History of Computing at the University of Saskatchewan: The First Three Decades, 1957-1987

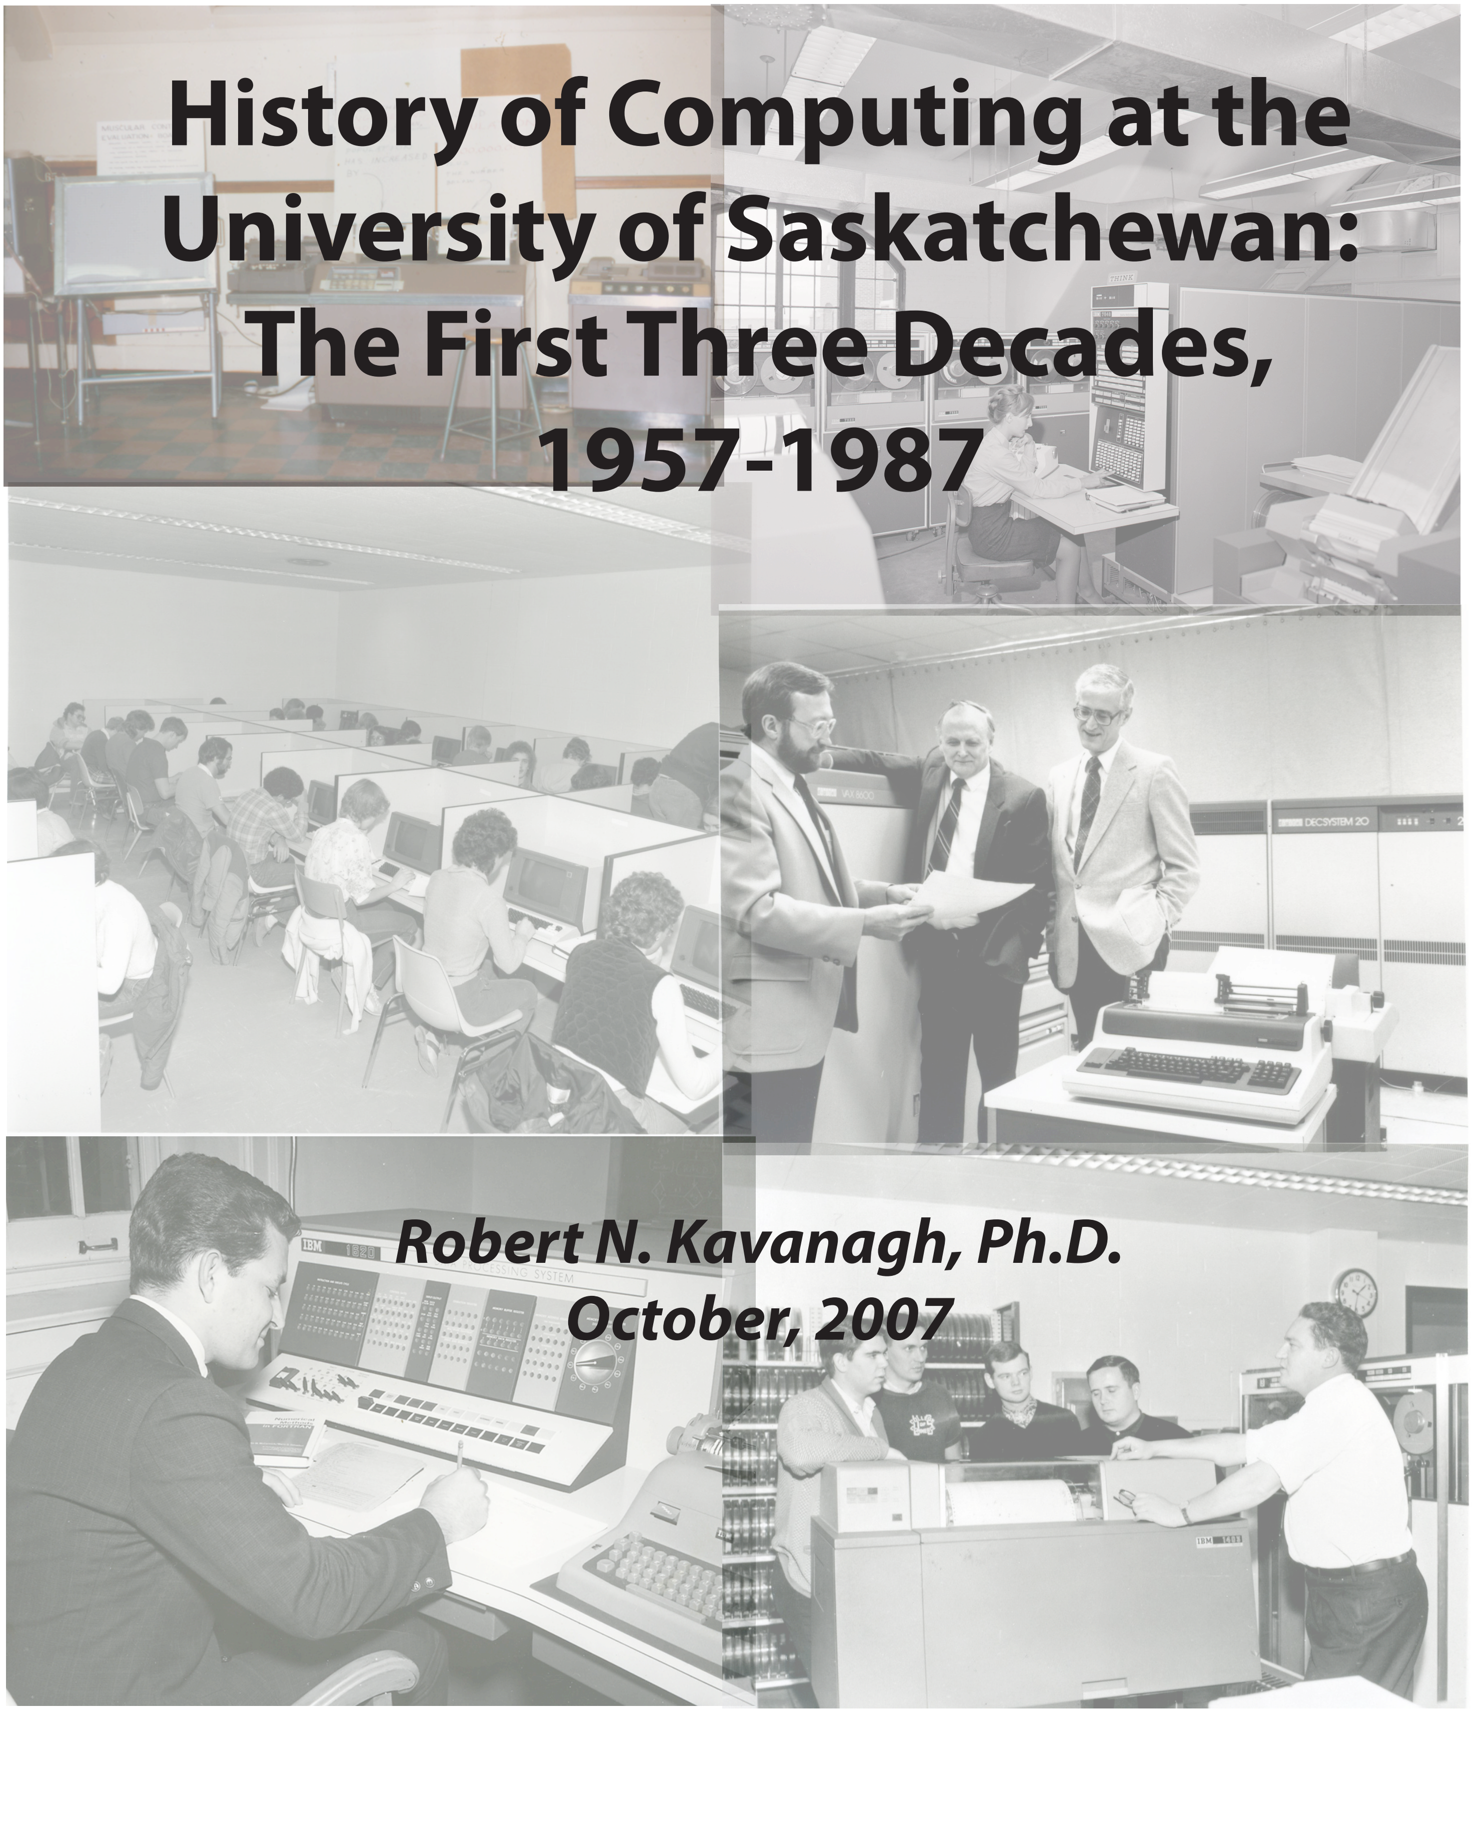

Robert N. Kavanagh, Ph.D.

October, 2007

1. Introduction

This document was created for the Technology Week events in November 2007, marking the 50th anniversary of computing at the University. Many more aspects of this history could have been included: computers in the Library, Administrative applications, technology used in energy conservation, etc. And, most noticeably absent from this memoir is the perspective of the students, faculty and staff who used the facilities and services as they evolved over time. Time constraints have not allowed the inclusion of these important sources in this mostly personal recollection of computing at the University. However, with this effort as a good start, perhaps we can look forward to a more complete version of this history, with the input of more voices.

2. Beginnings: 1957-1966

The University of Saskatchewan entered the computer age in 1957(1), with the installation of a Royal Precision LGP-30 computer in the basement of the Crop Science Building. This was a machine that was designed to fit into a Steelcase™ office desk, and had 4096 words of 31-bit magnetic drum memory. It was jointly owned by the Saskatchewan Research Council, the National Research Council Prairie Regional Laboratory and the University. The machine was used primarily by specialists in the three organizations. At the University, this was mostly faculty in the Mathematics department.

Contrasted with today's microcomputer chips that have tens of millions of transistors and millions of logic elements, the LGP-30 had only 15 flip-flop elements. This machine pre-dated transistor usage in digital computers, and its digital circuitry was built from vacuum tubes.

This era of digital computer design could be characterized in large part by the need to be very imaginative in the design and implementation of the machines, because of the physical restrictions on components available at the time. Designers were almost 3 decades away from the availability of low cost memory, chip technologies that consumed very little power, and components that were physically small. The design of the LGP-30 was, in many ways, a classic triumph over the primitive technologies available at the time. Its designer, Stan Frankel, a Manhattan Project veteran, set out from the start to house his machine inside an office desk from a well known office furniture maker. He achieved compactness of the computer by building a bit-serial architecture, getting each of the 15 flip-flop elements to do multiple tasks, according to which state of the 31-bit serial process the machine was in at a particular instant.

Input to the machine was through a Flexowriter keyboard and paper tape (ten 6-bit characters per second!). Output was primarily to the Flexowriter printer, but a paper tape punch was an option.

The machine was first manufactured in 1956 with a retail price of $47,000 US. It weighed 740 pounds. It contained 113 vacuum tubes and 1450 diodes, and consumed 1500 watts of electrical power. The machine installed at the University cost $30,500 (an early example of educational discount?). Each of the three founding partners contributed an initial

The first digital computer on campus: the LGP-30 machine in the basement of the Crop Science Building, installed in 1957.

$10,500 to acquire the machine and cover the costs of installation and early maintenance.

In 1963, an IBM 1620 machine was installed on the third floor of the Engineering Building. It was a punch-card oriented machine, operated by its users (one at a time). Some programming courses were given in Engineering, Commerce and Mathematics, but there was no full-time support organization for computing until 1965, when an IBM 7040 'mainframe' was purchased with funding assistance from the National Research Council. NRC at that time had a one-time program to 'kick-start' computing capacity and usage at Canadian universities, and for the UofS, this resulted in the acquisition of the IBM 7040, a significant step up from the model 1620. Along with the facility, the Department of Computing Services was created, under its first manager, Mr. Glenn Peardon.

The IBM 7040 was a 'smaller brother' of the IBM 7090/7094 machines, which were powerful mainframes of the day. They were among the first transistorized machines built. The 709/7090/7040 machines were 36-bit word length machines (before we had 8-bit 'bytes'), with an address space of 32,768 words of memory. The installation at the UofS, including some keypunch machines, tape drives and output devices was rented from IBM at a cost of $11,760 per month(2). The 7040 was installed in July 1965 and the 1620 was retired in October of that year.

A personal anecdote about the 7040 is too tempting to exclude. Some of us who were using the machine noted that from time to time our programs were inexplicably aborted, and we were quite perplexed about this circumstance. We went to the manager to see if we were doing something wrong (after all, this was a new, 'fancier' machine than the 1620, to which we were accustomed). He came back to us a few days later, somewhat embarrassed. He reported that

The IBM 1620, installed in the Engineering Building in 1963. Note the amount of switches and knobs that the user had to manipulate to get a program activated.

The IBM 7040, the first 'mainframe' used by the University, installed in the Engineering Building in 1965. The card reader is front right, and in the left foreground is the side of the line printer, where all output was obtained. Several tape drives are shown in the background.

one of the computer operators, who felt that she had a kind of energetic connection to the machine, would abort programs whenever, according to the patterns in the rapidly flickering lights on the control panel, she thought the programs were causing the computer some 'discomfort'. She was given additional training.

By the fall of 1966, (only a little over a year later) the demands for computing were exceeding the capacity of the 7040, and an IBM 1401 was added to handle more work. However, as a sign of things to come many times in the history of computing on campus, the addition of the 1401 was approved only if the Computing Centre could handle it within a reduced operating budget.

The most obvious characteristic of computing during this period is that users had to *walk to the computer centre*, develop programs on punch cards, submit the cards, get their output on paper (or 'intermediate' decks of punch cards) and return to their office or class. This was the era where the only form of computing was 'batch' computing. There were no interactive terminals, no graphical output devices (except primitive pen plotters) and no 'on-line' systems. During this time the first administrative applications were developed for Payroll, Financial Reporting, the Library, the Registrar's Office and University Hospital. Programs were run one-at-a-time; the technology of the day had not yet developed any form of multiprogramming. FORTRAN was the most common 'high-level' programming language, but many users prided themselves on their mastery of 'Assembler language', which was a symbolic representation of the underlying machine operations of a particular machine. Random Access Memory was expensive and small – the 7040 had only 32 K words of 36-bit memory. Programmers had to squeeze results out of very limited program address space and data storage.

The use of digital computers in this period was largely restricted to mathematics, science and engineering. Few, if any, applications to the humanities had been developed. Social science applications were in their infancy.

3. Protracted Decision-Making and the Creation of a Government Computer Utility: 1967-1973

In June of 1967, it was proposed (3) that a new machine be obtained, an IBM 360 Model 50. This was a mid-sized mainframe of that era, and it was expected to provide sufficient capacity until October 1968. It was installed in a new location in room 70 Arts Building, and although it was expected to be in service for only one year, it was still in place by July 1970.

This was a period of emergence of some important new features in the digital computing industry. First, we note that in 1968 the ARPANET was established, which was the experimental platform upon which digital packet-switching communication was developed. Most important to us today, it also developed the first generation communication protocols for the Internet, which came about some fifteen years later. Thus networking of computers was beginning to be a 'hot topic'.

Second, punch card input was no longer the sole means of getting programs and data into the machines. 'Interactivity', in the form of typewriter-sized keyboard/printers, and later, Cathode Ray Tube devices (CRTs), was starting to develop. Users could access the computers from 'terminals' in their office or in shared 'terminal rooms' remote from the 'computer centre', and they could input programs and data, control the execution of their programs, and get their results, without having to walk very far, and without having to punch cards. It is important to note that creating punched cards was not a process tolerant of making keystroke errors: if you mispunched a column, the card was ruined, and you had to start over again. When interactive terminals were introduced, there was at least some primitive form of text editing, so that keystroke errors could be corrected without having to re-type an entire line of your program or data file. Keypunching could be very frustrating, especially for a poor typist, and simple text-editing, even in its primitive forms at first introduction, seemed so much easier to create program and data files.

The systems that provided such interactivity were known as 'time-sharing' computers. Many users could simultaneously be working on terminals around the installation, and even the best of typists could not present much of an input challenge to a powerful digital computer. The operating systems on the machines (the programs that governed the operation of the machine, and allocated its resources among all of the work currently waiting to be performed) gave each user a tiny 'time-slice' of processing time, on a rotating basis, and thus gave the impression (at least when the overall system was not heavily loaded), that each user had the machine to themselves. However, having said that, it should be obvious that the operating systems that provided good time-sharing response had to be designed for that purpose: a processor scheduler for batch processing was an entirely different beast than one that allocated processor time for a time-sharing system. Alas, OS/360, the operating system for the 360 series of IBM machines, was never conceived as a time-sharing system. In fact, some would argue that even as a batch processing facility, it was far from being 'user friendly', because you had to learn Job Control Language, which were the cards preceding your actual program, just to get your 'real' program into the machine.

One step taken in Saskatoon was the purchase in July 1970, of the Hewlett Packard 2000A time-sharing system, which was a minicomputer capable of handling up to 16 simultaneous users (later expanded to 32). It was programmable in BASIC, a new language designed to allow quick creation of simple programs from an interactive terminal. It rapidly became a favourite service on campus, as it was so easy to use, and met the needs of many students and faculty. It was however, only a stop-gap measure to providing time-sharing for the Saskatoon campus, and the demands for more computing capacity, both in Saskatoon and Regina, were not going away.

Lastly, by the late 1960s, most universities had established academic departments of Computer Science, and the development of formal degree

Dr. J. Cooke (right) and students in Room 70 Arts, late 1960s. The IBM Model1403 line printer, one of which they are standing around, was the workhorse of printers for many years. It could print an 11" by 14" page every 2-4 seconds. IBM lore says that one of the design challenges for this printer was to ensure that in the unlikely event that the print chain broke, all the pieces would be contained by the surrounding cabinet, and not injure anyone nearby. The chain contained all of the character shapes, and as it rotated horizontally at very high speed, individual hammers in each print column would strike a character as it flew by, thus printing all of the characters on a line.

programs was well under way. In Saskatoon, the department was formed in 1968, under the headship of Blaine Holmlund, and was called the Department of Computational Science. It was initially created as a department of the whole university, i.e., it did not belong to a particular college, as it was felt that it would have strong programs in all of Engineering, Arts and Science and Commerce. It did not join up with a college until many years later. The number of students studying the discipline was growing, as was the number of students who were majoring in other fields, but wanted to take a course or two in computer programming as an aid to the work in their own disciplines.

Thus in 1967 there were some trends that were driving demand for new functionality, as well as steadily increasing capacity requirements, at the UofS and elsewhere. But that is not the whole story.

There was one other monumental issue at the time. Strictly speaking, when we say 'University of Saskatchewan' as the name of the institution to which this period of history refers, we are ignoring a critical fact. There was one 'University of Saskatchewan', but it had two campuses: one in Saskatoon and one in Regina. The former Regina College was renamed the University of Saskatchewan, Regina Campus, in 1961. The President's office of the University was in Saskatoon, and each campus was headed by a Principal and their respective institutional administrations. This two-campus structure prevailed until 1974, when, following a Royal Commission study (4), the University of Regina was created as a separate institution. Thus in 1967, the decision-making structures required collaboration between two institutions that found themselves often in disagreement. This context had a major impact on the direction of computing in the whole province in general, and the two universities in particular, for the next several years.

It is impossible to say what would have been different if the two campuses had been free to pursue their own computing directions at this time. However, what prevailed for much of the time from 1967 to 1970 was an acutely growing gap in capacity and functionality in Saskatoon, and frustration at trying to develop strategies that fit both campuses, where legitimate differences in needs prevailed. As a final factor, it must be noted that the cost of computing capacity in those days was proportionally much greater than today, so plans for expenditures in this area attracted a lot of scrutiny, especially from government.

As they neared the end of the decade of the 1960s, both campuses of the 'University' found themselves running out of computing capacity. A number of stop-gap solutions were attempted, but it was becoming clear that a major review of needs, and bold responses to the emerging industry trends, were needed. The two campuses conducted comprehensive assessments of their computing needs in 1970-71. Not surprisingly, the functional requirements strongly emphasized time-sharing as the preferred mode of accessing new capacity, on both campuses. In addition, there remained a growing requirement for more batch processing capacity, particularly as the administrative applications continued to be developed and expanded in role and functionality. In October 1971, the Computer Advisory Committee in Saskatoon (5) received a recommendation from the Universities Study Group (that had responsibility to the University as a whole), that the campus needs could best be handled by acquiring <u>both</u> an IBM 370/155 and a DECsystem 10. The former was the platform of choice for batch processing and administrative work, and the latter was de-facto the standard time-sharing machine within the ARPANET community in the United States.

What followed was prolonged inability to agree on strategies to meet the expressed needs by the two campuses. By March of 1972, there was a new recommendation that an IBM 370/155 be installed in Saskatoon and a DECsystem 10 in Regina, with networking connections between the two. The Saskatoon campus, now critically short of capacity, went ahead with the acquisition of the 370 in the spring of 1972.

The 370/155 had 1 Mbyte of RAM – about 1/1000 of the RAM that is installed in a typical desktop computer in 2007. It had 600 Mbytes of disk storage, again only a tiny fraction of what we have on our desktop machines today.

However, these planned investments were being scrutinized by government. As early as November of 1971, the Government of Saskatchewan stated that it was interested in "common solutions to the problem of providing computing resources to the Province". The seeds for a province-wide computer utility were being sown.

The argument for a utility took the approach that computing was expensive, and that economies of scale should rule the acquisition of computer capacity. In 1950 a computer scientist at IBM named Herb Grosch (who was born in Saskatoon) had formulated what became known as Grosch's law: computing capacity increases as the square of the cost of the system. By that logic, the overall (to the province) least expensive strategy to meet the needs of both campuses … and <u>also</u> the entire public sector, would be to buy the biggest machine you could afford, and hook everyone up to it.

This promise of cost-saving by centralization of what was deemed to be an expensive, specialized resource became an unstoppable train. In May 1973 the Government of Saskatchewan created a new Crown Corporation, the Saskatchewan Computer Utility Corporation (SaskCOMP), and the University, composed of two campuses, found itself as the customer of a new government agency. Other customers were the other Crown Corporations and the Government of Saskatchewan Systems Centre.

There were many voices at the Saskatoon campus opposing this development, and the fears of what would happen when the computing service was controlled by an outside agency were strongly stated. As events played out over the next several years, these concerns were significantly validated.

As a final blow, the strongly expressed needs in Saskatoon for time-sharing access were not to be satisfied during the next 6 years of having to be a customer of SaskCOMP. The decision to purchase the IBM 370 was taken, but no action was taken to providing time-sharing services. In hindsight, it might be said that the pressure to do something to add capacity was so great that such a decision at least addressed part of the needs of the campus. However, it is the case that IBM's best efforts to graft time-sharing functionality (via the 'Time Sharing Option', or 'TSO') on to an operating system best suited for batch processing was a source of much frustration, and considerable expense, in the years to come.

As an almost laughable oversight in this whole process, the Saskatoon campus had acquired the 370/155 from IBM at prevailing educational discounts. When it became known that SaskCOMP would assume ownership of the machine, IBM reminded everyone that it would not honour the educational discount for a crown corporation. This snag was resolved by the University buying the system and leasing it back to SaskCOMP! In the end, a facility that did not meet Saskatoon's needs for time-sharing was installed in room 70 Arts, but under new ownership by SaskCOMP, and a Xerox Data Systems Sigma-9 (a time-sharing system) was installed at the Regina campus.

In summary, the period of 1967 to 1973 was characterized by unsuccessful attempts to reconcile the differing computing needs of the two campuses, and woolly thinking about strategy for providing appropriate computing capacity and functionality. Government leaders were swept up by the notion of centralized control of computing, and despite opposition from the Saskatoon campus, the University capitulated. Few paid attention to the fact that Grosch's Law was almost 20 years old, and its validity was already diminished as a guiding factor in acquisition strategy. In essence, bureaucrats took the view that computing in the classroom or research lab was no different than data processing in government. The pain of this faulty thinking was to be felt in Saskatoon for several more years.

4. Restricted Computing Activity: 1974-1978

SaskCOMP assumed the management and operation of the facilities previously owned by the Saskatoon campus, but these facilities remained in the Arts Building, lower floor, for about 5 years. SaskCOMP also hired many of the former staff members from Computing Services. In the late 1970s, the main facility was relocated off-campus, first to temporary space in a warehouse in the north end of the City, and eventually SaskCOMP became a tenant in the new Galleria Building in Innovation Place in 1980.

During the years immediately following the formation of SaskCOMP, computing at the UofS (a separate institution after 1974) suffered continuing restrictions. During this period SaskCOMP service rates increased regularly, and sizably, and exceeded the increases to University budgets. Simply put, SaskCOMP, as a Crown Corporation, had much more expensive staffing and overhead costs than would have prevailed in a university operation, and those costs were passed on to its captive customers. The universities (both in Saskatoon and Regina) were not able to keep up with these increasing service rates. In other institutions computing, both as an area of study and as a tool used in most departments, was growing rapidly. But at the UofS, because of the growing gap between University budget resources and service rates from the provincial utility, we did less and less computing each year. It should be noted that this was at a time when the number of students wanting to study computing was steadily increasing.

To add insult to injury, it became evident that any use of TSO (the time-sharing function at SaskCOMP) was enormously expensive. Because students, who were, after all, trying to learn how to use the tool, often made mistakes in usage, it was not uncommon for a confused and innocent student to incur charges of over $100 in a few minutes of using TSO. Thus the use of TSO by students essentially became prohibited.

SaskCOMP did not see student computing service as important. As an example, Dr. John Cooke, then department head in Computational Science, recalled an event which was potentially catastrophic for student computing. At that time, since time-sharing on the IBM system was not available to students, they prepared their programs on punch cards, fed them to a card reader and waited for a few moments for output to return on a nearby high speed printer. This facility, located in the Arts Building even after SaskCOMP left campus, was called the HOTT Terminal, for Heavily Oriented Towards Turnaround. The system was set up so that student programs, which were usually small and required only tiny amounts of processing time, were given a high priority on the mainframe so that the student could submit their cards and receive their output almost immediately. When you were learning to program this was a great advantage, since student programs often required many changes (finding the errors and punching new cards) before they got them correct. However, one day the turn-around time on the HOTT Terminal, instead of being a minute or two, became 30 minutes or more. When Dr. Cooke complained to SaskCOMP, he was told "Surely you don't expect us to put the interests of students above those of a really important customer" (they were running a payroll for another crown corporation at the time).

While it might sound like an increase in delay from a minute or two to a half hour would not have serious consequences, you need to consider more closely how student work on programming assignments was done. The HOTT facility was entirely 'self-service' - there was no one to gather output and file it in any cubicles for later pickup. You watched the printer, waited until output emerged with your name on it, and went away to correct your program if necessary. With a turn-around time of only a minute or two, this was a task that could fit into time between classes: you might get in two or three 'runs' of your program in 10 to 15 minutes, and you could concentrate on your programming efforts in a very efficient way. At half an hour, the turn-around time prohibited that model of student work, and would require dedicating a large block of time (probably in the evening) to working on your programming assignments, with a lot of unproductive time spent waiting for the results.

This lack of appreciation for impact of service policy changes was typical of the causes for frustration with SaskCOMP. Happily, Dr. Cooke's appeals eventually fell on more sympathetic ears, and at least this dysfunction was corrected.

The period of frustration with being unable to access services that met the needs of the University came to head in 1977, when the UofS embarked on another major review of its computing needs.

5. Re-establishment of University Control over Computing, and a time of Aggressive Expansion in Functionality and Access: 1978-84

In 1977, the Advisory Committee for Academic Computing Services, under the Chairmanship of Dr. A. Wacker of the department of Electrical Engineering, conducted a comprehensive, college by college assessment of academic computing needs. This work re-affirmed the desire for time-sharing access, still unfulfilled after 5 years of making do with only the HP 2000 as a source of simple terminal access to time-sharing services. This work resulted in the release of a Request for Proposal to the vendor community in late 1977.

As a personal aside, I was at this time on sabbatical leave, working within Digital Equipment of Canada Ltd, on network communications products and strategies. For obvious reasons, I personally remained at arms' length from the work of the committee evaluating vendor proposals, and was not a part of any of the selection deliberations.

However, upon my return to campus in July 1978, I was appointed Director of Academic Computing Services, and found myself responsible for the final stages of acquiring a new time-sharing facility for the UofS.

As of 1978, SaskCOMP was still the primary supplier of computing services to the campus, and the Government still believed that it had to be in control of computing services for the University. As such, SaskCOMP was given a seat on the committee that was trying to choose a course of action for meeting future computing needs. While that might seem reasonable at face value, I came into my new job to join an on-going process that was in my mind curious at best and unfair competition at worst. The rules set out by SaskCOMP were that they could, as members of the selection committee, review all of the bids from computer vendors. Then, after the selection had been made, and the University had prepared a detailed plan and budget for acquiring and operating the new service as a university-run facility, SaskCOMP would be allowed to put in a bid to own and operate the facility and sell its services to the campus. The flaws in this process were obvious, but those were the rules.

As I started my tenure as the new Director of

Academic Computing Services, I spent the summer of 1978 finalizing the budget and recommendations for the acquisition of a DECsystem 2050 time-sharing facility, to be located in Arts 70, which had been unused as a computer centre since SaskCOMP had moved out several years earlier. When our work on detailed costs and plans was completed, we had to turn it over to SaskCOMP for them to use to prepare their own bid. As much as we wanted very much to see the university finally get some appropriate time-sharing service and to be able to control its policies of use ourselves, I could not see how SaskCOMP could fail to sustain its control, by virtue of being allowed access to all of the materials from the selection process and all of our budget plans. With all of the information at their disposal, it seemed to me that all SaskCOMP had to do was submit a proposal that was somewhat less costly, or perhaps even the same cost, as the University's own estimated costs. That would have put the Board of Governors in a difficult situation, as a decision by them to refuse such an offer from SaskCOMP would have 'political overtones', as the Government of the day still remained committed to the idea that the proper way to deploy computing resources provincially was through its own crown corporation.

What followed was almost comic. SaskCOMP used portions of our report, verbatim, in their formulation of a bid. However, in their use of portions of our budget plans, they sometimes used the wrong parts, or left out important parts. In essence, we had to coach them through the structuring of their response, as we had to ensure that if they were going to quote our estimates, they should at least quote them accurately.

In the end, SaskCOMP submitted a proposal that offered about 80% of the capacity of the system for about 120% of what the University estimated it would cost to operate the facility itself. Dr. Leo Kristjanson, President of the University at the time, went to bat with the Board of Governors in support of university control and operation of the facility, and he certainly took some political heat over this stance. In the end, the Board agreed, and we took back control of the main academic computing services.

The DECsystem 2050 was installed in early 1979 and was upgraded to a 2060 in 1980. It supported up to 80 simultaneous users, but at that level of activity it was frequently overloaded. It initially had 256K of 36-bit RAM, 450 Mbytes of disk storage and 2 tape drives for user data storage and backup of the system. It had a large number of programming languages including BASIC, FORTRAN, COBOL, SNOBOL, LISP, and others. It had libraries of common packages such as SPSS for statistical use, and graphics packages for plotting output in graphical form. Following its installation, there was a rapid growth in its use, as the needs first expressed in 1970 were finally met with some adequate functionality in 1979. As an academic facility, it was especially helpful for the scientists who had very large computing requirements, for hours of processing time, to do chemical modelling, structures analysis, etc. The faculty in these areas has been restricted by SaskCOMP service rates for several years, and the DEC-20 offered substantial

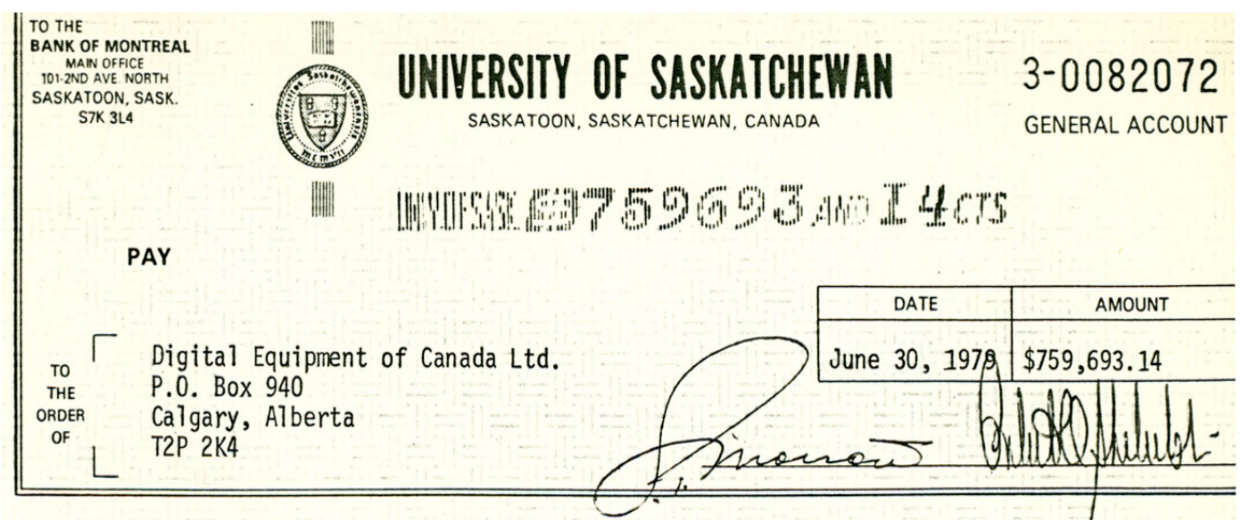
TO THE
BANK OF MONTREAL
MAIN OFFICE
101-2ND AVE. NORTH
SASKATOON, SASK.
S7K 3L4

UNIVERSITY OF SASKATCHEWAN
SASKATOON, SASKATCHEWAN, CANADA

3-0082072
GENERAL ACCOUNT

759693 AND 14 CTS

PAY

DATE	AMOUNT
June 30, 1979	$759,693.14

TO THE ORDER OF
Digital Equipment of Canada Ltd.
P.O. Box 940
Calgary, Alberta
T2P 2K4

The cheque which paid for the DECsystem 2050 in 1979.

capacity (for its day) to do 'number-crunching' at off-peak hours.

The DECsystem 2050 cost over $750,000 in 1979 dollars - about $2,105,000 in 2007 dollars. This underscores an important aspect of the history of the computing industry, not only at the UofS, but in general. In 2007, we are accustomed to being able to purchase truly remarkable amounts of processing and storage capacity for amounts easily handled by most personal credit cards. However, that was certainly not the case in the past. Consider the comparison shown on this page in disk storage, comparing 1979 with 2007. The DECsystem 2060 typically used disk units that held about 175 Mega-bytes, at a cost of about $55,000 each, in 2007 dollars. They were the size of a household washing machine, taking up about 1 m^2 in floor space. One can now purchase a 500 **G**iga-byte drive that will fit in the palm of your hand, for only about $370 (approximately a million-fold reduction in cost per unit of storage). The latter unit has the capacity of 2,857 of the drives on the DECsystem 20, and is smaller than an average textbook - as opposed to requiring almost 3,000 m^2 of floor space (about the space of fifteen-2000ft^2 bungalows, with every square meter of the floors covered with disk drives, and much of the basement occupied by air conditioning). The electrical consumption requirements would show a similar order of reduction as the price per storage unit.

Why is this important? It is important to understand that people using computers twenty years ago were constantly aware of the limitations on their efforts caused by restrictions on data and program storage space. ALL of the several thousand regular users of the DEC-20 fit their storage requirements in **less than 1Gbyte of user disk space**. Nowadays RAM and disk space are so inexpensive and physically small that hardly anyone ever worries about not having enough of either of them. The 'Year 2000 Problem' although mostly a non-event, had its origins in the high cost of digital storage in the 1970s and 80s - programmers simply could not design applications that anticipated the need to store the identification of the year in 4 characters, as opposed to only 2.

500 Gbyte disk drive in 2007. Equivalent to almost 3,000 RP06 drives of 1979. Model courtesy of NeuralNet Interactive, Saskatoon.

Three DEC RP06 disk drives. Each one of these units, about the size of a household washing machine, held 175 Mbytes of information.

The DEC-20 produced over 100,000 BTUs of heat out the rear of the machine. Normal practise of the time for cooling computer facilities was to control air temperature in a closed system: cool air came into the room, the machines heated it up, it was then cooled back down again and brought back in as cool air, with a bit of fresh air added in so that people working in the facility did not have to breathe stale air. In this facility, however, the heating engineers at Physical Plant decided on an open system. Rather than expend energy cooling the hot air back down after the machine heated it up, they captured the hot air from behind the machine by isolating the rear of the cabinets with an insulated curtain that hung from the ceiling and clung to the top and sides of the DEC-20. They then took that hot air into the Arts Building heating supply, and used it to supply some of the heating requirements of the building - an energy saving. Since the space above the computer room was used for a number of Psychology laboratories with experimental rats in cages, we sometimes referred to the DEC-20 as the rats' furnace!

Thus the arrival of the DECsystem 20, and the restoration of University control over its main academic computing facilities, were a welcome shift in campus computing. Although we did not 'catch up' with similar sized institutions for some time yet, we at least got on the same road.

There were other important developments that occurred at this time. First, although the DEC-20 provided for some welcome time-sharing capacity, the main form of student computing remained in the form of punching cards and submitting them through the HOTT facility. In a room on the main floor of the Arts Building there were about 12 key-punch machines which students could use (when they were all working - keypunches were notoriously prone to being our of order). Once programs were key-punched, students took their card decks to the HOTT facility one floor below in room 49, in the space between the Arts and Commerce buildings. The continuing use of keypunch machines was becoming a problem. First, they were a very inefficient form of program/ data entry - if you made a keystroke error, the card was ruined and you started over. Second, student numbers were rising and there simply was no more space for additional keypunches in the room. There was also a cost issue: each keypunch cost over $100/ month rental in 1978 (almost $300 in 2007).

The DEUS student lab in 1979. Fourty-four Cathode Ray Tube terminals were provided for students to create their programs and submit them to the SaskCOMP IBM mainframe off-campus. DEUS was designed and implemented by Mr. Peter Hardie of Academic Computing Services.

To address this challenge, a new system was locally developed. The selection committee of 1977-78 had seen an interesting system in use at the University of Waterloo when they did site visits as part of their vendor evaluations. The concept was to use a minicomputer to host a number of Cathode Ray Tube (CRT) terminals, where students could write their programs with the assistance of a simple text editor, and then submit the programs right from the terminal - no more punch cards and no more walking to another floor to get output! Output still came back in the same fashion as in the HOTT facility, but now the program entry would have a more user-friendly form of program creation, and a lot more terminals.

It was determined that while the Waterloo system (called 'Widget') had merit, it would not transport well to the UofS, so Academic Computing Services undertook the design and implementation of a system that was similar. This system was called DEUS, for Data Entry, University of Saskatchewan. It consisted of a DEC PDP-11/70 minicomputer and 44 CRT stations in room 145 Arts. The system was almost single-handedly designed and written by Peter Hardie of Academic Computing Services. It so happened that the Administrative Systems group already had a PDP-11/70 in the basement of the old Administration Building (now the College Building), and Jack Billinton, then head of Administrative Systems, generously gave us permission to use the machine on Sunday afternoons to create and de-bug DEUS. We rented a machine called a protocol analyzer, and used it to monitor the flow of communications back and forth to the SaskCOMP mainframe, where the student programs were still being run. Peter thus 'reverse engineered' the function of performing what was known then as Remote Job Entry to an IBM mainframe.

There were some initial teething problems in switching over to DEUS from keypunches. The greatest problem was that in the late 1970s, students were poor typists, and while this handicap also affected their keypunch accuracy, it really showed up once they were on a system that could measure keying speed and accuracy. Also the increasing numbers of students caused them to rapidly adopt a strategy that, once they got their hands on a free CRT, they stayed there, no matter how ineffective they were. This lack of 'throughput' became apparent within the first week of use, and it was addressed by imposing a 30 minute limit on how long a student could stay logged on to the system. After that, throughput increased to a satisfactory level, although the system was frequently running at almost full capacity.

Although it was not perfect, DEUS was a distinct improvement over punching cards. Among other benefits, instructors no longer had to carry boxes of punch cards to class to hand out one copy per student of the common 'data card decks', that were used as input to programming assignments. Now DEUS automatically appended the data file to each program sent over to the IBM machine, and the instructors had the added advantage of being able to make invisible any 'tricky' data that they wanted the student programs to handle properly. The system was exported to University of California Berkeley campus, University of Hawaii, Bell Northern Research and Cray Research Labs.

Now that the University had acquired a time-sharing system, a collateral challenge was to increase the access to it from all over campus. Otherwise, users would still be walking, possibly to other buildings, to find interactive terminals. Recall that 1978 was 5 years prior to the adoption of Ethernet as a standard, and the technology of the day for connecting remote terminals was the use of twisted-pair telephone wire, with devices known as 'Limited Distance Data Sets' or simply 'data sets' on each end. The serial communications port on each terminal device would be connected to the data set, and the data set at the other end would be connected to a time-sharing input port … and the user was in business. In 1978 we had two problems in trying to offer this access. First, we had no wire in the ground that connected campus buildings to the computer centre in Arts 70. This was tackled with the installation of a cable of 600 pairs of 26 gauge copper wire, starting in Arts 70 and fanning out across campus. Every building that it passed through would be allocated a number of the pairs, and the most remote building got what was left as the remaining pairs finally got to their premises. It took 2 pairs for each data set, so we had the capacity for 300 terminals (it seemed like an outrageously large number at the time), dropped off in clusters of a few dozen in each major building, as the cable got thinner and thinner the further it got from the Arts Building. It should be noted that the cable leaving Arts was almost 10 cm in diameter, filled with grease, and was extremely stiff. It almost entirely filled the conduit that at the time transited between Arts and Thorvaldson Buildings. Years later, as that conduit was needed to install Ethernet cables, the grease had hardened to a degree that Physical Plant staff could not pull the old cable out of the conduit. It remained as a 'legacy' component of campus infrastructure for several years, and Ethernet installation was accommodated as other needs for more conduit or tunnel capacity developed over time.

Conversations about allocation of these 300 access points were interesting. I recall in particular meeting with representatives of the Thorvaldson Building in 1978-79, to get their reaction to a suggested allocation of 24 terminals. The reaction was polite disbelief. I was told (I remember the words well) : "24 terminals!! Why, this building will NEVER need 24 terminals".

In any event the cabling project was completed. There were two main 'routes' of the cable. From Arts the cable went through Thorvaldson and on to Geology. At that point it spit into two: one route continued on through Physics, Administration, Crop Science (now Archeology Building) and finally to Engineering, while the other crossed the Bowl to the Library, McLean Hall and Health Sciences. The project marked the first significant (and successful) collaboration between Computing Services (now Information Technology Services) and Physical Plant (now Facilities Management). The staff of Physical Plant, although new to this form of campus infrastructure, responded with energy and competence. We could not have done this without their skill and knowledge – they knew where the tunnels were, and how to get around the obstacles!

So by roughly 1979 we had 'wired the campus' (tongue in cheek). However, we had an additional technical difficulty. The description above states that one end of a line was connected to a terminal, and the other end to a computer port. What if the user wanted to be able to sometimes connect to one computer, and other times to another? There already were several computers that were operating in time-sharing mode (DEC-20, HP 2000, Library catalogue and the first campus-wide word processing system, called Word-11). What we needed was a kind of 'phone exchange' for computer terminals: a user would turn on a terminal, connect to the 'exchange' and then ask (by typing the name of a service) to be connected to a particular computer. Academic Computing Services and Administrative Systems (Jack Billinton, Director) did some research on available systems to do what we wanted. We concluded that a Canadian company, Gandalf Technologies Inc., of Ottawa, was a good choice, with a device called a Private Automatic Computer Exchange (PACX). This system was well known and highly functional. We actually were within a day or two of ordering a Gandalf PACX when we received a visit from Mr. George Spark and Mr. Nigel Hill of Develcon Electronics, a Saskatoon company that was in the business of manufacturing and selling data sets for twisted-pair networks. "We've been thinking", they said, "about designing one of these switching things, and we hear that you are looking for one. Would you be willing to take a chance on us?" So, we did. We were using quite a number of their data sets, so we were familiar with the company.

We worked with Develcon for almost a year, talking through how the thing should work, what features it should have, how to manage it, and so on. Eventually we took delivery of Serial Number 000 of the Develcon 'Dataswitch', which went on to make Develcon a very successful company (until overwhelmed by Ethernet networking about 7 years later). We purchased additional machines for installation in Arts, Engineering and Administration buildings, as the twisted-pair network grew. It was a happy result: the Dataswitch was a rugged, functional and affordable solution to a common problem. We took a chance, and it paid off for both the University and for the local economy.

As we entered the 1980s, several new challenges and opportunities emerged. First, in 1981, IBM

Develcon Dataswitches in rear of photo, in Arts 70. Small boxes on top of VAX in foreground are first Ethernet concentrators to which terminals were connected to have access to the new Ethernet network, in 1985.

introduced the IBM Personal Computer. Nothing changed the world of computing so fundamentally. By putting their name on a desk-top device, IBM legitimized an industry that was in its infancy. For the first time, a user could have on their desktop a machine that could meet most of their requirements – sometimes all of their requirements. Early machines were relatively expensive, so they were not purchased in large quantities immediately, but as the 1980s progressed, the desktop computer moved from 'curiosity' to 'pretty good idea' to 'mainstay' for computing needs.

As these devices grew in number, they introduced a new connectivity challenge. Whereas the prior mode of working by walking to the closest 'terminal room' was acceptable, users who invested in personal computers to be installed in their offices expected to be able to connect to the campus network from those locations. Needless to say, the infrastructure that allocated 300 access points within common terminal rooms did not anticipate this requirement. Two measures were used to address this challenge. First, some users were able to purchase a 'modem' that allowed them to use their office telephone to dial into a bank of answering modems in the computer centre, and thus get connected through the Dataswitch to the service they sought. This was acceptable but limited in transmission speed to about 1200 bits per second (bps). Second, we rented a steadily increasing number of 'leased lines' from SaskTel, which were spare telephone pairs, throughout the campus, not needed as phone lines, but which could be used to provide a copper path to the terminations back in Arts 70. These leased lines would have data sets on either end, and could usually achieve higher speeds, up to 9,600 bps.

The introduction of desktop computers got an additional boost in 1984 when Apple introduced the first Macintosh computer. It set a new standard in ease of use compared to the IBM PC and its clones. Remember that this was years away from the availability of Microsoft Windows on the PC, and thus the first users in the PC world did not have anything like the ease of use introduced by Apple in the Macintosh.

Academic Computing Services took the position that personal computers were intrinsically good, and that our services should augment the functionality provided by the desktop machine. In the early days of desktop computers some university computer centres and many commercial organizations took the view that desktop computers were a threat to them, and fought against their use. We will return to the role of campus desktop computers within a discussion of campus networking, as it was the combination of better personal computers and improved network access that really caused the desktop machine to achieve pre-eminence as the most important single component of a user's access to modern computing functionality.

By 1984, the University had many challenges to face in the provision and support of its computing and networking infrastructure. First, the DEC-20 was now 5 years old, operating at full capacity and without a growth strategy from Digital. For some time the large computer customers of Digital Equipment Corporation (those that had DECsystem 10s and DECsystem 20s) had been pushing on DEC to announce a new higher-end system that would allow graceful migration from our current systems into new facilities with the order of 2-5 times the capacity of the old. In the meantime Digital had developed in 1978 the VAX 32-bit computer, and it received great success in the marketplace. By the mid-1980s, Digital announced the end of the line for their 36-bit architecture and consolidation of future planning around the VAX architecture. While this was surely the correct strategy for them, it caused a lot of frustration and delays among their customers who were trying to choose a replacement strategy for their DEC-10s and DEC-20s, the UofS included. Thus one challenge was simply providing more capacity in the shared systems for academic use.

A second challenge was to build a new campus network. By 1983 the IEEE had published the Ethernet standard, based on a specification coauthored by Digital Equipment Corporation, Intel Corporation and XEROX Corporation (the so-called 'DIX' proposal). This new technology promised network speeds previously unheard of (a shared common channel of 10 Million bps, with individual user connections of initially about 1 Million bps). In addition, the technology used new media: a shared coaxial cable, with published rules for installing the cable and tapping it for user connectivity. By fall 1983, early devices were available, and Dr. Paul Sorenson and I went on a field trip to visit some of the Ethernet players in the Bay area of California. We visited a company named Ungermann Bass, who were one of the first suppliers of Ethernet hardware, and from whom we soon purchased some equipment to be used to learn about this exciting new technology.

We also visited Robert Prentis at XEROX Palo Alto Research Center, the 'home' of Ethernet (Robert Metcalfe, a former XEROX employee, invented Ethernet). It was Prentis whose name appeared as the XEROX coauthor of the original 'Blue Book' from the three initiating companies. As we drove up to the building to meet with Prentis, I recall thinking that for sure we would be able to see answers to one of the more troubling early issues we were concerned about: how could you install this thick cable that didn't bend very well? Ethernet cable initially was the so-called 'thick Ethernet' which used a coaxial cable about 1 cm thick and a minimum bending radius of over 1 metre – it was physically difficult to install, particularly in the old buildings that made up most university campuses. Here we were at the place where it all started – surely they would show us exemplary ways to retrofit old buildings with this ugly orange cable. However, as we entered Prentis' office, we were shocked to see that hanging from the ceilings in the hallway were common cup hooks, probably purchased at the local hardware store, then somewhat haphazardly screwed into the ceilings and walls, which held up the Ethernet cables, in plain view. Further, each office door had been subjected to some crude work with a hacksaw, to cut a notch in the top of the door just at the place where the cable wanted to be as you closed the door. If the cable caught as the door was closed, they just sawed a bigger notch in the top of the door. I recall sitting in Prentis' office thinking "If Ethernet is what we want, how would we ever get Physical Plant to go along with this scheme?" In truth, none of us would have counseled this method of cabling, as it was totally insecure and rather ugly. But it did show us that we had some barriers to overcome if we wanted to install Ethernet throughout the campus.

In addition to Ethernet, there was one other candidate emerging as the technology of choice for building Local Area Networks (LANs). IBM had developed a technology called Token Ring, which shared almost nothing in common with Ethernet: different cables, different protocols, different device interfaces. The situation was to the LAN world what the battle between VHS and Beta was for the video tape industry. Both systems were backed by substantial supporters: IBM for Token Ring and the DIX consortium for Ethernet. There was no shortage of scholarly analysis of the two systems, and each side had its nay-sayers who predicted gloomy performance for the other contender.

A third challenge was the level of financial support for campus computing and networking. Table 1 is taken from an annual survey of university computing facilities in North America. The data for Canadian institutions represented 30 universities, including the UofS. The analysis shows that financial support for computing facilities and support staff at the UofS would have to more than double to come up to the Canadian average in 1983-84. It also shows that our 'strongest' performance was in squeezing out as many terminals and time-sharing ports as possible from our budgets: the data in those categories shows that available resources were expended as much as possible for user benefit, and in one category we actually exceeded the Canadian average.

This table typifies the history of the University of

Table 1. Comparing the University of Saskatchewan with other Canadian universities in several measures of computing support and services available. From 1983-84 Directory of Computing Facilities in Higher Education, compiled by C. H. Warlick of University of Texas at Austin.

	Support Staff per 1000 Students	Central Academic Computing $ per 1000 Students	Support Staff per 1000 Faculty	Central Academic Computing $ per 1000 Faculty	Time-sharing Ports per 1000 Students	Shared Terminals per 1000 Students	Time-sharing Ports per 1000 Faculty	Terminals per 1000 Faculty
Canadian Average	3.9	246,657	52.4	3,492,577	31.8	44.5	220.1	675.2
U of S	1.6	103,679	23.5	1,526,056	20.4	29.0	302.4	427.4
Change to equal Canadian Average	+143%	+140%	+123%	+129%	+55%	+67%	-27%	+58%

Saskatchewan: you could find similar comparisons between the UofS and other Canadian universities in so many areas: clerical support for faculty, Library resources, technical support staff, physical facilities support, administrative support resources – and so on. These comparisons bespeak the culture of this university to tenaciously do the best it can, providing as much service as it can, in spite of chronic under-funding. In 1983-84 one would have been reluctant to put money on the University being able to successfully move forward in what was a very challenging time in the computing world.

Thus as we approached the mid-1980s there were some very significant challenges facing the UofS in terms of its campus computing and networking environment:

- We had an aging time-sharing system for general academic computing, the DEC-20, for which there was no obvious expansion path
- The total connectivity provided by University-owned and operated network infrastructure was only about 500 terminal access points, and they were all low-speed (up to 9,600 bps) connections
- Emerging standards for Local Area Networks were on the way, but it was still not clear which of the Ethernet or Token Ring standards would prevail
- Campus financial and staff resources to support computing and network activity were seriously under-funded

6. Reaching for New Levels of Capacity and Networking: 1984-1988

There is a well-worn saying: "when the times get tough, the tough get going". Our strategy to meeting these challenges was to set out a bold, but achievable plan, with advanced functionality for computing and network users, and sell ourselves as partners to industry. The mid-1980s were a pivotal time in the industry. Whoever won the LAN race would have significant business advantage for some time, and whoever found ways to marry desk-top computers with shared resources in mainframes (they were not yet called 'servers') would deliver functionality in services that would give their users a competitive edge.

The result was a vision detailed in a document called **Project ACCESS**: The **A**dvanced **C**omputing and **C**ommunications **S**ystem at the University of **S**askatchewan, published in December 1984. This plan was released as part of a Request for Proposal for new academic computing and campus-wide networking facilities for academic users in early 1985. A Selection Committee, chaired by Dr. J.F. Angel, got to work throughout the first half of 1985, and evaluated vendor responses. A short list of IBM and Digital was chosen, and two site tours of 6 faculty and staff were conducted in May of that year. We toured corporate offices and university reference sites for both vendors.

In my opinion, the differentiation between the two companies came in the LAN area: they both had excellent options to replace the DEC-20. On the IBM tour, our exposure to Token Ring technology was largely in the form of talks from IBM scientists, in their headquarters, who were trying to convince us, on paper, that Token Ring was the superior

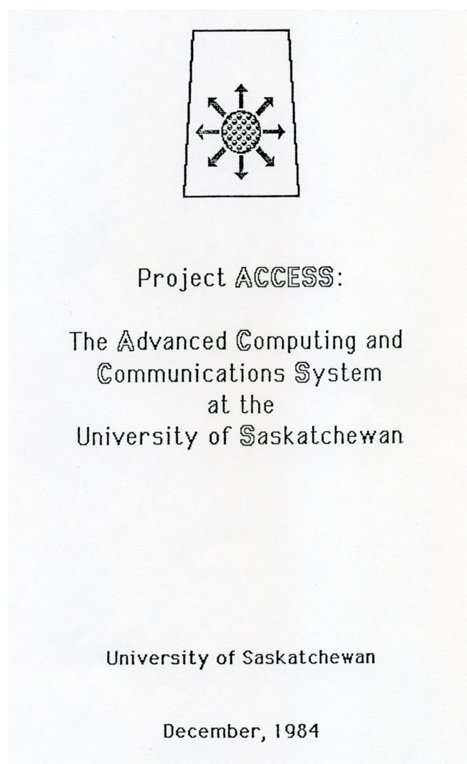

Project ACCESS:

The Advanced Computing and Communications System at the University of Saskatchewan

University of Saskatchewan

December, 1984

Cover page of the Project ACCESS vision statement. One of the early uses of document creation on an Apple Macintosh, first sold in February of 1984. The graphic was meant to convey a number of interconnected 'nodes' in a campus network, with province-wide impact.

Table 2. Recommendations of the Selection Committee in 1985 for replacing the DECsystem 2060 and growing capacity over the next three years. Note that one VAX 8600 had about twice the capacity of one DECsystem 2060.

	1985	1986	1987
Number of VAX 8600s to be installed	one	two	three
Location	Arts 70	Arts 70 (2)	Arts 70 (2), Engineering (1)
Approximate Capacity in Millions of Instructions Per Second (MIPS)	4	8	12
Total on-line Disk storage	3 Gbytes	5 Gbytes	8 Gbytes
Number of tape drives	2	4	8
Number of user ports	120	240	360

technology. We did not see a working installation of Token Ring in any customer site. When we visited the Digital sites at other universities, campus personnel (not Digital sales staff) would take us into the computer facilities, lift up a removable floor panel and point to 'the ugly orange cable' and say something like: "this stuff works well; we intend to wire the campus with it". I think the tour convinced us that Ethernet was in the lead – history would show that we were right!

So now we had a technology choice for creating one campus-wide network (it was common at the time for universities to have one network for academic computing, and a physically separate one for administrative purposes) Going to one common network was a bit of a 'hard sell' to the administrative users, who were worried about whether unauthorized users could access say, the financial or student records systems. Time and security technology eventually smoothed out this worry. We also had a migration strategy for replacing the DEC-20 with multiple VAX 8600 systems. Table 2 shows the Selection Committee's recommended plan for replacing the DEC-20. One VAX 8600 was roughly twice the capacity of the DEC-20, so this plan called for approximately a six-fold expansion in time-sharing access for academic computing.

However, although we had a plan, could we afford it? Recall that the **Project ACCESS** strategy was aimed at trying to build a partnership with industry, so that we could get some form of special treatment from one or more suppliers, to accomplish our vision. Basically our pitch to the industry was something like "notwithstanding the fact that we have very thin financial and human resources, we know what we are doing, and if we could get a little help, we could showcase your products to the world." In fact, part of our message to industry was precisely that we could show that an organization that was way below the average in financial and staff resources could develop a highly functional campus network and array of services, if we could work together and manage our precious resources carefully. A manufacturer would rather have a reference site that showed that you did not need large budgets to use their products than one that was awash in financial and staff resources. The UofS could certainly be an example of minimally-funded support for computing and networking!

As the summer of 1985 came and we were still trying to put some flesh onto a partnership with Digital, a very fortuitous development occurred. Digital, in its attempt to lead the pack in campus-wide networking, announced a special three-year program: The Campus-Wide Investment Program. They sought applications from North American post-secondary educational institutions for the program, and we were eager and ready to apply. By November, largely through the personal efforts of Vice-President Blaine Holmlund from the UofS and Ken Copland, President of Digital Equipment of Canada, we were selected to join eleven other institutions in the program. There were two Canadian participants; University of Saskatchewan and University of Waterloo. The other ten partners were in the United States, nine universities and one very large Community College, Maricopa Community College in Phoenix, Arizona.

The CWIP program structure was that Digital would provide a 45% discount (normal educational discount was 15%) on hardware and software for approved projects, to a maximum dollar amount over the three years of the program, 1985-1988. For the UofS, the total discount amount was

$4,500,000. Projects had to be approved by a project management committee at each partner institution. The project management committee was made up of representatives of both Digital and the partnering institution. So now we could purchase products for 55% of their list price, but where did we get that money? The Government of Saskatchewan had at that time a special fund (University Renewal and Development Fund) for university projects that they approved, and we successfully obtained $4,200,000 from that fund for the duration of CWIP. The remaining $1,300,000, to make up 55% of a total of $10,000,000 in projects, came from existing university budgets.

There is one amusing anecdote about the Digital representatives on the project management committee. There was a member representing Digital Equipment of Canada, Mr. Al Seamans and another who represented Digital corporate offices in Massachusetts, Mr. Andy Maisland. The committee met four times a year for three years, and that meant that a meeting in Saskatoon in the winter was unavoidable. It turned out that in fact we met in Saskatoon in January 3 times. When the project was over in 1988, Mr. Maisland confessed that when his Boston colleagues found out that he had to attend meetings in Saskatoon in January, they told him that likely there were not enough warm clothes in his closet to keep him comfortable on his visits here in the winter. But, as he told us at the project wind-up, it transpired that each of the three times that he came to Saskatoon in January, it was actually warmer in Saskatoon than it was in Boston!

Needless to say the CWIP program, and the extra capital funding from the Government of Saskatchewan, was a major boost to the University. It has to be stated, however, that there was at no time in the program any additional funding for more

Dr. Robert Kavanagh, Vice-President Blaine Holmlund and Dr. Joe Angel in the computer room in Arts 70, 1986. DEC-20 in the background and new VAX 8600 behind Holmlund.

support staff, so the existing staff had to absorb all of the new activity during this period without any extra help. There were 22 projects approved for the UofS over the course of the program, ranging from new central computing facilities in both academic and administrative areas, to college computer labs, to distributed academic facilities in the College of Engineering and the department of Computer Science to, of course, the building of a campus-wide Ethernet network.

All of the projects were challenging, but none more so than the building of a network. In 1985 hardly anyone really knew much about retrofitting old buildings with Ethernet cabling and distribution to offices and labs. As mentioned earlier, the cable was thick, it did not bend well and it could only be 'tapped' at certain points. Obviously, none of the existing buildings were designed to include LAN cabling, and new wiring conduits, wiring 'closets' to serve a particular area of a building, documentation to track the installations – all had to be developed as new infrastructure, within spaces that were in full use. It was much easier to install in new construction, so we seized every opportunity to piggy-back on renovation projects and new building design. We rapidly found out that this task was expensive: we could have spent all of the extra funding, and much more, on just building a campus Ethernet. Somewhere along the way the staff involved in building the network coined 'Kavanagh's Law': no matter what the technology, whether it be twisted pair copper with data sets on each end, or Ethernet, the cost was a constant: about $1,500 to add a connection to the campus network. That 'law' remained valid for many years. Considering there were thousands of connections to made across all of the offices, labs, student work rooms across the campus, there was a shortage of funds and certainly a shortage of human hands to build this new network.

In spite of these challenges, we built a very significant beginning of a campus Ethernet. In 1988, Digital Equipment of Canada surveyed corporate and university sites across the country and pronounced that the UofS had the largest Ethernet network in the country – about 2,000 connections by that time. As the special funding for CWIP expired, and we had to fall back on University capital budgets to advance the network, the next many years were persistently challenging for everyone concerned. Unfortunately priorities for new connections were often distorted by joining up with renovation projects for other purposes, which lowered the cost, but were not necessarily in the high demand areas. Also, it made sense to try to do a big area of a building at a time, because of some economies of scale in the location of wiring closets and doing the labour to install cables, so it was difficult to be precise about allocating new connections to the areas of greatest need. For several years we used the available capital funds to 'match' funding from the departments and colleges, to make the funding go further. It was a very difficult problem to address with restricted resources.

One of the CWIP institutions, Johns Hopkins University in Baltimore, Maryland, had a very unique solution to this problem. In the mid-1980s, they had raised about $500,000,000 from a funding drive with their alumni. All of the CWIP institutions met twice a year, at one of the partner sites, to share information about progress on their projects. When we met at Johns Hopkins in 1987, all the other eleven institutions were wondering why the hosting site did not seem to be having any problems with building the physical Ethernet network, as the rest of us were. The Johns Hopkins hosts said that as a result of the successful capital funding drive, they were in the process of replacing all major buildings on the home campus, and the problem of installing Ethernet had just been addressed by making it a part of each new replacement building!

One final story about trying to wire the campus. I had been preaching to our staff that we needed to make the network as 'invisible' as possible, i.e., users should not have to be technical specialists to use the network. I found out that 'invisibility' was not the right word, from Professor Marshall Gilliland of the English Department, who had been Chairman of the Advisory Committee for Academic Computing Services in the early1980s. Marshall was always on the lookout for information that he could feed back to the computing organization. On one occasion he reported a conversation with two of his colleagues. He had asked them how important was the campus Ethernet network. Their response was a vehement assertion that it was a "total waste of money", and an "administrative boondoggle that takes resources away from faculty". He looked at them in confusion and said "but I don't understand, don't both of you use your office desk-top computers to connect to the Library on-line catalogue?' "Oh, that", they said, "that is absolutely essential; we <u>have</u> to be able to do that." So I learned that perhaps 'transparency' is a desirable characteristic of the network, but ‘invisibility’ should definitely not be a goal!

The participation in the CWIP program was a great opportunity for the University. We significantly advanced the capacity and functionality of computing, both academic and administrative. Perhaps more importantly in the long run, we established the foundation for a campus-wide Ethernet network, which brought desktop computer users into a world of connectivity and function that was not even dreamed about in the days of low-speed network connectivity. The success of this period is owed to the staff in the departments and colleges who undertook CWIP projects, and the staff in Computing Services (Academic Computing Services and Administrative Systems were joined into one Department of Computing Services in 1986) who worked very hard and with great innovation to accept the challenge of the program.

7. Conclusion

This document has been assembled in a hurry, to partially provide a history of computing for the Technology Week celebrations in November 2007. As such, there a several glaring gaps in the reflection here, dictated by what could be easily accessed in a short time -- mostly my own recollection of events. For example, the stories told here apply almost exclusively to the academic component of computing at the University; there is a parallel, interesting history of the efforts of University staff to steadily evolve administrative applications, Library services, Facilities Management applications, etc., that needs to be told. The evolution from card oriented batch applications, from students standing in line to register for classes, from paper personnel records, to modern on-line services ... those are additional stories worth recording.

Further, this record concludes just as desktop computers were coming into widespread use. By 1987, there already were a few thousand personal computers on campus, and no history of computing would be complete without picking up the story of how those devices so fundamentally changed the infrastructure and services that provide information, data analysis, communication tools and administrative services. In the period concluding in 1987, we were just on the early edge of those developments.

Finally, this narrative is a little long on technological history, and short on the actual experience of those who studied with and used computers in so many interesting ways at the University.

However, having identified some important gaps that should be addressed in a more complete history, I hope that this document provides an interesting and informative glimpse into the first 30 years of computing at the University. It was my privilege to have served as a part of this history.

8. Acknowledgements

I want to thank Carol Kavanagh, my spouse, who edited and proof-read the work. Drs. Rick Bunt and John Cooke read the first draft and provided many helpful suggestions. Tim Hutchinson of University of Saskatchewan Archives provided much appreciated assistance in locating archival material. David Emigh of Information Technology Services provided photos of the equipment in Arts 70 from the 1980s. I received material from Dave Bocking of the Department of Computer Science and Frank Winter of the Library, but their work will have to await inclusion in a more complete history to be developed at some time in the future. My sincere thanks to all.

9. References

1. Minutes of the Board of Governors of the University, 1957, U of S Archives.

2. Minutes of the Computer Advisory Committee, January 1967, U of S Archives.

3. Minutes of the Computer Advisory Committee, June 1967, U of S Archives

4. Encyclopedia of Saskatchewan, Education, at http://esask.uregina.ca/entry/education_in_saskatchewan_timeline.html

5. Minutes of the Computer Advisory Committee, October 1971, U of S Archives.