

PRELIMINARY REPORT

The Future of Computing at Carnegie-Mellon University

The Task Force for the Future of Computing

28 February 1982

Allen Newell, Chairman

Thomas Bajzek
Brian Berry
Lorna Cheng
Richard Cutkosky
Steve Director
Charles Eastman
James Hoburg
Charles Kriebel

Jill Larkin
Tom Michalak
Dale A. Miller
John Pople
J. Ray Scott
Lee Sproull
Howard D. Wactlar
Lori Zuti

Staff

Susan Elster

Judy Hammer

Nick Pealy

Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

PRELIMINARY REPORT

The Task Force for the Future of Computing at CMU (TFFC) was created by President Cyert in September 1981 to formulate a view of what computing at CMU should be like in the mid to late eighties. In a two page Charge (reproduced as Appendix 1) he laid out the following reasons for establishing the Task Force. CMU is already heavily committed to computing, and planning is going on in several places that could substantially increase this commitment. Many groups on campus believe that the computer is an integral and important part of their own future. The effects of computer use are becoming large enough to affect all aspects of campus life, not just those that already use computers directly. Thus a global planning effort is in order.

We use the term *computation* broadly to mean all processing and communication of information by electronic means. Numerical computation, text processing, data base inquiry, image processing, electronic mail, symbolic computation, interactive graphics, electronic music, computer-aided design, electronic publishing -- all are computation. Correspondingly, we use the term *computational facility* to mean the hardware, software and human support that provide computation.

The goal of the Task Force is to attempt to determine the implications for the entire CMU community of a future of much enhanced computation, and to formulate what kind of future would benefit us most. This effort includes predicting bad effects, and proposing possible remedies.

The Charge requests a preliminary report by the beginning of the 1982 Spring semester. This is the requested report, and it is indeed preliminary. The main section describes a set of issues that we deem important. For each of these we set out preliminary positions or recommendations. To make these issues intelligible, we precede this section with three others: the first on the ingredients from which any computational future must be composed, the second on the current state of computing at CMU, and the third on basic positions underlying our discussion of the issues.

In formulating this report we have met intensively through the Fall (some ten 3-4-hour meetings). We have also met with all the departments and with several other groups (some thirty noon-hour meetings) to gain knowledge of the diverse ways computers are used on campus.¹

A central purpose in circulating a preliminary report is to promote discussion about the issues it contains. We will arrange public forums for discussion and feedback. We also encourage feedback to any member of the committee, or by sending computer mail to TFFC on any Computation Center TOPS system or on CMUA at Computer Science.

Ingredients of a Computational Environment

A modern computational facility is constructed from a few basic physical components: computers for processing, channels for communicating, memory for storing information permanently, interfaces for interacting in different modes (e.g., printers, graphics terminals). To these are added software components: an operating system and programming environment; an array of programming languages; and, increasingly, many tools (e.g., statistics packages, simulators). Finally, there is a human support organization to tie it all together. In the current era, all these components can be acquired and assembled from many independent sources, so that a facility is composed, rather than being bought whole from a single vendor. Thus, we need to know the broad

¹We will issue later a report that describes in more detail the current state of computing on campus, especially as we saw it through the eyes of those who discussed these matters with us. However, we want to acknowledge the large impact these discussions had upon us and this report, and to express our gratitude to the three-hundred-odd people who took the time to educate us.

options available for each kind of component in order to understand the options available to CMU in constructing its future.

Computers. Computers can be characterized by their processing power. However, the power of a given computer architecture (e.g., a VAX) can vary by a large factor (e.g., five) depending on the details of its construction, how much working memory is used, and what kind of data (e.g., floating point or integer) is being processed. Thus, the particular computer architecture (e.g., VAX vs DEC 2060 vs IBM 4341) doesn't matter much -- what counts is the effective power the computer delivers for the task at hand. As a result, we can adopt a familiar computer, such as the DEC 2060, as a useful, if rough, standard of comparison. Then a VAX can be described as a third of a DEC 2060. More accuracy would be provided by adding a separate ratio for processing other types of data, e.g., a VAX is about equal to a 2060 for floating point processing.

In general, computers sufficiently powerful to do important work are getting smaller and cheaper, at roughly a decrease of 20% per year in unit cost, which amounts to a factor of a third in five years. These decreases mean that increasingly useful-sized computers need not be owned by large units (e.g., a computation center) but can be owned by smaller units (e.g., departments, research groups, or even individuals). Given both the diversity of kinds of computers suggested above, and the lowering of costs, it seems almost certain that CMU's computational future will involve many kinds of computers owned by many different entities, including "personal" computers that might be owned by individuals.

The personal computer has been at the center of much discussion on CMU's computational future, largely because of the proposal to integrate them into undergraduate education. The defining characteristics of a personal computer are: A true computer, capable of independent standalone operation, that is used by a single person. Personal computers vary from very small (Apple I, about one fiftieth of a 2060) to small (the recently announced IBM machine, one fifteenth of a 2060) to moderate (Perq, one sixth a 2060) to large (Dorado, a Xerox research machine, about equal to a 2060). Correspondingly, current costs vary roughly between one thousand and one hundred thousand dollars.

As with other computers, for a given price, the power available increases substantially every year. A personal computer with about a tenth the power of a 2060 may reasonably be expected to be available for a few thousand dollars within the next few years.

As true computers, personal computers have varying amounts of memory and communication facilities, various word sizes and floating-point capabilities, and support a particular array of languages and tools. They all interface to the user through a keyboard and video display (the latter varying immensely in quality and rate of interaction). Personal computers also vary significantly in address space (from 12-bit up to 32-bit) which determines ultimately the possible complexity of programs. Address space is an important parameter of computers generally, but becomes critical in small computers.

Because a personal computer is dedicated to a single user, service is not degraded by other simultaneous usage, and there is freedom from many of the frustrations of time-shared systems. This standalone character is not, however, inviolate -- with communication links, a personal computer becomes an intelligent interface device. Personal computers are pleasant to interact with, because relatively more of their resources are devoted to the interface. Finally, a personal computer may be inexpensive enough to permit private ownership. For very small personal computers, however, these gains may not be useful because many tools (e.g., document preparation programs) are not available on very small machines, and if they are available their functionality may be minimal.

Memory. Primary (main) memory and a modest amount of secondary (disk) memory provide the working memory of a computer and simply contribute to its general computational power. The principal separate use of

secondary memory is permanent storage, which is a large and continually growing need. Today secondary memory means disks, ranging from a hundred thousand characters (floppies) to a billion characters, with the cost per character decreasing as size increases. Memory can be distributed and shared, if networking is good, and if the software on individual machines supports it.

Archiving memory (today, magnetic tape) is also needed. Although this is not a major item in CMU's current environment, more intensive computation might make it a more major issue and expense.

Communication. As the preceding comments suggest, increasingly a computational facility is made up many distributed computers. Thus communication to and among computers is a major part of a computing environment. Communication (or networking) affects at least four major aspects of an environment. First, as computer-implemented services become more useful, they are needed more continuously, and ready access becomes essential. Second, computer services are coming to include communications per se -- electronic mail, bulletin boards, conferencing. The value of such services increases radically as everyone in the community achieves ready access. Third, networking at high data rates allows the distribution of resources among computers on the basis of cost-performance tradeoffs. Resources can be shared and not duplicated. Fourth, reliability can be obtained by redundancy and communication to alternative resources.

The most common current networks connect terminals to a single computer, and involve hardwired lines to immediate locations (terminal rooms and local offices) plus use of the telephone system. The rate of terminal access is increasing from 30 characters/sec (more or less standard now) to 120 characters/sec (the best available with a modem on a standard voice-grade telephone line). With each increase in data rate, access increases and the computer becomes easier to use. No one willingly goes back to a lower data rate once a higher one has been assimilated.

Communication between small dedicated laboratory computers and larger central computers, a vital part of current scientific and industrial practice, involves both recording data and downloading programs to the small machines. Many groups on campus are doing this regularly; others are just beginning. The extension to undergraduate laboratories is just beginning, but will be pervasive.

Networks linking together computers in a local area are coming into routine operation. At CMU, an Ethernet and a DECnet link the respective computers of the Computer Science Department and the Computation Center. Separate networks can be connected through gateways. Local area networks can be fast enough that a computer can use memory elsewhere on the network as easily as its own, and that one computer can use another for special tasks or for sharing resources. Thus technology is achieving the ability to weld a collection of computers into an integrated system.

Long-distance networks (ARPANET, TELENET, TYMNET) allow information to be transmitted nation or even world-wide. These are significantly slower than local networks, but do permit remote terminal access, electronic mail, and rapid transfer of files, including documents and programs. The extent of such networks is growing continually, though still spotty.

Types of Data and Interfaces. Although all computation is done in bits, useful work is always based on data structures attuned to the relevant task. There are special structures for integers, floating point numbers, text, multifont text, symbolic structures, graphics, speech and images, as well as many more specialized data structures. Each structure presents a unique combination of computational requirements. Even integers and floating point numbers are quite distinct in their requirements and are embedded in data structures of various intrinsic sizes (which is why computer power is measured separately for each). Graphics, speech signals and visual images all have requirements for high-capacity communication and large memory. Color displays increase such requirements further. For these reasons the computational facility as a whole should be designed to support

some collection of data structures, with the appropriate computers, communication, memory and interfaces to create a balanced system. Failures in this coordination (e.g., thinking only of a particular device, such as a xerographic printer, without corresponding terminals for display), produces unbalanced systems that extract a real cost throughout the system.

In particular, data structures place strong and specific requirements on input and output interfaces. Multifont text is not possible without high quality displays and multifont printers. Speech provides another example --without microphones or loudspeakers what can be done with a computer system that processes speech? Though interfaces seem to present a bewildering variety, they can be understood in terms of what data structures they support.

The types of data available make an immense difference in what sorts of computing are possible and what groups in a community can use computing profitably. The provision of multifont text has opened up computing to the general office world. Graphics is just becoming widely available, and throughout the campus we found many groups for whom its availability would provide a major increase in the utility of computation. As with other basic types of data (e.g., multifont text, floating point numbers) the potential uses of graphics are very diverse. Examples include statistical displays, design applications (requiring high-quality graphics), artistic media, large dynamic classroom displays, and the special symbolic notations needed for mathematics.

Software. The situation is more complex than in the days when software was divided into operating systems and programming languages, and there is much inconsistency and awkwardness. A first feature of the current state is a variety of tools: statistical packages, linear programming packages, editors, electronic mail, simulations, computer-aided design systems, document production systems, and on and on. Although these tools are sufficiently complex and powerful that learning them is much like learning a programming language, once learned, a good tool allows one to bypass the phase of writing a program, turning the computer into a device one uses immediately to solve problems. Intensive and sophisticated user communities are emerging which do essentially no programming, but only use a wide array of tools.

A second feature of the current software situation is the welding together of tools, languages and operating systems into unified programming environments. These environments are complex and only gradually become comfortable places to work. The environment for the 2060s, for example, extends back fifteen years, while the less comfortable environment for the VAXes is much newer.

A third aspect of current software is the slow development of machine independence, started with programming languages, FORTRAN, COBOL, PASCAL, and to some extent LISP. Many tools, written in these public languages, are also system independent. Even operating systems are beginning to be machine independent (e.g., UNIX and C/PM). Machine independence has a long way to go, however. For one thing, the quality of implementations on different machines often differ significantly.

Support. A collection of hardware and software requires human support to to weld them into a usable computational facility. Support includes acquisition, installation and maintenance for both hardware and software; also user education, consulting, systems analysis, and preparation of documentation and directories. At one time a computation facility periodically acquired a total computer system from a single vendor. Now environments are built through small increments from many sources. Thus the flow of new things requires continuous support (e.g., new documentation, increased user consulting). Modern computer systems do have tools (on-line documentation, demand printing, on-line walk-through of problems by users and consultants) to help with some of these functions. However, the basic ingredient of the support system remains what it has always been -- a well-managed human organization.

The Current State

CMU is already one of the most intensive university computing environments in the nation. Figure 1 summarizes the computational situation as of Fall 1981 in terms of the hardware facilities and the numbers of users. It does not reveal the number and diversity of operating systems, programming environments, programming languages and software tools.

The total annual computing budget at CMU can be usefully lumped into three parts: Computation Center (\$2.5 million), CSD and Robotics (\$2.3 million), and departmental systems (Administrative Systems, Chemistry, Physics, Psychology, etc.) (\$1.1 million). Amortized equipment costs are included in these figures. However, departmental minicomputers are not included, nor is the small amount spent by individuals and individual research grants on terminals and small personal computers. Most important, the extensive support for departmental systems provided by graduate students and faculty is not included.

For the Computation Center, averaging over the past several years, about 35% per year is spent for capital acquisition and 65% for operations and support. Roughly, 20% of their budget goes to undergraduate education, 45% to research and graduate education (which cannot be separated), 25% to administration and 10% to external and commercial uses. Almost all of the CSD, Robotics and departmental budgets go to research and graduate education. Putting it all together yields a total annual budget of \$5.9 million, with 9% for undergraduate education, 68% for research and graduate education, 17% for administration and 6% for external and commercial. This provides about \$140 per undergraduate student per year and \$2000 per faculty/graduate-student per year. Average figures for computer usage must be interpreted with special caution, since usage is always distributed very unevenly, with about 90% of the computing being done by about 10% of the users.

Organizationally, there is a Vice Provost for Computing and Planning, who is responsible for all computing on campus. He directly supervises the Computation Center and Administrative Systems. The latter has responsibility for all administrative computing on campus; it is a software organization, working entirely on Computation Center systems. Major decisions about computing are made by the Computer Board, made up of the Provosts and President in addition to the Vice Provost for Computing. There is an advisory Computer Policy Committee, which is broadly representative of the whole campus. There is also a recently established Computer Education Committee whose function is to explore how to use computers educationally in the CMU environment. The departmental facilities are administratively separate from the above organization, though their acquisition decisions must be approved by the Computer Board, and several of them have their machines maintained by the Computation Center. They are mostly too small to have distinct organizational structures. The exception is the Computer Science and Robotics, which run a joint facility. There are about 45 full-time equivalent people in the Computation Center, 15 in the CSD-Robotics facility and 20 in the departmental facilities (including Administrative Systems), for a total of about 80 people.

Current Planning. CMU is not sitting idle with respect to computation. Numerous segments of the campus community already have or are laying plans to deal with or take advantage of increased computation.

The most important of these plans is the proposal to work toward providing convenient and round-the-clock access to personal computers for all undergraduates, permitting an integration of computation into the educational program. This would be attained at some time in the future, conventionally put at five years. This proposal was put forth by President Cyert in Spring of 1980, and there has been considerable discussion with sentiment expressed in several directions. The administration is actively working towards this goal by exploring possibilities for a joint venture with an industrial organization, that could provide resources for personal computers. A Technical Committee, whose members are computer experts drawn from several departments, is working on the technical details.

Meanwhile a wide variety of departments are engaged in individual planning and development of new computer facilities, including a network of powerful personal machines (computer science), a local center of research-oriented machines including several VAXes (psychology), new resources for undergraduate education (electrical engineering). There are many more examples.

The Task Force has reviewed with some thoroughness all these developments. It views its role as formulating positions and recommendations towards helping all these efforts work together.

Basic Positions Underlying this Report

Committees to study computation are a regular feature of university life, both here and elsewhere. Commonly, major concerns are first with equipment and second with economics -- with cost/benefit analysis in the narrow sense. The Task Force believes that a more useful focus for its own analysis is based on the following positions.

#1. The base-line is a substantially increasing use of computers

The starting point for analysis might be whether there should be any increase in computing. But the neutral assumption is, in fact, that computing will continue to increase substantially, with or without comprehensive planning. Both growth of computation in the external world and the already deep involvement of the university drive CMU along this path. The important issues are then how to shape this growth and whether to seize the moment to make a dramatic move to accelerate the increase in computational facility to attain some worthwhile goals.

We take the position that a substantial increase in computing at CMU is fundamentally good. By a substantial increase we mean that most of the community would no longer find computing a scarce resource; terminals would be available when needed, response times adequate to most tasks, storage sufficient, and printing quick and convenient.

We believe in the potential of this increased computation not only to help all manner of separate activities, each in their own way, but also to support new modes of integration and community. Although the ongoing computer revolution has often failed to achieve these goals, and we share with many in the CMU community concerns about the effects of letting the computer seep further into our lives, withal, we believe the promise is great and the benefits clearly outweigh the dangers. We believe our energies should go into finding the right way to proceed and to do a high-quality job.

#2. The analysis should focus on designing an environment and predicting impacts

Our goal is to set up guidelines for the growth of an appropriate computing environment. These guidelines address what should come first, major advantages that could accrue to CMU, impacts on all segments of the campus, and identification of negative side effects with ways to abate them.

Our guidelines do not prescribe total plans for action (e.g., what equipment, what vendor). But we believe that such specific plans should grow from the design of the environment we begin to outline.

We have also not addressed details of how the current environment should be modified in the short term. Some short term decisions are being made and implemented. Others should be formulated in keeping with the longer-range goals we address here.

#3. Analysis should be in terms of general types of computational facilities

Computation Center

Total user community: more than 3500 users from entire campus (except CSD)

5 DEC 2060 TOPS-20 Systems as follows:

TOPSA (2400 megabytes) for Administration, CC development

TOPSB (1400 megabytes) for GSIA & SUPA research & graduate education, external

TOPSC (1400 megabytes) for CIT research & graduate education

TOPSD (800 megabytes) for undergrads, H&SS research & graduate education

TOPSE (600 megabytes) for undergrads, MCS & CFA research & graduate education

1 VAX 11/780 (300 megabytes) for research & graduate education

2 DEC11/45 RSTS systems (80 megabytes each) for word processing

All systems networked together via DECnet (1 megabaud link) with connections to:

Psychology, Chemistry, Physics VAXes and CSD 2060 (also 1 megabaud)

Terminals (most hardwired, some dialup; most 1200, 2400 baud, some 300, 4800): over 550

1 Xerox 9700 Multifont printer (2 pages/sec max)

Computer Science Department and Robotics Institute

Total user community: 400 in CSD, RI, Psych, EE, Math for research & graduate education

1 DEC 1080 (similar to a 2060) (1400 megabytes)

1 DEC 2060 (1060 megabytes)

1 DEC KA10 (one fifth 2060) (40 megabytes)

5 VAX 11/780s (600 megabytes each)

6 VAX 11/750s (half 11/780) (400 megabytes each)

All systems networked together via Ethernet (3 megabaud)

Personal computers: 18 Altos (one sixteenth 2060), 44 Perqs (one sixth 2060)

1 Dover multifont Xerox printer (half Xerox 9700)

Terminals (half 1200, half 9600; most with access to all systems): 250

Departmental systems

Architecture VAX 11/780 (600 megabytes)

User community: 25 for research & graduate education, jointly with CSD

Chemistry VAX 11/780 (500 megabytes)

User community: 40 for research & graduate education

Electrical Engineering VAX 11/780 (160 megabytes)

User community: 70 for undergrad & graduate education

Mechanical Engineering VAX 11/750 (120 megabytes)

User community: 20 for research & graduate education

Mellon Institute Computer Engineering Center VAX 11/750 (180 megabytes)

User community: 25 for research

Physics VAX 11/780 (1700 megabytes)

User community: 60 for research & graduate education

Psychology 1 VAX 11/780 (400 megabytes) & 2 VAX 11/750s (120 megabytes each)

Total user community: 50 for research & graduate education

All systems connected to DECnet or to Ethernet

Figure 1: Summary of computer facilities at CMU, Fall 1981.

We believe that the important issues can be discussed better by avoiding details such as vendors or specific equipment. As discussed earlier, the important aspects of a computational environment are specified by general characteristics: processing power, addressing space, amounts of secondary memory, network bandwidth, etc. Discussing the issues in these terms permits focusing on essentials.

This view implies little attention to relative costs. While we have no quarrel with economics, narrowly defined cost-benefit analyses shift attention away from how large changes in computation will affect the entire campus community. Furthermore, if two pieces of equipment or software are sufficiently comparable to make a cost-benefit comparison valid, then choosing between them is probably an implementation decision better left until we have defined the total environment we want.

Behind this position is an assessment that large amounts of computational facility can be acquired by CMU within costs that are acceptable, due in large part to the radically decreasing cost trends in the industry. Thus, economics is not a primary concern in our discussion.

The Issues

We now present the issues we believe are the most critical to our computational future over the next decade.

1. Access to Computational Facility

Gaining access to computational facilities includes everything between deciding to use the facility and starting to accomplish something. Thus it involves making contact with hardware and software (e.g., a terminal and a port), discovering relevant tools, and learning how to use them. No matter how effective the computational tools themselves, gaining access is an impediment to the task at hand. If computation becomes a much larger part of our lives, good access, broadly defined, becomes absolutely critical.

The position of the Task Force is that effective expansion of computation must include generous resources and attention to making access easy and pleasant. This position has the following consequences:

1.1. The campus must have a high-quality local-area computer network

The network must provide every user with direct access, from his own local access device, to all resources on campus available to him. The Task Force believes this is the single most important technical step that can be taken to enhance the computational environment of CMU. The network should also extend beyond the campus boundary so that access does not decline or disappear when working at home. (We recognize the technical difficulties here.)

1.2. There must be good means for discovering available computational resources

A multitude of computational tools exist both in the Computation Center and scattered through the community. There needs to be an information system of sufficient scope to index all these tools. This information system itself must be universally accessible, presumably both on-line, over local and remote networks, and in hard-copy form.

1.3. There must be access to appropriate data bases

As with other organizations, CMU finds essential its accumulated information: libraries, accounts, inventories, space occupancy, student and personnel records, schedules of classes and events, alumni records, proposals, and more. Enhanced computation can provide access to such information for anyone with a legitimate interest (with due regard for privacy and security). The benefits to CMU in increased efficiency, accuracy, responsiveness and

community would be very large.

1.4. Good documentation and human consultation are essential for access

Interaction with knowledgeable people lies on a continuum with good documentation -- both are needed in some mixture to obtain good access. Documentation and user consulting are required for the entire computational environment, not just the part residing in the Computation Center. Likewise, the obligation to provide these resources is not solely the responsibility of the Computation Center.

Ideally, good documentation of all facilities should exist in a variety of modes (on the computer as well as in mass-printed or demand-printed form), and at a variety of levels, especially primers for beginning users in addition to exhaustive system descriptions.

Although the continuously changing array of tools makes full documentation beyond the resources of any single organization, we believe documentation requires substantial effort, creativity, and resources, so that the most essential jobs can be done. For instance, a highly responsive organization could produce, for those systems most in demand, brief documentation made permanently available through networking in a highly indexed on-line data base. This organization might usefully overlap with human user consultants, who could use their experience with users in producing user-oriented documentation.

Final comment

We have purposely placed access first on our list of issues. This is primarily because enhancing computation means enhancing computation usage, which means good access. However, we are also responding to the universal cry of frustration we have heard from the campus community about the current state of access, including all the issues discussed above. This frustrating current state was also the major cause of cynicism and doubt that CMU could implement a truly beneficial major expansion of computing.

In an effective computer-intensive environment, access must be easy, reliable and pleasant -- even gracious. This is an absolute requirement for any major increase of computation, far more important than, say, how to package computation (e.g., as personal computers). A good campus-wide computer network plus a good organizational framework for the other aspects of access would provide a basis for letting CMU grow into an immensely productive computational future. Failing to provide this will almost surely lace any other scheme with painful frustration.

2. Computation and the Educational Process

The computer is a tool -- a means, not an end -- to be used by the instructor when judged valuable. Like the book, it is powerful and has particular usefulness in the instructional setting. Again, like the book, it can be used to convey great wisdom or immense trivia. More than the book, computers can get in the way of education, if access is difficult and use is tedious and time consuming, or if the student comes to use the computer inappropriately.

Based on these views, the Task Force takes the following positions on educational computing:

2.1. Vastly increased computational facility can be a major aid to education

The uses of modern computation range from the management of education (grades, assignments, scheduling), to information display in the classroom (dynamic situations), to text-processing for writing (both generally and in composition courses), to problem solving in data-rich (social science) and computation-rich (engineering) areas, to tutorials in drill and practice domains (e.g., solfege in music), to simulations (engineering), to exercise checkers (proofs in logic), to design tools (engineering), and much more. The sophistication will increase with

time, and in the end it may rival the book in its impact on education.

No simple characterization of the use of computation in education is possible. The gains arrive along many different avenues. It is clearly a rich resource with possibilities for impact that keep expanding. It is also clear to us that an integrated facility (as in Issue 1) opens up these possibilities.

2.2. The use of computation will be diverse and extremely uneven

The diversity of potential uses implies that usage will also be uneven. Substantial segments of the university will use the computer very little. The computer is not relevant to some educational efforts; some faculty will legitimately choose not to use potential that does exist; and many uses are not cost effective because they require too much effort, either of faculty or students.

The preceding statement is important only because we are talking about a very intensive computer environment -- in other environments use is extremely uneven because most people do not use the computer at all. But greater availability of computers should not create expectations that everyone will use them. Our view is that the extent and variety of computer use for education will be impressive, but that certainly it will not be used everywhere.

Furthermore, new tools, however useful, do not change the fundamental role of the teacher in choosing freely how best to educate, without distracting pressures either from climates of opinion or from direct administrative expectations.

2.3. Access to robust computation is absolutely critical when computation is woven into education

Most educational facilities -- classrooms, books, paper and pen -- are extremely robust, and even when catastrophe strikes (the heat goes off, a book is lost), substitute arrangements are manageable. The computer is not so robust, and as described in Issue 1, access is a real problem. Locking into this technology means that failures in the technology can become failures in instruction. Thus, major expansion in computation for education must include generous support of robustness and access.

Low reliability and poor access are especially destructive for individuals being introduced to the computer, not only because they are unable to work around problems in the short term, but because their initial problems can affect their long term view of computation. Students strongly told us that initial bad experiences had really put them off, and argued strongly for better attention to new student users.

2.4. Computer literacy is an important educational goal

Literacy -- narrowly, being able to read -- can be made to carry a broader meaning: Being able to use a basic intellectual skill competently and intelligently, though not with mastery or scholarship. Literacy in some areas (e.g., writing, mathematics) is considered an essential component of becoming educated. We believe that computer skills should be added to this list.

Computer literacy is not equivalent to learning to program, but is different in at least the following ways: First, using a computer often means not programming, but using available tools -- editors, electronic mail, statistical packages, simulators, computer-aided-design systems, data base systems, and so on. Using these tools intelligently requires skill, together with knowledge of when to use them appropriately. Second, using a computer effectively involves more than just local skills for using particular languages or tools. It requires understanding the fundamental nature of the computer, what kinds of things it can and cannot do. Third, if the role of computation is to increase qualitatively at CMU, computer literacy must include competence in the local computational facilities. This part of computer literacy is in fact an important part of good access.

Computer literacy is not a separate educational objective. If computation becomes more intensive in our environment, then we must convey to our students, as part of their education in all areas, the difference between intelligent and inappropriate use of the computer.

Without entering into the design of curricula for computer literacy, we believe that many different devices are appropriate for different functions -- continuously available short courses to introduce the tools in the environment, diverse special-purpose courses offered by a variety of departments, computer-relevant topics integrated into other courses, on-line self-teaching aids for acquiring new tools, as well as regular courses in computer science.

Although computer literacy is most important in undergraduate education, a good literacy program would benefit others as well -- visitors, graduate students, existing faculty who had not previously found time or need to assimilate the new technologies. All of these people will find our environment impenetrable if there are not easy ways to assimilate available tools.

2.5. The social dimension of computation

Much speculation exists about the social effects of greatly enhanced computation, especially at the undergraduate level. Much of the speculation is evoked by the notion of each undergraduate having a personally owned computer. There is fear that this will isolate students, who will spend all their time interacting with their machines and not with other persons; also that officialdom (including faculty) will penetrate too deeply into students lives. But also envisioned are ways to use broadly available computation to make possible new social activities and to communicate and organize public events (social, intellectual, sports).

Attention must be paid to social questions, even though the nature of the questions will only gradually emerge. We make the following recommendations: First, the computer should not be restricted to the world of work; it is a general aid to living like the telephone and the television. As with these other aids, there will be a tendency towards overuse which must be counteracted not by restricting the use of computers, but by helping students understand the increasing societal role of the computer, including both its powers and its limitations. Second, much computing should take place in social settings. Terminal rooms are places of learning, communicating and socializing. Such places must be pleasant and conducive to such activities. When terminal rooms are crowded, noisy and laid out only for person-terminal interactions (all charges we heard repeatedly about the current scene), then the good things happen only fitfully at best.

Making computation readily accessible implies putting it near or in residences. But there are problems in placing devices, which may be hot, noisy and space-taking, into an already crowded residence world. Even beyond these mechanical aspects, computers may violate space that is now "where one gets away from it all." Solutions may exist, but they can be expected to be costly. We recommend that these problems be taken seriously and that, if necessary, resources be diverted from the computers themselves to support modifications to student residences or provision of auxiliary working spaces.

3. Management of the Computational Facility and its Expansion

As the community comes to depend increasingly on computation, the infrastructure that supports computing must become increasingly responsive to the needs of the community. We include here all the issues of who controls the various computational facilities, who decides on their growth, and what sorts of administrative and organizational structures are appropriate to a world increasingly saturated with computation.

Classically, computation is organized in a university in a small, centralized administrative hierarchy, abetted by an advisory committee, with essentially no integration of independently supported research computation. The proposed fully-networked CMU environment would be very different. The set of significant decision makers will be of the order of fifty to sixty, and will include members of almost every department. Already, there are

significant facilities in several departments, all of which are definitely part of the overall CMU computation system. They do and will continue to make demands on the general system -- for maintenance, software, compatibility, printers, storage, etc. Their quasi-autonomous growth, through independent entrepreneurial efforts, will provide a substantial fraction of the total resources of the system. We believe the large set of decision makers is a great strength of the system, though we acknowledge that it poses difficult management problems.

The position of the Task Force is that effective expansion of computation depends fully as much on enlightened management as on high quality technology. Based on this view, we believe:

3.1 Sufficient resources must be provided to build and maintain a first-rate infrastructure

In the recent past, CMU seems to have followed a policy of preferentially putting resources into equipment rather than support services. This imbalance must be ameliorated if a first-rate infrastructure is to be built.

There may be little romance in budgeting funds to support new hardware facilities. There may also be an underlying feeling that personal computers and networks should make support services unnecessary. Yet, if funds sufficient to the cause are not available, the equipment cannot be well-used. "Running lean" is not an appropriate strategy for building a high-quality infrastructure, particularly where concern already exists about the robustness of services that can be provided (see 2.3 above).

3.2. The Computation Center should play a strong role in creating an environment for decentralized systems

The Computation Center should encourage individuals or departments to develop and control local facilities, and to provide services to the community, as appropriate to their special research and educational roles. Similarly, the Computation Center should concentrate its energies on creating central facilities where central control and action seem most appropriate, e.g., networks, mass storage or large acquisitions of common equipment.

Encouraging decentralization does not imply a weak Computation Center. On the contrary, our view is that the Computation Center should create a strong central facility that permits the other facilities to flourish. A strong central system, with stable interfaces and support systems, provides a world to which the other facilities can adapt and exploit to everyone's advantage.

The overall goal should be to move control, expertise, and services as close to the user as is possible, without sacrificing the quality of important shared facilities. Achieving this goal will clearly not be easy, and will require both creativity and resources. However, if done well, our development of an innovative and effective support infrastructure would not only be of great benefit to ourselves, but serve as a model for others.

The development of such a strong but supportive central facility requires inventiveness plus aggressive planning and implementation. Thus, though we favor committees that aid and advise in this process (e.g., the current Policy Committee), the control and responsibility for the central system should be vested in the Computation Center.

3.3. The Computation Center should manage by incentives

Even with greatly enhanced computing, resources will always be limited compared to all the tasks it would be nice to do. A great deal of these resources will be controlled by the Computation Center. In so far as possible, management should be by incentives -- by making it pleasureable and profitable for people to use the facility in the ways that fit into the general plans of the center. This use of market mechanisms is a central means by

which goal 3.2 above might be satisfied, and is to be contrasted with management by edict, which simply specifies how users must behave.

3.4. There must exist a continuously updated public plan

Since decision making is dispersed, coordination requires an explicit and regularly updated plan. Such a plan should provide at least a common view for all decision makers of what computation currently exists and what is solidly planned. Even this minimal planning in widely accessible form would be an immense aid. The responsibility for formulating and updating this plan should fall equally on all of the decision makers, and not just on the Computation Center. Although there is an essential tension between commitment through plans and the freedom necessary to be appropriately opportunistic and flexible, the kind of plan suggested above restricts the freedom of entrepreneurs very little.

3.5. Expansion must be managed without major disruption

There is no predictable steady state towards which the computational environment is tending. The most accurate prediction is that change will be the order of the day in the foreseeable future. In this chronic expansion, we must be extremely careful not to sacrifice the present for the future. Fortunately, technology seems finally to support incremental evolution without massive disruption at every step. Management of a major expansion will itself be an important and substantial activity, which will require generous resources.

Certain guidelines are critical. Good organizational support should precede technological expansion. Publicized plans should be laid out to enable all to plan their own adaptations to change. Expansion should be conservative in preserving past facilities for backup, providing contingency resources to smooth over difficulties. High risks should not be taken without the participation of those affected.

4. The Rewards of Leadership

By advocating a substantial increase in computation at CMU, we are advocating that CMU lead universities into a world saturated with computation.

Even with this commitment, CMU has a range of options on issues of leadership. For many, the focus is only on internal goals and not on leadership, but others throughout the community also sense the wider opportunity. Here, we probe what is implied by leadership in the wider community, and what paths of leadership we recommend.

4.1. Large payoffs come from being a leader

Leadership has its price, but also its rewards. There are direct payoffs for those involved in leadership, but also rewards for the institution more generally. First, leadership engenders the ability to gather the resources and the talent to do great things in the area of the leadership itself. Second, leadership and excellence in one area often translate into a general ability to generate resources and aid for adjoining areas and for the institution as a whole.

Thus, if CMU's development of computation is sufficiently prescient and responsive as to attain a position of leadership, then many rewards follow. Specifically, many of the resources necessary to carry out development can come from external sources, interested in learning about effects of intensive computer use in society.

In this respect, CMU is currently exploring the possibility of a joint partnership with a vendor. The reward to any vendor is the opportunity to develop its systems in a prototype of future educational environments. This provides both a test bed and a living example for marketing purposes. The rewards to CMU are a major sharing of the costs of computational expansion, together with less tangible rewards such as direct impact on the

evolution of the educational computer industry. While the details of any such joint partnership are far from determined, it clearly would involve a long term relationship with a large industrial computer firm. Though part of a developing pattern in the 80s of university-industry relationships, such a partnership contains some risk -- as one would expect in an attempt to seize leadership. The Task Force believes such developments can be well worth the risk, although the details must be examined carefully.

4.2. Computation has some remarkable properties as a candidate for a leadership area

CMU, as every university, seeks areas in which it can be excellent and provide leadership. Excellence in the various distinct scholarly areas are largely the separate achievement of each (although excellence in teaching in one area certainly requires good quality teaching throughout the university). In contrast, computation, although peripheral to the main tasks of the university, supports all other areas. It enters into some much more than others, and in separate ways in each, but it enters directly. Thus all areas have a direct interest in the quality, diversity, and organization of the computational facilities. Leadership in the development of an outstanding computational environment can directly help other areas in their own search for leadership and excellence. Thus, development of outstanding computation has a unifying, synergistic influence on the university as a whole.

The number of directions in which a university might strive for excellence are limited. We are all familiar with the counsel to small institutions to build on strength. Computation is one of our strengths, one we can exploit. A desirable goal, and one that is within our grasp, is to be able to demonstrate to others the best way to use an excellent computational environment to further the traditional activities of a university. At the same time, we can take the lead in developing new areas which computation might open for universities.

4.3. How to deal with the costs of leadership

There are costs of leadership. First, there is intrinsic risk in exploring any new area. Thus, we should move cautiously and proceed in a way which maintains the greatest possible flexibility, and care should be taken to establish contingency plans to restrict this risk. Second, there are opportunity costs, for choosing to attempt leadership in one area means choosing not to attempt it in others. We believe there is no way out of some exacerbation of this problem in the present case. We would counsel a strategy of being sure the benefits of computing are truly widespread. We also believe that when a community chooses to place resources in one area, it accepts the obligation to be sensitive to the effects in other areas.

Conclusion

We stop at this point, not because we have run out of issues, but because attention needs to be focussed. Many issues of utmost importance have not been touched -- security, privacy, the relative needs of special sorts of computing, how we proceed from our current situation (with its problems), and others. Many of these issues can be dealt with only in the context of a specific total proposal for how the future should develop, which must include specific vendor offerings, performance parameters, costs and time scales.

This preliminary report has presented the issues we see as most crucial to setting the CMU community upon an appropriate path. We believe we have been able to reflect here, not only our own assessments, but some of the sense of the campus. However, much diversity exists on campus and we know we have missed important opinions. We wish this report to lead to further discussion so that a more adequate expression of the entire campus can occur. As mentioned in the introduction, we will be taking steps to hear that expression.

Appendix

Charge to the Task Force

Richard Cyert
29 October 1981

I am creating the Task Force for the Future of Computing at Carnegie-Mellon University to develop a comprehensive model for the role of computing at CMU during the decade of the '80's and guidelines for realizing that role. I see this committee as critical in further involving the entire campus in our planning for the future.

Such an effort takes place in a context in which the importance of computing on our campus continues to increase dramatically. This is occurring pervasively and not just along a single dimension and for a special group. The causes of this increase lie in large part in the exponential growth and ramification of computing in society as a whole, which impinges on CMU as it does on every university. But the causes lie also in our own repeated decisions throughout the last quarter century to involve ourselves with this new element -- using it in our research, understanding its nature, building it into the fabric of our scholarly lives.

These decisions continue in full force: the proposal for personal computers throughout undergraduate education; the proposal in MCS for a special experimental facility for large numerical computation; the decision of Computer Science to construct SPICE, its own powerful personal computer environment; the decisions to add yet more DEC 2060s to the Computation Center; and the acceptance of some Hewlett-Packard small systems to be used in undergraduate education in Electrical Engineering and in the elementary computer courses.

It is imperative that we attain a broadly-based view and a sense of perspective about where we want to go. Yet, no moratorium can be called on the existing activities while the Task Force deliberates. Current activities flow from real needs embodied in earlier plans and from substantial current planning and discussion. For some activities, time is of the essence. Moreover, we must recognize that planning and decision-making regarding computing is a continuous process, with each new plan applying, in effect, a course correction to the ongoing process. Thus, the Task Force should proceed with dispatch, so its view can begin to influence our course as soon as possible. Concurrent activities and plans about computing must proceed with caution because the Task Force's report may cause shifts in the direction that CMU takes. But they need not be suspended.

In the light of the above, I would like to have a preliminary report by the beginning of the coming Spring semester. The Task Force itself will need to determine what must be accomplished in toto, hence when a final report can be completed. However, that we face continuous incremental decision-making about the computer implies that the final report should not be a final view. Rather, it must be a structure that can operate as a planning framework for continual modification and updating. It will not do to have to evoke in 1983 a new Task Force for the Real Future of Computing at CMU, just because we have learned so much more and the options available have expanded so rapidly. Thus, I would hope the Task Force would address this issue of continuous planning.

I have personally committed myself to a major expansion of the role of computing at CMU. You certainly know of my public pronouncements on this score and the discussions that have been held this last year, mostly around the possibility of personal computers for undergraduates. This commitment has grown from my sense that CMU has already taken this direction by those quarter century of decisions and also from my own belief that it offers us an immensely exciting future.

It is clear that this direction implies substantial funds. My own commitment has included a belief that a large

fraction of such funds can come from external sources, if our plans are sufficiently exciting and realistic. But there is no doubt that a real fraction must also come from internal funds. I do not want the Task Force to be concerned primarily with the trade-off of funds for computational facility versus other activities. No fixed set of funds is available, so such a trade-off cannot be addressed except in the light of a model for what computing at CMU should be, hence what prospects exist for external funding. On the other hand, there is no planning without some sense of resources. I would hope the Task Force would find some way of addressing priorities on the desirability of various computational facilities and the interactions that couple various uses together, so that the trade-offs can be addressed more easily as information becomes available about external sources.

The Task Force should consider all factors relevant to the use and impact of the computer on CMU -- on education, research, campus life, recruitment, national image, type of student, administration, whatever. You may engage in any investigations you think appropriate. As you may know, several other standing committees have also been created to address various aspects of computation on campus: the Computer Policy Committee, the Instructional Computing Committee, and the Technical Committee. These committees all have immediate ongoing responsibilities and thus differ from you. However, you will all need to cooperate closely. To help this happen, the chairpersons of each of these committees are also members of the Task Force. I should note that the administrative structure for handling and planning computing is an appropriate concern of the Task Force.

The Task Force is advisory to me, Provost R. L. Van Horn, Vice Provost D. E. Van Houweling, and the entire University. Doug Van Houweling will provide staff support for the Task Force and help you with any problems you have.

I have asked Allen Newell to chair the Task Force. The membership has been drawn to be widely representative of the entire campus. This has made for a fairly large committee, but I am sure you agree this was the appropriate decision.

I wish you well on your task and trust that you will develop a view and planning guidelines for the role of computing at CMU that will exploit this revolutionary technology to our benefit and yet be both realistic and commensurate with our long-term goals.