8. Project 2000 logo

Several quotes at the announcement capture the spirit of P2K. From Fred Rogers, senior vice president and chairman of ADSPAC: “A number of colleges, departments, and central offices already have been rethinking or reengineering the ways administrative work is done in their areas, eliminating unnecessary steps and duplication of effort, streamlining their processes, and simplifying procedures. It is time for us to capitalize on these successes. By the year 2001, we hope to experience significant annual savings in administrative costs.”

Rawlings concluded: “Project 2000 will improve our day-to-day learning, teaching, and working lives at Cornell. Most important, it will allow us to focus our human, financial, and capital resources on those core missions that are most vital to the life of this university: teaching, research and service.”

A key to the savings plan and redirection of resources was the elimination of so-called shadow systems, all those home-grown information systems in operating units that helped them manage their resources and meet the demands of other offices at the university or external agencies. Often these systems, designed for the specific use of a department or individual, duplicated the information in the central systems, or had information that only the department or individual knew. Sometimes the local data contradicted the information in central systems, a difference

that could be sometimes (but not always) traced to timing—transactions that had not yet been recorded in a central system. Regardless of the condition, there was known to be considerable duplication of effort everywhere.

The early and most optimistic plans at this time called for the five core information system projects to be finished before the year 2000, with the rollout of the HR/Payroll system in 1997, Student Administration and Sponsored Programs in 1998, and Finance in 1999. The budget for Project 2000 was \$50 million, roughly one-half coming from new money and the other half from reallocation of existing funds, mostly staff salaries. The new money was for the purchase of new hardware and software and for consultants and other experts needed to train the Cornell staff.

The project team and staff that developed the first Project 2000 plan and documentation was led by Fred Rogers and Dave Lambert, with early participation by John Weisenfeld and then Ron Ehrenberg after Weisenfeld left Cornell. ADSPAC and its individual members formed the leadership group that argued out the different approaches that could be taken. The work on both technology and approaches was mostly done by Mark Mara, Rick Jones, Bob Cowles, and Tom Boggess. Mara and Jones had the expertise and experience with administrative systems and technology, while Cowles and Boggess had the experience and expertise in computing systems and advanced technologies. John Rudan, as acting director of Information Resources, carried much of the organizing and report writing load that was taken over by Helen Mohrmann when she joined Cornell as the director of administrative systems and distributed technology in February 1996.

Understanding that Cornell needed to upgrade most of its major business systems for a long list of reasons, the responsible data administrators, who were all represented at ADSPAC, took their own initiative to explore options for new systems. In part, they took these actions to get away from systems tailored to specific Cornell requirements and the VM-Adabas environment. They were very interested in taking advantage of the new emerging distributed technologies. This combination of factors led to further studies to determine the best possible future path.

¹¹⁶ “Rawlings unveils Project 2000, a strategy for change at Cornell,” *Cornell Chronicle*, March 7, 1996.

Project 2000: Best of Breed vs One-Vendor System Selection and Plan C

The two main approaches that ADSPAC considered could be classified as “best of breed” and “one vendor.” The best-of-breed approach was based on each responsible data administrator selecting the top-rated commercially available system for the unit’s specific needs. So, for example, after considerable study and evaluation and discussion with other institutions, a team led by Lyman Flahive and Glenn Beardsley selected the PeopleSoft system for Human Resources/Payroll. Public Affairs, in turn, promoted acquisition of the BSR (Business Systems Resources) system for Alumni Affairs/Development, as it was top rated by other institutions comparable to Cornell. Similar investigations were made for Finance and Accounting, Student Administration, and Sponsored Programs.

While there were advantages to this best-of-breed approach (generally the systems were in production at other sites and each vendor had a proposed or working migration strategy to the client-server model), there were serious perceived disadvantages. In particular, different vendors implemented their systems with different database systems and different technologies, all of which would take more effort to install, support, and use than the current well-honed single-technology VM-Adabas environment shared over all systems. The costs of all these disparate environments loomed larger than the advantages. The introduction of all these different systems to the clerical, supervisory, and management staff across the university, and all the training and support these systems would need, could not be overlooked. An important requirement for the new systems was that they had to be distributed systems so that input could come from all relevant points on campus and output could be sent to all relevant staff outside the central offices. Last, it was not clear how data could be shared between such different systems to provide an integrated, consistent view of university data.

Unfortunately, there were no real vendors in the “one-vendor” model. Vendors who offered comprehensive and integrated systems for higher education, SCT, for example, did not have products for large and complex research institutions; their market was largely with the smaller institutions. Two vendors were given serious consideration: PeopleSoft and TRG (The Robinson Group).

TRG was headed by John Robinson, who had a strong track record in the higher education market, having founded Information Associates and several other software companies. TRG had been working with a select group of financial officers from research universities, including Cornell, to define and develop

a financial system. As TRG considered this issue further, they decided to offer a more comprehensive package for large universities and started development of appropriate software.

PeopleSoft^w had enjoyed considerable success with their human resource (HR)/payroll system for commercial firms and was considering modifying this system for nonprofit organizations like Cornell. At about the time Cornell was seriously considering them as a vendor, PeopleSoft organized a higher education division and started aggressively promoting their offerings to this segment of the market. They grew quickly by acquiring smaller companies that had expertise in student systems in particular. So, while PeopleSoft only had an HR/Payroll system, they claimed to offer an integrated client-server solution for all Cornell’s needs. At about this time the industry “buzz word” was ERP (enterprise resource planning), that is, systems that provided information about the entire organization and which facilitated organization-wide planning and deployment of resources. Given PeopleSoft’s past and current successes and their future product plans, they looked like the better choice.

There also was a third plan—plan C as it was called—mostly promoted by the technical staff and the consultants who were involved in the early stages of the project. Although they wouldn’t put it in writing, consultants from Integrated Systems Solutions Corporation, IBM’s consulting arm, and Andersen Consulting, both employed at the start of the project, openly discussed this alternative when they talked with the technical staff. Plan C was to implement one system from either the best-of-breed group or the one-vendor group on a trial basis. That is, this would be the learning experience that would bring out all the implementation problems that Cornell would experience and all the shortcomings of the application package and the vendor before making a final commitment.

This idea had some real advantages. In particular, the widespread use of several new technologies would be tested: a new relational database management system and the client-server technology system from a vendor would be put to serious use across the administrative spectrum of users. Also, it was thought that bringing external software to Cornell would force a re-engineering Cornell’s business processes. At the time, Cornell business systems were seen as arcane and handcrafted to suit personal peculiarities, when basically Cornell was in many ways no different from any other commercial enterprise. If the implementation failed, the loss would not be great and the shortcomings of the vendor’s offerings would be exposed. If the implementation succeeded, continuing with that vendor would have a high probability of success. In either case, Cornell would have significant leverage with any

vendor that an early contractual commitment failed to secure. The executive staff thought otherwise, and so Plan C was not given serious consideration.

PeopleSoft Selected as Vendor for P2K Systems

Under pressure to make a decision on the one-vendor versus best-of-breed approach, a team of responsible data administrators (RDAs were members of ADSPAC) and technical staff from Cornell visited with TRG. They did not recommend selecting TRG systems. The team had concerns about TRG's choice of technologies, the amount of staff resources they had available, and their commitment to meet the schedule Cornell was planning. ADSPAC supported this decision. This was more of a blow to TRG than to Cornell, for it appeared that after failing to get Cornell's commitment, TRG closed shop. PeopleSoft became the only choice. Lambert and Rogers sealed the deal at a meeting with Dave Duffield, president of PeopleSoft, in California, and a formal contract was signed before Cornell closed for the winter break in December 1996.^x

Project 2000 Activities and Teams

Project 2000 was initially advanced under the sponsorship of the President's Council, with a leadership group, the Project 2000 Council, composed of faculty, deans, and executive staff representatives. The Project 2000 Council was to resolve policy questions, set priorities, and monitor progress of the project's Steering Group. The Steering Group, which had oversight over the various project teams, was to be managed by Cathy S. Dove, director of management services, and Helen Mohrmann. A total of 12 teams were organized to provide leadership and take responsibility for recommendations or actions. Five teams were responsible for the five system areas: Student Administration, Finance, HR/Payroll, Alumni Affairs/Development, and Sponsored Programs. Six teams provided leadership for projects dealing with issues common to the entire project and to the five other teams. These were re-engineering, communications, training, reporting, technical/design/integration, and evaluation and measurement. The project office itself brought the total to 12 teams. Generally speaking, Mohrmann had oversight of the technical teams and Dove of the other teams.

Early in 1997, Jack Freeman, who had successfully carried out a similar effort at the University of Pennsylvania, was appointed executive director of Project 2000. His responsibilities included working with and through the Steering Group and the individual teams.

Cornell MVS and VM Software Upgrades

Although all the work with P2K took a lot of the attention of technical staff supporting business systems and the new computing environment, it was still necessary to keep the Cornell workhorse operating in top efficiency, processing all the current Cornell business systems. In 1996, IBM marketing and pricing policies led to a total revision of the operating system and related software on Cornell in order to continue receiving technical support and avoid increased costs. A new version of MVS needed to be installed at this time, providing the opportunity to clean up much of the homegrown and commercial software that had been installed and modified for needs no longer relevant. In particular, since VM was now being used primarily as the launch platform for online Adabas transactions, it was no longer necessary to continue features that supported the previous research or student use. So, for example, the VMBatch service was eliminated, as was the locally developed VMTape that facilitated the use of magnetic tape processing in VM. All magnetic tape work would have to be done in MVS. As part of this same move, the Adabas nucleus was moved to MVS, a move that also improved the performance of all administrative transactions.

The installation of the new version of MVS/ESA, labeled OS/390 R1, was in itself a large undertaking; cleaning up the use of magnetic tapes and encouraging greater use of online storage presented some large problems that had developed over time. However, given the need to make those changes, the decision was also made to install the IBM standard products that would come with the systems software for tape management and access control (RMM replaced CA-1 and RACF replaced ACF-2). With these changes, the MVS systems software could be used unmodified, creating a simpler and more easily maintainable environment. To gain this simplicity, several practices had to change, including coding various features to distinguish expiration dates and monthly lease charges. These obstacles were overcome and the new environment was installed in September 1996 (personal communications with Mark Bodenstein and Mike Garcia).

With all these changes in software and the earlier installation of RAMAC disks and magnetic tape robotics, magnetic tapes were now to be used for importing or exporting data (such as direct deposit of paycheck amounts to participating banks) and for long-term storage of information. Magnetic tapes were not to be used for intermediate storage, as had been the practice, such as for temporarily storing print files so printing could be done offline and for reprinting reports if a print job failed. As has been noted earlier, starting in 1994 a whole host of changes had been made to improve the storage environment. Higher-

capacity RAMAC disks and a 3494 tape robotics unit had been installed to facilitate the storage of information, and new high-speed network-based laser printers had replaced all the older printers. In the end, this was just the first step in a transition to a different operating style that took advantage of technology and the changing environment and reduced costs for systems software in the order of \$150,000 per year.

1997 to 1998 at Cornell University and CIT

During 1997 there were a number of significant transitions and changes in the supercomputing program at the Theory Center as well as in CIT services and staff. The changes were both large and small in their effect on the campus and in their significance for the future. In the summer of 1997, there were three significant changes in CIT: a minor mainframe upgrade along with the elimination of staffing of the mainframes during weekends; the elimination of Docutech-based printing services; and the transfer of PC sales to the Campus Store. Data entry services ended. CIT moved its computers and the server farm and the Network Management Center to Rhodes Hall.

Several key staff transitions occurred: VP Lambert left Cornell, and Tom Dyckman took over in an acting capacity; long-time staff member Dick Cogger moved his network development programming group to Vienna Systems in downtown Ithaca.

A campuswide committee was appointed to deal with the Year 2000 potential problem. Plans continued to be made to implement new client-server business systems as part of Project 2000. By the end of 1998 a new HR/Payroll system, along with a new online time system, were installed and were in everyday use.

During this period, the entire campus network was upgraded to provide higher-speed transmission, and new services were introduced. Cornell was invited to join Internet II, a substantially higher speed network supported by NSF, and networking continued to be a regional interest.

Theory Center Loses NSF Support, Reinvents Itself

All the great and varied accomplishments of the Theory Center, and all the future plans based on the current model, came to an abrupt end in March 1997. Cornell was not selected by the National Science Board as one of the two institutions to be supported by NSF's new Partnerships for Advanced Computational Infrastructure program. Under this program, only two of the four original national supercomputer centers would be supported, and San Diego and Illinois were chosen while Cornell and Pittsburgh were not. The phase-out of the old program started, and the Theory Center looked for new ways to continue its very active programs.¹¹⁷

By April 1998 the re-invented Theory Center was announced, with an intention to retain its high-performance computing resources to sustain Cornell's prominence in research. Tom Coleman was the new director. It was announced that the center had an IBM RS/6000 Power Parallel system with 160 processors and would strive to maintain Cornell's prominence in supercomputing.¹¹⁸

Another CornellC Mainframe Upgrade

In July 1997 the CornellC mainframe was upgraded from the current System 390 model 9672-R32 to model 9672-R24 to achieve 50 percent more computing capacity and overcome the poor performance and overload conditions experienced at the end of the spring semester. This upgrade came just two years after the installation of the model R32 and numerous other performance improvements—the installation of new tapes and disk drives and new operating system software and changes in the application programs. The growth pattern presented as part of the 1995 upgrade continued. It was expected that this upgrade would be the last for CornellC, given that Project 2000 was under way to replace all the old Adabas interactive and batch applications with new client-server systems. Other users of CornellC were encouraged to look for alternatives to meet their computer needs before the year 2000.

CIT Electronic Publishing Phased Out; Use of CUPID Ends

With the departure of VP Lynn, the initiative to organize the electronic printing outlets on campus into a larger conglomerate, like the Cornell Electronic Publishing Working Group, stalled. By 1997 the Computer Resources division of CIT had worked out a cooperative agreement with Media Services to share workloads and charge the same rates. The driving force for CIT activity was no longer the Cornell Library and the reproduction of deteriorating books, but the publishing of its own documents, EZ-Publish activity, and course packs (preprinted course notes) for the Campus Store. The capabilities of the Docutech system to store documents and produce small volumes of material on demand fit with the unpredictable demand for course packs at the beginning of each semester. In fact, there was some discussion about the potential for producing the packs on demand—that is, allowing students to order a pack while in the Campus

¹¹⁷ "CU comments on NSB supercomputing decision," *Cornell Chronicle*, April 3, 1997.

¹¹⁸ "A reoriented Theory Center: Renewed focus on Cornell Research," *Cornell Chronicle*, April 9, 1998.

Store and having it printed on site and made available to them before they finished their other shopping. This, of course, would have required a Docutech in the store itself.

Given the presence of other print shops on campus, including the large University Printing Services operation, and that printing was not a core CIT activity, in mid 1997 CIT phased out the Docutech system, sold it to University Printing Services, and went out of the publishing business.

At about this same time, the efforts on CUPID (Consortium for University Printing and Information Distribution) to move “bits” and not books between cooperating sites also came to an end, according to Rich Marisa, who took over technical responsibility for CUPID (personal communication). CUPID was implemented at Cornell, and although CIT used it to publish materials at different print shops in Ithaca, and the code was installed and used productively at other sites as well as being licensed to Net-Paper, the momentum and energy to continue dissipated due to other priorities.

CIT Transfers Microcomputer Sales and Support to Campus Store

In mid-1997 the CIT microcomputer sales and repair division (CIT Sales and Support) and its operations were transferred from CIT to Cornell Business Services, and that division of CIT went out of existence. Earlier VP Lambert had appointed a blue-ribbon commission to decide on the future of the sales operation.¹¹⁹ The commission recommended that the CIT operation be moved to the Campus Store to simplify the user interface and to effect operational cost savings. The recommendation reversed the arrangement that had been in place since 1984, when the sale of hardware and software was deliberately split between CIT and the Campus Store, respectively. The feeling at that time was that the Campus Store was much better at selling off-the-shelf software commodities, while hardware was still very much a technical design issue of monitors, “mother boards,” and other ancillary features in which CIT specialized. For this reason, and to reduce redundancies, the Office Equipment Center was folded into CIT’s personal computer sales operation in 1987. Ten years later, however, when all kinds of personal computers were available as packaged commodities through a large number of sales outlets, continuing these separate hardware and software outlets no longer served users well in terms of cost or convenience.

After the commission’s recommendation was accepted and the transfer took place, the Campus Store opened their new Technology Connection department, which combined sales, sales consult-

ing, and repair and support services. The Technology Connection continued the tradition of having back-to-school fairs, although not at the scale of fairs held previously in Lynah Rink.

CIT Data Entry Operations End

On October 1, 1997, the CIT data entry operation was phased out. The last two CIT data entry operators were Donna Poole and Shirley Harders. Old timers still referred to this function as “card punching,” even though the physical card had long since been replaced by card images on magnetic media. This service, whose primary function was to transfer information from paper to electronic form, was no longer needed. It had a long history going back to the 1930s in the Department of Agricultural Economics in Warren Hall, to the 1940s in the administrative data processing Machine Records unit in Day Hall, to the 1950s in the Computing Center in Rand Hall. Data entry was now being done all over the campus directly from workstations into data collection systems by manual or automated means. A small vestigial group continued to perform this function in the controller’s office, pending the installation of the new payroll system that would completely distribute data entry all over the campus.

As an interesting counterpoint to the rate at which technology was becoming obsolete, the same year (1997) that CIT announced it would no longer distribute or support the Bear Access version for DOS, DOS had lasted at most 20 years on campus.

CIT Servers and Network Management Center Moved to Rhodes Hall

During 1997 CIT began planning to move its computers and servers from CCC to Rhodes Hall, because the downsizing of the Theory Center and supercomputing operations had resulted in unused computer room space. Not only would such a move provide more modern computer room space, but the synergies with the Theory Center, the Network Management Center, and CIT would reduce costs while improving service.¹¹⁹ Also entering into the discussion was the accepted view that space in CCC was now premium central campus space that should be put to other uses.

After the decision was made to combine operations, the move was planned in three parts: to consolidate Theory Center equipment, to move the CIT CornellC mainframe, and to move the servers and the Network Management Center. The first move was done in early November. The second move of CornellC was done

¹¹⁹ “Major computer, network resources to be consolidated,” *Cornell Chronicle*, October 30, 1997.

over the Thanksgiving break when students, faculty, and staff were away from campus.

The third and the most complicated move took place between December 25, 1997, and January 2, 1998, a period of low activity and traditionally the time when such major equipment moves were scheduled. Over 100 computers, mostly servers, had to be disconnected, boxed up, and transported and then reconnected in their new home on the seventh floor of Rhodes Hall. All the moves were completed with a minimum of service disruption. One of the interesting parts of this move was that one of the IBM 3494 Tape Robots was left in CCC and connected to CornellC in Rhodes Hall using the installed fiber on campus. In this way backups to CornellC data and systems could be easily executed off-site, providing an additional layer of security for the continuity of business operations.

Credit for this successful move—one in a continuing string of successes in relocating complex computer operations every few years since 1986—goes to the leadership team. Members of the team were Mariann Carpenter, manager of computing systems; Andrea Beesing, manager of network operations; Benjamin Brown, Theory Center manager of operations; Billie Dodge, special projects manager; Sanjay Hiranandani, network engineer; Peter Baker; Bill Biata; and Jim Doolittle. The work of hundreds of other CIT and Cornell and vendor staff also contributed to this success.

As part of the third move, the Network Management Center—the NMC/NOC—was moved to Rhodes Hall. When NMC was set up in 1990 in Caldwell Hall, it was largely concerned with monitoring CITNet and TheoryNet, Cornell's main central campus networks, and their external connections to BITNET and NSFNet, 24 hours a day, 7 days a week. It later expanded to monitor international connections for Sprint Corporation and other entities. It took on increasing responsibilities so that by 1997 it was monitoring key components of the campus telephone system, the central servers such as e-mail, local department network components, and CornellC, the remaining Cornell mainframe. Future responsibilities would likely continue to increase as the network, and network applications such as video, became standard practice and service reliability was critical. During its lifetime the NMC led an itinerant life, first located in the basement of CCC (1988), then at 120 Maple Avenue (1995), then back to CCC (1995), and finally to Rhodes Hall. These moves of the NMC alone were a good example of the increasing sophistication and capability of the Cornell technical staff to move such operations without major operational disruptions.

When all the computers and servers and the NMC were relocated from CCC to Rhodes Hall, the supporting staff members from Network and Computer

Services were located in nearby offices on the seventh floor of Rhodes Hall as part of the agreement. This move vacated the third floor of CCC, which was later converted to offices for the Office of the University Counsel. When this was completed, CIT occupied the first and second floors in CCC and a portion of the computer room in the new addition, roughly one-third of the computer room space. The rest of the computer room was not assigned for use.⁷

VP Lambert Leaves Cornell

The final and most significant transition in 1997 was the move of VP Lambert to Georgetown University at the end of the year. Lambert was appointed to the newly created position of vice president for information services and chief information officer.¹²⁰ Lambert had been at Cornell for nine years, serving as the first director of network resources for seven years and then as vice president for two years. Lambert played a leadership role in the significant network developments at Cornell, especially the installation of the Ethernet and FDDI networking technology early in the 1990s, followed by his encouraging and expanding the use of network applications, such as Bear Access and the web. Lambert also supported the creation of the Cornell Policy Law Institute, which positioned Cornell in a leadership role in addressing the legal issues increasingly impinging on IT. Responding to the concerns noted in the Merten Report, which recommended that academic and administrative units assume accountability for their IT needs and actions, Lambert took preliminary steps to organize the College Information Technology Officers (CIO) group.
^{aa} Several meetings of the group were held before Lambert left. Last, he played a key role in getting Project 2000 under way.

Tom Dyckman Appointed Acting VP

As an interim move, Provost Randel initially put in place a troika model where Helen Mohrmann, Ann Stunden, and Pete Siegel—“three gifted leaders,” in his own words—would constitute an executive committee with Randel as chair. This leadership arrangement would see CIT through a transitional period until a new vice president could be recruited. However, as the recruiting period stretched out, in September 1998 Randel appointed Thomas R. Dyckman as acting vice president for information technologies. A professor of accounting in the Johnson Graduate School of Management, Dyckman had expe-

¹²⁰ “H. David Lambert accepts post at Georgetown,” *Cornell Chronicle*, November 19, 1997.

rience in administration, having been associate dean and acting dean of the Johnson School.

CIT Staff Transitions; Cogger, Regenstein, and Siegel Leave Cornell

There was some significant staff turnover in CIT in 1998. Dick Cogger and his entire NCS Advanced Technology group left Cornell as a unit and joined Vienna Systems, setting up an office in downtown Ithaca. It had been announced earlier that Cogger would be retiring from Cornell; this move was a novel way for him to do so. As Siegel described it in an e-mail to all CIT staff, "This talented software development team has a once- or perhaps twice-in-a-lifetime opportunity to move together to create an R&D center as part of Vienna Systems,^{bb} at the same time staying in Ithaca." One could say that this was an excellent example of how telecommuting was changing the employment opportunities for talented technical staff. The individual staff members joining Cogger were Pete Bosanko, Larry Chace, Josie Cundy, Steve Edgar, Jianmei Li, Tom Parker, Melinda Shore, Joy Veronneau, and Andy Wyatt.

At about this same time Carrie Regenstein, director of academic technology in the Division of Academic Technology and Technology Services, joined the University of Rochester as director of the University Computing Center and assistant dean for educational technology. This move was a marvelous opportunity for Carrie but a significant loss for CIT and Cornell.^{cc}

In mid 1998, Pete Siegel, director of network and computing services, left Cornell and was appointed director for academic information technology at Iowa State University of Science and Technology.^{dd} This brought to a close Pete's long and distinguished career at Cornell, where he had played important roles in bringing high-performance supercomputing to Cornell, in managing the Theory Center and CIT, and in fostering a spirit of cooperation between different organizations. Jim Doolittle was appointed acting director pending a decision on the new VP, who would have the responsibility for hiring someone to fill this position.^{ee}

CSARS Use Ends

As part of the financial restructuring in this period, the last use of CSARS (Computer Services Accounting and Reporting System) for mainframe billing purposes occurred at the end of June 1996. CSARS was first put into production in July of 1985 to produce billing statements for the CIT mainframes. At that time there was a substantial amount of real-dollar income to be collected, but by the mid-1990s, most of the use of CornellC was for business systems of the university. In addition, many items such as

tape and disk storage were no longer chargeable. The importance of CSARS was in the collection of computer usage statistics, which were used as the basis for allocating total costs of computer operations across the administrative units. The fiscal year billing for 1995–96 was used for cost allocations until other methods were developed.

Business Systems

The year 1997 stands out for addressing the Year 2000 (Y2K) problem and organizing for Project 2000 (P2K). In preparation for implementing the systems that composed P2K, decisions were made on acquiring Informix, a new database management system, and providing a new data entry process for payroll records called COLTS (Cornell On-Line Time System). A production architecture was developed to guide the implementation of PeopleSoft systems, and a survey was conducted in late 1997 to solicit perceptions about Cornell's administrative processes and the readiness of the campus for the changes to come from new systems. In 1998 the P2K schedule was revised, and efforts were directed toward installing the PeopleSoft human resources/payroll system.

Year 2000 Committee Appointed

In 1997 an organized effort was under way in CIT to address the Y2K problem for Cornell-developed software (the so-called legacy systems), and CIT was preparing a larger university approach to minimize the potential problems. That year Marilyn Baxter, Y2K project manager in Administrative Systems and Distributed Technology, was advising users to check owner's manuals and to check with manufacturers to see if the products they were using were Year 2000 compliant.¹²¹ Campus users were advised to check web sites that were tracking the state of Y2K-compliant vendors about their particular products. Baxter was advising users to have their fix in place by the end of 1998 so that there would be a full year of usage before the rollover date.

However, in 1998 there was additional concern that the university was not adequately addressing the Y2K issue. In part this concern came from the Audit Committee of the Board of Trustees, many of whom were involved in the external business world where some efforts were perhaps better orchestrated and organized, and also from KPMG (the external auditors). Also, as it now seemed unlikely that all the central business systems would be replaced as part of the P2K project, those in continued use would have to be

¹²¹ "Protecting against the 'millenium bug,'" *Cornell Chronicle*, September 18, 1997

“remediated” to be Y2K compliant. Central Audit was also becoming involved with Y2K issues to address any of Cornell’s business practices that could be affected by the Y2K rollover.

Responding to all these concerns in the fall of 1998, VP Dyckman appointed John McKeown, a seasoned Cornell administrator, as Cornell’s Y2K project leader. With the support of President Rawlings and Provost Randel, McKeown alerted the campus to the need for Cornell to be fully prepared for Y2K, assuring that this was a university-wide priority issue and that a process was being developed to make certain all campus units would carry out their responsibility. The Y2K core project team consisted of McKeown as director, Steve Worona from OIT, Cecilia Cowles to deal with communications to the campus, Keith Boncek for Facilities and Campus Services, Dennis Butts for Central Purchasing, and Mary Bouchard from the University Audit Office. From CIT there was Ron Parks (replacing Marilyn Baxter) for central administrative systems, Dan Batholomew for desktop issues, Rich Marisa for research issues, Jim Doolittle for Network and Computing Systems (NCS) contingency planning, and Mark Bodenstein for mainframe, server, and network issues. This group worked with 30 senior administrators from key campus constituencies, such as colleges and major administrative operations, to make sure all possible problems were anticipated and resolved. In McKeown’s opinion, this was a cooperative venture where the campus units really worked together to make sure there were no surprises at the Y2K rollover.¹²²

Project 2000 Organizing Steps in 1997

Executive director of Project 2000, Jack Freeman, spent the early part of 1997 working closely with Mohrmann as the technical director of the project and with Robert DePalma from Andersen Consulting. Andersen had been hired to replace IBM’s Integrated Systems Solutions Corporation (ISSC), as they had more experience with installing PeopleSoft systems and had a larger seasoned staff than IBM. In addition, Cornell, as an early adopter of the PeopleSoft systems, took a leadership role in organizing other institutions that had already successfully installed a system or were considering PeopleSoft as a vendor into lobbying groups to get PeopleSoft to design systems to meet their needs. Fred Rogers was instrumental in creating a Strategic Advisory Council of CIOs and CFOs from other institutions to work with PeopleSoft on satisfying their collective needs.

At another level, staff from the business systems and technology sectors at these institutions formed a special-interest group of PeopleSoft users or prospective users to work with PeopleSoft staff at their level.

Mohrmann was chairman of this group for two years. PeopleSoft took notice of all these efforts, supported them, worked with individuals and groups to understand their needs, which were different from those of the commercial sector, and as possible, responded to those needs.

During most of the year a high level of activity by all the Project 2000 teams was directed toward preparations for installing their respective systems, while the cross-function teams continued working to meet their objectives. In many cases, the functional teams acted as beta-testers of systems or partial systems and features as PeopleSoft rolled out early versions of its products for higher education.

When he left in December 1997¹²³ six months ahead of his original plan, Freeman said, “Project 2000 remains on track and eventually will offer great benefits to the university.” Fred Rogers took over Freeman’s activities after he left, and no replacement was appointed, although the plan was to do so. By this time the project schedule had slipped almost a year. There were valid reasons for this situation. Uppermost was that as the teams became trained in, and more familiar with, PeopleSoft systems, the complexities were far greater than anticipated, and some of the shortcomings, like the lack of a web interface, were critical. At the same time, the fit between the capabilities built into the PeopleSoft systems and Cornell’s operating style grew larger rather than smaller, presenting serious problems. Also, scheduled upgrades to PeopleSoft systems, which provided all the maintenance changes in regulations, technology, etc., became a disadvantage rather than an advantage, because Cornell had to fit in with the PeopleSoft schedule and could not proceed entirely on its own.

Selecting a New Database Management System—Informix

One of the important projects that was going on in parallel with the decision on the software vendor for Project 2000 was the selection of a new database management system. A comprehensive set of criteria was developed, which included the support of large and small databases with equal ease, the support of a high volume of transactions, and queries from decision support systems such as data warehouses and data marts. The offerings further had to be compatible with a wide range of software vendors and a good fit with

¹²² John McKeown, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, December 20, 2000.

¹²³ “Project 2000 director will leave his post at the end of this month,” *Cornell Chronicle*, December 4, 1997.

the future directions of Cornell and the industry in general.

The principal choices given serious consideration were DB2 (a nonmainframe product) from IBM and products from Informix, Oracle, and Sybase. After listening to presentations and making corporate visits, checking with other customers and reliable sources, and using the evaluations from the Gartner Group, the finding of the technical team was that “either Informix or Oracle will provide an adequate relational database computing environment for Cornell’s needs.”¹²⁴ Although that was the technical conclusion, the team leaned toward Oracle for a number of reasons. Because any license would make the database management system available across the campus, a wide range of current and possible users were solicited for their input about such products. By this time a number of installations on campus were using Oracle as well as other products that met their need or came at the right price. At the time Oracle had a major share of the market for those products and pretty much commanded a price of their choosing. When negotiations with Oracle resulted in a proposed cost of almost \$2.25 million to acquire their license, the executive staff decided that the product from Informix, available at \$600,000, was a better choice. Given this large difference in costs and the sense that PeopleSoft preferred Informix as their vendor of choice for future development, Cornell acquired the Informix system.^{ff}

COLTS—Cornell Online Time System

One of the more interesting and successful projects that became associated with P2K was COLTS, the Cornell Online Time System, first introduced in mid-1997. As Donna Taber describes it, COLTS started earlier when Administrative Systems and Distributed Technologies (ASDT) was encouraged to “harvest” successful applications of technology developed by other campus units.¹²⁵ This was in keeping with the sentiments at the time that campus units had done interesting, innovative, and cost-saving projects using IT for administrative processes, and central units should look at these for wider applicability.

One such case was the work done by David Stull in Cornell Cooperative Extension. Stull built an online system for collecting time-keeping data for hourly employees. A project team of Donna Taber and Kevin Leonard from ASDT and Kelly Thompson and Ken Ceurter from Payroll worked with Stull to generalize his system for the entire campus. Basically, time-keeping data was entered online by the individual employee through a web browser interface and edited for various consistency checks. It was held in a database that was later fed into the HR/Payroll system to generate paychecks. The approval process could

be done online or on paper; COLTSI (the first version) could generate a voucher that had to be signed and returned to central Payroll if a supervisor so chose. However, a supervisor could also electronically approve the time for his employees. COLTS was the first Cornell business application to use the NetID and Kerberos for signature authority. Since there was a software-based limitation of 3,000 individual users, no strong action was taken to eliminate the paper process.

There was also a third mode of input using KRONOS, a separate electronic time-keeping system used by units with large numbers of hourly employees. Input to KRONOS was done by swiping employee ID cards, from which point the system took care of the time calculations and accumulations that, after approval, were fed into COLTSI to later go into the old payroll system. This multiple time-keeping entry continued until late 1998 when the new PeopleSoft HR/Payroll system went into production.

Development of Production Architecture

In late 1997 and early 1998, a Project 2000 production architecture was developed to define the components necessary for continued processing and support of the Project 2000 business systems that were to be installed.¹²⁶ An experienced team of Cornell staff (Rick Jones, Mark Mara, Gary Buhmaster, Robert Wilkinson, Dave Wakoff, Nancy Van Orman, Tom Boggess, and Art Wallace), representing different divisions and technologies across CIT, wrote this key report. The report focused on the PeopleSoft HR/Payroll system being implemented with the ability to accommodate the other systems in time.

In order to generalize the types of interactions between the systems and the end users, and so be able to develop the overall architecture plan along with the hardware and software components, the access levels were classified as shown on Table 10. The key differentiation between power access, occasional access, and self-service access was in the level of training required. It was acknowledged that a single user could interact with systems in different roles; for example, a manager entering personal vacation or sick time using self-service access could at other times make a salary change for an employee using power access.

¹²⁴ “Relational Database Technical Evaluation: DB2, Informix, Oracle, and Sybase,” Tech Team Analysis, November 15, 1995.

¹²⁵ Donna Taber and Kevin Leonard, “COLTS, The Cornell Online Timekeeping System,” Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, June 21, 2001.

¹²⁶ Rick Jones et al., “Cornell University Project 2000 Production Architecture,” July, 31, 1998.

Table 10. Project 2000, Application Access Types

<i>Access Classification</i>	<i>Training</i>	<i>Number of Users</i>	<i>Frequency of Use</i>	<i>Application Examples</i>
Power access	extensive	small 0–100	constant	Manage academic records
Occasional access	moderate	moderate 100–1,000	often	Student employment, COLTS approvers
Self-service access	none/limited	large 1,000–10,000	occasional to constant	Just the Facts, Employee Essentials, COLTS data entry

Once this architecture was stated and accepted, various technical solutions were proposed, which are discussed in great detail in the report itself.

Survey of Administrative Processes and Campus Readiness for Change

In November 1997, 2,544 surveys were mailed to all administrative staff on campus “to obtain a baseline against which to assess the perceptions of administrative staff regarding today’s business processes and administrative systems prior to reengineering and the implementation of PeopleSoft.” In particular, the survey was designed to obtain information regarding

- processes identified as most critical to Project 2000’s success in each principal area;
- outcomes expected to impact how easily administrative data may be accessed, updated, reported, integrated, and shared;
- perceptions of campus readiness for organizational change; and
- information regarding the extent to which administrative staff have access to technology resources.

Of the total sent, 772 surveys (30 percent) were returned, a number sufficient to make reasonable statistical estimates on the full sample.¹²⁷

There were many findings from the data but the most important were:

- Staff working in all areas reported that (1) providing customers with appropriate administrative information is a very important component of their jobs, and (2) falls substantially short of performance expectations. (There were many, some large, some small, gaps between how important the staff perceived and rated the process to be and how satisfactory they rated the process to be.)

- The data suggested that the process and systems changes anticipated from Project 2000 are perceived by administrative staff (1) to be critically important to the work they do, and (2) not to describe Cornell’s present environment.
- Overall, the results suggested that Cornell’s current work environment—at least as perceived by administrative staff—may not be best positioned to effectively manage change of the scale and complexity likely to be precipitated by Project 2000. (Some administrative units were better prepared than others, and the suggestion was made to address those units individually.)

In terms of access to technology resources, Table 11, taken from the survey report, presents the information very well. Besides directly reflecting the status of desktop systems in the administrative area, it is a reasonable representation of the general situation on campus. It also reveals one of the complications discovered as P2K unfolded—that the large number of Macintosh systems were not going to be well accommodated by the PeopleSoft HR/payroll implementation.

The report on production architecture and the results of the survey provided a good foundation to guide the implementations of the new PeopleSoft systems.

¹²⁷ Agelia Dumas, *Perceptions of Cornell Administrative Processes and Campus Readiness for Change: Report of a Survey Sponsored by Project 2000, Conducted November–December 1997.*

Kronos, Configuration Management and Security, and Database Administration.

The December 1998 deadline for HR/Payroll also became the deadline for COLTSII. This was a totally new system, designed in response to the on-campus experience with the original system. In addition to accepting individual input, it continued to be the entry point for data from the KRONOS time-keeping systems. The manual system for data entry was no longer possible, so unless the unit had a KRONOS system, employees and supervisors had to use COLTSII. This new version had no limitation on the number of employees and continued the internal checking and time and labor distribution calculations before feeding the information into the PeopleSoft system. Within a short time, COLTSII simply became COLTS.

As the project headed to the completion date, the project team was busy developing reports that would be available when the system went into production. It was well recognized that the self-service Employee Essentials would not work with the new HR/Payroll system, nor would the Data Warehouse be in operation for lack of a feed to the new system. The production architecture report had laid out different ways for what became known as data delivery, providing information to operating departments in printed or electronic form. While a number of technical solutions were discussed, when it became known that PeopleSoft intended to use Actuate, a report development and distribution application program from Actuate Corporation, as their preferred reporting tool, this product was acquired for the HR/Payroll system. After extensive consultation with departments, the Reporting Group (Graham Hall and Neil Belcher headed by Robert Wilkinson) developed 65 reports for the central offices, HR/Payroll, and different internal units such as Benefits Administration to use in the conduct of their work. The reports were made available in electronic form to central HR/Payroll for distribution on request to departmental offices for another year before web-based reports and data sets, with appropriate security, could be developed and installed to make the data directly available to those offices.

All told, there was reason to celebrate when the New Year came about on January 1, 1999. The new HR/Payroll system was in production (on December 16, 1998) and COLTS was generally free of problems. People just understood there were going to be glitches and worked diligently to report them, while the CIT and HR/Payroll staff worked equally diligently to overcome them. One of the outcomes of following this strategy was that at times it was simpler to program a solution into the PeopleSoft system, even if this created some short-term problems. While this action violated one of the tenets of the project, which was

not to make system modifications, this was at times, according to Mohrmann, the only expedient way to go. It was simply impossible to attempt policy or procedure changes across the university to address the issue.

At this point, as the reality of the complexities of PeopleSoft technology was fully exposed with the installation of the HR/Payroll system, Project 2000 was more or less put into hold status. Adding to this decision was the fact that the project costs had grown to almost \$35 million real expended dollars, and it was not clear that Cornell could afford to move forward to install the other systems. That is not to say that the other systems would be as expensive, because clearly a lot of the cost was in training, in hardware and software, and in other sunk costs that would not be repeated. Nonetheless, the question of continuing on the chosen path needed to be examined more carefully.⁸⁸

CUDA Fades Away

CUDA, the Cornell University Distributed Accounting system, was disabled on June 30, 1998, after being in productive use for 11 years. Yoke San Reynolds, university controller, made this decision in 1997 to provide sufficient lead time for users to consider other alternatives. Both the success of Accounting Data Warehouse, which satisfied most of the needs of departments for similar information, and the costs of maintaining CUDA entered into that decision, according to Michelle Reichert, CIT staff member who had supported CUDA users for the last eight years of its use. CUDA was developed in 1987 as the first distributed administrative computing application and had served departments admirably in providing access to the information in the central accounting system. When first released, CUDA started with 20 users; by the time it was disabled, there were about 120 users of the system in 87 departments.

Networking and Network Services

During the years 1997–98, the use of network services continued to grow. For example, by the end of fiscal year 1997, use of electronic mail had grown to over 100 million messages that year, compared with 69 million the year before, and CIT had to take measures to increase network capacity. New network services for printing were added, while other services such as Bear Access were upgraded to newer technology. Subscriptions to ResNet continued to increase, and EZ-Remote modems were upgraded to higher speeds. Cornell was invited to join a select group of institutions forming an NSF high-speed network, and an Office of Distance Learning was formed. Interest con-

tinued to be high in networking in the area, and the IthacaNet conferences continued. The CU-SeeMe consortium faded away, and CU-SeeMe became a total commercial product.

Cornell Network Upgrade

By 1998 the network technology first installed in 1991–92 had lasted just about six years. To address the problems of being at the limit of capacity, in 1998 a \$1.9 million project was approved to start making improvements to the backbone and to limited systems beyond the backbone, the so-called edge systems to buildings or to LANS, to maintain short-term network viability. Essentially this phase started to replace the IBM 6611 routers with a combination of Cisco 8540 and 1900 routers.¹³⁰

The network topology changed from the backbone being defined as the connectivity between the 6611 routers to the backbone being defined as a series of subnets with the 8540 routers operating at gigabit speeds connected to the 1900 routers operating at 100 Mbit speeds. However, all the 1900 routers were in the node rooms and not in the buildings, as in the older topology. For many LANS this meant simply connecting to the 1900 router ports without any interruption or change of service. In some cases, by using a Cisco 5500 router, a workstation would be changed from a shared to a switched Ethernet, providing 10 Mbits directly to the desktop.

This cutover eliminated the FDDI technology for routing IP to a switched Ethernet technology operating at gigabit speeds between the 8540s and at 100 megabit speeds from the 1900s. This approach was a change in direction from the proposed use of ATM technology to a more conventional and proven approach that left the workstations unaffected. No action was taken on upgrading the telephone services until after the decade ended.

Net-Print—New Network Printing Service

In April 1998, CIT announced a new Net-Print service that would be available at the start of the fall semester after a successful pilot.¹³¹ Net-Print let students print across the network on fast, high-resolution laser printers and bill the charges directly to their bursar accounts or to special cash accounts. This new system replaced the VendaCard system of paying for laser printing and at the lower cost of 10 cents rather than 15 cents per page. During the pilot phase as many as 13,000 pages were printed daily, a substantial increase over the 10,000 pages printed per month with the older VendaCard system.

Bear Access Upgraded

In June 1998 a new Bear Access was put into production. Bear Access changed from the Project Mandarin technology to a new architecture dubbed Project SALSA (Service and Licensed Software Acquisition), a Cornell-developed technology that was more powerful and flexible than the Mandarin version.¹³² Ron DiNapoli, Project SALSA technical lead, said that “Project SALSA has several important benefits. It can deliver services using installers instead of downloading individual files. It can support multiple viewers, like web viewers. It can keep track of a service’s location on the user’s computer, even if the user moves the service. And it makes it much easier for developers to build and maintain services.”

Earlier in the spring of 1997 CIT announced that it would no longer support the DOS version of Bear Access and would focus exclusively on Windows 95 and Windows NT platforms for the IBM-compatible PCs. CIT argued that it had to allocate its resources for the increased use of Windows 95 and Windows NT on these newer platforms and it was supported in this decision by the computing support providers across the campus. Users were advised to contact the Help Desk for assistance in moving to other services.

EZ-Remote and ResNet

Off-campus EZ-Remote services continued to be improved. By October of 1998, the EZ-Remote-HI fee-based service was offering 160 modems running at 56 Kbps with improved Bell Atlantic technology. As a result, those users with this high-speed modem would see improvement over the 28.8 speeds still being offered. The EZ-Remote LO (free) service continued to provide 100 modems at a speed of 14.4 Kbps.

By the fall of 1998, ResNet had 6,450 ports in operation, and 74 percent of those were subscribed. “Wintel” machines now accounted for 93 percent of the subscriber machines compared with 7 percent for the Mac. These statistics remained essentially the same until the end of the decade.

Cornell Awarded Grant to Join NSF High-Speed Network—vBNS

In March 1998 it was announced that Cornell was one of eight New York colleges and universities to receive a grant to support connections to vBNS, a

¹³⁰ Office of Information Technologies (OIT) Reports, “Next Generation Network,” www.cit.cornell.edu/oit/ebg/ngn.html.

¹³¹ “Net-Print service center—with enhancements—will be fully operational by fall,” *Cornell Chronicle*, April 30, 1998.

¹³² “Bear Access 98 is here; instructions for installing it are on the web,” *Cornell Chronicle*, June 18, 1998.

special high-speed network of NSF vBNS (very high performance backbone service) had been at Cornell since 1995 as part of the supercomputer initiative in the Theory Center. The grant would allow the service to continue and be more widely available on campus. When created, vBNS originally connected the national supercomputer centers, but it had expanded since then. At the time the grant was announced, the number had grown to 63 institutions with plans to add 29 more. At the announcement, Siegel noted that use of the vBNS would be restricted to “meritorious applications,” which might include the transmission of video and data as well as collaboration and data sharing among researchers. He also noted that “even now, Cornell has more traffic going over the vBNS than over the commodity Internet, because we are doing more collaboration with other vBNS schools.” The vBNS^{hh} also was seen as an early step in the development of Internet II,ⁱⁱ a project to provide a separate high-speed network for education and research.

Cornell Office of Distance Learning Formed

The Cornell Office of Distance Learning was organized in 1997, and David B. Lipsky, who had served as dean of the School of Industrial and Labor Relations, was appointed director reporting to the provost.¹³³ Creation of this office was one of the first steps in exploring the possibility of using communications technology to make Cornell courses available to people around the world. The office was to facilitate and coordinate the use of distance learning throughout Cornell through encouragement and assistance, research, and support activities. The goal was to stretch across the globe Cornell’s instruction in degree and nondegree programs, executive education, and international programs and strengthen its land-grant mission. This was clearly a response to the concerns expressed in the earlier FABIT report and in the Draft Planning report, both of which called for this kind of action.

In June 1998 the office received an anonymous gift of \$500,000 to accelerate the growth of programs that would allow students, faculty, and experts around the world to interact using a blend of computing and video technologies.¹³⁴ At the time over 100 Cornell faculty were offering or planning to offer distance learning courses.^{jj}

IthacaNet—Van Houweling Keynote Speaker

The annual Tompkins County networking conferences continued to be held in 1997 and 1998. Starting in 1997, IthacaNet partnered with Academic Computing Services at Ithaca College to sponsor Ithaca Technology Week at about the same time as

the networking conference, and they also held an Educational Technology Day at the college during this week. The 1997 conference—NTC97, Networking the Future—was held in Statler Hall on the Cornell campus. One of the keynote speakers was Doug Van Houweling, who entitled his presentation “When the World Is a Place, Where Is Tompkins County?” The 1997 IthacaNet award was given to the Electronic Futures Committee of Tompkins County, which was organized to consider this increasingly important topic on a regional basis. The chair of the Electronic Futures Committee was Charlie Evans, a member of the Tompkins County Board of Representatives and a long-time systems programmer at CIT.

The fourth annual Networking Tompkins County Conference was held in March 1998 at Ithaca College, with over 200 people attending. IthacaNet and Academic Computing Services at Ithaca College followed the same pattern of sponsoring Ithaca Technology Week in conjunction with the Educational Technology Day at the college. Starting in 1997 and continuing in 1998, the conference ran multiple tracks (technology status, legal issues, and future directions) to present different topics and allow attendees to attend the session of their choice. Steve Cisler, a former Apple researcher and now a full-time advocate of community networking, was one of the keynote speakers. He talked about the evolution of community networks since the 1970s. Homer Smith, who founded Lightlink, the area’s first commercial Internet service provider, was given the second annual IthacaNet Award for Achievement in Networking Tompkins County. Interestingly, almost half the people in the audience were Lightlink subscribers.

There was talk at the conference that Time-Warner Cable would soon be providing network services over its cable plant, and that Bell Atlantic would be offering ADSL services over the telephone network. Marty Luster, New York State assemblyman, closed the conference with concerns about the “technological haves and have-nots” and said he was seeking a substantial amount of state funding for computer technology in public and school libraries.¹³⁵ During this conference Mugwump, Ithaca’s still-functioning Macintosh users group, sponsored the session “Mac Inter-Networking: What’s in It for You?”

¹³³ “Cornell opens distance-learning office to expand university boundaries,” *Cornell Chronicle*, April 10, 1997.

¹³⁴ “Half-million-dollar gift accelerates CU’s distance learning efforts,” *Cornell Chronicle*, June 18, 1998.

¹³⁵ “Networking Tompkins County conference updates some Internet dreams,” *Cornell Chronicle*, March 26, 1998.

CU-SeeMe Fades Away

In 1998 Cornell sold the remaining rights to CU-SeeMe to White Pine Software. By this time both the Cornell version and the White Pine version offered both black and white and colored transmissions available on Macintosh and IBM-compatible systems. Given White Pine's greater interest and resources, Cornell let the consortium fade away, and after that point, the commercial version from White Pine became the product of choice.

Instruction

More FABIT-Funded Teaching Classroom Upgrades

In round one of the FABIT-supported teaching classroom upgrades taking place in 1997 and 1998, 50 new data/video displays were installed in 65 classrooms, which were then upgraded to a tier 1 technology level (an overhead projector with a liquid crystal display to project computer displays, and a network connection). The cost to FABIT for this round of first projects was \$735,000. In round two, which followed, the total FABIT funding allocation was \$1,265,500. This amount raised another 60 classrooms to the tier 1 technology level and funded several new innovative projects that were at the tier 2 level (tier 1 requirements plus phone jacks for teleconferencing and a satellite download). Every school and college was now participating in this push to upgrade classrooms for the use of IT in lecture and laboratory classes.

In the fall of 1998, the provost allocated FABIT a total of \$775,000 for round three projects to upgrade classrooms. Notable in this round was the roughly \$500,000 for upgrading instructional computing facilities in Carpenter Hall for the College of Engineering, many of CIT's facilities, and the advanced computing facilities in Uris Library. These upgrades included new equipment as well as improvements to the facilities themselves.

Computer-Based Training Introduced

Also in June 1998 a new training paradigm was introduced when Technology Training Services signed an agreement with CBT Systems to provide interactive, computer-based training for students, faculty, and staff members. Thirty end-user courses were offered on Internet skills, Microsoft Office 95 and 97, and Windows 95, and Macintosh courses for FileMaker, Microsoft, and Adobe products. All these could be accessed from the user's workstation over the network. There was no fee for these courses. In addition, there were over 200 professional course titles covering subjects such as Java and web administration, for which there was a nominal fee to recover costs.

1999 Highlights

We take a somewhat different approach to recapping the events of historical importance for the year 1999, since it is both the last of the decade and also the last to be covered in this manuscript. We list the highlights, besides the Y2K and P2K projects, and the goings and comings of staff who had influenced or would influence the future of IT.

Polley McClure Appointed VP for Information Technologies

The most significant event for CIT in 1999 was the appointment of Polley Ann McClure as vice president for information technologies in March. She arrived on campus later that year. Before coming to Cornell, McClure was vice president for information technology and communications and professor of environmental sciences at the University of Virginia. At Cornell, McClure also was appointed professor of ecology and evolutionary biology. McClure had been very active at the national level, serving as a member and chair of the board of directors of both Educom and CAUSE and in 1999 as the chair of the board of EDUCAUSE, the new organization that combined the two previous organizations concerned about IT in higher education. VP McClure's appointment brought a renewed sense of stability to CIT, which had been without top leadership for almost 18 months and had also undergone some significant losses in director-level staff.

Formation of the Faculty of Computing and Information Science

To emphasize the role information technology was taking across a broad range of disciplines at the university, Randel announced the new post of dean for computing and information science and appointed Robert L. Constable to the new position effective July 1, 1999. Constable had just finished a six-year term as chair of the Department of Computer Science. The action grew out of recommendations in the preliminary report of the Cornell Task Force on Computing and Information Science, which "recommended the creation of a new administrative structure for computer science, tentatively called a 'Faculty of Information Science.'"¹³⁶ The new academic term "faculty" was created to complement the traditional academic structures at Cornell and was roughly suggestive of an undergraduate version of the cross-college fields in Cornell's Graduate School.

During the early fall, there was considerable discussion about this new organizational model, especially

¹³⁶ "New CU dean named for computing, info science," *Cornell Chronicle*, September 2, 1999.

as it seemed that Randel had not discussed this move with the Faculty Council and other interested parties. There was some concern that the budget for the computer science department would now be removed from the Colleges of Engineering and Arts and Sciences, which shared this responsibility, and that the Department of Computer Science might dominate this new arrangement.

Given the strong support for the vision of this new unit, the task force and Constable recognized some of the concerns being expressed by individuals and groups and issued their final report in November 1999. The first sentence in the report, paraphrasing Ezra Cornell's statement, summarized their recommendations: "Cornell University should undertake to become an institution where anyone can bring ideas from computing and information science to bear on any discipline." The main recommendations of the report were:

1. A Faculty of Computing and Information (FCI) should be created, although the actual name may differ from this.
2. The faculty should be the focus of significantly expanded educational and research activity in computing and information.
3. The FCI-affiliated faculty should reflect a balance of faculty on campus with expertise in computing and information areas.
4. Certain academic disciplines are critically important to the success of the proposed faculty, and the computer science department (CS) should be administratively in the FCI from the start.
5. An advisory board representative of the computing and information constituencies on campus should be formed to advise the dean on matters related to the faculty, such as new courses and programs, appointments, and other academic and research program matters.
6. The faculty should develop and oversee a new university-wide undergraduate computing program modeled after the Knight Writing Program.

After acceptance of the report, actions following the recommendations were put into place.^{kk}

CIT Staff Transitions and Turnover

Staff turnover and organizational transfers continued to impact CIT during the second half of 1999. In July Helen Mohrmann^{ll} assumed the position of director of Networking and Computing Systems. This

was an excellent move for both Mohrmann and CIT, as networking would continue to be the most critical campuswide issue for the future, and she could bring her talents to bear on these issues while the future of Administrative Systems was resolved. No interim appointment was made in ASDT, the division Mohrmann had previously directed, and the management staff reported to VP McClure.

In November, Ann Stunden left Cornell to assume the position of director of the Division of Information Technology and chief information officer at the University of Wisconsin at Madison. During her three-year tenure at Cornell, "Annie" had brought a new vision and new ideas to supporting academic computing at Cornell. She left Academic Technology Services a stronger and more focused organization, dedicated to service excellence and the fostering of partnerships.

In December long-time senior administrator in CIT, Alan Personius, took early retirement. For the last five years of his tenure Alan had become the chief financial officer in CIT and had been involved with the then most recent financial restructuring, trying to develop an understanding of and rationale for CIT's budgeting and charging policies and practices. At the time, VP McClure announced some interim moves pending further organizational changes and assignments.

Fred Rogers Leaves Cornell

Also in December, Fred Rogers^{mm} announced he was stepping down from the position of senior vice president and chief financial officer to assume the presidency of a new Internet company to serve the education community.¹³⁷ He expressed the desire to work full-time in the development of innovative business solutions for higher education, an issue high on his priority list during his tenure at Cornell. Other responsibilities in his roles in finance and university operations had detracted from his ability to focus on this keen interest. His vision for administrative computing and business systems had been on the mark, but perhaps too bold for Cornell to accomplish at the time. Neither the technology nor the people were ready. Nonetheless, Rogers deserves a large amount of credit for trying to improve administrative processes on campus, for improving cooperation between central and college units, and for trying to reduce the costs of operations.

In short order, President Rawlings announced that he would recommend the appointment of Harold D. Craft Jr. as vice president for administration and chief

¹³⁷ "Senior Vice President, Frederick A. Rogers will leave Cornell to head new Internet company," *Cornell Chronicle*, December 9, 1999.

financial officer effective February 1, 2000, and that action was completed.

CornellC Mainframe Upgraded to Model 9672-R25

Despite a May 1998 upgrade to the VM operating system to capture efficiency improvements and to remove the last of the local system modifications, performance continued to deteriorate. In addition, by late 1998 it was clear that Project 2000 was well behind schedule and the “legacy” systems would have to continue to provide services for the foreseeable future. Given this situation, aggressive testing was needed to ensure that all such systems would be Y2K compliant. Such testing and regular production could not be done on the existing system. As a result, CornellC was again upgraded in February 1999 from the installed model R24 to a model 9672-R25 to provide a 20 percent to 30 percent increase in computing capacity.ⁿⁿ The system was leased for a period of three years. The system was configured to operate as two logical partitions, or LPARS, one for the Year 2000 project, the other for continuing use by online or batch production business systems, so that they did not interfere with each other.

One of the key applications still using CornellC at this time was the NOTIS system for the Cornell Libraries. It was expected that this system would be converted to a client-server in mid-2000 using the Endeavor system running on SUN servers.

Local Cable TV Network Services

In 1999 Time-Warner Cable, holders of the local cable TV franchise, began offering their Road-Runner^{oo} network services to homes. Using a cable modem attached to the broadband TV cable, homeowners could get shared Ethernet services to their home computer and out to the Internet and other services. The \$40 per month charge competed favorably with the alternative services offered at \$20 per month from other Internet service providers plus the cost of a separate telephone line. In a way, this could be seen as a culmination of the activity that IthacaNet began in 1995 with their conferences on Networking Tompkins County.

IthacaNet—Schrader Keynote Speaker

The fifth Annual Networking Tompkins County Conference was held in March 1999, organized by IthacaNetPP and Ithaca College Academic Services following the same format as in the past except that the conference was held in Statler Hall on the Cornell campus rather than at Ithaca College. William L. Schrader, Cornell '74, the founder, chairman, and CEO of PSINet, Inc., was the keynote speaker.

Schrader had been to Cornell earlier in 1996 when he was the keynote speaker at the 1996 Cornell Society of Engineers Annual Conference discussing “Advances in Information Technology and the Future of the Internet.” Schrader had also been featured in Time magazine as an “Internet mogul,” whose company PSINet was growing worldwide. Schrader was predicting a whole new Internet and web-based business model for industry.

The third IthacaNet Award for Achievement in Networking Tompkins County was given to Road-Runner, Time Warner Cable, Syracuse Division, “in recognition of high-speed online service to residential customers in Tompkins County.” As it turned out, this was the last of the Annual Networking Tompkins County Conferences; the “bubble” burst on the dot-com revolution, with far-reaching consequences for the network industry and Internet-based services as well.^{qq}

Network Security

Security issues continued to be an important and everyday concern during the year. In May, Kevin Unrue was named security coordinator in OIT with responsibilities for overseeing security of the university’s computers, networks, and data. This important position had remained unfilled for over a year, and it was expected that now increasing attention would be paid to this critical issue.

During that year hackers succeeded in breaking into Traveler’s Mail service, and NetID passwords may have been compromised. Users were advised to change their passwords just in case.

The HAPPY99.EXE virus came to Cornell as an e-mail attachment, and users were alerted that they should delete the file and not click on it as they would be infected. This Windows 95/98 computer worm was spreading around the Internet and was seen at Cornell. These were but a few of such incidents reported.

Software Piracy

Software piracy, the illegal and unauthorized distribution or use of computer software, continued to be an issue of concern. In February 1999 the *Cornell Chronicle* published the article “Don’t Be a Software Pirate: The Dos and Don’ts of Copying Software,” which advised students, faculty, and staff about this serious issue. The article was done in a question-answer format and covered the full range of legal and illegal uses of software, the potential consequences of actions that violated Cornell and national policies, and sources to get more information.

More Network Upgrades—On and Off Campus

Activity to upgrade the network and the “edge” systems continued throughout 1999. A proposal asking for \$2.6 million to complete the backbone upgrade was approved in mid-year. To make sure that the currently installed non-Y2K-compliant IBM 6611 routers were totally removed before the end of December, \$1.6 million was advanced earlier in the year. All 108 such routers were replaced during routine morning maintenance and without affecting network services. Once all the new Cisco routers were installed, there were selective upgrades to building or LAN connections that had performance problems. To rework the entire network, the building wiring, the node rooms, the entire cable plant, and Resnet still had to be upgraded. These projects would take place in the year 2000 and after. Assuming all those projects were implemented, the service improvement would be noticeable, as the network speed would be 100 Mbits to every desktop. However, the roughly estimated cost of \$80 million would probably require some selectivity in scope and project timing.

After the 1999 upgrades were completed, the Cornell data network can be depicted as shown in Figure 8.

EZ-Remote

EZ-Remote-LO services were discontinued in July and replaced by Express Lane services, which continued the free modem connection to Cornell but limited an individual to four 15-minute connections per day. The modem speeds were at 56 Kbps, much faster than the previous 14.4 speeds for this service. Users were encouraged to switch to EZ-Remote, now the name for the fee-based service, which also offered connections up to 56 Kbps.

E-mail Activity

During the fiscal year ending June 30, 1999, approximately 160 million e-mail messages were sent and received by the Cornell mail system, compared with 41.2 million messages in 1994–95, the first full year of client-server e-mail services.¹⁷

Central Calendaring Service

In July 1999 the first steps were taken to develop and operate a central calendaring service as a campuswide resource for scheduling meetings. The funding model would make this service available to faculty and staff in all colleges and units at no charge, the same as for e-mail. Up to this time many units were using Meeting Maker for their own local purposes, but this product would not scale to meet the larger campus needs. A

campus task force selected Corporate Time as the best system for Cornell’s needs. Using Corporate Time, users would be able to maintain secure online calendars to schedule meetings, appointments, tasks, conference rooms, and equipment from their Mac OS, Unix, or Windows platforms. Gail Honness, the project leader for CIT, prepared a rollout plan that called for a pilot phase to begin in March 2000 and then increasing the user base to full campus use.⁵⁵

New CIT Publications/Computing at Cornell Web Site

In the fall of 1999 CIT published its first tabloid newspaper, *Computing at Cornell*. Recognizing that information development and distribution have to be priorities for an IT organization to be successful, CIT continued this evolution of its publications to reach interested audiences. With an initial run of 38,000, the tabloid was distributed to all incoming students during fall registration and inserted as a supplement in the *Cornell Chronicle*. The tabloid replaced the aging Bear Access and EZ-Remote brochures. It also referred to the Computing at Cornell web site, where more information could be found. The “@cornell.edu” column continued to appear in the *Cornell Chronicle* as a monthly technology news update.

Computer Science Course Statistics

The Department of Computer Science continued to teach more courses, with an accompanying increase in student enrollment, during the decade as displayed in Table 12.

Table 12. Department of Computer Science Course Statistics during the 1990s

Year	Number of Courses	Enrollment in Courses	Total Cornell Students
1990–91	74	3,431	18,742
1998–99	86	5,678	19,021

FABIT-Supported Teaching Classroom and Laboratory Upgrades

FABIT was allocated another \$750,000 in the fall of 1999 for what was now called FABIT2000 projects. Uris Library and Mann Library were allocated over \$300,000 for computing facilities and, for the library, a laptop loaner program.

By this time FABIT had spent or allocated \$3.5 million on upgrading classroom and facilities, resulting in a total expenditure of \$7 million counting matching funds. Out of the 314 identified for such upgrades, 108

DATA Backbone

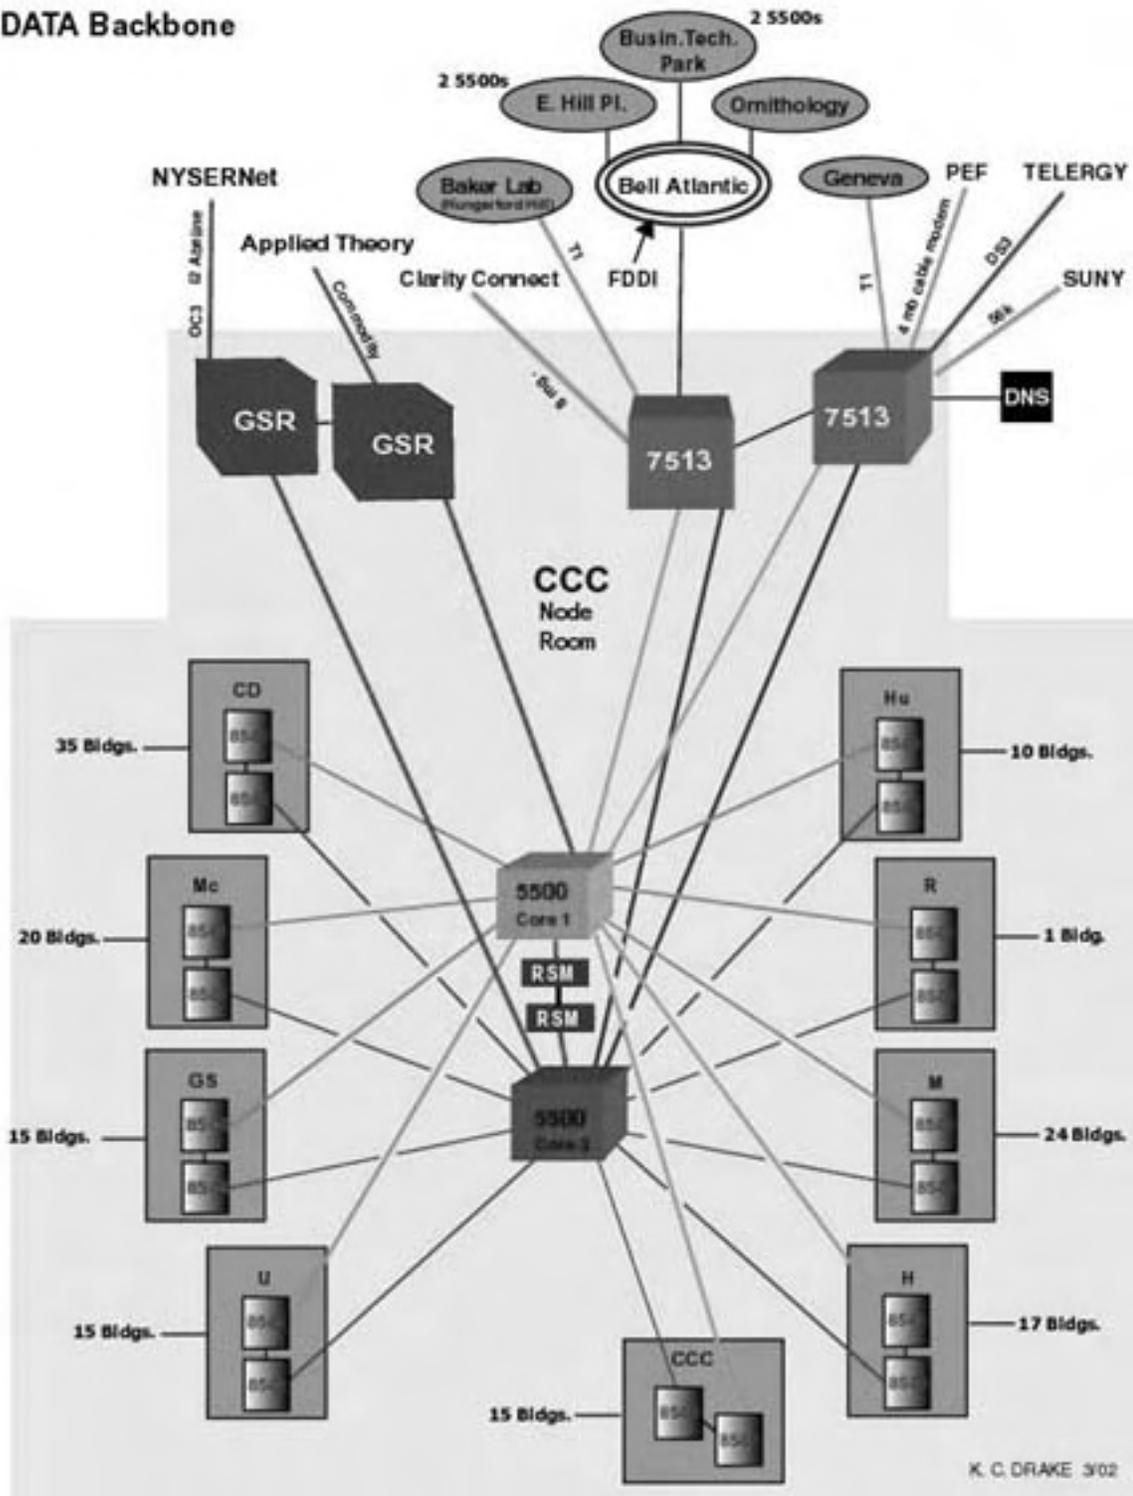


Figure 9. Cornell data networks after 1999 upgrades

classrooms were brought to tier 1 level (an overhead projector with a liquid crystal display to project computer displays, and a network connection). Several academic units, notably the Johnson School and the Hotel School, had brought 100 percent of their classrooms to tier 1 level, while the ILR School had over 80 percent. This was quite an accomplishment over three years. With the interest now shifting to the use of wireless technology in the classroom, the library,

and student lounge areas, future plans would have to consider how to use this technology in classrooms and laboratories.

Bear Access with Runway

Bear Access 99 was released with a new feature: Runway. This feature allowed users to customize the Bear Access window with buttons for applications that

the user preferred. Runway restored the feature to customize windows, available with the first Bear Access, but not available in the first use of Salsa technology. Also in this release, users were encouraged to use web e-mail for those times when they were away from their workstation; later that year Traveler's Mail, previously used for this service, would only be available through EZ-Remote connections or other off-campus access.

EZ-Backup Growth

EZ-Backup was serving 2,000 computers and holding 3,500 gigabytes of data by the end of the decade. This was a spectacular growth of 20 times the number of computers and 80 times the number of gigabytes of storage since the service started in 1995.

Net-Print Activity

By the end of 1999, the daily volume on Net-Print was peaking at about 17,000 pages per day. One day it hit a record of 61,000 pages. Various small service improvements were made, such as printing watermarks and adding small-print identification on pages (as done with faxes) to avoid a separate banner identification page. The main improvement was that the service was extended to more printers across the campus.

Information Technology and the Libraries; HeinOnline

The number of IT projects that involved the Cornell Library or library technology exploded during 1999. Workshops and forums as well as library staff were concerned with topics such as digital libraries of the future, managing digital collections, preserving institutional electronic records, examining interfaces to visual collections, and last, increasing the number and variety of online databases. This progress was quite an achievement given the starting position of the libraries at the beginning of the decade when the first digital library project got under way as a joint study using the Xerox Docutech printer.

One of the CIT success stories, based on CIT's experience of working with digital libraries since early in the decade, was HeinOnline, which provided online access to historical legal materials.¹³⁸ The success of the Making of America Project drew the attention of William S. Hein and Co., Inc, the world's largest distributor of legal periodicals, which was interested in placing its collections on the web. The availability of Dienst, the Cornell-developed digital library protocol, and the expertise of the staff in CIT, the Cornell law library, and the computer science department, led to the success. According to Rich Marisa, who developed the interface and continues to support the service, the location and retrieval of documents is simple

and straightforward. When first put into service, HeinOnline comprised 30 journals and over a half a million pages, presenting the scanned images of the original documents. By 2004, the collection comprised over 12 million pages and was subscribed to by almost every law school in the country as well as many overseas (personal communication from Rich Marisa).

Theory Center Switches to Dell Equipment

The Theory Center ended its long relationship with IBM and started a new relationship with Dell Computer, Intel, Microsoft, and Giganet. The new supercomputing system was called an AC3 Velocity Cluster and was made up of 64 rack-mounted Dell Poweredge 6350 servers, each incorporating four Pentium II chips and running the Windows NT operating system.¹³⁹ Giganet provided a switch that allowed the servers to communicate with each other at 100 megabytes per second. This new system was rated to run at 122 gigaflops, much faster than the 76 gigaflops of the IBM SP that it replaced, and was much less expensive as a result of using off-the-shelf components.

New DataMart, Old Data Warehouses, Mark IV Fades Away

In December 1999 the Payroll DataMart was released, providing additional information and corresponding tools for access to data in the new HR/Payroll system. At this point, the storage of a much larger number of data elements in the new HR/Payroll system and the supporting technologies started achieving some of the objectives set forth when Project 2000 was first announced.

The use of data warehouses that drew their data from the continuing legacy systems, Finance and Student Administration, for example, continued to enjoy expanded and increasing use.

Mark IV, the first high-level query-language/report generator at Cornell, put into use in 1972, was deleted from the system. It had served its purpose for many years, and it was now time to use more contemporary technologies.

¹³⁸ "HeinOnline brings historical legal materials to the web," *CIT News* at www.cit.cornell.edu/computer/news/news00/heinonline.html.

¹³⁹ "At Cornell, a cluster of Pentium processors becomes a super-computer," *Cornell Chronicle*, October 5, 1999.

The Millennium Arrives

The clock ticked down toward December 31, 1999. The Y2K Committee continued to work with college and administrative units and its own constituent members to make sure that Cornell was totally Y2K compliant. Further, it wanted to make certain that Cornell would experience no harmful effects to its physical plant, to its information technology infrastructure, or to individual computers as well as to experiments that needed temperature control or continuous power. All possible media—including newsletters, e-mails, web site references, meetings, and road shows to departments or interdepartmental groups—were used to keep the issue in front of the campus.

As the year progressed, there was increasing confidence, according to John McKeown, who was directing Cornell's efforts, that Cornell would experience few or no problems.¹⁴⁰ But, as he described it, the campus could not relax entirely. Every one of his conversations with NYSEG, the sole provider of electricity to the campus, would end with their disclaimer that while they were ready, there could be problems ahead of them.

McKeown's own words from his oral history capture the time of rollover:

When we got to the very end of the year we decided that if we were encouraging lots of people to have a strong presence on campus, so would we. So our entire committee was around on New Year's eve just to make sure things were going well and to provide whatever help we could. More significantly though, there were hundreds of people from the facilities department who were going around checking on various buildings.... So, there were lots of facilities people around and lots of individual department people around looking after their specific needs. We ended up hanging out in the NOC, the Network Operations Center in Rhodes Hall. Of course there were loads of CIT people around. We purposely planned to take down some of the services because it's not a busy time at the university and people were convinced that if we had some machines in "rest" mode there was less likelihood of any problems if there were power surges or whatever else might happen. So there were lots of us, our entire committee including Polley McClure was around,

¹⁴⁰ John McKeown, Oral Histories—Cornell Computing and Information Technology, www.cit.cornell.edu/computer/history/, December 2000



Figure 10. The Network Operations Center at the Y2K rollover



Figure 11. VP Polley McClure and CIT staff at Y2K rollover



Figure 12. Action in the NOC during Y2K rollover

and we all counted down and looked out the window as the Ithaca College lights changed to '00' and held our breath a bit as we just didn't know what to expect. The bottom line of the whole experience was that we sailed right through it with very little disruption.

The biggest IT event of the decade—the millennium—was a big nonevent. At 2:00 a.m. the campus servers were brought back online and by 12:00 noon all essential services were up and running as usual. This same situation was repeated the world over. While a few individuals who expected the worst were surprised, most technologists were ready and viewed this as just another new year. The way things turned out provided endless opportunities for pundits and technology experts to speculate as to why things happened the way they did and perhaps what would have happened if everybody had not done the kind of work that Cornell did.^{tt} In Cornell's case, and CIT's case in particular, everybody could now focus on the future and all the opportunities to use information technology to further education and learning and research at Cornell University.



Figure 13. CIT Logo, 2003

Endnotes

^a It was a sad commentary on the times that CIT was accused of having too flashy a publication, unbecoming of the organization, even though it was produced at low cost by using advanced technology resources such as the Docutech Publishing System. *CITNews* was toned down, but in time it ended up where all such previous attempts to spread the word and promote the organization also ended: namely, they were abandoned.

^b Reading the copies of *Inside CIT* provided an excellent source of information on CIT plans, priorities, and difficulties for the years 1990 to 1994.

^c It would appear that the Committee on Instructional Computing appointed in 1991 did not come up with any documented recommendations or a report.

^d The December 24, 2000, Sunday edition of the *Syracuse Herald American-Post Standard* had an update on Robert Morris. In March 1999 he received a Ph.D. degree from Harvard. "Morris went on to make millions in the dot-com boom with an Internet startup that was bought by Yahoo. He joined the faculty at Massachusetts Institute of Technology, the nation's top-ranked computer science department. Along the way, he got married in a storybook-style wedding in a mansion overlooking the sea."

^e According to Scott Brim, he and Mark Fedor, another Theory Center employee, were peripherally involved in the development of the SGMP, the Simple Gateway Management Protocol, and the precursor to SNMP. SNMP itself, developed by the IETF (Internet Engineering Task Force), was adopted as a TCP/IP-based management standard in 1989 and has gained widespread acceptance since 1993. According to Brim, the real work on SNMP was done at RPI and Carnegie-Mellon.

^f Project Mandarin was organized as a separate corporation, with VP Lynn as the chairman and Mara as the president, until it disbanded in 1997 and work ended. According to Mark Mara, all or part of the technology was still in use in the year 2000 at Brown, Michigan, and Arizona, while Stanford and others such as Cornell used parts that fit in with their own migration to new software providing the same or equivalent features.

^g Exemplar faded away in 1996–97 when the interest in many places shifted to new client-server-based systems and away from Adabas/Natural mainframe systems

^h According to Cecilia Cowles, a rump group that met at the Statler bar on Fridays developed the format for the NetID. The group included Cecilia, Barbara Skoblick, Tom Young, Larry Chace, Steve Worona, and Stuart Lynn as more or less regulars, but others also joined the discussion. Barbara Skoblick gave the napkin on which the format was scrawled to Stuart Lynn when he left Cornell.

ⁱ Marc Andreessen, who wrote Mosaic, was then a student employee at NCSA. It was the success of this browser that led to Andreessen's leaving NCSA and starting Netscape, which then led to all the "browser wars" when Microsoft integrated Internet Explorer with Windows. This was in part the reason behind the U.S. Justice Department's taking Microsoft to court and Justice

Johnson rendering the judgment to split up Microsoft. The issue was resolved in 2002; Microsoft remained intact but was held to some obligations to allow vendors to pre-install non-Microsoft software with new systems sales.

^j The appointment of the Merten Committee did not go over well with CIT management and staff. While all the interesting projects and accomplishments at CIT produced a great amount of satisfaction and positive feedback from customers, the staff felt they were unappreciated for their hard work and long hours. Many felt they had to go off campus and talk about their success stories to get the positive feedback so important to their technical and personal satisfaction. Although one could look at some of the recommendations as leading to new cooperative working arrangements between the academic and administrative units, it was not clear that the complaining units would change their opinion regardless of what CIT did. There was definitely a loss of morale until a new sense of direction was established.

^k Dave Pulleyn was one of the first employees of the Cornell Computing Center when he joined the staff in 1957 after his first employment with the electrical engineering department at the university. In 1969 he was promoted to assistant director of operations, three years after OCS was formed by combining the Computing Center and Machine Records operations. He worked for many years at Langmuir and was a key player in designing and overseeing the new computer rooms in CCC and the Theory Center. He was the first CIT employee to retire at age 55, although he has continued to work part-time on CIT building and space renovation projects, where his expertise has been extremely helpful.

^l Koehler was at Stanford until late 1995 and from there went to Princeton University, where he was director of information systems. In 2002 he returned to Cornell University as director of business information systems. He is the only member of the IVY+ group of information systems directors to retain continuous membership by working only at institutions that were part of the group. He also holds the record for working at the most IVY+ institutions.

^m While the public reason for this reorganization was that Distributed Technologies (DT) had a broader mission than just the university's business systems, and DT was in fact supporting the development and maintenance of Bear Access, Project Mandarin, and related application programs, this also was a strategic political move. At the time, Fred Rogers, senior VP, was very much making the case that IR should report to him. If this was to take place, the separation ensured that Rogers would get the responsibilities and staff dedicated to supporting the business systems and not the staff involved with the newer technologies of broader interest to Cornell.

ⁿ After his return to California in 2001, Lynn was elected president and CEO of ICANN (Internet Corporation for Assigned Names and Numbers), a position he held until 2003.

^o SCAMP simplified the mailing and improved the quality of information that new undergraduate students received from Cornell between their acceptance and their arrival on campus. The project consolidated over 20 separate mailings into 3, saved on production and mailing costs, expedited the delivery of information, and reduced confusion among the students, parents, and staff. Over 200 people were involved with this project.

^P Worona left Cornell at the end of May 2001 and joined EDUCAUSE as the director of policy and networking in their Washington, D.C., office. Steve had a long and productive tenure at Cornell, starting as a student in 1965. In those 35 years, he produced such innovations as the electronic mail services on mainframes and CUINFO and all its related services, and he helped convince Microsoft that the network was the future of IT. He was vocal in arguing that CIT had to keep ahead of the technology curve, and he took actions to make that happen. One of his last accomplishments was to start the Cornell Computer Policy and Law Institute, and he became a well-known and effective spokesman for those concerns. He continues his association with Cornell through the institute.

^Q Dave Koehler deserves credit for first expressing this concern in 1990–91, as noted in the minutes of CIT senior management meetings at that time. Since the Y2K event was nine years away, the most reaction he got was a friendly smile or a “not my problem now” response.

^R The genesis of CourseInfo appears to have had several roots. It was the interest of Cindy van Es, a senior lecturer in the Department of Applied Economics and Management, to create a web site that would facilitate communication with students in her course. She interested Daniel Cane in taking on this project, which he did. According to the *Cornell Chronicle* of October 17, 1996, he developed his idea during an independent study during his sophomore year. The experience of Margaret Corbit of the Theory Center is somewhat different. She recalls working with Cane and Tim Chi in the fall of 1995 on developing a web site for the Theory Center’s online science book *Explorations*, which was still in use in 2003. That perhaps gave the group the experience they needed. Also, she and Dan worked on creating a virtual exhibit for the Syracuse Museum of Science and Technology as part of their cooperative process. Corbit and Cane continued to be in touch with each other after Cane left Ithaca to pursue his career with Blackboard.

^S Other Cornell notables in the information technology field are Jay Walker, ILR ’78, and Jeffrey C. Hawkins, ENG ’79. Walker is the founder of Priceline.com, one of the enduring “dot.com” companies that rose to prominence in the late 1990s. Hawkins is co-founder of Handspring, Inc, makers of the popular Visor handheld personal organizer. Hawkins invented the PalmPilot products and founded Palm computing in 1994 but left there to start Handspring in 1998. Hawkins was Cornell Entrepreneur of the Year in 2000.

^T While the Merten Committee review was under way in 1993, two divisions in CIT—Information Resources and Computer Resources—took steps to reorganize internally with the intention of improving services. Both divisions tried to leverage the team-based movement introduced by the QIP program that was then generally advocated in the industry and the country as a way to improve staff productivity. In IR it was estimated that productivity would double. It actually regressed because staff were preoccupied with managing themselves and their team-based resources when supervisors were eliminated. This new operating style, combined with a new scheme proposed by consultant John Hupp for paying for IR staff resources, was a near disaster. In CR the reorganization was not much more successful. It worked reasonably well for the shift teams, which tried to eliminate the job distinctions between operators and production controllers in the interest of all working together for customer satisfaction. It failed with the systems programmers, because there was no good focus to forming teams. Nonetheless, in both cases the experience was worthwhile, for it gave staff opportunities for expressing themselves, communicating with their supervisors, and understanding their own accountability.

^U The planning group of Rogers, Lambert, Mara, Cowles, Jones, and Rudan brainstormed about what to call this project. Some leaned toward using the year 2000, as there was some sense that the project could be completed by this year or soon after. I credit Lambert with first suggesting Project 2000!, with the exclamation mark as a sense of the excitement. This was later toned down to just Project 2000. Much later it could be said that the proper label for the project was Project 2000@#%&. Nonetheless, in keeping with the terminology of the times, it became known as P2K, just as the year 2000 issues became known as Y2K issues.

^V In the *Cornell Chronicle* of February 8, 1996, there was an article written by Dean Alan Merten entitled, “Will it happen at Cornell University?” Basically, he covered the difficulties of change in universities in general and Cornell in particular, the reasons for change, and the conditions for change. He concluded with the questions: “Should Cornell change? Yes. Can Cornell change? Yes. Will Cornell change? Frankly, I’m not sure. But I do know that if we combined the talents of our alumni, faculty, administrators, and students, we can not only make Cornell a better place, we can make it the best place.” The following month it was announced that Merten had been appointed president of George Mason University and would be leaving Cornell before July 1.

^W John Robinson and David Duffield were co-founders of Information Associates (IA) in Rochester, N.Y. After selling IA to GE, they went their separate ways. Duffield then founded Integral Systems, another successful company, the vendor of the first DB2-based human resource and accounting system, before starting PeopleSoft. Duffield was a Cornell graduate, Class of ’63 in electrical engineering, and received an MBA in 1964. He was voted Cornell Entrepreneur of the Year in 1996. He has shared his success with Cornell with endowments to the Veterinary College and for the building of Duffield Hall on the Engineering campus—the site of Cornell’s nanofabrication research.

^X Completing the PeopleSoft contract was one of those experiences one never forgets. I ended up as the go-between, alternating with Rogers, Pat McClary in the legal office, and PeopleSoft offices in Boston and California, trying to settle on the contract terms. The dollar amounts had been fixed. It was December 22nd, as I recall, and PeopleSoft had to have a signed contract so they could book the income for that fiscal year. This was not the first time that Cornell, or that I personally, was involved in this kind of year-end crisis. It didn’t help that this was the last working day of the year for Cornell and there was quite a bit of holiday spirit and socializing going on around Day Hall. Amongst all this confusion we did succeed.

^Y No copy of the “blue ribbon” commission’s report could be found to get more information. Various documents prepared before and after the transfer of PC sales to the Campus Store clearly support the outcome.

^Z In 2003 the College of Engineering /Theory Center asked to reclaim a portion of the space in Rhodes Hall assigned to CIT staff. Some of the staff were relocated to 110 Maple Avenue, while those remaining were squeezed into smaller quarters. In 2003 the CCC basement was refurbished for CIT staff space. All the windows removed in the 1985 renovation (in the interests of increased security for the computer room) were restored.

aa In 2004 a new group was formed called the Information Technology Managers Council (ITMC), composed of representatives from each Cornell unit and members of OIT/CIT senior management. While this group was formed as an outcome of the Workforce Planning effort in previous years, it seems very similar to the CIO group proposed 10 years earlier.

bb At the time that Cogger and his group left Cornell, Vienna Systems was an affiliate of Newbridge Networks based in Kanata, Ontario, just outside of Ottawa in Canada's "Silicon Valley." In time and after a couple of corporate changes, Vienna Systems became a part of Nokia Corporation, worldwide leaders in cell telephone technology. The group was phased out in 2002 with the downturn in the telecommunications industry.

cc In 2002 Regenstein took the position of associate CIO/director of the Division of Information Technology at the University of Wisconsin at Madison, where she joined the staff of Anne Stunden, CIO.

dd In August 2000 Siegel was appointed the first permanent chief information officer (CIO) and associate provost at the University of Illinois at Urbana-Champaign.

ee Jim Doolittle retired in 2002 from the position of associate director of Network and Computing Systems after a career of 33 years with Cornell.

ff In 2002 Cornell acquired the Oracle DBMS at a cost of \$1.2 million, about half the price first proposed in 1995, and projects were under way to phase out the use of Informix.

gg In 2000 the EOG (the Executive Oversight Group) was created to govern the processes for evaluating and implementing new administrative systems and enhancing existing systems. The EOG membership included the provost; the deans of the Colleges of Arts and Sciences and Human Ecology and the Law School; the vice president of administration and chief financial officer; the vice president for financial planning and budget management; and the vice president for information technologies.

hh The vBNS program was scheduled to end in March 2000 with the end of NSF funding. To maintain this high-speed networking service, CIT worked with NYSERNet to form NYSERNet2000 to continue high-speed access between universities and research sites in New York State and to provide two gateways to national vBNS networks. One was MCI Worldcom's vBNS+ and the Abilene network from the University Corporation for Advanced Internet Development, a national collaboration of over 170 universities that has partnerships with government and industry. The estimated cost was \$150,000.

ii Douglas Van Houweling was appointed chief executive officer and president of University Corporation for Advanced Internet Development (UCAID) in 1997. UCAID is the not-for-profit corporation that leads higher education's continuing role in Internet development, its major current project being Internet2.

jj In the year 2000, vice provost Mary Sansalone announced the formation of eCornell, a for-profit corporation to distribute Cornell course content via the Internet and other modes of "technology-mediated learning." The corporation would be controlled by Cornell but financed by private capital, with startup costs estimated to be \$50 million or more. This was a model advanced by other universities on their own, as at Cornell, or by forming consortia with other institutions. In effect, eCornell would become the distance learning outlet for Cornell. Most faculty did not readily accept the model at the time.

kk The dean of the faculty, Robert J. Cooke, called a meeting of all faculty members to present more information on this new Faculty of Computing and Information Science, the name used at the time. I attended that meeting and was fascinated by the different opinions, not necessarily objecting to the formation, but strongly objecting to the process by which Provost Randel had made this decision. From a purely personal point of view, I was pleased to see the study of computer science move to the forefront of concerns in the undergraduate program at Cornell and be as strong as the graduate program in the field. Soon after the final report was issued in July 2000, Provost Randel left Cornell to assume the presidency of the University of Chicago.

ll In January 2000 Mohrmann was appointed executive director of CIT, reporting to VP McClure. In this new position and role Mohrmann was to oversee CIT's academic and administrative computing services as well as the campus telephone and networking infrastructures. In the announcement, McClure emphasized that this was a new position to unify CIT's operations. In 2002 McClure and Mohrmann agreed that Mohrmann would take the position of director of administrative systems planning to bring her skills to assist this important activity and to fill a gap during a job freeze that prevented the hiring of a new director for Business Systems. Other temporary assignments were made in CIT at this same time to accommodate this change and as a result of the freeze.

mm With the resignation of Rogers, only Mara, Rick Jones, and Tom Boggess of the original planning team were still at Cornell. Lambert had earlier left for Georgetown, and Rudan retired in 1997. In the fall of 1997 Bob Cowles was dismissed along with six or so others in a rapid one-day downsizing of ASDT staff, a dramatic break with past termination practices. Seeing a 25-year employee terminated in this way caused a lot of anguish and loss of morale for CIT staff. In time Bob took a position with the Stanford Linear Accelerator as head of security for their PeopleSoft systems. Tom Boggess left Cornell in 1998 and joined PeopleSoft as a trainer for about a year but returned to become a member of the student area technology team for implementing a new student system. In 2000 Mara was promoted to director of distributed technologies architecture, and Jones was promoted to associate director, business information systems management.

nn This was the third mainframe upgrade during this decade despite VP Lynn's strong commitment to having the 1995 upgrade from the 3090-200J to the 390 9762-R32 be the last one at Cornell. Even more interesting is the fact that there was another upgrade in 2002, which improved performance (Adabas use still increasing at 15 percent to 20 percent each year) at a reduced cost as a result of technology and marketing practices.

^{oo} The use of the TV cable for computer networking only took about 15 years to come to Ithaca. I recall having discussions with Ken King and Mike Withiam, manager of the local cable service back in 1984 or so, about using the cable as a higher-speed alternative to telephone modems. This was about the time the Theory Center was coming into being and Ken Wilson was raising this same issue. Two-way communication was a strange concept to American Community Cablevision, or ACC, the franchise owners at the time, as they thought in terms of broadcasting signals out from their site to individual TV sets or homes.

^{pp} Acknowledgments should be given to the individuals and organizations that were part of IthacaNet and who sponsored the five conferences held in Ithaca. According to several sources, the core group included Mark Anbinder, representing several organizations over the period; Jean Currie, from the South Central Regional Library Council; Dan Dwyer, from the Cornell Theory Center; Bill Kaupe, from the City of Ithaca; representatives from BOCES and Ithaca College at different times; and Steve Worona, from Cornell Information Technologies. Jim Harper, initially at Cornell but later as an independent consultant, set up and took care of the web site materials and conference programs.

^{qq} In 2001 William Schrader was replaced as chairman and CEO of PSINet; although his contributions to the company and visions for the industry as a whole were lauded, the company was facing an uncertain future and possible bankruptcy. In 2001 Cogent Communications acquired the major U.S. assets of PSINet and continued operations under that name after integrating the joint assets. PSINet continues to operate in different countries around the world as independent companies.

^{rr} Cornell processed 360 million e-mail messages in the fiscal year 2002—a million a day—a growth of 2.25 times compared with the volume in 1999.

^{ss} Gail Honness reports that Corporate Time became a production service in June 2000, and in January 2003 it was being used by 6,500 staff and faculty.

^{tt} In an informal postmortem it was roughly estimated that the Y2K project cost about \$8 million in effort and real dollar expenditures across the entire university. This amount included \$2.5 million for the Medical College. In CIT, for example, it is estimated that 18 person-years of effort were expended. In addition, \$750,000 in mostly real dollar funds was budgeted to buy software and hire consultants for remediation efforts on 4.5 million lines of legacy systems code. Also, additional funds were spent to upgrade the mainframe so that Y2K testing was not impeded by production during 1999. That same postmortem documented about 16 “bad date” problems of little or no consequence as they were mostly display errors, and perhaps another dozen problems with malfunctioning central systems software. These date-related problems were quickly fixed on discovery.

Ron Parks notes that the one incident that “burned” him was when students enrolling for the spring semester 2000 using CoursEnroll in the fall of 1999 discovered that their display read the spring semester 1900. This got some splash coverage in the *Cornell Daily Sun*, which perhaps did more good than harm because it highlighted the problem. In defense of Parks, CoursEnroll had been scheduled but not remediated—it just hadn’t been on the schedule soon enough to stay ahead of the students. In any case, this was not a serious problem—just a matter of fixing a display of a date, which did not affect the basic student record.

Appendix

Computing at Cornell—Timelines

Contents

Organizations and Leadership	266
<i>Organizations</i>	266
<i>Leadership</i>	266
Governing and Advisory Boards	267
<i>Campuswide Committees/Boards</i>	267
<i>Year 2000 Project Committee</i>	269
<i>Academic Advisory Boards</i>	269
<i>Administrative Advisory Boards</i>	270
Locations for Central Computers and Satellite Stations	270
Computer Information	271
<i>Growth in Computing Power at Cornell</i>	271
<i>Central Computers—Timeline and Configurations</i>	271
Network Information	273
<i>Telecommunications and Networking Timeline</i>	273
<i>Abbreviated Cornell Networking Timeline</i>	274
<i>Detailed Cornell Networking Timeline</i>	274

Organizations and Leadership

Organizations

The names of the central computing organizations evolved over time. In the early days the name reflected the technology or the target population and service mission. With time, as computing devices dealt less with computational problems and more with information processing, the use of the term information technology became the popular reference.

1947 to 1966

Machine Records (Administrative)

1953 to 1966

Cornell Computing Center (Academic Research and Instruction)

1966 to 1980

Office of Computer Services (OCS)

- Combined academic/administrative center
- Merger of Cornell Computing Center and Machine Records

1966 to 1980

Computer Activities Group (CAG)

- Academic computing support for the statutory colleges

1980 to 1988

Cornell Computer Services (CCS)

- CAG merged with CCS in 1980.
- Office Equipment Center merged with CCS in 1987.

1988 to 1999

Office of Information Technologies (OIT)

- Cornell Information Technologies (CIT)
- Telecommunications merged with CIT in 1989.
- Microcomputer Sales/Support transferred to Campus Store in 1997.

Leadership

Until 1980 when Ken King was appointed the first vice president for computing at Cornell, the computing service organizations were led by a director reporting to an executive-level position. The organization charts presented are typically for when the organization had stabilized after new leadership.

1953 to 1964

Richard C. Lesser, director, Cornell Computing Center (reporting to T. P. Wright, vice president for research)

John W. Rudan, assistant director

1955 to 1966

Dominic Bordonaro, director, Machine Records (reporting to A. H. Peterson, controller)

1964 to 1966—John W. Rudan, director, Cornell Computing Center (reporting to F. A. Long, vice president for research); David H. Bessel, assistant director

1966 to 1967

David H. Bessel, director, Cornell Computive Center

1966 to 1968

Richard W. Conway, director, Office of Computer Services (reporting to F. A. Long, vice president for research)

Jeremy E. Johnson, associate director

Sid Saltzman, assistant director for academic services

Bill Worley, assistant director for systems programming

Dom Bordonaro, assistant director for administrative services

1968 to 1970

Erik D. McWilliams, director, Office of Computer Services (reporting to F. A. Long, vice president for research)

Richard C. Cogger, assistant director for systems

David W. Pulleyn, assistant director for operations

Erik D. McWilliams, acting assistant director for academic services

Peter Shames, manager of academic services

David Jennings, assistant director for administrative services

1970 to 1977

John W. Rudan, director, Office of Computer Services (reporting to S. A. Lawrence, vice president for administration)

Richard C. Cogger, assistant director, Systems

David W. Pulleyn, assistant director for operations

Douglas E. Van Houweling, assistant director, Academic Computing

1971 to 1977

Henry G. Vaughan, director, Management Systems and Analysis (MSA) (reporting to S. A. Lawrence, vice president for administration)

Judy Campbell, associate director; Edmund V. Hollenbeck, assistant director

1977 to 1978

Robert R. Blackmun, director, Office of Computer Services (reporting to S. A. Lawrence, vice president for administration)

Douglas E. van Houweling, associate director,

Academic Computing

Edmund V. Hollenbeck, assistant director,
Administrative Programming Services
Richard C. Cogger, assistant director, Systems
David W. Pulleyn, assistant director, Operations

1979 to 1980

John W. Rudan, acting director, Office of Computer Services (reporting to J. R. Cooke, special assistant to the provost)
Douglas van Houweling, associate director and director, Academic Computing
Edmund V. Hollenbeck, assistant director, Administrative Programming Services
Richard C. Cogger, assistant director, Systems
David W. Pulleyn, assistant director, Operations

1980 to 1987

Kenneth M. King, vice president, Cornell Computer Services
John W. Rudan, assistant vice provost, Operations Services
Russ Vaught, director, Administrative Computing
Gordon Galloway, director, Academic Computing

1987 to 1988

Norm Scott, acting vice president for computer services
John W. Rudan, assistant vice provost, Operations Services
Russ Vaught, director, Administrative Computing
Gordon Galloway, director, Academic Computing

1988 to 1994

M. Stuart Lynn, vice president, Information Technologies
Ruth Sabeau, director, CIT Services
Larry Fresinski, director, Workstation Resources
Dave Lambert, director, Network Resources
John Rudan, director, Computer Resources
Dave Koehler, director, Information Resources

1996 to 1997

H. David Lambert, vice president, Information Technologies
Ann Stunden, director, Support Services and Academic Computing
Helen Mohrmann, director, Administrative Systems/
Distributed Technology
Peter Seigel, director, Network and Computing Systems

1998 to 1999

Thomas R. Dyckman, acting vice president, Information Technologies
Ann Stunden, director, Academic Technology Services
Helen Mohrmann, director, Administrative Systems
Peter Seigel, director, Network and Computing Systems

1999

Polley A. McClure, vice president, Information Technologies
Helen Mohrmann, director, Administrative Systems Planning
Tracy Mitrano, director, IT Policy, and co-director, Computer Policy and Law
Steve Schuster, director, IT Security
Robert (Dave) Vernon, director, IT Architecture
Rohit Ahuja, director, Administration and Finance
David Koehler, director, Business Information Systems
Jason Rhoades, director, Network and Communication Services
Charles “Wes” Kahle, director, Customer Services and Marketing
Eric Fredericksen, director, Distributed Learning Services
Mark Mara, director, Integration and Delivery
Rick MacDonald, director, Systems and Operations

Governing and Advisory Boards

Campuswide Committees/Boards

During the existence of central computing at Cornell, starting with the 1950s, there were advisory boards of faculty and staff who advised the provost or other responsible senior university administrators on the delivery of computing/information technology services to the campus. They advised on organizations that were thought best able to deliver services as well as the computing devices and their deployment for those services. Financial issues were an important part of the deliberations and recommendations of these groups. At times there were separate boards to make recommendations for the entire campus as well as for the differing academic and administrative needs and customers.

The descriptions that follow are from documented membership lists found during the examination of different records in the computing archives; names and titles were taken from those records. These are not year-by-year listings of members.

1953 to 1961—Advisory Board, Cornell Computing Center

T. P. Wright, Vice President for Research, chairman
Robert J. Walker, Mathematics

J. Barkley Rosser, Mathematics
W. T. Federer, Biometrics Unit, Department of Plant
Breeding

**1961 to 1962—Executive Committee, Cornell
Computing Center**

J. Barkley Rosser, Mathematics, chairman
Richard W. Conway, Industrial Engineering and
Administration
Walter T. Federer, Biometrics Unit, Department of
Plant Breeding
S. S. Atwood, provost (*ex officio*)
Richard C. Lesser, director (*ex officio*)
Seymour V. Parter, associate director (*ex officio*)

**1963 to 1967—Executive Committee, Cornell
Computing Services**

Franklin A. Long, vice president for research, chairman
Richard W. Conway, Industrial Engineering and
Administration
Walter T. Federer, Biometrics Unit, Department of
Plant Breeding
R. J. Walker, Mathematics
J. S. Harding, Child Development and Family
Relationships
Richard C. Lesser, director (*ex officio* to 1964)
John W. Rudan, director (*ex officio*, 1964 to 1966)

**1967 to 1970—Advisory Board, Office of Computer
Services**

Franklin A. Long, vice president for research,
chairman
U. Bronfenbrenner, Child Development and Family
Relationships
D. A. Edwards, Physics (1974)
A. G. Feldt, Sociology (1975)
E. L. Gasteiger, Physical Biology
R. Henderson, Animal Science (1974)
R. Hoffman, Chemistry
W. R. Lynn, Civil Engineering
W. L. Maxwell, Industrial Engineering
M. S. Nelkin, Engineering Physics
A. H. Peterson, university controller
G. Salton, Computer Science
W. G. Tomek, Agricultural Economics
R. J. Walker, Mathematics
R. L. Walsh, Finance and Business Office, statutory
colleges
Later additions were
R. McGinnis, Sociology
S. Smidt, Business and Public Administration

**1971 to 1988—University Computing Board
(UCB)**

1971 (when formed)

Samuel A. Lawrence, vice president, Administration,
chairman

Nyle Brady, College of Agriculture and Life Sciences
Richard W. Conway, Computer Science
H. Justin Davidson, Dean, Graduate School of
Business
Arthur H. Peterson, controller
Geoffrey V. Chester, Nuclear Studies
Wallace B. Rogers, University Support Services

1975

Samuel A. Lawrence, vice president, Administration,
chairman
Richard W. Conway, Computer Science
W. Donald Cooke, vice president for research
A. Schultz, Operations Research
W. Keith Kennedy, dean, College of Agriculture and
Life Sciences
Arthur H. Peterson, controller
D. E. Van Houweling, Government
G. V. Chester, Faculty Council of Representatives
Jack Lowe, Sponsored Programs (Information Systems
Advisory Board)
J. H. Williams, Computer Science and OCS
(Academic Computing)

1977

Samuel A. Lawrence, vice president, Administration,
chairman
Richard W. Conway, Computer Science
J. Robert Cooke, Agricultural Engineering
June Fassenden-Raden, vice president for undergradu-
ate education
Ivor Francis, Industrial and Labor Relations
R. E. Hughes, Chemistry
J. S. Ostrom, controller
Andrew S. Schultz, Jr., Operations Research

1978

Thomas E. Everhart, dean, College of Engineering,
chairman
G. V. Chester, Nuclear Studies
Juris Hartmanis Computer Science
Timothy D. Mount, Agricultural Economics
J. S. Ostrom, controller
Don M. Randel, vice provost
Douglas K. Reece, Business and Public Administration
Paul F. Velleman, Industrial and Labor Relations
Richard N. White, Civil Engineering

1983

J. Robert Cooke, Agricultural Engineering, chairman
G. V. Chester, Nuclear Studies
Juris Hartmanis Computer Science
Timothy D. Mount, Agricultural Economics

J. S. Ostrom, controller
Don M. Randel, vice provost
Douglas K. Reece, Business and Public Administration
Paul F. Velleman, Industrial and Labor Relations
Richard N. White, Civil Engineering

1985

J. Robert Cooke, Agricultural Engineering, chairman
Hal Craft, Facilities
Jane Hammond, Law
Juris Hartmanis, Computer Science
John Lewkowicz, Veterinary Medicine,
Charles McClintock, Human Ecology,
Robert McGinnis, Sociology,
Richard Moore, Hotel School,
J. S. Ostrom, controller
Chris Pottle, Electrical Engineering
Vithala Rao, Graduate School of Management
Sidney Saltzman, City and Regional Planning
Paul F. Velleman, Industrial and Labor Relations
Richard N. White, Civil Engineering

1990—Cornell University Board on Information Technology (CUBIT)

1990

J. Morley, senior vice president, chairman
Mal Nesheim, provost
Fred Rogers, vice president for finance and treasurer
John Weisenfeld, vice president for planning
Norm Scott, vice president for research
Alain Sez nec, director, Cornell Library
Alan Merten, Johnson Graduate School of Management
Stuart Lynn, vice president for information technologies

Note: The CUBIT charter called for the creation of a University Advisory Committee for Information Technology (UCIT), whose chair was to serve on CUBIT. There is no documentation about UCIT and if it was ever formed.

Year 2000 Project Committee

1998 to 2000—Year 2000 Project Committee

John McKeown, director/project leader
Steve Worona, Office of Information Technologies
Cecilia Cowles, CIT (communications to campus)
Keith Boncek, Facilities and Campus Services
Dennis Butts, Central Purchasing
Mary Bouchard, University Audit Office
Ron Parks, CIT (central business systems)
Dan Batholomew, CIT (desktop issues)
Rich Marisa, CIT (research issues)
Jim Doolittle, CIT, (network and computer contingency planning)

Mark Bodenstein, CIT (mainframe, server, and network issues)

Academic Advisory Boards

1970 to 1972—Academic Computing Advisory Board

1970

Goeffrey V. Chester, director, LASSP, chairman
Erik D. McWilliams, OCS
Roald Hoffman, Chemistry
Ronald B. Furry, Agricultural Engineering
Robert McGinnis, Sociology
Howard L. Morgan, Operations Research
D. Hywel White, Physics
Chris Pottle, Electrical Engineering

1971 to 1980—Advisory Committee on Instructional Computing (UCB subcommittee)

1971

S. Saltzman, Policy Planning and Regional Analysis, chairman
S. Edelstein, Biochemistry and Molecular Biology
I. Francis, Industrial and Labor Relations
Ronald B. Furry, Agricultural Engineering
N. Lyons, Business and Public Administration
J. Maas, Center for Improvement of Undergraduate Education
C. Paul, Basic Studies, Engineering
D. Van Houweling, Government
J. H. Williams, Computer Science
J. W. Rudan, Office of Computer Services

1973

D. van Houweling, Government, chairman
G. V. Chester, Physics
Ronald B. Furry, Agricultural Engineering
Susan Gold, Business and Public Administration
J. Maas, Center for Improvement of Undergraduate Education
D. P. Greenberg, Architecture
Paul Velleman, Industrial and Labor Relations
J. H. Williams, Computer Science
J. W. Rudan, Office of Computer Services

1971 to 1978—Subcommittee on Research Computing (UCB subcommittee)

H. A. Scheraga, Chemistry, chairman
D. Hywel White, Physics (acting chairman)
D. P. Greenberg, Architecture, Art, and Planning
S. R. Searle, Biometrics Unit
H. E. Aldrich
W. L. Maxwell, Industrial Engineering and Administration
R. McGinnis, Sociology

1988 to present—Faculty Board on Information Technology (FABIT)

1993

Ronnie Coffman, associate dean, College of Agriculture and Life Sciences, chairman
Stan Bowman, associate dean, College of Architecture, Art, and Planning
Steve Ealick, Biochemistry
Geri Gay, Communications
Bruce Halpern, Psychology
Dan Huttenlocher, Computer Science
Alan McAdams, Johnson Graduate School of Management
Charles McClintock, associate dean, College of Human Ecology
John McRae, Asian Studies
Walter Mebane, Government
Tim Mount, Agricultural Economics
Anil Nerode, Mathematics
Tom Owens, Plant Biology
Robert Thomas, Electrical Engineering
Nancy Tuma, Sociology
Note: There were three FABIT subcommittees: Faculty Support, Student Access to Information Resources, and Classroom Facilities.

Administrative Advisory Boards

1967 to 1970—Administrative Systems Planning and Control Board

Note: No other information or other references to the board's existence were mentioned in available reports and correspondence.

1981 to 1988—Administrative Priorities Committee

W. G. Herbster, senior vice president
James Spencer, vice provost
Jack Ostrom, controller
Ken King, vice president for computing

1989 to 1990—Administrative Systems Steering Committee (ASSC)

Robert Barker, provost
Jay Morley, senior vice president
Mal Nesheim, vice provost
M. Stuart Lynn, vice president for information technology
Non-voting members: G. Mueller, university auditor; David M. Koehler, director, CIT Information Resources; Mike Whalen, University Budget

1990 to 1999—Administrative Data and Systems Policy Advisory Committee (ADSPAC)

1992

Fred A. Rogers, treasurer and vice president for finance, chairman
Richard W. Banks, Public Affairs
Harold D. Craft, Jr., Facilities and Business Operations
Allan A. Lentini, Human Resources
M. Stuart Lynn, Information Technologies
Larry I. Palmer, Academic Programs and Campus Affairs
Yoke San Reynolds, controller
Norman R. Scott, Research and Advanced Studies
John R. Wiesenfeld, Planning
Non-voting members: David W. Koehler, IR; Michael B. Dickinson, Internal Audit; Mike Whalen, Financial Planning and Budget

1994

Fred A. Rogers, treasurer and vice president for finance, chairman
Richard W. Banks, Public Affairs
Harold D. Craft, Jr., Facilities and Business Operations
Beth Warren, Human Resources
H. David Lambert, Information Technologies
Susan H. Murphy, Academic Programs and Campus Affairs
Yoke San Reynolds, controller
Norman R. Scott, Research and Advanced Studies
John R. Wiesenfeld, Planning
John W. Rudan, IR, non-voting member

1996

Fred A. Rogers, treasurer and vice president for finance, chairman
Ingeborg T. Reichenbach, Public Affairs
Harold D. Craft, Jr., Facilities and Business Operations
Beth Warren, Human Resources
H. David Lambert, Information Technologies
Susan H. Murphy, Academic Programs and Campus Affairs
Yoke San Reynolds, controller
Norman R. Scott, Research and Advanced Studies
Ronald G. Ehrenberg, Planning
Joycelyn R. Hart
Helen T. Mohrmann, ASDT, non-voting member

Locations for Central Computers and Satellite Stations

The computers and equipment supporting both academic and administrative efforts were moved to different locations every few years, beginning in 1953 when the first computer was installed on campus until 1967, when the Langmuir location was then used for almost 20 years. The establishment of the Computing and Communications Center in the former Comstock Hall in 1986 was to have provided a more permanent home

for these systems. However, in 1997 when changing circumstances in the Cornell Theory Center presented an opportunity for cost sharing of facilities and staff, all CIT computers and the Network Operations Center were moved to Rhodes Hall. In 2003 there were over 250 CIT servers supporting different applications and services.

1948 to 1967	Day Hall, Machine Records
1953 to 1956	Rand Hall, Cornell Computing Center
1956 to 1959	Phillips Hall, Cornell Computing Center
1959 to 1967	Rand Hall, Cornell Computing Center
1967 to 1986	Langmuir Laboratory for OCS/CCS
1986 to	Computing and Communications Center (CCC) (formerly Comstock Hall; remodeled and renamed)
1997 to	Rhodes Hall; location for CIT servers and Network Operations Center (NOC)

Satellite Centers—Academic Computing

Upson Hall, the first to be established, 1967
 Clark Hall, 1967
 Warren Hall (Computer Activities Group), 1967
 Riley-Robb Hall, 1967
 Uris Hall, 1972
 Martha Van Rensselaer Hall, 1982
 Carpenter Hall, 1982
 Uris Library, 1984
 McFadden Hall, 1985
 Clara Dickson Hall, 1985
 McGraw Hall, 1986
 Pleasant Grove Apartments, 1986
 Sibley Hall, 1983

Satellite Centers—Administrative Computing

Day Hall, 1973 to 1986
 Warren Hall, statutory business office, 1964
 East Hill Plaza, endowed and statutory accounting office, 1980

Computer Information

This chronological listing of the large computers that were installed at the computing center at Cornell comes largely from the “Computer Chronology” compiled by David W. Pulleyn, longtime director of computing operations. It has been supplemented with materials provided in various reports from the Cornell computing archive and from correspondence with Mark Bodenstein, keeper of the records after Pulleyn retired in 1993. It reflects configurations at points in time, largely when new systems were installed, and does not fully convey the dynamics of changes to the

computer itself and to the supporting equipment during the period that system was in use at Cornell.

Growth in Computing Power at Cornell

Table 1 is an attempt to provide some sense of the growth of the power of the different computers that were installed at the central site. The rating used is MIPS—millions of instructions per second. Experts argue that this is not a useful indicator of computer performance, but it is an available and consistent indicator over the time period. The table is intended to display the relative, not absolute, growth for the IBM systems only.

Table 1. Relative Power of IBM Computers Installed at Cornell

<i>Year</i>	<i>Computer</i>	<i>MIPS</i>	<i>Relative MIPS</i>
1956	IBM 650	0.001	1
1967	IBM 360/65	0.68	680
1974	IBM 370/168	2.3	2,300
1983	IBM 3081-K	13.5	13,500
1986	IBM 3090-200	28	28,000
1990	IBM 2090-200J	45	45,000
1995	IBM 390: 9672-R32	58	62,000
1997	IBM 390: 9672-R24	88	62,000
1999	IBM 390: 9672-R25	117	117,000

By way of comparison, the 2MHz 8080 chip from Intel used in microcomputers in the middle 1970s was rated at 0.64 MIPS, roughly comparable to the IBM 360/65 in the above table. In 1989 the Intel 33 MHz 486 chip used in IBM PCs was rated 27 MIPS, roughly comparable to the IBM 3090-200 installed at about the same time.¹

Central Computers—Timeline and Configurations

1953

IBM 605 Card Programmed Electronic Calculator (Research/Rand Hall)

Card input: IBM 418 Accounting Machine (tabulator)
 Card output: IBM 527 Reproducing Punch
 Printed output: offline IBM 418 Accounting Machine (tabulator)

¹ Ken Polsson, 2003 Chronology of Personal Computers, www.islandnet.com/~kpolsson/comphist/

1956
IBM 610 Calculating Punch and Card Equipment (Administrative/Day Hall)

1956
IBM 650 Electronic Computer (Academic/Phillips Hall)
Memory: 2,000 10-decimal digit words on magnetic drum
Card input: IBM 533 Read-Punch Unit
Card output: IBM 533 Read-Punch Unit
Printed output: offline IBM 402 Accounting Machine (tabulator)

1959
Burroughs 220 (Research and Instruction/Rand Hall)
Purchase cost: \$601,000
Memory: 5,000 10-decimal digit words of magnetic core
Card input/output: IBM 087 collator @ 240 cards per minute (cpm)
Printed output: IBM 407 printer @ 150 lines per minute (lpm)
Magnetic tape drives: four units, using one-inch-wide tape
Paper tape input/output: part of teletype console printer

1961
IBM 1401 Electronic Computer Card System (Administrative/Day Hall)
2,000 storage positions (but no multiply/divide commands)

1962
Control Data 1604/160A (Research and Instruction/Rand Hall)
The 160A computer was a small support computer performing card to tape, tape to print, tape to plot operations.
Purchase cost: \$1.065 million
Memory: 32,768, 48-bit binary words
Card input: 1,200 cards per minute
Card output: IBM 523 punch at 50 cards/minute
Printed output: 1,000-lines-per-minute drum printer
Magnetic tape drives: 8 units, one-half-inch-wide, recording at 200 bits/inch (bpi)
Plotter: Calcomp three-color ink
Paper tape input/output

1964
IBM 1401 Electronic Computer Tape System (Administrative/Day Hall)
Removed in 1971-72

1967
IBM 360/65 Computer (leased) (Combined Academic/Administrative System/Langmuir Lab)
Memory: 262,144 bytes of main memory, increased to 524K in 1968 Ampex core was used to upgrade memory to 2 Mbytes from 1 Mbyte in 1970-71.
Eight 2314 disk drives (25.7 Mbytes/drive, transfer rate of 312K bytes/sec)
One 2301 drum (4 Mbytes with transfer rate of 1.2 Mbytes/sec)
Eight 2402-2 tape drives (800 bytes/inch, 60 Kbytes/sec transfer rate); four with seven-track recording, four with nine-track recording
IBM 360/20 for card reading, tape preparation, printing from tapes; two 1403-N1 printers (1100 lpm, 132 characters/line); one 2540 card reader/punch (reads at 1,000 cpm, punches at 300cpm)

1974
IBM 370/168
Purchase cost: \$3.72 million
3 Mbytes of main memory eventually upgraded to 6 Mbytes
Upgraded with VMA (Virtual Memory Assist) in 1976
Processor upgraded from Model-1 to Model-3 in 1979
Supplemented with two Floating Point System Array Processors (FPS 190-L in 1978 and FPS 164 in 1982)
Converted to 1600/6250 bpi from 800 bpi magnetic tapes in 1975
Converted to IBM 3330 disk drives from IBM 2314 disk drives

1980
DEC System 2060
Main memory: 1 Megaword (1024K) of 36-bit words
Three RP06 magnetic disk drives and two TU77 magnetic tape drives
Upgraded to a model 2065 in 1984
Removed in 1986

1981
IBM 3081D
Single processor, water cooled
Purchase cost: \$2.5 million
Estimated to provide three times compute power of the 370/168 it replaced
Estimated to cost \$400,000 less per year to operate
High-performance option (HPO) installed in 1983
Performance machine assist (PMA) installed in 1983
Upgraded to 3,081K in 1984
A separate IBM 4381 installed for instructional use

1986

IBM 3090-200

2-processor system, bipolar technology, water cooled
Purchase cost: \$4.1 million
64 MB of main storage; 64 MB of expanded storage
32 channels, 81.7 Gigabytes of IBM 3380 DASD (at removal)

1990

IBM 3090-200J

2-processor system, bipolar technology, water cooled
Purchase cost: \$3.1 million
Estimated 50 percent increase in processor speed
128 MB of main storage; 256 MB of expanded storage
Life extended by software improvements in 1992

1995

IBM 390 Parallel Enterprise Server Model 9672 - R32 (leased)

Three-processor system, CMOS technology, air cooled
Estimated 30 percent to 40 percent increase in processor speed
512 Mbytes of main memory
40 channels, mixed parallel and IBM ESCON fiber channels
First IBM computer with IBM ESCON fiber channel
IBM RAMAC II DASD array subsystem replaced all 3380 DASD
IBM3494 tape robotics system installed earlier that year

1995

IBM 390 Model 9672-R24 (leased)

Two-processor system, CMOS, air cooled
Estimated 50 percent increase in processor speed
RAMAC and tape robotics systems expanded
When the CIT Computer Room was moved to Rhodes Hall, one 3494 tape robot was left in CCC and connected by fiber channel; the other was moved.

1999

IBM 390 Model 9672-R25 (leased for three years)

1 Gbyte of main memory
Two logical partitions (LPARS): one for Year 2000 Project; one for production use
16 ESCON channels, 18 parallel channels
2 RAMAC 9394 subsystems; 240 Gbytes using 3390 emulation
22.5 Gbytes of 3390 disk storage (DASD)
Two 3494 tape robots: one in CCC, one in Rhodes Hall.
In the year 2000, MVS was managing a total of 1.1 Terabytes of data; 363 Gbytes on active disk storage, 154 Gbytes in migration under storage management (HSM), and 648 Gbytes as back-up copies.

Network Information

Telecommunications and Networking Timeline

Remote communications, that is, between a user at a site different from the site of the central computer, began at Cornell in 1965² when the IBM 360/65 was located at Langmuir Laboratory, five air miles from campus. That required the use of special and ordinary telephone lines for transmitting data to and from the 360/65 to slow-speed typewriter-like terminals and high-speed card-reading and printing equipment at campus locations. By the late 1970s, connections were made to national networks (Tymnet and Telenet), but connections between Langmuir and the campus still used individual telephone circuits. Plans were made to use one of the national network technologies to build a campus network, but this failed to materialize.

The first major network innovations at Cornell occurred during the 1980s. X.25 multiplexers were installed to permit the sharing of high-speed telephone circuits between Langmuir and campus (to gain economies), followed soon after by the use of higher-speed coaxial cable provided by the local TV cable company. This development was followed by the installation of a coaxial cable-based network—Sytek—running between select campus locations and Langmuir, providing faster and less expensive services for direct-connected slow-speed terminals. In 1982 Cornell was an early participant in BITNET and later played a key role in routing traffic to Canada.

At the same time the Computing and Communications Center was established on campus in the mid 1980s, Cornell installed its own telephone system and wiring plant, using a combination of optical fiber for the backbone and twisted pairs of copper wire for secondary connections. The telephone system was operational in 1986 and was operated by Cornell Telecommunications.

During that same mid-decade time frame, Cornell played a key role in the creation of NSFNet and NYSERNet, acquiring the expertise to deal with TCP/IP network technology. In this same period, CCS installed Pronet using the TCP/IP protocols to create two network backbones, TheoryNet and CCSNet, connecting the computers at Langmuir and campus. At this point, Cornell was connected to the then-forming Internet. When the computers were moved

² The first known digital computer-to-computer (network) connection at Cornell was an innovative project carried out by Dave Bessel in 1963–64 when he transmitted information between the Control Data 1604 and the Burroughs 220 computers using a very long length of paper tape that snaked perhaps 50 feet across the floor between the paper-tape readers of both computers in the two different rooms. Both Dave and Tom Dimock recall the episode in their Oral Histories available at www.cit.cornell.edu/computer/history/

to the new Computing and Communications Center on campus in 1986, all the network equipment moved accordingly. To provide better support to network operations and respond to problem reports, Network Operations Centers (NOCs) were established in the Theory Center and CCS.

Given the distance limitations of Ethernet and the speed limitations of LANs such as Appletalk, Cornell built its own TCP/IP network connections that could accommodate multiple LAN technologies. Based on the AT-Gateways (so named for using the IBM PC-AT as the network router), they started to be deployed in 1987–88 and lasted through the decade.

The most significant developments occurred in 1990s when, in a joint project with IBM, Cornell installed a campuswide high-speed network using 100 Megabit FDDI protocol as the backbone and 10 Megabit UTP Ethernet for LANs and workstation connections. This major project was facilitated by combining the Network Communications units in CIT and the Theory Center and the Telecommunications (telephone system) unit in Facilities into a single organizational entity, Network Resources. At this time the separate NOCs for the Theory Center and CIT were combined into a single unit as well.

In the mid-1990s, several projects were carried out to investigate the use of ATM network technology to increase network speeds (“cells in frames”) and also for workstation telephony, but these did not lead to any changes at Cornell. Later in the decade several upgrades to this network were carried out, changing vendors (CISCO routers for the IBM ones) and technologies (Gigabit Ethernet for FDDI and switched Ethernet to the desktop) to increase network capacity. These upgrades allowed CIT to deal with increased use coming not only from increased use of e-mail, for example, but from video and image traffic. When the CIT computers were moved to Rhodes Hall in 1997, the NOC operations were combined with mainframe/server operations to improve coverage of all systems.

The timelines that follow include first a simplified format (Table 2), then a more detailed (almost annual) description of the network technologies at Cornell. They are reconstructed from multiple sources, such as annual reports, weekly reports, newsletters, technical summaries, and oral histories. As a result, they are incomplete and subject to information gaps and errors. The help of Mark Anbinder, Dave Auerbach, Dan Batholomew, Dave Bessel, Mark Bodenstein, Scott Brim, Dick Cogger, and Ken Downey was invaluable in creating them. In particular, special thanks are due to Dick Cogger, who, over much of this period, played such an important role in both the selection of network technologies and their development and implementation and without whose help the timelines would have been more incomplete.

Abbreviated Cornell Networking Timeline

Table 2 only shows five-year increments and presents the information from the perspective of an individual user at a slow-speed typewriter-type or video-based terminal.

Detailed Cornell Networking Timeline

This more detailed timeline presents the information in the year or the period of years the events occurred; changes to networks were not recorded in a consistent manner and, in most cases, were made within a calendar or fiscal year.

1965 to 1969

Dial-Up Connections

- Acoustic coupler modems (110 baud) were used to connect typewriter-based terminals with the IBM 360/65 computer using dial-up telephone connections.³

Fixed Connections

- Telpac (40 Kbits/sec) was used to connect IBM Model20 RJE's at Clark, Upson, and Warren Halls with the 360/65 at Langmuir.

Telephone System

- New York Telephone Company (later NYNEX) provided a Centrex system for campus, local, and long-distance services.

1970 to 1973

Dial-Up Connections

- Acoustic coupler modems became available at 300 baud.
- Direct-connection modems at 300 baud became available from New York Telephone and from others as the telephone companies were required to permit the connection of equipment from other vendors.

Fixed Connections

- Unitech RJE terminals replaced the Model 20s, and the expensive Telpac circuits were replaced by 4,800 bps and later 9,600 bps non-Telco modems on four-wire leased circuits.

³ 110 baud and 110 bps are equivalent ways of referring to the same underlying signaling rate. The name baud came from the pre-existing Teletype practice, while bps, Kbps, and Mbps are usages developed in the computing era. In all cases, the actual data throughput is less due to overheads and packing densities. For example, at 110 baud, the maximum net data throughput would be 70 bits/sec while with 10 Mbit Ethernet, the maximum net data throughput is approximately 9.5 Mbits/sec.

Table 2. Abbreviated Networking Timeline

Year	Dial-Up Speed (Home Access Speed)	LAN Speed	LAN Technology	Backbone Speed	Backbone Technology
1965	110 baud or 0.110 kbps				
1970	300 baud or 0.30 kbps				
1975	300 baud or 0.30 kbps				
1980	1,200 baud or 1.2 kbps	9,600 baud shared with 12 to 24 terminals; 9.6 kbps shared	ENA multiplexer using 9,600 baud shared		
1985	2,400 baud or 2.4 kbps	360 kbps (shared) 9,600 baud shared with 2 to 24 terminals; 9.6 kbps shared	Appletalk LAN Sytek Broadband Network (coaxial cable)	10 Mbps backbone	Pronet Token Ring supporting TCP/IP
1990	14.4 kbps	1–2 Mbps (shared); 1,000–2,000 kbps 10 Mbps (shared)	Omninet LAN with AT-Gateways UTP Ethernet LANs	100 Mbps	FDDI supporting TCP/IP
1995	28.8 kbps	10 Mbps (shared)	UTP Ethernet LANs	100 Mbps	FDDI supporting TCP/IP
2000	56 kbps cable or DSL at 0.5 to 10 Mbps	10 Mbps (individual) 10,000 kbps	UTP Switched Ethernet	1,000 Mbps	Gigabit Ethernet supporting TCP/IP

1974 to 1979

Dial-Up Connections

- These continued to be a mix of acoustically coupled and direct-connected modems at speeds up to 300 baud. Upgrades to 1,200 baud and full-duplex began to appear.

Fixed Connections

- A large number of leased circuits from Langmuir to points on campus, including terminal clusters and individual sites, used Vadec and other modems at speeds up to 1,200 bps. Initially, these modems required relatively expensive four-wire leased circuits, but technical advances and regulatory adjustments in the period allowed a migration to much less expensive alarm-grade circuits. Over 100 such circuits were in place.

Off-Campus National/International Networks

- Tymnet and Telenet public data networks were installed. Tymnet service had 30 to 60 ports, Telenet, 8 ports. Intended for slow-speed terminals, both networks operated at 110 or 300 baud initially, later upgrading to 1,200 baud.

Campus Networks

- A proposal was made to evaluate Tymnet and Telenet for the campus network. A prototype Tymnet network was installed and tested.

Telephone System

- A WATSBOX was installed for long distance dialing to reduce the cost of this service.
- A Centrex system from New York Telephone Company (later NYNEX) provided campus and local services.

1980 to 1981

Dial-Up Connections

- The service was upgraded to 2,400-baud modems, full duplex.

Fixed Connections

- ENA multiplexers were installed using X.25 packet switching technology. One 9,600-baud line from campus to Langmuir (alarm-circuit quality) carried 12 to 24 terminal connections.

Off-Campus National/International Networks

- Tymnet and Telenet services continued.

Campus Networks

- Sytek broadband network (using a combination of packet switching and broadband transmission technologies) was installed, connecting Langmuir with selected buildings on campus. A coaxial cable from Warren Hall to Langmuir was leased from the local cable TV provider. This was a special construction, not connected to their local cable TV system.
- Connections to Sytek could be made anywhere the coaxial cable went using a box that accommodated two terminals at 9,600 bps transfer rate.

1982 to 1984

Dial-Up Connections

- 1,200-baud modems were upgraded with flow control and full duplex transmission.

Fixed Connections

- IBM 3270 terminals were installed. The controller on campus was connected to Langmuir using 9,600-baud bisynchronous lines. By the end of the year, 26 controllers were on campus supporting 335 terminals and 35 printers.
- Synchronous links at 4.8, 7.6, and 9.6 kbps were migrated to broadband synchronous modems using the cable to Langmuir to replace many leased-line facilities from New York Telephone. These links served RJE terminals, ENA multiplexors, and the growing numbers of 3270 clusters.

Off-Campus National/International Networks

- Cornell was an early member of BITNET with leased lines running at 9,600 baud to CUNY in New York City and McGill University in Montreal and across town to Ithaca College.
- Cornell provided the BITNET link to Canadian NetNorth sites through the University of Toronto (UTORVM).
- The computer science department was connected to ARPANet with a 9,600-baud link.

Campus Networks

- An Ethernet network in the Engineering College connected several buildings with classic yellow cable.

1985

Dial-Up Connections

- Dial-up to Sytek was made available at 300, 1,200 and 2,400 baud, and from there to Tymnet or CIT mainframes. This replaced separate dial-up pools to the different services and computers.

Optical Fiber Infrastructure

- As part of the construction of the incoming new telephone system, new optical fiber was installed to approximately 20 buildings on campus. Roughly half of these fiber runs were for future data needs beyond the needs of the telephone system.

Off-Campus National/International Networks

- The computer science department was connected to ARPANET (the then Internet), to CSNet (Computer Science Network), and UUCP-Unix to Unix Copy, a utility protocol for computer-to-computer file transfers that was superseded by FTP.

Campus Networks

- Pronet (Token ring product from Proteon supporting TCP/IP protocols running at 10 Mbps) was installed for two interconnected campus backbones, CCSNet and TheoryNet. On campus, Pronet used the installed fiber runs. To reach computers at Langmuir and the computer science department computers in Upson, specially modified broadband modems from Fairchild were used at those locations and in Uris Hall. This 10-mile-long round-trip run gave Cornell the distinction of having the longest Pronet ring ever installed.
- The connection of the campus networks to the computer science department provided the campus's first access to the developing Internet via the department's ARPANET connection.
- Sytek, Pronet, and the bisync links shared the coaxial cable from campus to Langmuir, but used independent channels.

1986 to 1987

Dial-Up Connections

- The Sytek rotary was expanded from 64 to 96 ports just before the move from Langmuir to CCC in December 1986.

Off-Campus National/International Networks

- The Theory Center assumed operational responsibility for NSFNet, initially linking the five national supercomputer centers and the National Center for Atmospheric Research with 56 Kbps telephone lines.
- Cornell was a forming member of NYSERNet, supported in part with a \$1.2 million grant from NSF. Initial transmissions were at 56 Kbps, providing a much more robust connection for the campus to the Internet.
- Connections were now available to ARPANET, CSNET, NYSERNet, and NSFNet—all TCP/IP based—and also to BITNET.
- Cornell was an early participant in BITNETII (BITNET over TCP/IP) using VMNET software from Princeton.

Campus Networks

- The Theory Center organized the Network Information Support Center (NISC) to monitor and support Theory Center networks.
- AT-Gateways (developed at Cornell) used the installed copper wire for the Cornell telephone system to connect campus LANs to the Pronet backbone. One AT-Gateway could connect up to four LANs

using Ethernet, AppleTalk, or Omninet (for mixed Macs and PCs) LAN technologies. Eventually, 150 Appletalk networks were connected to Pronet using the AT-Gateways.

- TN3270 (locally developed) for the Mac and the IBM-PC compatibles was first used for access to main-frame-based administrative business systems via the campus TCP/IP network infrastructure. In the following years this technology allowed Mac and PC users all across campus to replace the previously growing network of expensive bisync 3270 clusters. Unlike other campuses, Cornell was able to use the same network for administrative and academic uses.

Telephone System

- On March 3, 1986, the new Cornell telephone system (PBX), System 85 from AT&T, went into service, providing campus, local, and long-distance services. The WATSBOX for long-distance services was removed.

1988

Expanded Role of Networking

- A Networking Task Force was appointed by Norm Scott, interim VP, to develop a plan for linking all computers on and off campus. The final report, concluded by incoming VP Stuart Lynn, recommended establishing the position of a director of networking.

Off-Campus National/International Networks

- Tymnet service was discontinued; Telenet services continued.
- The newly engineered NSFNet backbone went into operation at the end of July with a 1,544 Mbps T1 line and new switching systems, replacing the old 56 Kbps backbone.
- Merit Computer Network in Ann Arbor, Michigan, took over management of NSFNet from Cornell.
- The Theory Center served as a node for the Ontario universities' TCP/IP network to connect to NSFNet.

Campus Networks

- There were 100 AT-Gateways in operation; Sytek use was diminishing.
- The CIT Network Operations Center (NOC) opened in the summer.

1989

Off-Campus National/International Networks

- Campus users could access HEPnet (High Energy Physics Network) through Wilson Lab, and SPAN (NASA's Space Physics Analysis Network), through Space Sciences.
- Cornell and Princeton led an effort to make BITNETII into a set of fully connected core nodes plus many other nodes to accommodate increased traffic. Cornell served as an "INTERBIT gateway," routing e-mail between BITNET and the Internet.

Dial-Up Connections

- In June, 2,400-baud dial-up to Sytek was available using modems that cost \$500 to \$700.
- In July, 300-, 1,200-, and 2,400-baud dial-up service was available.
- Demand for better than dial-up access from off campus was increasing. In subsequent years, a variety of initiatives, including use of the TV cable and services such as ISDN and DSL, were investigated.

Telephone System

- AUDIX (Audio Information Exchange) voice mail began in October.
- The Department of Telecommunications became Telecommunications Services in the Network Resources Division of CIT, putting responsibility for the Cornell communications wire plant, used for both voice and data, under one management.

1990

Network Infrastructure Upgrades

- Negotiations for a major joint study with IBM led to a two-part project that both extended the optical fiber infrastructure and created and deployed a new level of backbone technology on campus.

Off-Campus National/International Networks

- The NSFNet backbone was upgraded from 1.5 Mbps to 45 Mbps.
- Cornell was awarded a contract for providing and managing all international NSFNet connections from INRIA in France and NORDUNet in Sweden, working in cooperation with SPRINT.
- PSINet replaced Telenet services in February.
- The first digital videoconferencing link between the ILR Conference Center on campus and ILR Extension in New York City was put into service.

Campus Networks

- UTP (Unshielded Twisted Pair) Ethernet from David Systems was selected as the technology for providing "commodity-level," low-cost data communication services to the desktop using wire already installed in campus buildings.
- The CIT Network Management Center (NMC) was formed by combining the Theory Center Network Information Services Center (NISC) and the CIT Network Operations Center (NOC).
- A pilot project was started to provide network services in residence halls with the installation of UTP Ethernet in 328 rooms in Clara Dickson and Mary Donlon Halls.

Dial-Up Connections

- Modem connections were upgraded to provide 9,600-bps service. Even so, increasingly very few dial-ups were used on campus, and this service was mainly used by off-campus users.

1991 to 1992

Off-Campus National/International Networks

- A fractional T1 link was installed between Ithaca and Cornell University Medical College in New York City
- 56-Kpbs dial lines were installed between the ILR Conference Center in Ithaca and ILR Extension in New York City.
- CSNet ceased operations.

Campus Networks

- As a result of the IBM Joint Study, a new FDDI backbone was installed (100 Mbps) using Cornell routing software on IBM 6611 routers.
 - The continuing expansion of the UTP Ethernet resulted in the installation of UTP connections in more than 1,000 offices across the campus (starting in 1991).
- Bear Access, the Cornell-developed launch pad for network services, was put into use.

1993 to 1997

Rapid Growth and Deployment

- All the major service offerings grew rapidly in this period using the technology base laid down at the beginning of the decade. These offerings included UTP Ethernet service fed by the FDDI backbone, which expanded the modem pool, the interconnections to external networks, and the number of telephone users. At the end of this period, intensive investigations were under way to determine the next technology level in all these areas.

Off-Campus National/International Networks

- The first international videoconference was held using CU-SeeMe software for the Global School House Project sponsored by NSF.
- The first IthacaNet conference was held.
- Lightlink, the first commercial Internet service provider (ISP) in Ithaca, was founded by Homer Smith.
- Three ISPs now served the Ithaca area: Lightlink, Baka, and Clarity Connect.

Campus Networks

- Research to investigate the use of ATM (Asynchronous Transfer Mode) was conducted to see if this technology could be deployed using “cells in frames” to both increase network speeds at a modest cost and support use of the network for telephone transmissions.
- RESNET services (shared Ethernet) to all campus residence halls began in September; there were 540 subscriptions out of 3,000 available.
- BITNET service ended at Cornell.
- EZ-Publish and EZ-Backup network services began in 1995.

- Appletalk forwarding ended; LANs now used Ethertalk.
- LAN Ethernet connections could use switched rather than shared connections.

Dial-Up Connections

- In 1993 EZ-Remote service was made available. EZ-Remote, offering a variety of improvements to dial-up service, focused on pooled resources for off-campus network access. EZ-Remote-LO was a free service; EZ-Remote-HI was a chargeable service at higher speeds. Both offered two ways to access network services—serial terminal access or serial line Internet protocol access (SLIP). SLIP was recommended for use with Bear Access.
 - Toward the end of the period EZ-Remote-HI modems were upgraded from 14.4 kbps to 28.8 kbps, while EZ-Remote-LO continued at 2,400 bps with plans to upgrade to 14.4 Kbps.

Telephone System

- The 10th anniversary of the installation of the Cornell PBX was celebrated in 1996. The system now provided service to 16,000 telephones.
- An investigation into the deployment of workstation telephony was carried out in conjunction with the ATM network proposal.

1998

Off-Campus National/International Networks

- Cornell joined vBNS, the high-speed research network sponsored by NSF, as one of eight New York colleges.

Campus Networks

- The Cornell backbone upgrade started to replace all IBM 6611 routers with CISCO routers and with Gigabit Ethernet.
- Net-Print service was introduced as a pay-as-you-go laser printing service for students with a direct charge to their bursar bill.

Dial-Up Connections

- EZ-Remote-HI provided 160 modems with 56 Kbps transfer speed.
- EZ-Remote LO provided 100 modems with 14.4 Kbps transfer speed.

1999

Off-Campus National/International Networks

- Time-Warner Cable provided local Ethernet network services to the home using cable modems.

Campus Networks

- The Cornell backbone upgrade was completed. All IBM 6611 routers were removed before the Year 2000 rollover. Switched Ethernet was now available to the desktop, with 100 Mbps to the desktop planned to come later.

- Redundant Cisco routers in the nine node rooms replaced the David Systems LAN switches with Cisco LAN switches. **[this doesn't make sense?]**

Dial-Up Connections

- Express Lane service replaced EZ-Remote-LO, providing free access, but it was limited to four 15-minute connections per day at 56 kbps.
- EZ-Remote (previously EZ-Remote-HI) was now a fee-based service, also at 56 kbps.