THE INFORMATION TECHNOLOGY CENTER

James H. Morris and Douglas Van Houweling
Carnegie Mellon University

Keith Slack, IBM Corporation

Confluences

The Information Technology Center is an organization developed to mediate and serve several purposes. Essentially, the ITC was formed by merging concurrent trends in diverse technologies and institutions.

Two significant technological trends - one in computing, the other in communications, are influencing the creation and development of distributed systems. The trend in computing is the phenomenal reduction in the cost of the digital computer's heart: processors and random access memory. Very powerful computers are currently available to nearly everyone who has a need of one. Indeed, the question has become: What can the average person do with a powerful computer? The trend in communication is the increase in the importance of communications and sophistication of media in our society. Although dramatic cost reductions have not been observed, both the market for television and the growing cost of transportation are promoting great growth in the cable television and telecommunications sector. These two trends answer one another's needs: the average person can use a computer to enhance his ability to communicate in a world where the amount of information available is growing exponentially. The growing shortcomings of today's communications, the intrusiveness of the telephone and the massifications of television, will be greatly offset by the use of computers to mediate communication.

Two important institutional trends - one within CMU, the other at IBM, have created a project and an organization -- the Information Technology Center. CMU is a small private university with a strong commitment to research on and use of computers. The administration of CMU decided a few years ago that a

heavy commitment to the uses of computers in education was it's top priority for the decade of the 1980's. In roughly the same time period, top executives at IBM decided to increase IBM's presence in universities. This was not merely a desire to sell more equipment to the university market; the major objective is to enhance the knowledge and decision-making skills of the future university-educated work force. In the fall of 1982, CMU and IBM decided to jointly develop a distributed computing system for the university. The agreement calls for a five year project, with a re-evaluation at the end of three years. The Information Technology Center is responsible for developing the system. The ITC consists of 30 to 40 people: ten of whom are IBM software professionals, the rest CMU employees. The organization is administered by CMU, and is primarily a software design and implementation group. The basic presumption is that the ITC will use IBM hardware and network technology whenever possible. The system developed by the ITC will belong to IBM; CMU will have a royalty-free license. Both organizations intend for the system to become an IBM product. As yet, there are no contractural commitments to deploy or purchase equipment.

The Product

The ITC's objective is to implement a distributed campus-wide network consisting of personal computers and various types of server machines. The community it will serve consists of about 4,000 undergraduates, 1,500 graduate students, 450 faculty, and 1,000 staff. The University consists of six very distinct Colleges and Schools. The computation requirements of the university are very large and diverse. Students and faculty use computers for programming courses, engineering and scientific computation, statistics, business simulations, and laboratory simulations. Administrative computing is done in the same context. Eventually, most of the actual computation in these areas will be performed on personal machines on a distributed network. This will free the large computers currently used for these computations for use in activities such as those that require massive database manipulation. (There are basically only three justifications for a computation to be performed centrally: it is so massive that it requires a very large computer; it is tightly intertwined with a central database, or it depends upon a system that has not been implemented on the personal workstation computer.)

The campus is already breaking new ground in the use of the computer as a communications media. High quality document preparation is a major service in great demand. Several sub-communities make heavy use of electronic mail and bulletin boards. The library is proceeding to computerize its card catalogue. Our project is destined to greatly accelerate and broaden these kinds of communication services. The significance of a communications system grows with the square of the participants. The large number of people coupled with the power of their personal computers and their need to communicate along many dimensions is the most novel, challenging, and potentially significant aspect of the project.

The most over-worked piece of equipment at CMU is the Xerox 9700 laser printer. The reason is quite simple: by using it and a document preparation system called Scribe, one can produce a good-looking document that looks like it might have been professionally printed. Anyone who writes - and there are a lot of them in a University - likes to seek their work presented in the best possible way. Automatic typesetting, demand printing, and other tools like spelling correctors can make writing more effective and enjoyable. In the short term, quality printing will be the strongest draw of the system; people's appreciation of electronic mail and other things takes longer to form.

This improvement of quality in communication does not stop with multi-font printing. That is only the tip of the iceberg. People also communicate with pictures and sound. Here, however, the technology is far beyond the printed world. There are not many graphic artists who would pay to use what is commercially available. We can expect; however, that simple kinds of illustrations will be easy for novices to include in their written documents.

The possibilities for communication do not end with words, pictures, and sound, no matter how powerful. A computer program is a form of communication we have just begun to explore. The programmer/sender of a message actually creates an interactive mechanism that the user/receiver can interact with to comprehend. Construed as art, compilers and operating systems are like houses people called users live in; the builders of these things must employ the same kinds of judgements and inspiration and understanding of their clients

needs that real architects do. The Adventure game is a rudimentary form of a new kind of art form: the interactive puzzle/mystery.

Unifying the current concepts of electronic mail, data bases, and information retrieval is another intent of the ITC. Electronic mail is essentially a one sender, one receiver, non-updating sort of system. Information retrieval systems are one sender, many receivers with very occasional addition of information. Data bases, as typified by an inventory control systems are many senders, many receivers, dynamic systems. Where should the ITC system fit in this multi-dimensional spectrum?

Electronic mail is a very powerful communication tool. Like the telephone it allows immediate communication over long distances. Unlike the telephone it does not require simultaneous sending and receiving. This crucial improvement makes it a far more useful tool than the telephone as anyone who has played "telephone tag" will attest. By allowing communication to be asynchronous, electronic mail increases the number and kinds of conversations a person can have. In particular, it permits an incredibly enhanced version of the time-honored dormitory bull-session. It is not uncommon for a discussion carried out via an electronic bulletin board to go on for several weeks and involve hundreds of participants. These discussions can be like a dormitory bull-session in that their direction is entirely random, people are free to wander in and out contributing as they wish, and they are spontaneous and fun. Also, the enthusiasm generated by these free form discussions generates brainstorming sessions to solve academic problems.

It is possible to have less free form, more productive and directed electronic discussions with the same properties. One can easily imagine that a high proportion of class room discussions will be carried out through electronics. The teacher and students will all belong to a message group that exchanges messages on chosen topics throughout the semester. In many ways this is a more reasonable medium for thoughtful discussions than face to face group meetings. People have more time to think about what they are saying: in fact, they are forced to because they must type it. When a teacher wishes to base grading on participation, electronic discussions provide a written record he can draw from.

Communication across time might seem to be a different topic from electronic mail - information retrieval or data bases perhaps - but we believe it should be looked at in the same context. To oversimplify: an electronic mail system is merely a distributed data base system in which individual recipients of mail search the data base for items with their names in the "To" field, in which recipients of bulk mail search the data base by topic, and so on. It is already commonplace for people to accumulate all the messages of a particular electronic discussion into a single message file for latecomers to the conversation. Unfortunately, this process is somewhat hit-or-miss and does not produce good information for the long term. The purely mechanical task of collecting everything on a common subject should be part of the system. The thing that will always require intelligence and attention is the task of editing this information in the light of subsequent events and information.

The concept of a growing, democratically created, intelligently edited data base will greatly amplify another venerable, if dubious, campus tradition: the fraternity files. If, year after year, you capture all the written communication associated with a course - assignments, tests, papers, etc.- the potential exists to change the nature of how people learn the body of material. It will no longer be just the nominal teacher of the course who does the teaching; it will be generations of unseen students transmitting their solutions to problems and views on issues. This potential can be realized, however, only if this data is suitably culled for what is truly worth studying.

The Project

Our efforts are divided between two major areas: The Storage and Transport System and the Work Station.

The major challenge in the Storage and Transport System is to provide an absolutely reliable system that can store and distribute information for the community. This part of the system bears the brunt of the huge size of the community. The basic image of the storage and transport system is that of a giant time-sharing file system. In fact, it will be implemented by many computers distributed in a network, but the user will see only one file system and need be only occasionally aware of the physical distribution of the sources. It must be expandable without limit. Mechanisms for security, accounting and metering will be provided.

For the Work Station system, the challenge is quite different: to develop a user-friendly system that utilizes graphics and other aids to assist non-expert users in achieving their various objectives. While we shall be very conservative in the storage and transport systems, favoring reliability over functionality, in the workstations area the issue of functionality will be paramount. The workstation will provide a high quality raster-graphic issue of display on which multiple windows will give access to concurrent activities. A high quality "what you see is what you get" editor will be provided. A major challenge is to develop a programmable, yet mouse-driven interface.

Immediate Plans

1983: Explore various technologies and implement a small, 40 person network based upon interim Work Stations. Also, help to specify an advanced Work Station.

1984: Deploy a small network to a community of about 200 users, and assess its ability to expand.

1985: Deploy a large system to about 1,000 users, based upon an advanced workstation.

Challenges for the University

If one assumes the technical success of the ITC's project narrowly construed -- thousands of people equipped with personal computers, able to communicate through a dependable storage and transport network -- even greater challenges and questions remain.

The major one is that of using the system in meaningful ways to enhance the educational process. In the coming years, many educators will devote large portions of their time to rethinking how to teach given these new tools. Exactly how this effort is to be fostered is unclear.

The whole question of the organization that develops and distributes software needs to be re-examined: Should it be subcontracted like the Cafeteria or student-sponsored like the campus Radio Station?

It seems clear that the users of this system will be drowning in information. Nobody will have the time to look at everything, and only a small percentage of it will be worth looking at for any particular person. There will be a desperate need for the electronic equivalents of the newsroom editor, the literary critic, the art reviewer, the computer program reviewer, etc. These people will become more influential than the individual producers of information and services.

There was once a mechanical engineering student at CMU who maintained his position at the top of the class because he knew how to search the literature. He took every problem to the library and searched until he found a good solution. Occasionally he had to so some invention, but he always began with that he could find from the past. One has the vague feeling that this was cheating, but there is no doubt that he was going to be an incredibly effective practicing engineer as long as he had access to this information base. We had better come to terms with this method of problem solution, because our tools are about to make it much more attractive and effective. Our general idea that students must learn to "think for themselves" is in danger. We could find ourselves in the situation that the only way to have them do it is always to present them with entirely new problems.