

## NIC COLLECTION

# The Evolution of ARPANET



ARPANET, the world's first packet switching network, has been a vital element in the nation's computer and communications research activities, but its managers determined that it was technologically obsolete. The ARPANET spirit lives on, however, in the networks it has spawned, now known as the Internet, which will supplant ARPANET for IS-critical research.

BY BRAD SCHULTZ

**A** torch in data communications research is being passed from ARPANET, the world's first packet switching network, to the Internet. It is on the Internet that most of the nation's important research about computers is done, with profound implications for commercial IS managers.

"The Internet" is the colloquial name for the largest operational superset of research-oriented internets. Its three major subnetworks are ARPANET (Advanced Research Projects Agency Network), CSNET (Computer Science Network), and NSFNET (National Science Foundation Network), each of which is a network of computer networks.

### Allowing Researchers To Interact

According to Vinton G. Cerf, vp at the Corporation for National Research Initiatives, Reston, Va., ARPANET and its successors and progeny, "such as the NSFNET, CSNET, public data nets, local area networks, and regional networks, have formed the basis for interaction among a very large fraction of the computer science and engineering research community." He adds that "this infrastructure has been adopted by other researchers not directly involved in computer communications research."

The size and complexity of the Internet are awesome. Some university computer centers within it have local

area networks larger than the ARPANET, says internetworking consultant Einar Stefferud of Network Management Associates, Huntington Beach, Calif. It also dwarfs some commercial packet nets.

As Stefferud tells it, the commercial


networks—some wide-area, many more LANS—is unfathomable in such terms.

As evidenced at the Enterprise Networking Event in Baltimore last June (see "Enterprise Networking: MAP/TOP Clears Its Biggest Hurdle," June 1, p. 19), commercial IS managers have a big stake in the research on the Internet, which began with the development of the Internet Protocol (IP) and the Transmission Control Protocol (TCP) on ARPANET.

### Eliminating Costly Gateways

The IP, which made the Internet feasible, allows networks that are proprietary—at the levels of links and interfaces between computers and links—to be networked without costly gateways (to the extent that IP modules run at the nodes of those networks). The TCP/IP protocol suite is eliminating the need for such costly commercial IS procurements as software systems, which track the flow of individual packets across a network, and gateways, which bridge mutually incompatible networks (see "A Close-up of TCP/IP").

Slated for termination, ARPANET is now only a small subnetwork of the Internet. Last December, U.S. Army Maj. John Mark Pullen, the Defense Advanced Research Projects Agency (DARPA) program manager responsible for ARPANET's \$10 million budget, and USAF Maj. Brian Boesch, who also manages DARPA distributed systems programs, decided that ARPANET had grown technologically obsolete as a high-risk



CNRI'S CERF: Replacement of ARPANET is an essential step.

IS manager who boasts of managing tens of thousands of terminals would probably provoke blank stares from internetworkers. Who could count terminals in the Internet? Who would want to? The Internet's nesting of networks within

PC  
N that

is new  
turning  
in, on

X is  
k for  
work  
stret-

copy

INK

3-604

## A Close-up of Transmission Control Protocol/Internet Protocol (TCP/IP)

The entire Internet complies with the Internet Protocol (IP), which was developed along with the Transmission Control Protocol (TCP) to meet the Defense department's especially tough internetworking requirements. IP is implemented as a switching/routing module, usually software, on computers at an internet's nodes.

Connected by trunks, IP modules provide services for use by modules that implement higher-layer protocols, such as TCP. This is the definitive characteristic of layered protocol networks. Each layer (except the lowest) uses the services of lower layers and each layer (except the highest) offers services to higher layers.

The Internet implements transport protocols besides TCP, but TCP has certain advantages, given the needs of many Internet users, and generally runs with IP.

A considerable number of TCP/IP-conformant networks, some of them outside the Internet, now operate in the commercial sector, academic institutions, and civilian government agencies. Meanwhile, TCP-4 has emerged as an alternative to TCP, and the Connectionless Network Service (CLNS) has emerged as an alternative to IP. Both alternatives conform to the ISO's Open Systems Interconnection Reference Model.

The following sketch of TCP/IP builds on a physical mail metaphor and draws on advice from internetworking consultant Einar Stefferud of Network Management Associates, Huntington Beach, Calif. It also serves to describe TCP-4 and CLNS.

In a strictly TCP/IP internet, data travels between sender and receiver applications inside segments, created by TCP modules, which in turn travel inside packets of a special format called IP-datagrams. These IP-datagrams allow TCP segments that comprise a long applications datastream to take different routes through the internet toward their common destination.

TCP modules are active only at the communication path's end points, where applications data are inserted into or removed from TCP segments, which in turn are inserted into or removed from the IP-datagrams. IP modules are active at the end points in the layer below the TCP modules and at the switching/routing nodes across the internet.

Envision the full complement of data a user is trying to transmit across an internet as a document consisting of white typed pages, so numerous that they must be segmented in a series of separate interoffice envelopes, each inscribed with a segment number. Each of these interoffice envelopes is then inserted into a postal envelope, which is inscribed with the destination office address for mailing. Following the metaphor, the white pages represent applications data, the interoffice envelopes represent TCP segments, and the postal envelopes represent IP-datagrams. There are also yellow pages that represent data about white pages.

The sending system's TCP module accepts a set of white pages from an application, counts how many vowels are typed on each, and stuffs a certain number of white pages in each interoffice envelope along with a yellow page that states the

vowel total, the number and order of white pages in that envelope, and how those white pages should be concatenated with the white pages of other interoffice envelopes to reconstitute the original complete document.

The TCP module hands the interoffice envelopes to the IP module, which stuffs them into postal envelopes for mailing and inscribes on each the appropriate return and destination office addresses, as well as some damage detection indicators and the time when the envelope's value expires.

At the IP switching/routing nodes, IP modules route the postal envelopes to the designated offices on the basis of the inscribed postal addresses, but these modules ignore the interoffice envelopes until the postal destination is reached. IP switches also check the condition of envelopes and destroy damaged or overage envelopes.

At the receiving postal address, an IP module removes an interoffice envelope from each postal envelope and hands it to the indicated office (the appropriate TCP module) for processing. This office processing includes detection and correction of errors and disorder. To resolve errors, the receiving TCP module re-counts the vowels on the white pages in each interoffice envelope, compares that total with the yellow page total, and may ask the sending system to resend any segment for which the totals conflict.

To resolve disorder, the receiving TCP module compares what the yellow page says the order of white pages should be with their arrival order. The yellow page also guides the concatenating of those white pages with the white pages of other interoffice envelopes. Protocols of a layer higher than TCP are concerned with such issues as the format of the white pages, and the style and grammar of what is typed on them, whereas protocols at the IP layer and below are concerned with postal carrier activities, such as bundling, unbundling, and sorting postal envelopes, and handling envelope bundles at airports and on aircraft.

The essential problem IP solves is how to build an internet from nodes and links that comply with a wide variety of lower-level protocols (such as X.25, Ethernet, and token passing). Because nodes and links differ as to the speed and method in which envelopes are bundled, unbundled, sorted, trucked, and flown, it is impractical to standardize exhaustively at the lower protocol layers in internets.

Internets always will need to satisfy the lower-layer protocol requirements peculiar to different generations of technology, let alone those peculiar to different systems architectures. So, IP was contrived to at least end divergence on postal envelopes. Without something such as IP, two networks that differ on envelope standards can only exchange data via a gateway, a system that opens and re-orders envelopes, and converts address information.

Gateways impose extra overhead and risks. IP obviates the need for them, while TCP provides the required higher-layer function of monitoring the condition of the data as the data arrive at the destination, as well as the state of the processes that move the data through the network.

### TCP-4 HAS EMERGED AS AN ALTERNATIVE TO TCP.

rese:  
occa  
PANF

ARPA  
mov  
thei  
Inte  
isfac  
prov  
kille  
will  
that  
ARP  
win  
Inst

ARP

pris  
cre:  
gy l  
link  
dat:

ope  
178  
spe  
er  
swi

atic  
seg  
atu  
tro  
fro  
ing  
be  
ain  
alt  
ne  
us

m:  
bo  
bu  
sp  
1C  
w:  
m  
sf  
to  
in  
ty  
pi

C:  
ti  
30  
at  
e  
a  
c  
b

## The Evolution of ARPANET

research resource. Minor user activities occasionally were swamping major ARPANET arteries, Boesch says.

So, Pullen and Boesch notified ARPANET users that they would be moved, without serious disruption of their work, to a new Defense Research Internet (DRI). After those users are satisfactorily ensconced in DRI, and DRI proves to be stable, ARPANET will be killed. Boesch estimates that probably will occur in four years' time. He muses that some hardware that played a role in ARPANET's historic achievements may wind up on display at the Smithsonian Institution.

### ARPANET's Honorable Discharge

ARPANET's obsolescence is no surprise, Cerf says, since it originally was created for 50Kbps lines and "technology has reached the point where 1.5Mbps links are needed to support the mass of data exchange now required."

With the growth of high-speed LANS operating at rates from 10Mbps to 178Mbps, he adds, it is evident that speeds of 100Mbps to 1.8Gbps and higher will be needed in the backbone, with switching capacity to match.

DRI will have experimental and operational subnetworks, says Boesch. The segregation is intended to insulate operational activities, such as routine electronic mail and computer conferencing, from the risks of experimental networking. But there is a catch-22, Boesch adds, because experiments on networking, aimed at verifying the value to users of alternative approaches to networking, need to involve a reasonable sampling of users doing what users normally do.


Trunk link speeds on DRI's experimental and operational subnetworks will both begin at 1.544Mbps, Boesch says, but the experimental subnetwork's link speed eventually will be phased up to top 1Gbps. "Although not without its awkward moments," Cerf says, "the replacement of the ARPANET with a much higher speed capability is an essential step toward the creation of a new information infrastructure for the research community and should be viewed positively by all parties interested."

The Internet, however, lives on. CSNET and NSFNET serve computer scientists and other scientists who need to use advanced data networks. CSNET was created mainly to let academic computer scientists share computational resources and exchange mail, explains one of its creators, David J. Farber, professor of both computer and information science

and electrical engineering at the University of Pennsylvania.

CSNET inspired NSFNET, which ties together several supercomputer centers and is also aimed at academia. The NSFNET, however, is directed at researchers interested in computers in any scientific discipline. Some researchers use what was the first major academic network, BITNET, Farber explains, but BITNET mainly serves academic administrators and falls short as an internet.

The value of research on the Internet is well exemplified by MEMNET, a fledgling project that Farber and his colleagues have in the works, which could eliminate the need for network protocols altogether and, in turn, a lot of products that commercial IS managers now must buy. MEMNET is concerned with the opportunities posed by ultrahigh-speed



### MEMNET SHOWS THE VALUE OF RESEARCH ON THE INTERNET.

links, which can move at least a few hundred million bits per second.

A computer network configured as a MEMNET essentially operates as a single giant computer. How big can a MEMNET be? The size of a college campus? A county? Texas? The world? Farber and his colleagues plan to find out, implementing it initially on a slice of the Internet that spans the Middle Atlantic states.

MEMNET runs a high risk of failing to meet expectations and of posing unforeseen problems. But not to worry. The Internet, and especially ARPANET, throughout their histories have been intended as havens for high-risk research.

It was in that same spirit—to explore the then unknown merits of networking computers—that the agency now called DARPA hired Bolt, Beranek & Newman (BB&N), Cambridge, Mass., to create ARPANET, which opened in 1969.

Just after ARPANET opened, one of its principal designers, Robert E. Kahn, left BB&N for DARPA and began some high-risk research with packet switching ideas, from which the TCP/IP packet switching protocols ultimately emerged,

Cerf says. Last year, Cerf was made an IEEE fellow for having led the TCP/IP development. The TCP/IP protocol suite was first standardized by the U.S. Defense department and later was adopted as a de facto industry standard.

At DARPA, Cerf says, "Kahn started the packet satellite, packet radio, and internetting programs. Respectively, these looked at sharing a common satellite channel by multiple ground stations [dynamic allocation], sharing a common high-speed radio channel by a number of mobile digital radio units, and the establishment of a protocol architecture and specific protocols enabling heterogeneous packet networks to be linked and operated as if forming a common internet."

The research culminated in the first public demonstration of packet switching in 1972. "For the first time," Cerf says, "the technique of packet switching was demonstrated in a hands-on fashion. Especially impressive was the [ARPANET] network's ability to support echoplex operation [the echoing of characters back to the user for display on his or her video screen] in real time... [because] each character was packaged up as a packet, sent store-and-forward style through the packet net, processed by the host, re-packetized, and sent back for real-time display.

"Apart from demonstrating the feasibility of the packet switching concept," Cerf continues, "the ARPANET stimulated the development of a variety of new computer-based applications, such as electronic mail exchange, among hundreds of host computers and thousands of researchers. Before, the only electronic messaging available was found in much less responsive military systems, such as the AUTODIN network. With the introduction of electronic messaging in the highly heterogeneous host environment of the ARPANET, the value of layered protocols and common standards for computer communication was made dramatically apparent."

### ARPANET's Ties to Commercial Ventures

The demonstration led BB&N to launch Telenet, the first commercial packet net, in 1975, with consequences that transformed the role of computers in commercial communications. BB&N's Telenet subsidiary went public, diminishing BB&N's share in its equity, and it later was acquired by GTE. Telenet Communications Corp., Reston, Va., is now a subsidiary of U.S. Sprint, a joint venture of GTE and United Telecommunications.

## The Evolution of ARPANET

ARPANET was also the focal point for training two or three generations of graduate students in the difficult art of protocol and network design, Cerf notes.

"The many new startups arising in the mid-to-late 1970s and early 1980s in this field were often the direct result of ideas evolved from personal experiences with the ARPANET," Cerf says.

Examples include 3Com Corp., Mountain View, Calif., which was begun by Robert Metcalfe, the co-inventor of Ethernet; Bridge Communications, Mountain View, Calif. (now a division of 3Com), founded by Bill Carrico and Judy Estrin; and Advanced Computer Communications of Santa Barbara, Calif., founded by Roland Bryan.

ARPANET's progeny, the Internet, has no overarching formal management structure, although its operations are guided to varying degrees by a few boards and committees with members in common, generally organized by the federal government. Chief among them are the Federal Coordinating Council on Science, Engineering, and Technology (FCCSET, pronounced "fix it"), the Feder-

al Research Internet Coordinating Council (FRICC), and the FRICC's Internet Activities Board (IAB).

The FRICC was created to coordinate the development of a new major

### ARPANET IS NOW A SUB-NETWORK OF THE INTERNET.

subset of the Internet, tentatively called the Interagency Research Internet (IRI), which would be the largest operational superset of civilian federal research internets. R&D-intensive civilian agencies such as NASA, and the companies that do business with them, would use IRI.

IRI was proposed about a year ago by the Network Working Group of the FCCSET's Committee on High Perform-

mance Computing. A spokesman for the group, Barry Leiner, assistant director at the Research Institute for Advanced Computer Science in Moffett Field, Calif., tells DATAMATION that "things are moving forward rapidly" on IRI.

"The various agency networks are being interconnected," says Leiner, "and agreements [are] being initiated to allow appropriate sharing of resources. NSFNET has played a key role in all of this by providing a coherent method for providing access for the university community."

"There will be some short-term rough spots," Leiner adds, "as the agencies and the FRICC try to cope with the massive explosion in the use of the networks, compared to the available resources [network and funding], but that is called being a victim of success, something we shouldn't really complain about." IRI's total operations budget would begin at around an equal level with ARPANET's (\$10 million per fiscal year) and would reach \$13 million by fiscal year 1991.

#### Internet's Boundaries Defy Definition

What complicates attempts to manage or even analyze the Internet is that its boundaries defy definition. Any computer in the world capable of IP-conformant communications (Berkeley Unix systems, for example) can connect to the Internet, in principle, though administrative policies exercised on both sides of the connection would determine availability of the Internet's resources. Access to the Armed Forces' MILNET, for example, is extremely restricted in order to preserve national security, yet the IP-conformant MILNET might be considered part of the Internet.

Because BITNET does not conform to IP, it can only exchange data with the Internet via a special gateway at the University of Wisconsin. For the same reason, commercial companies and individuals lacking IP modules need gateway services to exchange mail with Internet addresses.


One such gateway on the market is the DASNET gateway service of DA Systems, Campbell, Calif. A long list of commercial electronic mail services, including facsimile (Group 3), MCI Mail, Telemail, Telex, and Easylink, can front-end to DASNET and exchange electronic mail with ARPANET, CSNET, NSFNET, and most other Internet addresses. ■

*Brad Schultz is a freelance writer based in New York City.*

## Mainframe Interconnection Is Our Line

### Our Products Include:

- Blue Cable "New"
- Bus/Tag Testers
- And Much More



### Our Pledge:

- Zero Defects
- Overnight Shipment
- Technical Support

### Mainframe I/O Bus/Tag Cable Assemblies

For Pricing And Information Call

**REN Electronics Corp.**

(A Major OEM Supplier)

755 New Ludlow Road, South Hadley, MA 01075

East of Miss. R. & MN. 413-536-1800 West of Miss. R. 719-275-7595

# MATION

## NEW PRODUCTS

### 50 CIO: Misfit or Misnomer?

BY RALPH EMMETT CARLYLE  
What is the true status of the much vaunted chief information officer? The mythology has it that this bridge between the executive suite and the computer room is a vital piece of the management mix, and perhaps the breeding ground for future ceos. But a survey conducted by DATAMATION and Coopers & Lybrand finds few in its even hold the title, and those who perform the function find they have little corporate bite. With:

### 53 Kodak's Katherine Hudson: Clearly a CIO Model

A conversation with an IS director who is also a corporate vp.

### 65 Managing a System Transformation

BY FREDERIC G. WITHINGTON  
What do you do when business requirements mandate that your entire information processing environment be transformed? Panic? No need. New products and the experience of others will make the task easier.

### 71 The Evolution of ARPANET

BY BRAD SCHULTZ  
ARPANET, the first packet switching network, had hosted U.S. datacom research activities well, but technology has passed it by. Its progeny, the Internet, is ready to assume the helm for IS-critical research. With:

### 72 A Close-up on TCP/IP

How this Internet transport protocol works.

### 77 Hardware

Compaq brings 386 power to its low end, as IBM launches three versions of the PS/2 Model 70-386. In Trends: what does the V.42 modem error-control standard promise users?

### 80 Software

IBM debuts an optical-based document management system. In Trends: U.S. user companies are placing off-site operations in Ireland.

## DEPARTMENTS

### 4 Letters

### 82 Calendar

Even if the Cards aren't in the pennant race, St. Louis still has the ShowCASE Conference to look forward to in September.

### 83 Career Opportunities

### 88 The Marketplace

### 88 Advertisers' Index

## INTERNATIONAL 48-1

Does not appear in all copies.

### -1 The European 25: Smashing the \$50B Barrier

BY PAUL TATE  
1987 saw key shifts in IS development and purchasing patterns that will affect European vendors' fortunes in the '90s. Meanwhile, revenues for the top 25 IS firms in Europe were up 26.5%. With financial data on:

### -2 The DATAMATION European 25

### -6 The 25 Leading European-Based IS Companies

### -8 IBM Europe's Major Markets

Cover illustration by Min Joo Hong

DATA MATION

INTERNATIONAL EDITION

AUGUST 1, 1988  
VOLUME 34  
NUMBER 15  
THIS ISSUE, 181,247 COPIES



1988 JESSE H. NEAL AWARD



1987 JESSE H. NEAL AWARD

F  
R  
I  
If  
inf  
me  
Ne  
the  
act  
co  
for  
ne:  
su  
acc  
off  
Ba  
dir  
wo  
for  
cor  
exp  
Aln  
bel  
pos  
amb  
bus  
of t  
the  
be t  
chi  
are  
just  
De  
tecl  
to t  
poli  
ma  
sign  
hov  
sof  
coc  
ma  
tow  
drc

5 AM: "Why is logging on so slow?" the early-bird operator mumbled. "I'm the only one on the system." Later that morning the system operator explained the problem to him. Someone had slipped in a Trojan horse. The perpetrator's program was scooping up all the permissions and access rights of people when they logged on, then turning things over to the real exec. The intruder's program was rigged so that if he executed it, he could read or write any of the files the others had rights to!

The account being used by the alien belonged to an honest Joe who had left the company a while back. Guessing Joe's password was easy, which seemed to explain how the alien could get in to use the account.

The system operator doctored the program to announce any log-on by the alien. When he logged on again via a network, the net connections were traced to a guy at a university. The traced connection revealed that the man had graduated and moved on. The connection to that particular network was shut down, hoping the alien would think the system had crashed.

The alien logged on again in 30 seconds through another network, using a lady's name! A fast trace revealed the woman's identity. She was contacted immediately. She was not on-line. The system operator promptly changed her password and killed the alien's browsing.

The perpetrator was back again in 30 seconds, using a new name...

Hardcopy  
Dec. 1984  
v 13:12

## Exploring Issues In Multi-Level Security

Military downfield blocking is helping the private sector develop security systems that curtail some pretty tricky tampering

by Harry T. Larson

Illustrated by Bill Buerge

NIC COLLECTION



**W**e have all read about computer crime, now growing to unnerving proportions. And we are only seeing the tip of the iceberg. The reported cases reveal an extraordinary variety of ways people commit crimes in systems that employ computers. We can no longer keep our heads in the sand assuming that the problem exists elsewhere.

Something is being done about the problem. A fundamental body of work spanning 15 years is approaching practical use. It is a story of battles of wits, massive failures, hard-won gains, and solid science with cloak and dagger overtones.

Perhaps you have heard distant mutterings about computer security, replete with new jargon and acronyms. Now that security-enhancing systems are appearing on the market, we would do well to check out the fundamentals, that we might make knowledgeable decisions when we adopt such systems.

Few computerized systems have been developed with the bad guys in mind. Loosely handled IDs, passwords, and daily code numbers have permitted capers ranging from invasions by hackers to the transfer of \$10.2 million to a Swiss bank account.

We have not devised systems to prevent disgruntled or greedy employees from lifting valuable programs or tinkering with the files, nor have university systems been designed with unscrupulous students in mind.

We have not designed to prevent computer sophisticates from taking over the operating system, proceeding to do whatever they choose — rummaging through files, corrupting data or programs, or introducing the tools of moles such as "trap doors" or "Trojan horses."

The list of methods for misusing systems is long, interesting and sobering. It is sufficient to say that we have a massive multi-faceted problem on our hands. It will come as no surprise that the defense establishment has long been concerned about this class of problems and has been sponsoring work to correct the situation. It may be a surprise, however, to learn that new methods which are now emerging are probably more relevant to security and privacy in non-military environments such as business,

industry, services, education, R&D and non-military governmental functions.

While the Department of Defense (DoD) properly protects its computerized systems, DoD has been aiming higher, for superior and less burdensome techniques that protect against malicious insiders as well as outsiders. Foundations for new methods for achieving computer security have been developed by the joint efforts of DoD, its contractors, the National Bureau of Standards (NBS), academia and computer companies. As might be expected, the trail through uncharted areas has been bumpy, strewn with instructive failures and partial successes. Although we have not yet reached the final des-

---

*"...necessary levels of security cannot be achieved by modifications or add-ons to existing software."*

---

tinuation, worthwhile concepts and methods are taking root in both hardware and software design.

#### **Security's Plucky Past**

A brief history is the best way to identify and explain today's activities. In the period preceding the early 1970s, DoD sponsored a series of major efforts aimed at developing new types of built-in security in computer systems. These projects typically focussed on changes in the (mostly commercial) software, capitalizing on protection features available in the operating system and supported by the hardware, e.g., memory protection and multiple levels of privilege.

Despite the ingenuity and capability of the architects, all known efforts fell short of the extraordinary goals that had been set. At the completion of each project, skilled individuals or tiger teams were assigned to try to break in, while working inside the overall external protection for the system. Time and time again the friendly teams succeeded.

With hindsight it is easy to say that probably few if any computerniks understood how tough this problem is, nor anticipated the many paths of

entry available to knowledgeable attackers. External security was okay. Internal security was not.

The time had come to re-examine the situation, to build foundations for a more promising way to devise smart built-in armor for computerized systems. A time for ferment in the late 1960s and early 1970s produced several significant concepts that have influenced computer security developments to this day.

One of the fundamental actions of early efforts was to define several modes of operation to provide an overall framework for how much and what kinds of security protection are required in various situations. At the most demanding extreme, beyond the state of the art at the time, was "multi-level security" mode, now entrenched in the jargon of the computer security community as MLS. In this mode of operation the computerized system contains information at several levels of sensitivity, in the military case called classifications Unclassified, Confidential and Secret Personnel using the system have several levels of access rights (called clearances), say Top Secret, Secret and Confidential.

The hardware and software required to exercise control are analogous to the way military security is maintained with people, paper, locks combinations, guards, procedures accountability, etc. A key role of the hardware and software is to provide internal mechanisms to control access to stored classified information or programs based upon a person's clearance and "need to know."

The need-to-know principle simply says that people who have, say a Secret clearance, don't have a right to see all information classified Secret — they have the right to see only the Secret information needed to do the job. This mode of operation applies to systems where multiple users are sharing the same computer resources and includes systems accessed from locations remote from the CPU. Using precautions like these, it is hoped the modern counterpart of the spy in the trenchcoat with a miniature camera will be denied from viewing the files.

Further, it is assumed that it is a "open" system, meaning people are present who have full knowledge and documentation about the system, including the design of security

hardware and software mechanisms.

The policy in the MLS mode requires that, even in an open system, the computerized security mechanisms must: prevent unauthorized access, disclosure, modification or destruction of data or programs; prevent unauthorized execution of programs; and prevent unauthorized viewing or monitoring — whether any of this be overt, covert, or accidental. A bit of challenge. The military had placed formidable demands on itself.

As you might imagine, all this triggered questions about feasibility and sparked lively debates: paranoia vs. realism. Are the extreme precautions essential? In one sense, they are not, because DoD employs many ways to protect its systems, even where they require layered protection and great expense. However, as in the intrigue of breaking secret codes, in the final analysis it is essential to make it too difficult, or cost too much, or take too long to compromise a system. The goal appears to have been to address all these dimensions of MLS with the very real possibility of providing markedly safer systems at a lower overall cost.

#### Avenues Of Intrusion

Virtually every aspect of hardware

and software was considered to be an avenue to penetrate or tap a system, including avenues for a malicious user. These are summarized in a 1979 report published for the Office of the Secretary of Defense, originally published in a classified report in 1970 and declassified in 1975. (Figure 1)

DoD had imposed on itself an extraordinarily difficult and worthwhile requirement for this most demanding end of the spectrum. It turns out however, that the developments designed to meet this challenge are paying dividends for nearly all computer users.



*This was no casual hacker. He clearly had lists of account names and passwords. It was beginning to look like a well organized crime. The FBI was called in.*

Cumulative experience, combined with the new insights force-fed by the successful break-ins, led to the inescapable conclusion that the necessary levels of security cannot be achieved by modifications or add-ons to existing hardware or software. The at-

tempts at providing integral arm left many cracks and holes through which intellectual swords easily penetrated. Program patches were disasters, typically introducing new ways to pierce the system. It became increasingly clear that high levels of built-in security can be attained or by designing it into the equipment and software from the outset. "Security first" became an established principle.

As the calendar clicked through the turn of the 60s into the 70s, a remarkable series of developments led to a third concept which has been a strong guiding force to this day. A planning study sponsored by the U.S. Air Force assembled a panel of representatives of the computer security field. In addition to knowing security requirements and having participated in the attempts to satisfy those requirements, some of the members served as a conduit for new technology, including the work of those who were establishing foundations for a theory of computer programming.

The deliberations of the study team produced a simple and powerful conceptual model. Drawing from good works of many and adding their own thinking, the team adopted

Ed. From "Security Controls For Computer Systems," Defense Science Board Task Force report, Willis X.

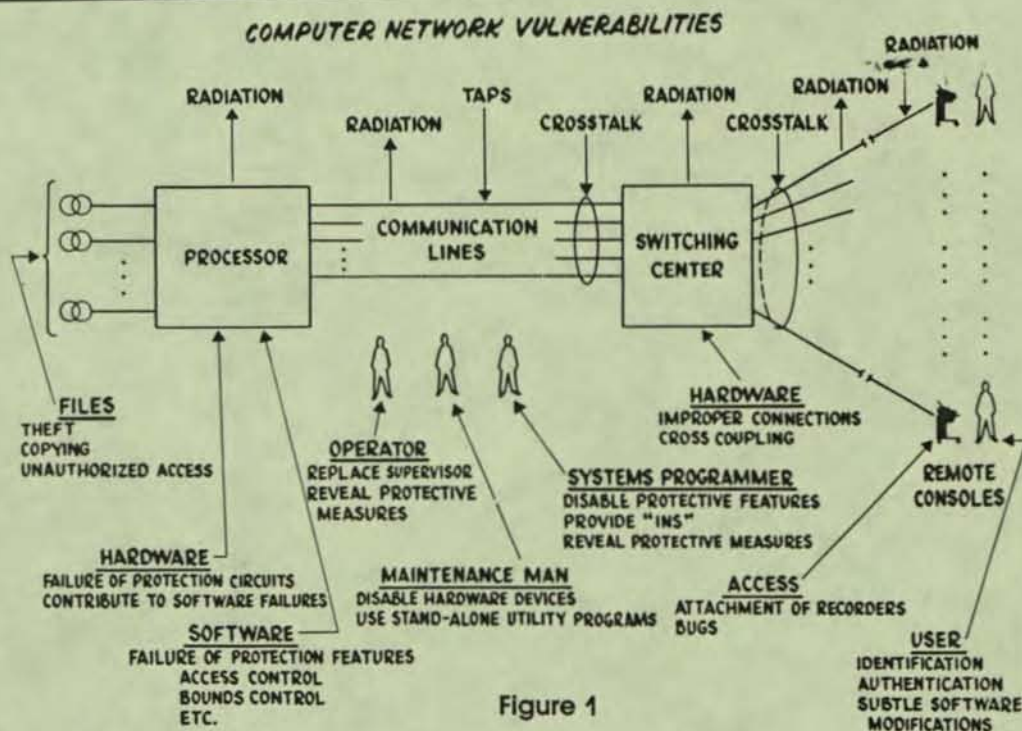


Figure 1

**Network Vulnerabilities**—By 1970 every aspect of computerized systems had proven to be a weak link in computer security.

Courtesy of Rand Corp.

concept of two entities: subjects and objects (Figure 2). Subjects include users, or programs or processes working on behalf of users. Objects are the things the subjects seek to employ to get the job done, including files, main memory, peripheral devices or other programs.

For the military example, the subjects have various security clearances and need to know levels, while ob-

depicted as an authorization database. These access relationships can be one set of rules for the military case, a different set for a business, and yet another for a university.

The study produced a set of principles characterizing the reference monitor validation mechanism. It is always invoked when a subject reaches for an object, a process which has become known as *complete medi-*

and completeness, the reference monitor was envisioned as being small and subject to exhaustive analysis and test, the completeness of which can be assured — verified. This small entity was dubbed the "security kernel." Today the term is sometimes used to mean the reference monitor alone, and sometimes includes the authorization rules which the reference monitor enforces.

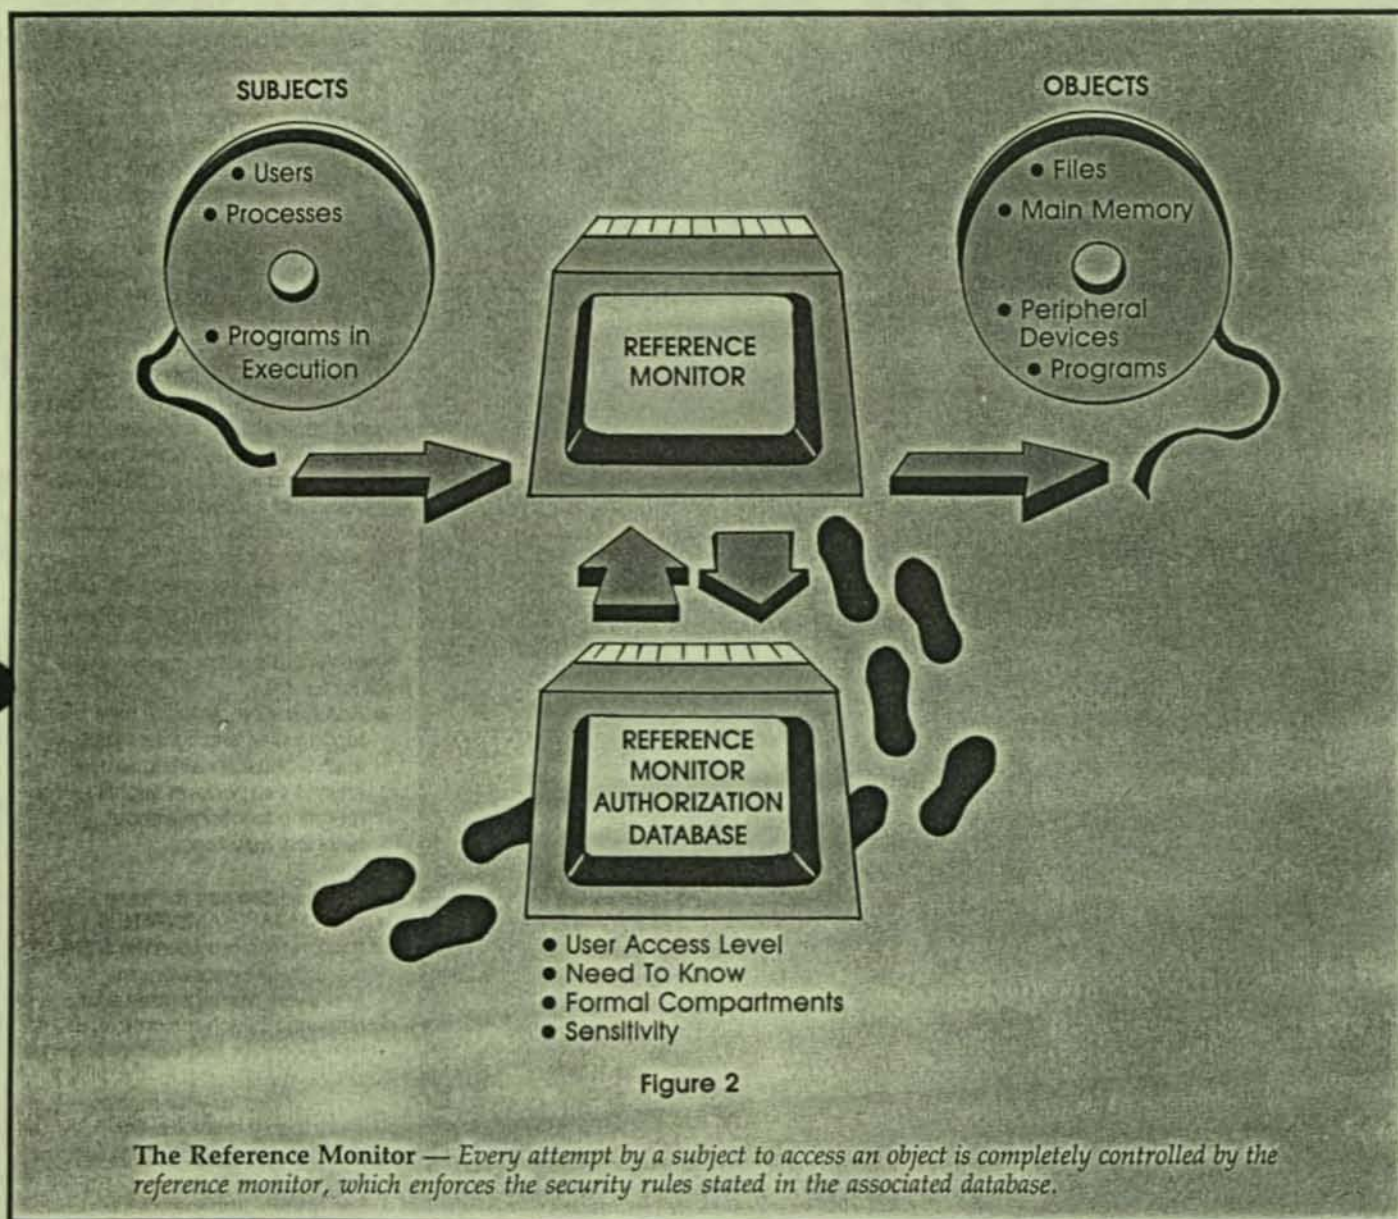


Figure 2

**The Reference Monitor** — Every attempt by a subject to access an object is completely controlled by the reference monitor, which enforces the security rules stated in the associated database.

jects carry varying levels of security classification. Between the subjects and objects stands a "reference monitor" which controls attempts by subjects to access objects. The reference monitor is a validation mechanism that enforces a set of rules that state authorized access relationships; the rules were later sometimes

ation. The reference validation mechanism must be tamperproof — isolated and self-protecting. It is to be verifiable, formulated in a way that can be analyzed and tested to assure its accuracy, self consistency and completeness.

To assure that it would be feasible to do the essential proof of accuracy

Although the small size is certainly desirable, it hasn't worked out that way to date. Nonetheless, the verification principle stands: there shall be no doubt that the reference monitor will enforce the security rules, no more and no less.

The study suggested that the authorized access relationships be

looked upon conceptually as an access matrix (Figure 3). The subjects label the rows of the matrix, and the objects, the columns. The elements of the matrix at each intersection of row and column state the access rights of that subject to that object.

For example, a subject with a Secret clearance will be blocked from accessing any object having a classification higher than Secret, and will also be blocked from accessing objects that fall outside that subject's need to know. For objects that the subject is permitted to access, the rules will also state whether the subject has read-only, write or deletion rights.

The study said in effect that no one knows how, but let's set out to implement the concept, step by step. That is a good summary of much of the research and development during the dozen years since the report of this seminal planning study.

#### Managing The Matrix

One class of development has literally implemented the access matrix, but the matrix tends to be very large and sparsely populated, so it uses computer resources inefficiently. More implementations tend to be built upon the equivalent of rows or columns of the matrix. For example, with each object, a list can be maintained stating which subjects may ac-

cess it, and their access rights. Terminology varies, but implementations along these lines are generally said to use the "access control list" approach. Or, with each subject a list of allowable objects can be stored stating access rights to each object. Implementations organized along these lines are sometimes known as *capability list systems*.

Although these structures do not exhaust the variety of methods developed to date, they are a useful

framework for understanding security products. These alternatives and variations on them, including distributed kernels, are in various stages of development by computer companies, universities and government contractors, some having reached the test and evaluation phase. New ground has been plowed by more than a decade of projects devoted to implementing secure operating systems supported by appropriate hardware, including developments aimed at "kernelized" operating systems and "provably secure operating systems." The latter do not necessarily use the kernel notion, but are intended to include security specifications in the system specifications and in every step of subsequent design of the software and hardware.

It became increasingly clear that standardized structured approaches were needed for development and testing. Formal mathematical methods were evolving for stating security policy, and rigorous mathematical methods were being developed for designing systems to assure accurate implementation of the policy.

Having seen everything else fail, DoD had reached the point that it required unequivocal mathematical proof of the design of systems which



*The FBI and the telephone companies finally put traps on the dial-in lines. It was tricky to walk the fine line, between keeping the alien out of the files and keeping him interested enough to stay on the line long enough to trace the call.*

Although the outlines of what was needed were becoming clear, it was equally clear that a considerable body of work had to be done to complete the development of the new methods, work which would require a large scale effort employing the talents and resources of the computer industry.

So the government sought to interest computer manufacturers in designing this new form of security into their systems. With one or two exceptions the computer industry didn't exactly jump on the bandwagon. Why go to all the trouble of undertaking an uncomfortable new design methodology and developing new architectures for hardware and software, all aimed at a goal not yet clearly defined, to satisfy a problem peculiar to the relatively small DoD market?

Ah, but the matter isn't confined to military security. The computer security community had developed insights that showed that the new methodologies would be valuable, if not essential, for a wide array of non-military uses of computers.

How do we prevent a terminal operator from creating fictitious suppliers and then sending company checks to them? How do we control access to company plans with money and off-color gratuities being dangled to employees to turn over such information? How do we protect against misuse of debugging aids or the maintenance man's utility routines?

The new techniques could clearly help to prevent many of the growing abuses of computers, meet business's needs to properly protect sensitive information, and help business and government meet increasingly stringent privacy laws.

#### Narrowing Avenues

To fill the need for cooperation be-

		OBJECTS						
SUBJECTS	Subjects			Files			I/O Devices	
	S <sub>x</sub>	S <sub>y</sub>	S <sub>z</sub>	F <sub>1</sub>	F <sub>2</sub>	F <sub>n</sub>	D <sub>1</sub>	D <sub>n</sub>
S <sub>1</sub>		Block Enable		Read Write				
S <sub>2</sub>			Stop		Update			
S <sub>3</sub>				Delete	Execute			
S <sub>n</sub>								

Access Matrix — Authorized access relationships between subjects and objects are defined in terms of "privilege." Figure 3

are to achieve the highest levels of security. The new formal tools, though promising, were still best considered to be research vehicles. They are, nonetheless, considered by DoD to be the only route that will assure reaching the peak they seek.

Not all programmers believe in these formal provable methods. The rather strident voices of a few veterans are heard in opposition, in words reminiscent of the antipathy between art and science.

Not all programmers believe in these formal provable methods. The rather strident voices of a few veterans are heard in opposition, in words reminiscent of the antipathy between art and science.

tween government and the private sector, NBS and several parts of DoD mounted an unclassified Computer Security Initiative, the first part of which was a technology transfer of computer security developments to the general computer community.

Launched in the late 70's, the Initiative has successfully made its case that the fundamental security requirements are indeed widely applicable. The Initiative's goals are to achieve widespread availability of "trusted" systems, and to develop technically sound procedures for evaluating security systems.

To serve as a focal point for this activity, a Computer Security Center has been established in DoD. A key effort of the Center recognized that MLS is not a binary matter, it is not a question of whether you have it or you don't. Environments and needs differ widely. Accordingly, the Center is working with computerdom to develop definitions of levels of security, ranging from no security features (for potential use in benign environments) to the strictest requirements for systems containing extremely sensitive information in high-risk environments.

The Center has also started a process intended to lead to an evaluated products list. To implement the security-first concept, a process is planned in which the design of a system is submitted at the conceptual stage and at succeeding design stages of development. At each step the Center will provide its evaluation and suggestions for strengthening the "trustworthiness" of the system. This process is in its infancy; few systems have been submitted at the conceptual stage.

Meanwhile, firms that have developed systems that offer security protection are submitting them for evaluation. Best guess is that at least a dozen commercial software and hardware/software systems have been presented, which have overloaded the Center. The Center is still coming up to speed and gaining appropriate staffing, and in concert with industry is still developing its methods for performing this new kind of evaluation. By October of this year, two commercial products had been rated and three preliminary evaluation status reports had been is-

sued with additional complete evaluation announcements in the works.

Do all these security requirements really apply in non-military environments? Yes — perhaps more so than in the military. The need is clearly established by the lengthy list of computer crimes in the private sector. But one more point merits attention.

Though the military has long had effective systems for controlling the use and dissemination of classified information, computerized or not, in the final analysis it all comes down to depending primarily upon the integrity of individuals.

Dependence on individuals is, of course, the reason the government goes to some lengths to evaluate a person's background and character before granting a security clearance. Once in the system, people are trained in the processes for maintaining security.

Few if any such formalities exist outside the military-industrial complex, nor is anyone likely to seriously recommend them. Such investigations, clearances, etc., are difficult but acceptable conditions in the military arena, but would be repugnant in other dimensions of our society.

Without the processes imposed in the military world, how is sensitive or private information to be protected from people acting out of greed, retribution, cheating for personal advantage, apathy, ignorance or the desire for a little fun and games?

These are indeed "open" systems, often open to use by people throughout the organization, and open to the sophisticates who operate the computer system. Such systems are, obviously, more vulnerable than those in the military environment. Thus in this sense it can be argued that the methods of providing MLS are more relevant here than in DoD and its contractors.

#### Safeguard Systems Forecast

Although progress is being made toward the highest provable levels of built-in security along the lines we have indicated, MLS does not appear to have been achieved in this strictest form — proven by formal mathematical specification and design. Nevertheless, we have come a long way. Systems have long been certified in the military for the various necessary levels of security, including the highest level and multi-level systems.

Many types of protective measures are employed, including power techniques that spring from evolution described herein, augmented by whatever it takes to do the job — often ponderous, brutally forceful, and expensive.



*Three Days Late  
The intruder was  
identified. A  
disgruntled former  
employee, his  
computer, floppy  
disks and related  
paraphernalia were  
promptly confiscated as evidence.  
He was apprehended—this time.*

The exact state of progress on the highest-provable front is unclear. Achieving proven MLS is proving to be a difficult problem which may take years to lick. On the other hand, what is currently under wraps in a number of places may be announced at any time.

When the formally proven products do appear, they will undoubtedly reduce throughput. Hardware and software developments are easing this burden, but it will remain a fact. There is no free lunch — how much are we willing to pay for security?

Meanwhile, useful aspects of the concepts reviewed here have been incorporated to a degree in products now reaching the market, and helping to provide intermediate levels of protection. Other security-enhancing products that use different methods have been available for some time. These offerings may well be adequate to do the jobs for which they are designed, but their level of security cannot yet be stated in a technical way because the measurements and evaluations are still in making. For the time being such systems will have to be evaluated by paranoid thinking, by how difficult they are to penetrate, and by successful usage. Meanwhile, keep passwords dry.

*Harry T. Larson is a senior engineer at a Southern California-based aerospace company, and a freelance writer.*

REPORT III

MARKET AND ECONOMIC  
IMPACT STUDY

GENERAL OVERVIEW AND  
VOLUME I: THE SUPERCOMPUTING COMMUNITY

**NYSERNet**

NIC COLLECTION



New York State Education  
and Research Network, Inc.

1095 Avenue of the Americas  
Room 1702  
New York, NY 10036  
Tel. 212-395-4480

REPORT III

MARKET AND ECONOMIC  
IMPACT STUDY

GENERAL OVERVIEW

Prepared By:

Alan K. McAdams  
Cornell University

For

MEIS TEAM

### REPORT III: GENERAL INTRODUCTION:

During the period of the last twelve months, the members of the NYSERNet Market and Economic Impact Study (MEIS) Team produced two prior reports that identified many of the needs and problems of individuals using data telecommunications networks for education and research at universities, a paper on the relationship between, "The Economic Benefits from and Public Support of a National Education and Research Network", plus the current report. In the prior reports, the team identified specific groups of individuals who are involved in education and research activities and whose needs required additional study. The several volumes of the current report are devoted to four of those groups:

- (1) Volume I, The Supercomputing Community
- (2) Volume II, The Biomedical Community,
- (3) Volume III, Education: K-12
- (4) Volume IV, Industrial Clients.

Generally speaking, these volumes review the needs of a specific segment of the education and research community, point out current problems associated with providing network service for that segment and outline recommendations for action that can be taken by an organization such as NYSERNet to serve that group most effectively. Taken together, these volumes present a broad spectrum view of the opportunities wide-area networking organizations such as NYSERNet have for making a positive economic impact on their regions--and on the nation as a whole.

Many readers have a primary interest in one or more of the segments covered in the current report, but not all of them. In order to facilitate the distribution of the report to those individuals, we have produced each volume under separate cover.

## GENERAL OVERVIEW: REPORT III

### INTRODUCTION

In this general overview section we evaluate the initial strategies incentives of NYSERNet in relation to the apparent strategies and facing other major actors such as IBM, AT&T, NSF and others. Then we augment this analysis with a discussion of the major recommendations which arise out of our studies to date. These recommendations are presented in two parts: those from earlier work and those enhanced by the current study.

The two fundamental strategies of NYSERNet have been:

(1) to focus on connecting NY State education and research institutions to the supercomputer that has come to be known as the Cornell National Supercomputer Facility, and

(2) to adopt the TCP/IP protocol suite for the network.

As we see below, these two strategies have proven to be prescient. Before evaluating them further, we review background materials including the sets of major recommendations that emerge from our studies.

### II BACKGROUND

In our initial market analysis we identified a number of uses for the supercomputer at Cornell by researchers throughout the country. We also identified a number of problems encountered by those users of the CNSF. In our second report we analyzed four types of networks including one that currently serves the CNSF. A series of recommendations for enhancing and balancing the current uses of NYSERNet as a mechanism for access to the CNSF and the other NSF supercomputer facilities, are included in the prior studies. We also contributed a paper that incorporates an analysis of the economic benefits from and public support of education and research networks. Several of the most important recommendations from those studies are paraphrased below.

- Network management at NYSERNet must be proactive if NYSERNet is to make an optimal contribution through education and research to the economic development and international competitiveness of its client region. Some of the areas in which a proactive stance is required are incorporated in the set of "Global Recommendations" in the next section. Others are incorporated in the recommendations immediately below.
- Optimal use of networks will require fully "transparent" network connection--connection that is no more difficult to use when accessing resources on the network than it is when the desktop system is used as a local system only. Transparency is to be sought from the desktop system through to target resources such as the CNSF, other supercomputers, remote data bases such as BRS, and other remote resources. NYSERNet must work to facilitate the development of this high level of transparency. This implies an ability to transport data initiated through multiple protocol suites such as APPLETALK, DECNET, XNS, etc., as well as TCP/IP.
- It is necessary to distinguish among three categories of users who can be connected to the network in differing ways. The three categories are:
  - Category I Researchers at major research universities and other research institutions at primary network nodes;
  - Category II Knowledge workers at public schools operated by local school systems; smaller educational and research institutions at other levels in the system; and
  - Category III Knowledge workers associated with category I and II institutions acting as individuals from their homes.
- NYSERNet must facilitate the interconnection of users from their homes (Category III users) to the main NYSERNet nodes. A technology which can be immediately implemented as a first stage to achieve this is a data over-voice facility available through the region's public telephone company, NYTEL.

- NYSERNet must catalyze the creation of new node "magnets" including sought-after databases to help generate traffic and to balance throughput on the network.
- NYSERNet must facilitate the rapid development of new and emerging technologies, such as facsimile, medical images, CAD/CAM, etc.

The conclusions from which these recommendations are drawn are re-enforced by the studies that have just been completed. These studies together with the earlier ones result in the following more global recommendations.

### III GLOBAL RECOMMENDATIONS

The recommendations below incorporate the new findings and conclusions from our current study. There are ten recommendations in all.

Recommendation (1): NYSERNet is well served by, and thus should continue its commitment to users of supercomputers as a major client group. These leading edge users represent a large and growing portion of the users of NYSERNet. They are establishing the requirements, approaches and procedures appropriate to supercomputing today that will be appropriate to more generalized, network computing of the future. Within a decade or less the power of today's supercomputers is expected to be available to users in their desktop systems. The experience gained by NYSERNet in dealing with the high transfer rates, graphics and other requirements of today's supercomputer users will stand it in good stead in the immediate future and throughout the decade to come as the network evolves from a packet switching data network to one that deals simultaneously with data, voice, and digital video.

Recommendation (2): NYSERNet is well served by, and thus should continue its focus primarily as a TCP/IP packet-switching network. To the extent economically feasible in the short to intermediate run, it should also transport other protocols such as DECNET, APPLE TALK, XNS etc. and do so transparently. One factor that confirms the primary focus on TCP/IP protocols is the current emphasis on UNIX by developers of diverse computer systems from desktops

through to supercomputers. The fact that UNIX is being used as a common, vendor-independent operating system through which to establish a unified operating environment for computing is the key. Each of the two consortia that are currently competing to provide a "standard" version of UNIX calls for the use of TCP/IP protocols (to evolve to OSI standards when these eventuate) for network communications. The widespread use of TCP/IP protocols assures that these protocols will have an important place in the evolution of the future universal standard protocol suite.

Recommendation (3): As shown in Volume III of this report, NYSERNet has an opportunity to contribute substantially to enhancing the productivity of the primary and secondary education system in NY State. NYSERNet's contribution can be made in a manner complimentary to efforts currently underway in the State through the NY State Department of Education. Many of the State's existing efforts are focused on enhancing its management and control functions using IBM systems and its SNA network protocols. NYSERNet is a TCP/IP network system and thus currently is incompatible with SNA networks. As highlighted in Volume IV dealing with industrial clients, this incompatibility can be used in the intermediate term as a basis for achieving the security of administrative functions as compared to operational functions. This dichotomy of function is consistent with our continuing focus on the classification of information into "current" information, that used to manage existing institutions and "new" information, that necessary to bring about change in institutions. NYSERNet can contribute significant value added--in part through "distance learning"--to the education and research process as it applies to grades K-12 in NY State.

Recommendation (4): NYSERNet must focus on New York State as a priority source of funding in its role as access vehicle for New York State users of CNSF. The analysis in Volume I of this study demonstrates that the direct interests of the various parties, other than New York State, in the CNSF can be met through means other than NYSERNet. NY State must look to the needs of NY State users in this arena, as well as in those that apply to the K-12 education process.

Recommendation (5): To enhance its value added for users, NYSERNet should focus on providing connectivity to those unique resources of New York State and as discussed in Volume II, to nearby areas of Canada that are available solely or most effectively through NYSERNet. The objective should be to make these resources available also to the rest of the country, especially through NSFNET and its successor networks. For example, the Brookhaven National Laboratory's unique facilities are available at broad bandwidths only through NYSERNet. Also, as discussed in Volume II many of the internationally renowned medical data banks and other medical resources of southern Canada would be most efficiently tapped for U.S. users through extensions of NYSERNet.

Recommendation (6): NYSERNet should be expanded to access major nodes throughout New England. Since many of the New England states are establishing major networking capabilities for their universities, they represent natural partners for NYSERNet. Such partnerships would become increasingly valuable as NYSERNet enhances its value-added capabilities.

Recommendation (7): The combination of supercomputing and high performance workstations is resulting in the rapid growth of computerized graphics and image processing, as discussed in Volume I of this report. These areas will continue to grow exponentially. NYSERNet must be prepared to facilitate and service this growth throughout the network under its direct control and proactively into the other two networking tiers, whose performance must be coordinated with its own if transparency for the user is to be achieved.

Recommendation (8): NYSERNet must be prepared to facilitate and service education and research related to medical imaging, a field which is developing rapidly (see Volume II). The appropriate initial step for NYSERNet's proactive efforts to facilitate the development of this field must be carefully evaluated. One approach could be to catalyze a jointly sponsored conference of expert medical practitioners, researchers, technicians and government regulators. Electronic transmission of medical images is an area that requires an in-depth understanding of the incentives faced by practitioners, the practices they currently

follow, the status of current and future technologies and the rate of change of costs for implementation.

Recommendation (9): NYSERNet can exploit a significant market potential in the research and education activities of industrial firms in its service area. In meeting this potential, however, it must insure the separation of the education and research services from the proprietary on-going business activities of these industrial firms. Both NYSERNet and the industrial firms must maintain a clear-cut separation between "new" information that can be carried over the network and "current" (commercial) information that cannot. In many industrial organizations commercial activities are implemented through SNA networks with which NYSERNet's network is incompatible. As suggested above in Recommendation (3), this incompatibility can serve as the basis for insulation of the two types of activities and the security of one from the other. Such a secure separation would permit a number of goals to be achieved simultaneously, at least in the intermediate term.

Recommendation (10): NYSERNet must focus significant efforts on bringing about federal funding for regional education and research networks. This recommendation grows out of the analysis in the MEIS paper, "Economic Benefits from, and Public Support of a National Education and Research Network", presented originally at the "NationalNet '88" Conference in Washington, D.C on March 29-30, 1988.

#### IV NYSERNET IN THE CONTEXT OF THE STRATEGIES OF OTHERS

As suggested in the Volume I of the current report dealing with "CNSF in the future of NYSERNet", we see that the choice by NYSERNet to focus on supercomputing coincides with two of IBM's current strategies that deal with areas in which IBM is not yet well established: supercomputing and networking for education and research. The choice by NYSERNet to adopt TCP/IP protocols, coincides with an AT&T strategy to focus on the UNIX operating system which deals with an area in which AT&T is not yet well established: computer systems..

One of the most important recommendations, both implicit and explicit in our earlier studies, is that to be most effective NYSERNet

must facilitate the achievement of true network transparency for its users. On this dimension NYSErNet is less fortunate. The strategy choices made by NSF, the major player in education and research networking at the federal level, are antithetical to the objective of true transparency for users. NSF conceives networking in this arena as composed of three tiers:

- (A) the backbone,
- (B) regional networks and
- (C) local networks within universities.

NSF sees its role as one of supporting solely tier (A), the backbone.

Also as detailed in Volume I of this report, the likely outcome of such a strategy is wide disparity in both regional and local network expertise and performance. This is likely to result in an uncoordinated system that is difficult, if not impossible to integrate. Much more will be required by the federal government and by implication by the NSF, to bring about transparency in networking for the user. The system of networks must be conceived in an integrated fashion; sufficient resources must be made available to achieve effective implementation of the system. At a minimum this would imply the creation of an institution to identify the "best practices" currently in use as of a given time frame. In turn these can be used as the basis for developing and then promulgating standards that all NSF-funded institutions must adopt in implementing their networks. Constant re-evaluation and update of the process will also be required.

#### The Focus on Supercomputing

As stated above, the initial impetus for the creation of NYSErNet was to permit the major research universities and other research institutions of NY State to have high performance network access to the CNSF. The supercomputer at Cornell has served as an important magnet for traffic over NYSErNet. It continues to be the focal point for the development of new services likely to be offered over NYSErNet. The capabilities of the supercomputer are sufficient to stress the networking schemes at each of the three tiers identified by NSF. Each tier must function effectively and each must interrelate with the others in order for transparency to be delivered to the user.

Integrated performance of all three tiers will be a continuing requirement as demands at each level of network become greater and greater.

Given that most expert observers forecast that the power of today's supercomputers will be incorporated in the user's desk-top systems within a decade or less, it is imperative that network providers become acquainted fully and early with the requirements that such systems can and will place on them. By being required to serve the supercomputer at the CNSF locally, and to provide access through NSFNET to supercomputers elsewhere in the country, NYSERNet is gaining the experience that will serve it well as the network evolves from data only to data, voice and digital video.

As documented in Volume I, IBM is moving rapidly to reestablish itself in the realm of high performance scientific computing. IBM's approach to supercomputing based on the 3090 system with vector facilities. The CNSF is the showcase, demonstration site, and performance safety valve for IBM users of these systems. As we noted, since the CNSF happens to be the site for which NYSERNet was established to provide access, the success of the CNSF is likely to be consistent with an enhanced performance by NYSERNet. This is not necessarily the case, however, as the analysis in Volume I demonstrates. On the other hand success for NYSERNet is almost certainly positively correlated with enhanced performance for the CNSF.

Also as documented in Volume I, IBM is moving rapidly to establish itself in the realm of high performance networking in support of education and research. IBM is a partner in the Merit Consortium (Merit, IBM, MCI) for installing and managing the enhanced NSFNET backbone, tier (A), that NSF is committed currently to support. Here NSF's priorities and those of IBM reinforce one another. An effective high performance backbone will provide IBM with experience vital to its corporate objectives. It will also greatly facilitate NYSERNet's ability to provide service to its region: many of today's service problems for NYSERNet have their origin in the inadequacies of today's NSF backbone. IBM's strategy and that of NYSERNet reinforce one another.

In summary, IBM's two strategies are congruent. It is focusing on reestablishing itself in supercomputing; its premier site is the CNSF. It is focusing on establishing itself on high performance

research and education networking; the premier network in this field is soon to be the NSF backbone. And the CNSF is currently the most active node on the NSF backbone. We see further that these IBM strategies are consistent with the interests of the NSF and of NYSERNet.

### The Focus on TCP/IP

NYSERNet's initial decision to focus on the TCP/IP protocols is consistent with the main thrust in networking for research and education institutions. TCP/IP is an outgrowth of an experimental protocol suite established by the Department of Defense as a temporary, experimental expedient in the absence of "something better". TCP/IP has been widely adopted by government networks. It is here to stay, at least as a stepping stone to the widely discussed and advocated OSI, the Open Systems Interconnect protocol suite. Fortuitously, in its UNIX operating system AT&T adopted TCP/IP as its preferred telecommunications protocol. UNIX has been widely adopted by manufacturers of computer systems of many diverse types to provide their users with a standardized, vendor-independent operating system. This has greatly broadened the base of TCP/IP users.

Recently AT&T established agreements with computer systems companies to develop a standard version of UNIX for use with various types of systems. The first such agreement was with SUN Microsystems, the manufacturer of high performance workstations. SUN and AT&T are developing a standard version of UNIX for use on workstations. The second is with UNISYS, the second largest mainframe manufacturer. UNISYS and AT&T are developing a standard version of UNIX for use on mainframes. To date AT&T has not developed a similar relationship with a manufacturer of supercomputers. In Volume I, however, there is a recommendation for establishing a four-way relationship among AT&T, U.C. Berkeley, the Lawrence Livermore Laboratories and a manufacturer of supercomputers.

This effort by AT&T has spawned a counter effort by a consortium of other firms led by firms such as IBM, DEC, APOLLO, HP to develop their own "standard" version of UNIX. This consortium has created a significant budget for the joint effort. We note again, however, that this group does not involve a manufacturer of supercomputers. Given the rapid evolution of workstation and other

desktop stations with the expectation that soon desktop systems will have the power of today's supercomputers, this appears to be a significant oversight. Given the importance of networking to the future of computing, and the focus of US National Laboratories on the networking of supercomputers, the absence of a National Laboratory from both competing groups appears to be a major oversight.

The intimate relationship between UNIX and the TCP/IP protocols provides the link between NYSERNet and AT&T's strategy. UNIX has been the operating system of choice for the research community. It is because of the commitment of R&D personnel to UNIX that the pressures exist to create the unified, vendor-independent operating environment extending from the workstations to the supercomputer. All of these efforts reinforce the decision by NYSERNet to focus on the TCP/IP protocols.

The absence of supercomputer manufacturers and the National Laboratories from the UNIX standardization effort is a problem that will effect the speed and direction and development of network computing. This implies that it is in the interests of NYSERNet to take what actions it can to overcome this deficiency in current standardization efforts.

To the extent that this deficiency is not overcome, a major opportunity will be missed. The US has as a great national resource that can be used to create comparative advantage--the experience and expertise of its National Laboratories. US international competitiveness needs demand that optimal use be made of these resources. NYSERNet can help to catalyze this outcome.

REPORT III

MARKET AND ECONOMIC  
IMPACT STUDY

VOLUME I: THE SUPERCOMPUTING  
COMMUNITY

Prepared By:

Alan K. McAdams  
Cornell University

For

MEIS TEAM

### REPORT III: GENERAL INTRODUCTION:

During the period of the last twelve months, the members of the NYSERNet Market and Economic Impact Study (MEIS) Team produced two prior reports that identified many of the needs and problems of individuals using data telecommunications networks for education and research at universities, a paper on the relationship between, "The Economic Benefits from and Public Support of a National Education and Research Network", plus the current report. In the prior reports, the team identified specific groups of individuals who are involved in education and research activities and whose needs required additional study. The several volumes of the current report are devoted to four of those groups:

- (1) Volume I, The Supercomputing Community
- (2) Volume II, The Biomedical Community,
- (3) Volume III, Education: K-12
- (4) Volume IV, Industrial Clients.

Generally speaking, these volumes review the needs of a specific segment of the education and research community, point out current problems associated with providing network service for that segment and outline recommendations for action that can be taken by an organization such as NYSERNet to serve that group most effectively. Taken together, these volumes present a broad spectrum view of the opportunities wide-area networking organizations such as NYSERNet have for making a positive economic impact on their regions--and on the nation as a whole.

Many readers have a primary interest in one or more of the segments covered in the current report, but not all of them. In order to facilitate the distribution of the report to those individuals, we have produced each volume under separate cover.

## VOLUME I: THE SUPERCOMPUTING COMMUNITY

### I INTRODUCTION

This volume of our paper deals with an assesment of the likely role of the Cornell National Supercomputer Facility (CNSF), first in the immediate, and then in the longer term future of NYSERNet. The supercomputer at Cornell has been a major traffic generator for NYSERNet and is likely to continue to be. But the future has been clouded by the recent Gramm-Rudman budget cutbacks in the federal budget for fiscal 1988. In investigating the issue of the impacts of these budget cutbacks, one is led very quickly also to issues of the intermediate and longer term that provide the context within which the more immediate questions must be answered. In part II we deal with the questions of the short-term impact of the budget cutbacks and related developments on users, on the CNSF, and on NYSERNet. In Part III we look at the significance of these and related developments in the intermediate term. In Part IV we move to the longer term issues.

The analysis of the intermediate and longer term issues strongly reinforces the most important conclusions reached from considering the short-term factors alone. That analysis also provides the basis for additional insights about the role of NYSERNet itself in relation to apparent, rational strategies of other major players such as IBM, AT&T, NSF and others in the arenas in which NYSERNet functions.

#### Background

The following background facts documented in our earlier studies are important for placing the task faced by an entity such as NYSERNet in perspective. NYSERNet was created as a two year experiment and has functioned in the context of a large number of other simultaneous innovations and experiments. The CNSF, like all the NSF supercomputing facilities, is a relatively new installation. It is optimized to perform multiple, medium scale supercomputing tasks, especially those that are data intensive, either in input or output, and thus require large amounts of disk space. NYSERNet is an experimental wide-area network (WAN), that provides access to the CNSF by NY State institutions including the major research universities, Brookhaven National Laboratory and the research-arms

of several important industrial firms. In this wide-area application, NYSERNet uses Proteon router technology that was originally designed for use in local area nets (LAN's). NYSERNet is a packet switching network that uses TCP/IP protocols. TCP/IP protocols themselves were intended to be experimental only, but in the absence of "something better", have been pressed into production use and have been widely adopted, especially by U.S. Government agencies and universities. These protocols are also the protocols of choice for users of UNIX operating systems, the operating system originally developed by Bell Laboratories and now widely adopted (in various versions) by a large number of different manufacturers of computer systems. Users with needs for large scale scientific computing facilities--the most frequent users of supercomputers--are very likely to be familiar with, and prefer to use UNIX operating systems on the desktop systems they use to access the supercomputer of their choice. To date IBM has not introduced UNIX for use with its 3090 computer systems. The CNSF computational facility, however, is based on the IBM 3090-600 system.

In sum, the users of CNSF prefer UNIX. The NYSERNet network for access to CNSF is based on TCP/IP protocols, the protocols of choice for UNIX users. The supercomputer production system and its associated development system function under the VM/CMS operating system of IBM, do not have UNIX implemented on them and have had VM/TCP/IP software available for users for less than a year.

From this litany one can see the potential for problems in implementation in the circumstances faced by these various organizations and by NYSERNet. Also documented in the prior studies, however, are the substantial successes that have been achieved to date by all of these experimental and early implementation entities in overcoming the difficulties they face. Further, in preparation of this study we learned that the CNSF personnel are especially up-beat. They see their facility to be poised for a quantum jump in service and performance. In support of their optimistic view is the fact that Cornell's share of the total packets transmitted over the NSFNET backbone is the largest of those for all six supercomputer centers on the net (this includes NCAR). It has ranged from 21-24% of that total over the last six months.

## II IMMEDIATE SHORT TERM ISSUES

We begin this section with a discussion of the cutbacks in the NSF budget which were passed-on, across-the-board to all parties supported by NSF. The three questions investigated are, what is the significance of the cutbacks for:

1. the budgets of users of CNSF?
2. the budgets of CNSF itself?
3. the budgets of NYSERNet?

All indications are that one answer applies to all three questions: the cutbacks will not be immediately disabling. Both current users and the CNSF appear to be able to absorb these cutbacks in the short run without being significantly impaired. The cutback will significantly interrupt growth and development, however. The direct effect of the NSF budget cutbacks on NYSERNet has been severe, but through two emergency measures it appears to have been contained. First, NYTEL/NYNEX have agreed to forgive significant payments for operating costs by NYSERNet as a short-run adjustment. Second, NYSERNet has received an emergency appropriation from New York State at least as a partial offset to the NSF cutback.

Although the short-run situation appears to be tolerable, it is clear that guaranteed, longer-run funding from the federal government is a necessity, both for the Supercomputer Initiative of NSF and for the networks required to support this Initiative.

### The Issue of the Short-term Survival of CNSF

A related question for investigation was, can we expect the CNSF to survive into the future? The answer emerges immediately. There is no question of the short-term survival of CNSF. The facility is too important to the major players that have a stake in its future for it to be significantly endangered. The major players are:

1. IBM
2. Cornell University
3. New York State

None of these entities could permit CNSF to fail.

The most obvious case is that of IBM. The CNSF is the worldwide showcase and capstone external facility for its approach to production supercomputing through use of the 3090 system. As discussed further below, IBM cannot permit this facility to fail; it serves too many corporate purposes.

Cornell University has featured the Cornell National Supercomputer Facility in much of its public relations and recruiting literature. The Supercomputer is truly integral to the research efforts of the Cornell faculty. It has become an integral part of the self-image of the University. The University cannot permit the CNSF to fail.

New York State has featured CNSF in much of its industrial recruiting and public relations literature. The facility has become an integral part of the infrastructure for economic development in the State. It has become an integral part of the educational infrastructure of the state. It has become an integral part of the international competitiveness strategy of the State. New York State cannot permit the CNSF to fail.

These judgements apply to the short-term. There are also factors of the intermediate term supporting these conclusions.

### III FACTORS IN THE INTERMEDIATE TERM

In this section we look at the intermediate term outlook for each of the three groups, the users of CNSF, the CNSF itself, and NYSENet. We investigate the effect on users of the pending cutbacks in their ability to use 800-dial-up access to CNSF. We follow this analysis with a look at the impacts on users of cutbacks by agencies other than NSF. Finally we look at the impact on users of developments of a general nature.

In the case of CNSF, we examine the other side of the 800-number access cutoff. Then we investigate the intermediate term significance of CNSF to its three major sponsors, Cornell, IBM and New York State. The discussion of the significance of CNSF to IBM opens-up a number of longer term issues. These are introduced in this section and examined in greater detail in the final section. Finally, the significance of CNSF to New York State is examined. In light of the enhancement of the NSF backbone network currently in

progress, a paradoxical implication emerges. This factor serves also as a transition to the discussion of the intermediate term impacts of the various factors upon NYSERNet.

NYSERNet plays a very important role for New York State institutions in relation to CNSF. As shown in our earlier studies and the preceding discussions, NYSERNet also has performed an important role in IBM's development of supercomputer capability.

#### CNSF's 800-Number Cut-Off

With NYSERNet becoming operational it became feasible to cut off "800-number" access to CNSF users within New York State. This took place on April 1, 1988 and, as of this writing, is scheduled for the rest of the country for July 1, 1988. NYSERNet does not serve the rest of the country. Other regions are served by regional nets that connect through the NSF backbone to the five NSF supercomputer centers including CNSF.

There appear to have been three impacts. First, there has been increased traffic on NYSERNet from New York State sites to CNSF. Second, there has already been some impact on SURANET users (users who represent the bulk of the remote users of CNSF). SURANET has a direct 56K connection to CNSF through the Maryland/Cornell link. This has permitted institutional users from SURANET-institutions to have direct access to CNSF without the necessity of going through other legs of the NSFNET backbone. This permits them to continue their access to CNSF and in addition, as compared to dial-up, it provides them with greater bandwidth. At first glance this would appear to be a total solution for SURANET users, but in reality it is not. Knowledgeable observers estimate that at least half of the total use from SURANET locations was taking place from the users' homes. (Though it has not been explicitly studied, this mode of work is believed to be followed generally by persons associated with major research and education institutions. We know it to be true for Bell Laboratories personnel and for personnel of the General Electric Corporate Research and Development Center.) Given that many of these users lack local access to their institutions from their homes, access will no longer be possible--except through a long-distance toll call.

The third result grows out of the reactions of these users to the pending 800-number cutoff. Our contacts at CNSF report that users

who anticipate losing 800-number access to the CNSF are putting pressure on their local institutions to provide funding for dial-up to CNSF from their homes. An argument used has been that since the local institution in the past provided neither the budget for the computer time at CNSF, nor the funding to provide dial-up access to CNSF; the least the local institution could do now would be to pay for telephone access. Some limited funds are also available from CNSF for users who are in particular need.

A positive effect on CNSF of its cut off of the 800-number access is the saving of approximately \$50,000 per month in phone charges. This saving is one step in the adaptation by CNSF, partially offsetting the NSF budget cutbacks.

#### Effects of General Budget Cutbacks on Users:

Budget cutbacks, in response to Gramm-Rudman, have been government-wide. As suggested above, the National Science Foundation dealt with these cutbacks through across-the-board cuts to all the people and institutions it supports. In contrast, the approach chosen by DARPA has been to cut out whole operations. For example, on the ARPANET, all nodes in the state of Texas have been discontinued. Similar actions have been or will be taken in other locations on the ARPANET where given nodes are known to have access to alternative networks.

In making its decisions, DARPA does not appear to have taken account of the likely impact of its actions on the various other networks. Problems of many sorts have arisen. For example, in the case of Texas University Net, the protocols used are DECNET. Thus, technically, universities in Texas are interconnected and do have alternate access to the INTERNET. Since their traffic is DECNET traffic however, the inter-connection is at best, complex. A second problem arises from the fact that Texas institutions are part of SESQUINET that connects to the current NSF backbone only at NCAR in Colorado. The shift of traffic from ARPANET to SESQUINET has resulted in a heavy overloading of the NSF backbone routers at NCAR. Similar impacts are occurring, or are expected to occur, in other locations where prior ARPANET service is discontinued. We see below that as currently envisaged, NSF's longer term priorities for support of networks are likely further to compound the problems of regional networks, and of virtually all users.

These problems are occurring at the time of changeover of management for the NSF backbone network from the current Cornell Theory Center management to management by the new MERIT-consortium (Merit-IBM-MCI). As an interim step during the transition, the contract to Cornell has been extended for three months through July 1, 1988. Since funding for the extension has been at only 75% of the prior level, however, a number of services have had to be curtailed for the existing backbone. For example, since April 1, it has not been possible to connect new nodes to the NSF backbone. Rather, new institutions must await new funding to be connected under the new management of the backbone.

As discussed earlier, one response by CNSF to the cutbacks in its budget has been to suspend 800-dial-up access to the facility. Some users whose local institutions do not currently have network access to NSFNet, have long distance toll dial-up directly to CNSF as their only option. The CNSF personnel are encouraging users who previously used the 800-dial-up to achieve local access to the INTERNET through their institutions. There are two problems with such a linkage. The first is the one just discussed, the possibility that the connection to the INTERNET is already congested. The second involves the typology of the local networks themselves. For example, it is often the case that faculty members' dial-ups to their own institution come in through particular nodes which are different from the nodes ordinarily used in their research. It may also be that a third node may be required for access to the INTERNET. In such instances, access to CNSF can become extremely cumbersome.

An additional problem has occurred in at least one instance. A faculty member at the University of Minnesota is able to access the University through dial-up, but achieves connection to the INTERNET through a computer system in a colleague's laboratory. That laboratory levies local charges on this faculty member for use of its inter-connection to the INTERNET. Under existing budgets there are no funds for the type of payment being required by this laboratory (or for any other data networking charges).

CNSF experts state that there are approximately 10 universities that are not connected to the INTERNET, but do have researchers currently using CNSF. These include the following five:

University of Connecticut  
University of Rhode Island  
University of Maine  
University of Nevada  
Tarleton State College, Texas

The only alternative for these researchers is long distance dial-up.

The number of dial-up lines maintained by CNSF has not been cutback. There are still 24 lines, half of which support 9.6K. (The other half are 2400 Baud). The question of dial-up access, then, resolves into, who pays for the phone call?

#### Related Matters:

Several organizations maintain leased-line connections to CNSF. These include:

University of Massachusetts  
SUNY Albany  
SUNY Buffalo  
Corning Glass Works.

The two SUNY campuses also have connection through NYSErNet. The NYSErNet connection uses the TCP/IP protocols. The reason for maintaining the leased-line connection despite the presence of NYSErNet is that it permits an IBM-to-IBM connection using the "PASSTHRU" software that is preferred by long-time IBM users, but is incompatible with TCP/IP.

CNSF personnel report that there has been little apparent cutback in the use of the facility following the cutoff of the 800-number access from New York State. It appears that virtually all the prior traffic has been shifted to NYSErNet. Here, however, the limited data collection and analysis tools prohibit greater depth of analysis. The information that is available relates directly to the use of the CNSF machines, not to the network itself. Use of CNSF is attributed to particular projects through the accounting software available for the CNSF machines. It is these data that show little, if any dropoff in use of New York State-based research projects during the month of April.

#### IBM, Supercomputing and CNSF in the Intermediate Term:

The conclusion that the CNSF is vital to IBM was driven home by a conference held in March, 1988 at Cornell entitled, "Super". This was a conference on the use of IBM 3090 systems with vector facilities for high speed scientific processing--i.e., supercomputing. The conference was jointly sponsored by IBM and the CNSF. Its participants were mainly from universities in the U.S., but also included 12 representatives from foreign countries. A main objective of the conference was for participants to share experiences in the use of vector-equipped 3090's in high performance scientific computing. Personnel from CNSF were clearly in the role of mentor, since many more resources and greater expertise have been assembled at CNSF as compared to the other installations

The CNSF is intended to be, and is a show case and demonstration installation. It also serves as an upward growth path and a safety valve for scientific computing by users of the 3090 at installations with fewer resources. At the conference users became acquainted with approaches to use of the 3090 by the CNSF that they otherwise might not have considered. They came to understand the types of scientific computing that can be achieved with increased resources. Users from other installations also learned that if the necessary resources are not available to them locally, they can apply to use the CNSF itself.

Users achieving remote access to the CNSF also automatically become a part of a virtual "networking research" project as they use networking resources--resources that are as yet in their development stages. All of these elements contribute directly to the commercial objectives of the IBM Corporation.

#### Background on IBM's Reentry Into Supercomputing:

The market for supercomputers is growing rapidly as these giant computer systems are being used over an ever-broadening range of applications. In recognition of these facts, IBM is moving to reestablish its role in this field.

The introduction of vector facilities with the 3090 systems and the extension of shared-memory-parallel processing across up to six processors in these systems, represent two steps in the direction of supercomputing. IBM is now advertising its 3090 systems as

"supercomputers", and it is taking a number of additional steps to establish itself as a significant participant in this market.

At least two major development efforts aimed at making supercomputers commercially available are taking place within IBM. At its Kingston Laboratories IBM has a sister installation to the Cornell National Supercomputer Facility. The Kingston Laboratories have IBM mainframe systems with vector facilities in a complex system with 21 Floating Point Systems (FPS) array processors. By comparison, the Cornell installation involves one IBM 3090-600 (expected soon to be two 3090-600's) plus seven FPS array processors. Paradoxically, the Cornell installation has greater IBM computing power than the IBM Kingston installation, while Kingston has greater FPS computing power than Cornell. In any case, these two production facilities, the internal one at Kingston and the external one at the CNSF, are providing IBM with substantial experience in interacting with production users of high performance scientific computers.

IBM's commitment to supercomputing is further illustrated by a second internal effort that involves the design of a supercomputer from the ground up. Not much is known about this effort. Recently IBM has also announced its support for the new venture by Steve Chen, formerly of Cray and the designer of Cray's Y-MP supercomputers.

#### Increasing Software Portability, an Additional Factor:

The potential reentry of IBM as a significant player in the supercomputer field is facilitated by a number of timely developments, the most important of which is that software is becoming increasingly "portable" from the supercomputer of one manufacturer to that of another. The four elements which are facilitating software portability are:

- (1) Standardized versions of FORTRAN.
- (2) Improved FORTRAN optimizing compilers.
- (3) Broad-spectrum libraries of algorithms optimized for each vendor's systems.

- (4) A common operating system--UNIX--now being implemented on most supercomputers.

As a recent entrant, IBM will benefit enormously from the increasingly easy access to the vast libraries of software that have been developed for the supercomputers of other manufacturers, especially Cray. These developments enhance IBM's likelihood of success in its efforts to reenter this segment of the computing market.

#### IBM's Interaction with Los Alamos National Laboratories:

A further indication of IBM's interest in supercomputing has emerged from a related quarter. A number of items have appeared in newspapers and trade journals about a cooperative venture between IBM and the Los Alamos National Laboratories (LANL). It is well known that managers in LANL's computation centers are urging IBM to become an active participant in the market for supercomputers. But IBM's UNIX operating system for the 3090 system, is not yet operational. This limits the applications that could be achieved by the 3090 as a production engine in the LANL networked-configuration which (as we see below) focuses on UNIX. Current levels of TCP/IP connectivity for the 3090 (VM/TCP/IP) are insufficient to permit the 3090 to function within the LANL supercomputer network. Nonetheless, rumors persist that a 3090 will soon be installed at LANL, at least as a test installation.

#### IBM, CNSF and NYSERNet: The Incentive Structure

If one cold-bloodedly examines the incentive structure of the current situation, one sees that for IBM it is important that the CNSF be successful; it is not necessary that the success result from NY State access to CNSF. Success for the CNSF, therefore, can not necessarily be equated with success for NYSERNet, though the reverse is more likely to be true. Institutions with a more direct interest in NY State must pay close attention to see that NY State is not disadvantaged by the presence and success of CNSF. That is, if through the augmented NSF backbone, research institutions outside NY State achieve highly effective access to CNSF, while NY State institutions do not, the presence at Cornell of CNSF would represent a net disadvantage to economic development and industrial competitiveness for New York!! It would enhance out-of-state organizations and firms and not those in the State (except Cornell itself). Note that NSF plans call for

priority of funding to the new NSF backbone; it is soon to be operational in its augmented form. It will have resources. As a major NY State industry, IBM would experience some pay-offs from development of NY State. These are minor, however, compared to those that accrue to the NY State Government, and to the utilities, especially the information utilities whose fortunes are enhanced by development, growth, and international competitiveness of firms and other organizations in the State. These entities must see to it that NY State institutions achieve at least comparable advantage to that achieved by out-of-state organizations.

NYSERNet is in being. It is one of the more effectively managed regional WANs. It is in the interests of NY State Government to support NYSERNet. By doing so, the State can assure that NY State is not disadvantaged by progress in, and development of NSFNET. NYSERNet's effective development is a goal greatly to be desired for the good of NY State and for the State's international competitiveness.

#### IV LONGER TERM MATTERS

In this final section, we look at some of the longer range issues related to NSF's funding priorities, to the development of supercomputers and to IBM's likely role in this development. Then we examine the significance of alternative approaches to networking supercomputers for networking in general. The discussion at the end of Part III touched on the fledgling interaction of IBM with LANL. Key elements in that interaction relate to the focus on UNIX-TCP/IP. These are symptoms of a broader debate currently raging in the field of supercomputing. These matters are discussed further at the end of this section.

#### Longer Term NSF Funding Priorities for Networks:

NSF's longer term strategy for funding networks is important to NYSERNet and to CNSF. NSF officials have publicly identified an hierarchical networking scheme and stated their intended priorities for supporting the scheme. The three levels of the hierarchy are:

1. The NSF backbone,
2. Regional wide-area nets (WANs),
3. Local-area nets (LANs) at node institutions.

Given current and projected funding levels, NSF will fund only the expansion and operation of the backbone. The implications of the NSF backbone development could go either way for NYSERNet. NSF will not fund either the creation or the operation of regional networks beyond what it has already done. NSF will not provide funding for local-area nets. Both of these policies are detrimental for NYSERNet, and under particular conditions, should be reversed.

NSF's position is that the regional WANs must be self-supporting through user fees. They envision two possible mechanisms to achieve this. First would be through a change in funding regulations to make charges for use of a regional WAN or WANs acceptable as direct costs in budgets for research projects supported by the federal government. Alternatively, the research institution could pay these costs for its users and include them as part of the indirect (overhead) charges applied to federal government sponsored research budgets in general.

Neither of these approaches is without major problems. First, the regional WANs themselves are in a start-up phase. Few activities in either the public or the private sector can be expected to, or are able to, cover their costs during the start up phase. Charging rates that might be reasonable in the long run won't cover costs in the short run. Charges that would cover short run costs are likely to be too high to attract users. Further, direct charges for network use sufficient to cover the full costs of the creation and implementation of a network, especially during a start-up phase, would violate the economic principles for efficient resource allocation. (This issue is discussed at length elsewhere). Second, existing research budgets are not permitted to include use of WAN's (other than the telephone) as direct cost items. What would sustain these regional WAN's during the transition period until the rule changes become effective? Third, in many instances with current network accounting mechanisms it is not even possible to attribute costs of WAN use to particular users, especially individual users within a research institution. The data collection and analysis software for use with current-day networks is rudimentary at best. It is difficult for network suppliers to analyze the usage of their systems, much less identify costs to individual users. For example, at CNSF, data collection software has permitted the development of an 18 megabyte database related to network use for the first quarter of 1988, but review of this data is just beginning and it's not automatic. It requires a very complex analysis.

The indirect cost route for covering the costs of regional networks is no better as a solution. Points one and three also apply to indirect cost recovery (the WAN's are still start-ups, and it is still not possible reasonably to measure usage in order to establish the charges). Even more important, research universities already are feeling pressure from research agencies and from OMB because of the dollar amounts involved in their existing cost-recoveries. Already many universities are foregoing recovery of large amounts to which they are entitled under existing rules. Significantly increasing the amount of allowable recoveries by changing the rules to permit WAN-charges to be recovered as indirect costs is unlikely to have a real effect. Universities are likely to feel the need to "eat" these costs along with those that they forego currently.

The argument put forth by NSF officials in support of their priorities is that direct funding by NSF for regional WANs would result in a clearly identified target for "budget cutters" within the Congress. That is, if there existed a substantial, explicit line item in the federal budget for the support of regional networks, it is likely that one Congressional committee or another would seek to "save" money by cutting that item. The fallacy in this argument becomes apparent in light of two facts. First, whether the costs of use of wide-area networks were included in research budgets as direct costs or as indirect costs, they eventually must be paid from (and thus they add to), the NSF budget itself. Second, Congress this year cut the NSF budget by approximately 20%! To make the paradox explicit, the NSF approach is to try to avoid creating a new line item which might be cut; instead they intend to include these new charges in their own budget--which was just cut for fiscal 1988.

What is required instead of the proposed NSF approach is the mobilization of a sufficient constituency in support of the creation and operation of research-oriented networks such that at least minimal funding needs are met and cutbacks become impossible. Anything less is a mere palliative.

The question of funding for local-area networks has similar dimensions. To carry on research efficiently under today's high-technology conditions requires that internal networking in universities also be of high quality and reliability. Yet, approaches to the internal use of networks in universities are in their infancy. Much experimentation needs to be done. Much needs to be learned

before such networks can be efficiently implemented and used. In other words, much research needs to be carried on in relation to the internal use of research networks by universities. This is an appropriate mission for NSF. The "spillover" effects from good research done to establish efficient approaches to network schemes within universities can be enormous. These would accrue to the institutions implementing the results of that research since such benefits though very real, are difficult if not impossible, to capture by the institution doing the initial research. This is a case in which government support and funding for research is fully appropriate to assist in getting the research done. Without it the research is unlikely to be done, or if done, unlikely to be done well. The appropriate agency of the federal government to provide such support is NSF.

#### Problems for Users Caused by Hierarchical Networks:

As mentioned above, NSF views networking for education and research as a three-layer operation. The existence of three uncoordinated layers can prove to be a significant problem for users. A network which functions well is a network which is transparent to the user. When a network does not function well, the user must be able to find a way to overcome problems. An uncoordinated, hierarchical networking scheme presents numerous impediments to easy solutions for users' problems. There is no single authority to whom the user can appeal. There is no single authority integrating the access sought by the user. This situation is likely to be one of the major impediments to the creation of networks that are effectively transparent in the fields of education and research.

When problems arise in the three tiered network, then, the point of contact for a user must be the support group at the local institution. Unfortunately, that support group faces the same problems as the user. Merely to identify the source of a problem, the support group often must trace back through all three levels of the hierarchy. On numerous occasions, the problem arises from interaction between levels in the hierarchy. There is often great difficulty in determining the cause of the problem--and more in finding persons with both the competence and authority to solve it. While the support people accept the need to understand each of the layers in the network and interactions among them, end users have no interest in either--nor should they. Nonetheless, given the degree of training of some support personnel, often they must.

The various regional networks are of different sizes and different levels of sophistication. For example, MIDNET, a small, simple, regional network is controlled from the University of Nebraska and is reported to be very effectively managed by a reasonable, individual as manager. Things function quite well on MIDNET. Other networks are much larger and more complex and thus face problems of much greater magnitude. An example is SURANET which is huge by comparison to MIDNET. It connects 32 nodes and is expected to add 15 more in the immediate future. Only one full-time person is assigned to the management of this network. Additional part-time people from various institutions help out, but only on an "as available" basis. Merely maintaining connectivity to 32 nodes can be a full time job. A staff of one cannot be expected to provide optimal levels of training and maintenance trouble-shooting. It can't be done. An additional problem is also exemplified by SURANET. It has but a single connection to NSFNET, that at the University of Maryland (which also serves as the connection to NYSERNet). If this node goes down, the entire network is disconnected; if it becomes congested, all nodes are effected. As we have seen, the latter has already occurred in the case of SEQUINET at its NSF-node at NCAR.

Similar problems exist with the local area networks at the individual institutions. There is a great disparity in the level of expertise available at these institutions. There are great differences in the complexity of the LANs at these institutions.

From the user's point of view, all three levels of the hierarchy must function, function together and well in order for network access to be "normal". Problems arise at each of the levels of the hierarchy. No coordinating institution exists. Users currently are faced with enormous difficulty in solving problems of access through the network to the resources they seek.

#### Supercomputer Networking and the Future of NYSERNet

Another longer-run factor important to the future of NYSERNet is the increasing prominence of UNIX as the operating system of choice to unify diverse computing environments that include desktop systems and workstations all the way up to supercomputers. UNIX uses TCP/IP telecommunications protocols. This augers well for NYSERNet as a TCP/IP packet-switching network.

Los Alamos National Laboratories (LANL) is the world's largest user of supercomputers. The Lawrence Livermore Laboratory (LLL) is a renowned development site for specialized supercomputer software, especially that used in networking. Both these national laboratories expect that in the immediate future, the user will have on the desktop, computing power equivalent to today's supercomputers.

These two national laboratories represent a truly unique resource for this country. Their experience with supercomputers is an area of unquestioned U.S. leadership. The experience which currently resides in U.S. National Laboratories--that of networking multiple computing systems of supercomputer power--must soon become commonplace if tomorrow's desktop systems are to be used effectively. The experience of these national laboratories must be harnessed and exploited as a building block for U.S. international competitive advantage. What they achieve at a single site must soon be replicated for effective interaction between and among multiple sites.

The outcome of the competition in approach by these two key national laboratories is likely to determine the outcome of supercomputing networking. Given the rapid pace of development, it is likely also to have a strong influence on the direction of networking in general--and thus the future of NYSERNet--over the next decade.

### Two Contrasting Strategies In Supercomputer Software Development

Lawrence Livermore Laboratories and the Los Alamos National Laboratories have quite different strategies for providing network support to their supercomputers. Stated in the broadest terms, Livermore attempts to provide customized networks for their systems while Los Alamos seeks to make maximum use of commercially available networking products and systems from workstations through to supercomputers. To this observer a hybrid approach making use of the best features of each could prove to be a desirable strategy.

LLL has consistently sought to develop its own systems software. Its supercomputer network currently uses CTSS, the Cray Time Sharing System, which LLL developed for networking

supercomputers at its installation, whether or not these were manufactured by Cray. It is currently working on its next generation of systems software for networks, NTSS, its "New" Time Sharing System.

The developers of NTSS advertise that its great advantage comes from its being designed from the ground up as a networking system for supercomputers and thus avoids many compromises that must be made when tying together preexisting elements of systems software for that purpose. The fundamental approach of NTSS is to make possible both a hardware and a software "bus-structure" with standardized interfaces to which the respective hardware and software elements of any product can be plugged-in. Once this dual bus structure were implemented, then additional systems could be added, or full systems or particular products could be removed from the network--without disruption to the operation of those that had not been changed. Each manufacturer would only have to insure that its products could plug-in and function at the appropriate hardware and software interfaces.

Backers of the LLL approach identify two major problems with the use of the UNIX operating system for networking supercomputers. First, at its initial design stages, UNIX was never intended to be used with supercomputers. UNIX is an operating system which deals with "bit-streams", streams of individual bits; while supercomputers use words of 60- and 64-bit in parallel. Second, UNIX was initially designed for stand-alone systems, and not for networks. These observers view the extension of UNIX to networked supercomputers as, at best, a design afterthought.

A question which immediately arises is why would anyone want to use UNIX with supercomputers? The answer leads us directly to the strategy of LANL. The computing at both LANL and LLL is high performance scientific computing. Users of supercomputers at both laboratories prefer to work from workstations networked to the main systems. Most such users are also experienced with and knowledgeable about minicomputers, especially VAX minicomputers. The vast bulk of these users use UNIX on their workstations and many also use UNIX on their VAX and other minicomputer systems. If the UNIX environment can be extended through to the supercomputer itself, then a single unified operating environment would become available to the end user for all computing.

The LANL proponents of UNIX argue that multiple advantages accrue from such a strategy. First are the obvious implications for the end user in having to learn only one major operating system. Second, given that multiple software vendors have developed ingenious, innovative software applications for workstations and for minicomputers, the enormous power of these independent development efforts would become available in the supercomputer environment as well. One of the major, current draw-backs to supercomputer use is the limited degree to which modern software applications have been implemented on supercomputers, since supercomputers have to-date represented such a restricted market. Through UNIX, the software developments from the other stages in the job process would be made available to the supercomputer field. At this juncture, it appears that IBM is casting its lot with the LANL approach to the networking of supercomputers.

On the other hand, under current institutional arrangements, when organizations such as DOD and NASA require vendors who wish to bid for government contracts for supercomputers to supply UNIX-based operating systems, the manufacturer of the supercomputer must develop the UNIX implementation for its systems. The resources that the manufacturer can bring to bear for this are likely to be no greater than the resources available to LLL. In any case, the expertise present in the software teams at LLL is generally recognized to be among the most sophisticated in the world.

LLL has a well recognized need to be able to upgrade their supercomputers as increasing computing power becomes available. Their concepts and approaches to systems software for networking supercomputers are quite well accepted. Yet, there is a great disadvantage to use of LLL-developed software: the only development/support teams available for users of such software come from LLL itself. It is not their mission to provide software support to a broad group of unaffiliated users.

#### A Recommended Partnership For Supercomputer Networking

These facts together suggest the desirability of a coordinated partnership effort. In essence, the objective of such a partnership would be to build on the expertise of the LLL systems software developers, but to achieve vendor support for the installation and maintenance of a hybrid system. This could provide the best of the

two worlds: the development of an integrated networking scheme by this highly regarded development team, plus support from the expertise, resources and user-outreach of a commercial vendor. This latter role would appear, at least on the surface, to be a major attraction for the developer of UNIX, AT&T.

Given the recent establishment of partnerships between AT&T and computer manufacturers for development of standardized versions of UNIX, it does not appear to be too great an extension of logic for AT&T and LLL to work together to extend UNIX into networks for supercomputers. Yet, this latter step appears not to be seriously under consideration. If it were done, the enormous expertise of LLL could be incorporated as an extension of UNIX and made available to a broad base of users of supercomputers, thus capturing the best of the two strategic approaches. If this can not be worked out with AT&T perhaps it could prove attractive to the newly formed consortium of other vendors (also see below) to develop their "standard" version of UNIX. Of course, ideally both groups would be well served also by including one or both of the US supercomputing manufacturers in their groups.

Part of the difficulty for anyone in achieving partnership with LLL, is LLL's long history of independent action. Seldom has LLL sought outside expertise when attempting to deal with their internal computer/networking problems. Paradoxically, LLL is managed by the University of California at Berkeley while the most widely used version of UNIX is "Berkeley UNIX", from the computer science department at the University. Since the two entities are managed within the same institution, U.C. Berkeley one would imagine that some coordination between the two might somehow become possible. On the other hand, what would really be implied by a more effective partnership would be a four way collaboration: U.C. Berkeley, LLL, and AT&T, plus a supercomputer manufacturer.

Assuming that the two entities within the University framework are unable to achieve a joint development effort with AT&T for a supercomputer version of UNIX, the LLL developers, nonetheless, could attend the UNIX "developers meetings" which occur periodically around the country. By pointing out possible shortcomings of proposed vendor-developed versions of UNIX for supercomputers and the advantages of the "twin bus" approach advocated by LLL for providing systems software to network supercomputers, the LLL people might be able to influence the

direction of the development of UNIX from supercomputer vendors. Given the well recognized, enormous expertise of the LLL group, there must be "something" that they are doing right. Given also the current focus of the federal government on technology transfer from the National Laboratories to the industrial world--in order to promote the international competitiveness of American industry--it would appear to be timely and highly desirable for such activity to take place. Something of the sort must come about in order for true network transparency to be achieved.

### IMPLICATIONS OF "STANDARDIZED" UNIX AND TCP/IP FOR NYSERNET

NYSERNet's initial decision to focus on the TCP/IP protocols is consistent with the main thrust in networking for research and education institutions. TCP/IP is an outgrowth of an experimental protocol suite established by the Department of Defense as a temporary, experimental expedient in the absence of "something better". TCP/IP has been widely adopted by government networks. It is here to stay, at least as a stepping stone to the widely discussed and advocated OSI, the Open Systems Interconnect protocol suite. Fortuitously, AT&T adopted TCP/IP as its preferred telecommunications protocol in its UNIX operating system. As we have seen, UNIX has been widely adopted by manufacturers of computer systems of many diverse types as a step toward providing their users with a standardized, vendor-independent operating system. This has greatly broadened the base of TCP/IP users.

The agreements AT&T established recently with computer systems companies to develop a standard version of UNIX are important. The first such agreement was with SUN Microsystems, the manufacturer of high performance workstations. SUN and AT&T are developing a standard version of UNIX for use on workstations. The second is with UNISYS, the second largest mainframe manufacturer. UNISYS and AT&T are developing a standard version of UNIX for use on mainframes. To date AT&T has not developed a similar relationship with a manufacturer of supercomputers or with a National Laboratory. Above, however, we make the recommendation for establishing a four-way relationship among AT&T, U.C. Berkeley, the Lawrence Livermore Laboratories and a manufacturer of supercomputers.

The efforts by AT&T have spawned a counter effort by a consortium of other firms led by IBM, DEC, APOLLO, HP, etc. to

develop their own "standard" version of UNIX. This consortium has created a significant budget for the joint effort. We note again, however, that this group does not involve a manufacturer of supercomputers or a National Laboratory. Given the rapid evolution of workstation and other desktop stations and the expectation that soon desktop systems will have the power of today's supercomputers, this appears to be a significant oversight. Given the importance of networking to the future of computing, and the focus of US National Laboratories on the networking of supercomputers, the absence of a National Laboratory from both competing groups appears to be a major oversight.

In sum it is the intimate relationship between UNIX and the TCP/IP protocols provides the link between NYSERNet and AT&T's strategy. UNIX has been the operating system of choice for the research community. TCP/IP has been the protocol of choice with UNIX. It is because of the commitment of R&D personnel to UNIX that the pressures exist to create the unified, vendor-independent operating environment extending from workstations to supercomputers. All of these efforts reinforce the decision by NYSERNet to focus on the TCP/IP protocols.

The absence of supercomputer manufacturers and the National Laboratories from the UNIX standardization effort is a problem that will effect the speed and direction and development of network computing. This implies that it is in the interests of NYSERNet to take what actions it can to overcome this deficiency in current standardization efforts.

To the extent that this deficiency is not overcome, a major opportunity will be missed. The US has as a great national resource that can be used to create comparative advantage--the experience and expertise of its National Laboratories. US international competitiveness needs demand that optimal use be made of these resources. NYSERNet can help to catalyze this outcome.

## SUPERCOMPUTER INTERVIEW SERIES

Date of Interview: 4/11/88

Company: CNSF/IBM

Persons Interviewed: Henderson Cole  
IBM Site Representative  
CNSF  
Cornell University  
Ithaca, New York

Elizabeth Schmerhorn 4/12/88  
Assistant Director of Outreach  
Cornell Theory Center, CNSF  
Cornell University  
Ithaca, New York

Additional interviews were conducted with:

Scott Brim 4/17/88  
Manager of Networking  
Cornell Theory Center  
Cornell University  
Ithaca, New York

Martine Haldran 4/17/88  
Director, Networking, CNSF  
Cornell Theory Center  
Cornell University  
Ithaca, New York

Steven Wolf 3/28&30/88  
Director of Networking  
Scientific Computing  
National Science Foundation (NSF)  
Washington, D.C.

Interviewed by: Alan Mc Adams

I discussed the conference "Super" with Henderson Cole of IBM. The conference took place March 28-30, 1988. It involved persons with IBM 3090 mainframes, generally with the vector facility. The participants included twelve people from foreign countries. This group now is planning for the second meeting to take place at the University of Kentucky next year.

Major Conclusions:

1. The CNSF serves as a capstone installation for all of these users. It has greater capability than any of the other sites. As an NSF sponsored facility it has developed great education and training capabilities. It has developed a great deal of teaching and training materials that can be made immediately available to others.
2. It is anticipated that the CNSF will serve as a "relief valve" for the other sites when their user needs exceed local capability. As such, it represents a growth path for scientific computing for all of the other sites.

3. The degree of vectorization achieved at CNSF is approximately 25% across all six processors. In accord with the algorithm for equating these matters, this implies that the FORTRAN code is 60% vectorizable at CNSF. This degree of vectorization implies (on the assumption of a three-fold speed-up through vectorization) the equivalent processing power to that of an additional six processor 3090. Cornell's degree of vectorization compares with approximately 10% or less of vector processing at the non-Cornell sites.
4. The CNSF 3090 system is considered to be "balanced". It supports a huge real memory, huge virtual memory, and an even more huge disk "farm". Users are able to maintain their programs and data resident indefinitely on disk. This represents a substantial advantage over most Cray installations.
5. Telephone dial-up still appears to be the preferred access mechanism for remote users. Henderson Cole describes the "weakest element" in access, to be the management of the networks.
6. The supercomputer is Cornell's "F" machine. The "C" machine is an IBM 3090-200 with two vector facilities. The main use of the "C" machine is for administrative purposes. The degree of vectorization on this machine is less than that achieved by 3090 users at other non-Cornell sites. Henderson Cole has offered the hypothesis that a significant portion of administrative computing is in fact vectorizable. He has suggested that experimentation be done to establish this point. There is no better place for this than at Cornell.

#### General Observations:

Henderson suggested that we speak with Geraldine McDonald at SUNY Binghamton. She was one of the presenters at "Super". Her talk focused on the myriad networks in New York state that someone in her position must pay attention to. We will follow up.

He discussed at some length the steps that must be taken to vectorize code. This involves the use of the vectorization compiler as an initial step. However, this is generally far from adequate. The next step usually involves identification of "hot spots", those areas of code that require the greatest number of machine cycles. When such "hot spots" are shown not to be vectorized, further investigation is desirable. Diagnostic tools are available to identify the factors that block vectorization for such code. Often the problems are relatively routine and easily overcome through minor changes to the code. Where this is possible, it should be done. Often it is possible to improve the degree of vectorization further by scanning for problems that are more difficult to resolve. Of course, there are codes not susceptible to vectorization and here nothing can be done.

Users, as they become more sophisticated, seek interactive diagnostic tools through which they can improve the degree of vectorization of their code. These tools are being tested and will become available more broadly in the near future.

Interview with Elizabeth Schermerhorn, 4/12/88:

Betsy pronounced the "Super" conference a great success. The major benefit from the conference was to get people to talk to each other and share experiences. Many people were quite surprised at the degree to which vectorization could be and is achieved on the 3090 at CNSF.

Given that the function of CNSF is "facilitating science", the demonstration that substantial degrees of vectorization could take place was a substantial achievement.

Cornell's support resources greatly exceed those available to other sites. A main focus of the efforts at CNSF have been on training and outreach. The "smart node" approach has proven to be extremely successful. The Cornell team effectively trains the trainers who then return to their individual sites to serve as "smart nodes". This has proven to be very efficient from the points of view of all parties. The Cornell personnel prepackage training and education materials deliver these to the "smart nodes" who in turn use them to train persons locally. When a new user at a "smart node" site contacts the CNSF, they are immediately referred back to their local "smart node" for introduction to the system. This permits Cornell consultants and trainers to focus on problems and questions whose solution requires a higher degree of sophistication.

The CNSF will provide training to any site that wishes to establish a "smart node". All costs for transportation, local expenses, and materials are covered by the CNSF. Each "smart node" is provided with an allocation of free computer time plus several pre-established small accounts which then permit new users to become acquainted with the system. Both the "smart nodes" and the users are also provided by the Cornell staff with a series of on-line teaching aids.

Major Conclusions

1. It is Betsy's perception that the CNSF is poised for a major upward move. A great deal of spade work has been done on which such an upward move can be based. The current activities are now stabilized and are running quite smoothly. The new elements that appear to be just on the horizon are the IBM version of UNIX, a jump to increase parallel processing (with the installation of the second IBM 3090-600), and the introduction and increasing support for work stations.
2. Betsy sees the CNSF at the center of the users group for IBM 3090s at universities. It has developed expertise in managing large numbers of scientific users. A significant portion of this expertise is transferrable to users at other sites. Foreign users are currently far behind users in North America as a result of what is available at CNSF. Europeans are anxiously seeking affiliation with the CNSF. A great deal of the technical support, on-line and through other means, can be made available to these foreign users.
3. Users are "getting stuff done". Many of them have "gotten hooked" on supercomputing, because it works. Betsy finds it exciting to visit with users today because they are excited about what they are able to do.

### General Observations

Betsy was very upbeat. In prior discussions the emphasis appeared to be on the limited resources available in relation to the enormous demands. Perhaps because of the even-more-limited resources available at other sites, the emphasis in our discussion of "Super" was on the advantages the CNSF has over "normal sites" which are not able to focus on supercomputing per se. Those sites must make decisions on how to allocate their 3090 mainframe among administration, teaching-support, and scientific computing. By contrast, the CNSF is able to focus the bulk of its resources on scientific computing. Further, those resources greatly exceed what is available elsewhere. This puts the CNSF in a very strategic position.

The conference offered an opportunity for the CNSF to acquaint other users both with what is possible and what is available. It is anticipated that as a result of the conference, very substantial improvements in scientific processing, at sites other than CNSF, will take place.

### Additional Interviews

The interviews with Scott Brim and Martine Haldran were incorporated directly into the body of the report. They are referenced as "CNSF contact persons". These experts provided the detailed, factual materials about networking that are incorporated in the report at various places.

Dr. Steven Wolf was a speaker at the meeting of the Supercomputer Subcommittee of the Committee on Communications and Information Policy of the IEEE, on March 28, 1988. He was also a speaker at the conference, "NationalNet '88", at the L'Enfant Plaza Hotel in Washington, D.C. on March 30, 1988. Materials in the text of the paper are based on his remarks at these meetings plus private interviews which took place before and after his public presentations.

## SUPERCOMPUTER INTERVIEW SERIES

Date of Interview: 2/19/88

Organization: LAWRENCE LIVERMORE NATIONAL LABORATORIES  
LOS ALAMOS NATIONAL LABORATORY

Persons Interviewed: John Rannelletti, Director  
Livermore Computer Center  
Lawrence Livermore National Laboratory  
Livermore, CA.

Dieter Fuss  
Livermore Computer Center  
Lawrence Livermore National Laboratory  
Livermore, CA.

(Presentation and meeting of the Supercomputer Subcommittee of IEEE)

Date of Presentation: 3/25/88

Norman Morse  
Director of Computing  
Computing and Communication Division  
Los Alamos Laboratory  
Los Alamos, CA.

Joseph Thompson  
Assistant Director  
Computing and Communication Division  
Los Alamos Laboratory  
Los Alamos, CA.

Interviewed by: Alan McAdams

Los Alamos National Laboratory (LANL) and the Lawrence Livermore Laboratories (LLL) are "rival laboratories". Each of them is managed by the University of California at Berkeley, but they operate very independently. Los Alamos has the largest number of supercomputers in a single installation anywhere in the world. LLL is a close second. Los Alamos now has nine Cray supercomputers with two very powerful eight processor Y-MP's on order for delivery at two or three month intervals later in this year. LANL intends to introduce the Y-MP's directly into their networking scheme.

#### THE LIVERMORE APPROACH

At Los Alamos I discussed the comparison between the Los Alamos and the Lawrence Livermore Laboratories. Their philosophies of computing are quite different. Lawrence Livermore is developing a distributed network system, set up in such a way that one only needs to know two interfaces--a

hardware interface, and a software interface, each on a "bus"--in order to put a new product into the network. All other elements of the network will continue to function undisturbed. For instance, when one puts a new product, such as a new supercomputer, into the network, one merely creates a table of its capabilities and requirements and plugs the system into each interface. The network operating system will use this table to treat the new supercomputer as a resource, just as it treated prior supercomputers or other devices as resources on the network. The hardware interface has been implemented already. The software interface, however, is more difficult to achieve. What Livermore must do is to completely rewrite the operating system of each type of supercomputer they have on their network because those operating systems were designed for stand alone (though front-ended) supercomputers, not for systems as part of a network. Furthermore, they were never intended to be interfaced directly into the general kind of sophisticated network that Livermore has designed.

The Livermore approach is very attractive because it makes possible rapid changeover from one kind of system to another. If widely adopted, such a network could greatly simplify the job of creating operating systems for designers of new systems: the network itself would have tremendous intelligence and thus would be able to perform many of the functions which are currently being built into the complex operating system for each type of new system. Therefore, it appears as if one should be able to accomplish much more rapid and significant changes on the hardware side of things and just substitute them into the network. The most difficult task is to provide the software to make the new approach operational.

If a system of radically different design, e.g., a largely parallel supercomputer, were plugged in to the network it becomes rather difficult to see how an existing network operating system could take care of that completely different architecture. There would have to be a mechanism for interfacing the parallel system into the network.

All observers with whom we talked believe that networks are going to be much more important in the immediate future to permit access to different types of systems. One cannot put every type of computing capability into one system. Therefore, the user must have access to different types of systems in order to do different kinds of problems. In order to have access to different systems, one needs excellent networking capability. The user seeks relative ease of access to remote resources, just as if they were on the desktop. If Livermore has a better conceptual approach to the use of networks for access to such diverse resources, then they may have moved out a decade ahead of others and may well have attacked problems that other people are going to face in the future. If they have solutions to those problems and those solutions are effective, then people should be adopting them.

Today's supercomputers are about a thousand times more powerful than the first supercomputer, the Cray I. Everyone expects that in the near future a person can have a desktop system that costs in the range of \$5,000 yet has the power of a Cray I. Since the Cray cost about \$10 million new, these desktops would represent approximately a 1,000 fold decrease in cost. These relationships seem to be rather consistent.

If one has a network of desktop computers all with the power of Cray I's, and one has National Laboratories that are the most experienced users of Cray and other supercomputers, then the experience of the laboratories becomes directly relevant to the activities of the people using desktop computers. With the equivalent of what were supercomputers on the desktops of every individual, users will require both very powerful networks and assistance from people who have real experience in using powerful networks. It seems as if the National Laboratories are a proper test bed for the new era of computing.

Livermore has conceived what it considers to be the ultimate networking scheme. They do all of the development work themselves. They have rewritten the operating systems, and created the network and the protocols that go into the network. They have set the network up as a messaging system, in what they consider to be an integrated, slick design such that all elements fit together. They do have all of the modern elements in the system: supercomputers, mainframes, banks of minicomputers, and large numbers of high performance work stations. The hierarchy of systems functions in the following way: a number of work stations interconnect into a bus, and series of workstation busses then interact into busses that tie together minicomputers. Then, those cascade up into the main hardware bus which permits the connection of multiple supercomputers. So, at least in concept, it is possible to sit down at a work station and have six different windows open and processing simultaneously one on each of six supercomputers. Livermore has approximately 3,000 potential users with this kind of capability.

#### THE LOS ALAMOS APPROACH

At Los Alamos they have an entirely different approach. They try not to write anything that they do not have to. They try to use commercially available software and utilize their highly skilled computer and communications people to tie together products that others produce. They do not want to create nor maintain new software themselves. They want to expand the effectiveness of what others do.

There is one area especially where they must do things in a much more sophisticated way than others outside the deep national security arena, however. They need to have assured security on their network since they are a major weapons laboratory. They must provide total security for network computer resources that are decentralized for interactive use by their scientists and engineers. When Livermore builds its own products, it inserts security measures directly into their system as an integrated package. However, Los Alamos must provide security on top of services and products of others with lesser security needs. In essence, it must retrofit security on systems that were developed with lesser security in mind.

At Los Alamos they have nine Crays of various generations (some 1a's, 1s's, XMP's, Y-MPs). When they bring in the Y-MPs, they plan to retire some of the older systems. In addition to the supercomputers, they also have 151 VAX computers. As one can see, they have a huge array of VAX's, all with DECNET. Digital Equipment Corporation always advertises that, "DEC has it now", meaning that they have a uniform architecture from top to bottom of their

product line, as well as the ability to interconnect all of these systems in their product line with a single uniform set of protocols, called DECNET. All of that is true. The problem, however, is that the range within which the DEC systems function is narrow relatively. They have no intention of providing supercomputer capability, and apparently perceive no comparative advantage in extending their expertise into the supercomputer area. To date, DEC has not been a significant player with desktop systems, and is not a lead player in the work station market even though they have microVax's which operate on the desktop and with work stations. DEC is mainly in the minicomputer area.

What Los Alamos plans to do is to provide a "super set" of DECNET to provide added security. (Besides all of the security that they provide, they will also use what DEC has provided with DECNET.) An irony of the current world is that it appears that DEC's protocols are closer to meeting the international standards (OSI) than anyone else at the moment. DEC has committed itself to modifying DECNET so that as OSI becomes better clarified, by the actions of International Standards Organization, DEC's protocols will coalesce to those of OSI. That is exactly what Los Alamos is also doing. They are standardizing on standards provided by others, and moving with DEC toward OSI simultaneously. Since the Department of Defense has also developed another set of protocols called TCP/IP (which itself is evolving toward OSI), they are doing all their internal networking beyond the DEC subset through TCP/IP. When TCP/IP and DECNET coalesce with OSI, then the two will become merged.

The workers at Los Alamos are not adopting what Livermore is advocating because of the following reasons. First, Livermore has not completed what they are advertising. Second, the things they have completed are stated to be "inefficient" in that they require twice as much overhead, such that they must use a lot of computing power just to run the network system. Third, its approach is unique. It does not build on what exists otherwise in the world.

Certainly Livermore is investing tremendous resources to develop its new approach. The Los Alamos people appear to agree that, in concept, this approach makes sense. However, one of their strong objections is that the only support source for Livermore products is Livermore itself. Since Livermore's resources are limited, the workers at Los Alamos are not certain of how soon their new approach will become available. Furthermore, they also pointed to the inefficiencies in its current design.

(Additional interview notes are being edited and will be incorporated into the published volume. These interviews served as a basis for the observations in the overview section and volume.)

## SUPERCOMPUTER INTERVIEW SERIES

Date of Interview: 4/25-26/88

Company: MICROELECTRONICS CENTER OF NORTH  
CAROLINA (MCNC)Persons Interviewed: Alan Blatecky,  
Director of University Relations, MCNC  
Research Triangle Park  
North CarolinaPat Richardson  
Communications System Data Network Supervisor,  
MCNC  
Research Triangle Park  
North CarolinaPeggy Montgomery  
Video Operations Supervisor, MCNC  
Research Triangle Park  
North CarolinaSamuel Averitt  
Manager, Communications and Hardware  
Computing Center, NCSU  
Raleigh, North CarolinaDr. Jim M. Fikry  
VBEE Administration  
Industrial Extension Service  
School of Engineering, NCSU  
Raleigh, North Carolina

Interviewed by: Alan Mc Adams

There are two important networks at the Microelectronics Center of North Carolina (MCNC). The first is the high bandwidth of two-way, full motion, video network. The second is the 45 megabit high-speed data network. Each of these networks is intended to connect six major educational institutions in North Carolina. The six are :

University of North Carolina at Chapel Hill  
University of North Carolina at Charlotte  
Duke University  
North Carolina Agricultural and Technical College  
North Carolina State University  
Research Triangle Institute

The networks do perform their intended function of providing a sufficient critical mass for establishing significant microelectronics design capability in North Carolina.

On the educational side, the television network permits the offering of a broad range of graduate courses in microelectronics. The objective eventually is to integrate the course offerings of the various institutions to permit graduate students the richest possible educational experience. To date, the offerings have been approximately 12 per semester. However, with the exception of the originating institution, the numbers of students who make use of courses through the network is quite limited. For example, at Duke, the usual enrollment is 1-3 students. The modal enrollment at any remote site is 3.

There are a number of advantages to all of these institutions from the educational network interaction. Probably the most significant is that it clearly differentiates the universities in this region from peer institutions. The presence of courses from multiple institutions of higher learning is a powerful selling point for prospective students and faculty. Second, in situations such as occurred in the semester just completed, it is often possible for a given institution to substitute for a course offered on the network for a local offering that was or is difficult to staff. This occurred in the case at Duke in that the computer science department at Duke accepted a course on Construction of Compilers offered by Chapel Hill to fulfill the requirement for such a course at Duke. Since the course was advertised to graduate students at Duke as fulfilling this requirement in the department of computer science, 12 graduate students took the course for credit through the network. Since there was a large number of students taking the course at Duke, the department assigned a course T.A. for the Duke students. The T.A. worked intensively with the students at the local site to the great satisfaction of all concerned. This is an example of what can be achieved through the two-way television educational process.

From the point of view of the computer science department at Duke, the compiler course was a great success both as an educational experience for the students and as a resource conservation step for the department. The faculty member who ordinarily would have offered the course locally at Duke was on sabbatic leave and did not need to be replaced.

This instance illustrated the factors which optimized the effectiveness of the real time television educational opportunity. 1) The receiving school had a real need for the course through the network. 2) The transmitting school had highly qualified well accepted instructional capability for the course being offered. 3) The students were provided by the receiving institution with the appropriate "signal" that the course was of a quality appropriate to their educational experience, i.e., Duke indicated that the course would be accepted as meeting distributional requirements for its computer science department. This instance contrasts with what to date has been the usual situation. Usually students are made only generally aware of the availability of courses on the network by the institutions that could receive

them. Courses are generally not included on the menu from which courses meeting distributional requirements must be selected. For both students and faculty, the network course is looked upon as a possibly interesting elective. Given that the total graduate student enrollment in computer science at Duke is 60, it is not surprising that the number of students selecting such an elective is small.

The above contrast in response rate by students to the network offering provides a powerful lesson for those who seek to optimize the contribution of the network to the educational program in microelectronics. The three key elements required to achieve this are now clearly understood. When they are present, the contribution of the network can be optimized. When they are absent, small numbers of students at receiving institutions should come as no surprise.

### The Quality of the Educational Experience of Courses Through the Network

A course offered throughout the network is different than a course offered through a standard classroom situation. It appears, however, that the educational experience can be quite comparable. "Something" is lost with the loss of direct human contact. There is a significant difference in the way in which written material is presented during the classroom portion of the course. Nonetheless, the educational experience is comparable.

What are these differences? First, there is no blackboard. It has proven virtually impossible to transmit information to remote site through standard approaches as use of a blackboard. For both the local and the remote sites, an overhead camera is used to replace the blackboard. That is, the instructor will present written material on 8.5 by 11 inch paper, photographed from directly overhead and displayed on the TV monitors. Note that this occurs both at the local and the remote site. What students get at both locations is identical. This requires some adjustment by both the instructor and the students. Once again, the medium is different but the results are comparable. What is lost is the record on the blackboard which is continually available to the students (until it is erased). What is gained is the ability to present written material of much greater sophistication and precision. One instructor commented that the freedom to select complex diagrams and equations from multiple text sources and to have them immediately available through the overhead camera proved to be a great educational boon. A second major advantage identified by this instructor was that material presented live by creating notes on 8.5 by 11 paper constituted a permanent record of what had been transmitted. (It is not erased at the end of the period--or before) This instructor is so pleased with his permanent record that he tends to make it into a formal textbook over the summer for one course and into a technical note for the other.

There are additional advantages/disadvantages from the differences between the network course and the course in the usual classroom. From end-of-course questionnaires, the reaction of a number of students to the network

course is that it is much better organized and much more formally presented. The instructors appear to be better prepared. Students report a downside to all this however. They state that "things sometimes go to fast" and that instructors appear to jam a significantly greater amount of material into a one semester course.

Students at remote sites did not appear to be inhibited from asking questions during a classroom session any more than were those at the originating site. Since the network provides two-way video communication and requires that the receiving site "sponsor" the course locally, including providing local camera operators. The questioner at the receiving site can be seen by both the instructor and the students at the sending site. Figure 1 shows a diagram of a typical TV classroom with the instructor at the front of the room, flanked by two 42 inch projection monitors. Three small monitors are available to the instructor on the podium; one to show what is going out over the network, one to show what is available through the overhead camera, and one to show students at the remote site(s). Any of the screens can be split four ways. Also facing the instructor is the main front view camera plus a large monitor, generally showing students at remote sites. From our discussions with a number of instructors, we learned that each has "tricks" for involving the students at remote sites. Often the remote sites will be polled for questions, with polling taking place alphabetically. As in normal classrooms, it is also possible for all sites to have persons speaking simultaneously, and this sometimes happens.

From the point of view of the instructor there are also additional advantages to teaching on the network. First is the feeling that one is "doing one's part", by originating a course from one's own institution in exchange for those that are received. ( There is no financial incentive provided.) Second, this is a new, challenging and exciting experience. Third, this may be the only way that a critical mass of students can be assembled to make it economic to offer the given course. Fourth, on another dimension, the instructor is provided with greater flexibility. If a student has missed a class, the instructor is not called upon to provide "private instruction"; the student can be directed to a video tape. Similarly, if the course had previously been offered on the network and the instructor finds himself/herself unable to meet a scheduled class, it may be possible to substitute a video tape from a prior period offerings. ( In fact this was done by a young instructor who had assisted one of the giants of computer science, Fred Brooks, in offering a course the prior year. The young instructor was able to substitute a video tape of a Fred Brooks class for one of his own and Professor Brooks agreed to answer questions in person for the fifteen minutes at the end of the session.) Fifth, it may be possible for colleagues to provide assistance even in real time. For example, a young instructor found himself in difficulty with a complex concept during a network transmitted class. A colleague who happened to have his video monitor tuned to the network, observed the problem, descended to the stairs, joined the class, and resolved the problem to the great satisfaction of all. In fact, a similar incident occurred at a later time in the same course. (We were fortunate enough to have heard about these incidents both from the instructor in question and from others who were aware of them. All parties spoke of these instances in very positive terms.)

There are also disadvantages. First, it is a great deal of work. There is no doubt that offering a course through the network requires significantly greater preparation. Second, there is some risk since this is a new medium. Even experienced instructors are not assured of success in presenting a course this way. Third, the excitement wears off relatively quickly. Fourth, you really never know who actually is watching the network and where; there is the potential for "big brother" to be watching. Fifth, the problems of coordinating the delivery of materials to multiple sites in a timely fashion are significant. (MCNC is assisting the local sites financially to permit them to have a network educational coordinator to see to the myriad of problems and details that arise.)

### Video Entrepreneurship at NC State University

While the other sites on the network deal generally with only the real time broadcast of course, NC State has an extension program that is highly dependent on video tape distribution. The School of Engineering has its own "industrial extension" program that deals largely with video tape. This group attempts to achieve multiple goals through its television activities. It participates in four different programs:

- 1) MCNC
- 2) The National Technical University (NTU)
- 3) Instructional Television Fixed Service (ITFS)
- 4) Video Based Courses

The above discussion applies generally to MCNC type course whether originated at State or at another site. NTU offers graduate education in engineering nationwide, generally, to businesses on their premises. This is done through satellite transmission. ITFS instruction is essentially a subset of the NTU activity. It is generally limited to small, medium size companies in North Carolina and thus is not directly competitive with NTU offerings which are generally made under contract to national organizations. State also offers credit courses and short courses to organizations of all sizes, directly through its own efforts.

In developing new courses for credit, State tries hard to make these offerings available through more than one of the above channels. For example, it will select approximately 24 courses to be offered for credit through video tape, 12 in each of 2 semesters. From these offerings, it will identify 5 or 6 each semester to be offered to NTU to be used on the national satellite network. In turn, NTU selects from the offerings of its various participating institutions, those which it decides to broadcast. For the courses it selects, NTU will provide an honorarium of approximately \$800 to the instructor and then beyond a certain minimum, will provide an additional fee per student enrolled. In addition, NTU makes similar payments directly to State, also with additional fees per student beyond a given minimum.

State will not offer its short courses to NTU for a very good reason. These short courses are a gold mine. They support the the for credit offerings by generating revenues in excess of their costs. The short courses are developed in the following way. First, a topic of interest is identified. Then a

person competent to offer such a short course is also identified (not necessary from NC State itself). Once the instructor has agreed to offer the course, then State will provide financial and technical support for its development. The usual initial stipend for development of a short course is \$2000. The instructor must develop a complete study guide for the course and make that available to the producer of the TV presentation. Jointly, they will select those elements which are most appropriate for the video portion of the course, leaving others for inclusion only in the study guide. The production team is also prepared to provide "field shots" as appropriate. If this is a first time video experience for the instructor, the production staff.....The instructor is required also to prepared and present a 2-3 minute promotional "plug" for the course. This serves several purposes, not the least of which is to permit the instructor to see himself/herself on tape. Usually this approach is all that is required to acquaint a first time video presenter with enough information to overcome any minor problems of presentation. The instructor is compensated further by 35% of the revenue net of production costs that are achieved from the short course. The more successful of the short courses generate up to \$10,000 to the instructor. NC State retains the other 65% and recoups its full production fees. The video section generates approximately \$600,000 per year to NC State. These funds are generally plowed back into new production efforts. The value of the raw tape in inventory is also approximately \$600,000 without consideration of any value added.

NC State uses professionals in the preparation of its videos. Now there is a role for students but only under the supervision of full time professionals. The staff that supports these operations numbers 10, 4 of whom are exempt personnel. With the exception of approximately half of the total personnel and the physical space provided, the operations are self-funding. This includes the highly sophisticated production equipment that is required for the operation. The physical facilities involve two video classrooms, plus a video conference room. Each of the classrooms is configured generally as described above. One classroom seats approximately 40, the other approximately 60. There are occasions when each is filled to capacity. However, given the small, average number of students who generally enroll in networked courses at receiving sites, State has found that the small conference room is greatly to be preferred for use by these students. There are a number of plans for additional small video classrooms to accommodate up to 12 or 15 persons. Funding is available for a number of these rooms which are referred to as "satellite" facilities.

#### Live Broadcast vs. Videotape

Depending upon the site of the discussions and the persons involved, we heard different perceptions of the relative merits of live network courses as compared to course work provided through videotape. After attempting to sort out the conflicting appraisals, we offer the following hypothesis. Each is more effective in particular circumstances. The case for videotape is perhaps more easily made. For the types of students who are the clientele of the engineering industrial extension service, videotape is greatly to be preferred. This clientele is made up of full time employees who have professional responsibilities competing with their educational efforts. For such people, to arrange to be present at a fixed time and place to participate in a live broadcast of a classroom activity can be extremely confining and may often be

impossible. Such clients greatly prefer the flexibility provided by the videotape. For these persons the substance of the materials on the tape is what is important and not the presentation. They will find a time to view the tape and absorb its message.

On the other hand, full time students have as their highest priority the pursuit of their course work. They are accustomed to being at a particular place at a particular time for their courses. They also have become acclimated to interaction with instructors and their classmates. This latter point is very important. Often this can be the most rewarding portion of the classroom experience--and its extension to the preparation and follow-up to class. For these people the videotape can be perceived as barren while the ability interact in real time with an instructor and with classmates provides a significant added dimension of stimulation. For such people, there is no substitute for the live classroom experience.

### The Data Networks

The data networks tie together the major research campuses identified above. The development of this network has been gradual leading to the 45 megabit pipe size among the main campuses. The timing of the development is as follows:

January, 1985	12 megabit
June, 1986	25 megabit
March, 1988	45 megabit

As the bandwidth was increased, equipment required for the lesser level of bandwidth was shifted outward in the network. The current 45 megabit pipe size provides multiple T-1 lines. To date, only a modest portion of the available capacity is in use, and some T-1 lines are used exclusively for development work.

MCNC has been serving in a catalytic role. It has had an enormous advantage for this function in that it has had significant funding available for it to permit it to provide resources directly to the universities. This has been a mixed blessing. The universities actual resent the fact that MCNC has resources available in a manner that is totally outside of their control. On the other hand, MCNC has found itself powerless to "require" the universities to take any given action. It has consistently functioned through persuasion and example. The MCNC network is a star network. The center has been able to provide the capability in-being to each of the universities only to find the universities were not in a position to interconnect. One university questioned the need for the network at all. The situation has changed very rapidly, however. The university just mentioned now has the largest number of LAN's in place as Ethernet.

The objective of the data network is to permit joint research to go forward among these universities. This is only beginning to happen. However, each of the universities has excellent communication now with MCNC itself. Faculty members of each university commute to and from MCNC and are now able to function effectively from either end of the network connection with their home institution.

We visited three campuses and were able to discuss their networking activities in some detail at each campus. Our impressions necessarily were fragmented and new information provided new perspectives throughout the visit.

Perhaps the most impressive single installation is that of the computer science/computation center building at UNC Chapel Hill. Norman Vogel of the computer science department was our host. He is a truly dynamic and impressive person. We joined a discussion by Dr. Vogel with a group from the SRC, the semiconductor research corporation. Vogel was providing a detailed demonstration of the communications scheme for the building. Copies of his transparencies are included as Appendix Y. In a nutshell, the building is "wired" to a higher degree than virtually any other facility with which we are familiar.

Vogel is attempting to make use of Ethernets to a degree and for purposes beyond which they were designed. As a result, his installation is experiencing some unusual and transient difficulties. The resolution of these difficulties is expected to advance the state-of-the-art.

From later discussion with technical persons elsewhere in the system, we learned that the computer science building is "an oasis" on the North Carolina campus. So far it appears that there is little coordination of communications function by persons in various authority positions.

This situation contrasts rather strongly with that which we discovered at NC State University. We had an extended discussion with Sam Averitt who is manager of communications and hardware at the Computing Center. Averitt is a no-nonsense guy with an excellent reputation among his peers. From his position he has been able to bring about some order on the campus of NC State. Virtually the entire campus is now wired through broad-band cable. This was installed only after an experimental facility had been established and tested within the Computing Center. Most of the current needs of the campus, including the production origination, reception and up-link of televised courses can be met through the broad-band. Averitt is currently designing an initial experimental fiber-net again to be tested over a significant period prior to any final decision on the introduction of fiber elsewhere on the campus.

(The state of North Carolina has made available a site adjacent to the current campus with acreage greater than that of the current campus for the development of the next stages of NC State.)

3071-88 vol.2

REPORT III

MARKET AND ECONOMIC  
IMPACT STUDY

VOLUME II: THE BIOMEDICAL COMMUNITY

**NYSERNet**

NIC COLLECTION



New York State Education  
and Research Network, Inc.

1095 Avenue of the Americas  
Room 1702  
New York, NY 10036  
Tel. 212-395-4480

REPORT III

MARKET AND ECONOMIC  
IMPACT STUDY

VOLUME II: THE BIOMEDICAL COMMUNITY

Prepared By:

William L. Dougan  
Cornell University

For

MEIS TEAM

### REPORT III: GENERAL INTRODUCTION:

During the period of the last twelve months, the members of the NYSERNet Market and Economic Impact Study (MEIS) Team produced two prior reports that identified many of the needs and problems of individuals using data telecommunications networks for education and research at universities, a paper on the relationship between, "The Economic Benefits from and Public Support of a National Education and Research Network", plus the current report. In the prior reports, the team identified specific groups of individuals who are involved in education and research activities and whose needs required additional study. The several volumes of the current report are devoted to four of those groups:

- (1) Volume I, The Supercomputing Community
- (2) Volume II, The Biomedical Community,
- (3) Volume III, Education: K-12
- (4) Volume IV, Industrial Clients.

Generally speaking, these volumes review the needs of a specific segment of the education and research community, point out current problems associated with providing network service for that segment and outline recommendations for action that can be taken by an organization such as NYSERNet to serve that group most effectively. Taken together, these volumes present a broad spectrum view of the opportunities wide-area networking organizations such as NYSERNet have for making a positive economic impact on their regions--and on the nation as a whole.

Many readers have a primary interest in one or more of the segments covered in the current report, but not all of them. In order to facilitate the distribution of the report to those individuals, we have produced each volume under separate cover.

## INTRODUCTION

There are those who have argued that in several hundred years the history of science will label the period of time corresponding to today as the "Age of Biology." Recent advances in medical imaging, genetics, biochemistry, microbiology, pharmacology, virology, surgery and the practice of medicine in general have led to tremendous advances in our expectations about our survival and quality of life. Very new fields like genetic engineering and magnetic resonance imaging promise to improve things even more in the near future. The situation is not without its problems, however. The explosion in information, the increasing strictness of malpractice laws and the advances in medical technology threaten to destroy the productivity of the average researcher by forcing him/her to spend large quantities of time just staying current with scientific developments. Data telecommunications networks (and the computerized resources to which they give access) show promise in reducing some of these problems.

The success of networks will depend heavily on how they are designed. The success of the designs will depend heavily on maintaining close contact with the groups of individuals who will use the networks to get things done. Current problems are not caused so much by the unavailability of technology to meet demands. They are caused by the inability of individuals to organize themselves to make coordinated, organized use of this technology. In order to achieve this, practicing physicians, medical students, students in the allied health fields, biomedical researchers, faculty members in medical schools and representatives of non-profit and government agencies need to be consulted constantly in order to determine how networks can be improved.

This paper is an attempt to relate some of what is happening in the medical and biomedical fields. It was written to inform the providers of network service in New York State just what medical and biomedical students and researchers want and need. There are sections outlining (1)the history of computing and networking in medicine, (2)the characteristics of the populations of individuals doing medical research, (3)technological trends in medicine that affect the demand for network services, (4)needs for network services expressed by members of the biomedical community, (5)problems with providing this service and (6)recommendations for

services that will serve the needs of the medical and biomedical population.

## HISTORY

### MEDICAL EDUCATION AND RESEARCH INSTITUTIONS

The history of networking in medical education and research institutions generally parallels that of post-secondary institutions. In the middle and late 1960's, the little computing that was done within the biomedical research community was done on centralized mainframes, usually those belonging to the university with whom the medical school was associated. As the prices of computing machines dropped and as minicomputers were introduced as an inexpensive alternative to mainframes, medical institutions began to acquire their own computing facilities. Typically, the introduction of computing resources into clinical institutions was led by hospital administrators. These individuals wanted the speed and efficiency that computers could provide for working with patient accounts.

The differences in needs among groups of users often led to the acquisition of completely different (i.e. - incompatible) types of facilities in different segments of the biomedical community. Different philosophies of computing and consequently different equipment was purchased: (1)for universities, (2)for the medical schools associated with those universities and (3)for the teaching hospitals that were associated with the medical schools. Often, the three divisions of the teaching hospitals themselves--research computing, clinical information systems, and administrative information systems operated with independent and incompatible systems. These separate systems were often isolated from facilities outside the hospitals. To varying degrees, this isolation still exists in many medical institutions today.

With the introduction of the personal computer, larger numbers of medical personnel became familiar with computing. Many individuals took the leap into the electronic age because they could afford their own computers. Others (particularly medical students) were willing to try computing solutions for biomedical projects because they had become familiar with computing earlier in their lives. More individual research projects could also afford their own computational facilities, and more individual physicians began

to acquire computing machines for their private practices. This created a higher level of awareness of and demand for computing on the part of the medical population. It also pushed the development of intra-facility networking (which grew along with the demand for interconnection of the individual machines). Packer (1987), citing a survey of 2500 hospitals produced by Shared Data Resources, a consulting firm, stated that the use of personal computers in hospitals increased 39.7% during 1986. The same survey revealed that the 2500 hospitals added 5379 new positions for DP professionals during 1986.

Packer (1987) also states that during the early 1980's, most hospital managers who purchased PC's did so independently (that is, without consultation with central hospital management). During 1984-86, these managers began trying to network the menagerie of different PC systems in order to gain access to shared data resources. Today, purchasing of PC's used for hospital management is often centralized in order to make networking easier. Though there is currently much interest in networking PC's to centralized networking resources, the facilities to do this are still being developed in the bulk of hospitals. There was little development of inter-facility networking until the early to middle 1980's. Newald(1987) states that a survey of 80 hospitals conducted by the National Research Corporation, a consulting firm, indicated that 23% offer computerized linkages with physician offices. Another 32% say they plan to offer such linkages in the future. Newald also cites a study by Sullivan and Frost, another independent consulting firm. She states that the percentage of physicians automating their offices will go from 35% to 80% between 1986 and 1990. According to Newald, spending on information systems by physicians is expected to go from \$546 million in 1986 to \$1 billion in 1990.

In the early to middle 1980's, two developments led to the consolidation of university-affiliated medical institutions into campus networking schemes. First, due to the deregulation of the telecommunications market, university administrators began to consider the economics of providing their own local telephone service. Because data services could also be provided using the same technology, university administration considered the strategic importance of networking facilities for research and teaching. Some universities began to provide for connections between main campuses, associated medical schools, and teaching hospitals. These connections generally took place at locations where the medical

school and teaching hospitals were located in close proximity to each other and to the main campus. (In cases where medical schools are located at some distance from teaching hospitals and from university facilities, none of the organizations are adequately linked.)

The second reason why medical schools were connected to main campuses was that connections between different types of systems became possible. A variety of new products that permit the interconnection of networked computing systems from different vendors was introduced by third party vendors. These products are still being introduced at an increasing rate and with an increasing amount of functionality. They are instrumental in providing for at least basic connections between the facilities on campuses and those in teaching hospitals.

## LIBRARIES

Historically, many of the advances in wide-area network services at medical institutions took place in medical libraries. Consider for example, the National Library of Medicine. In 1960 the National Library of Medicine (NLM) transferred some of its *Index Medicus* to a mainframe computer. (The *Index Medicus* is a cross-referenced index of articles from a number of medical journals.) In 1964, all references in the *Index Medicus* were available in machine readable form, and the Medical Location and Retrieval System (MEDLARS), a computerized bibliographic search service, was created. From as early as the 1970's, librarians have had access to MEDLARS through a well-known and widely-used network service developed at NLM. This service is called MEDLINE, and it has grown into the most widely-used data search service in medicine today. The Long-Range plan of the NLM predicted in 1987 that users had performed more than three million searches on NLM computers in 1986. When the service began, it relied on direct-dial telephone access and 300/1200 baud modems. Currently, it relies heavily on the value-added carriers like GTE-Telenet and TYMNET. It is also available through a number of commercial vendors of network services. The available transmission rates have risen to 9600 baud, providing the user has equipment to maintain such a connection.

In 1981, a program called Paper Chase became available for use with NLM resources. Paper Chase simplified the process of using NLM resources by providing an easy-to-use interface for

unsophisticated or unschooled users. This trend continued with the distribution of Grateful Med (a program with similar aims) for home computers in 1984. Today, libraries are moving more slowly. A JAMA (1987) article recently attributed the relative slowness of advances in networking libraries to the lack of funding available, the fact that individuals use medical libraries less than other libraries and the restriction created by a concern for confidentiality on the part of other divisions of medical institutions.

### DIVISIONS OF THE USER POPULATION

Users of networks and network services in the biomedical community can be divided in four different ways. The first set of categories separates different types of organizations: (1) medical schools, (2) teaching hospitals, (3) other post-doctoral teaching organizations (e.g. blood banks, government agencies, clinics and service organizations, etc.) (4) universities and (5) research institutes. The second set of categories separates administration, clinical operations and teaching and research activities in each of the previously mentioned organizations. The third set of categories distinguishes between professional staff and non-professional staff in each of the previous categories. The final set of categories distinguishes between types of individuals involved in medical education and research activities: faculty, research staff, clinical staff and students. Together, these divisions create about forty potential user groups (ignoring any field-specific differences between individuals). Current research-oriented network services tend to be concentrated on serving faculty members at medical schools, research staff at universities and research institutes, and to a lesser extent, administrators at universities, hospitals and medical schools. Commercial network services tend to be focused on serving practicing physicians in their private practices. This leaves a great number of individuals unaccounted for under current networking schemes. A discussion of the approximate size of some of these divisions is made below.

### MEDICAL SCHOOLS

In August of 1987, the Journal of the American Medical Association (JAMA) listed 126 accredited schools in the United States offering the M.D. and 16 in Canada. The State of New York was the

home of 12 of these schools, which means that nearly 10% of accredited medical U.S. medical schools fall within the geographic area served by NYSERNet. A listing of New York medical schools is given in Appendix A, Table I. Of these 12 medical schools, 11 are located near current NYSERNet nodes in Albany, Buffalo, Syracuse, Brooklyn, Stony Brook or Rochester. Because of this proximity, it would be reasonably easy to extend NYSERNet's coverage to include some or all of these organizations.

In addition, of the 16 accredited medical schools in Canada, 9 are located in regions which border on New York State. A list of these institutions is given in Appendix A, Table II. Some of these institutions are also close to existing NYSERNet nodes and hence might be good prospects for connection at some later time.

According to AMA (Jolly, Tasket and Beran, 1978) the 126 U.S. medical schools spent \$10.281 billion in the 1985-86 academic year. Of this expenditure \$2.55 billion was spent for research-related activities and \$531 million was spent for academic support including audio-visual services, libraries and computer centers. If we assume that all schools spend in equal amounts, it would mean that 1.08 billion was spent by New York State medical schools; that \$255 million was spent on research and \$53 million was spent on academic support by those institutions during the 1985-86 academic year. Given that a number of the medical schools located in New York State are both large and prestigious, the figures in each of these categories are likely to be larger than the amounts given.

#### POST-GRADUATE TEACHING ORGANIZATIONS

In a review of graduate medical education in the U.S., Crowley, Etzel and Shaw (1987) state that there were 6,332 accredited clinical post-graduate programs in the United States in January of 1987. Of these, greater than 1/5 (23%) were located in the mid-Atlantic region. Mid-Atlantic programs enrolled approximately 18,000 residents, which accounts for 23.5% of all post-doctoral medical residents. The State of New York had the largest percentage of programs of any state (13.7%) and trained the largest number of residents (14.4%).

Crowley, Etzel and Shaw identified 1,570 institutions and agencies as "offering a significant portion of graduate training in

1986-87." These institutions included 1,343 hospitals (with a median size of 378 beds) and 227 non-bed institutions, such as offices of state medical examiners, blood banks, ambulatory clinics, mental-health agencies and pathology laboratories. The percentage of each type of organization is given below:

Local government or hospital district authority	11%
State controlled	15%
Federal institutions	10%
Church-related institutions	14%
Controlled by non-profit organizations or agencies	49%
Miscellaneous	2%

These institutions frequently offer more than one program of post-doctoral education. In this study, the number of programs per institution varied from 1 to 41. Two-hundred of the sponsoring institutions sponsor 4,358, or 69%, of the programs. These 200 institutions are responsible for 70% of the residents on duty. Ninety percent of the programs have a formal relationship with an accredited medical school.

## FACULTY

In 1986-87, there were approximately 64,000 full-time faculty at medical schools in the United States (Crowley, Etzel and Peterson, 1987). Approximately 14,500 were in the basic sciences (e.g. - microbiology, biochemistry, etc.) and approximately 49,500 were in clinical fields (e.g. - pediatrics, orthopedics, obstetrics/gynecology, etc.). There also were 8,230 part-time faculty and, most interestingly, 122,149 volunteer faculty at the medical schools in the U.S., for a total of approximately 194,000 faculty. Assuming that New York accounts for 10% of each of these groups, we can project that there are approximately 6,400 full-time, 800 part-time and 12,000 volunteer faculty in New York medical schools alone. It is most likely that nearly all faculty have some clinical responsibility and that the full-time and part-time faculty have research responsibilities. Often, faculty have joint appointments in teaching hospitals and in the research institutes or universities with which medical schools are affiliated. From this information, we can estimate that the number of faculty members involved in medical research-related activities in New York is about 7,200 persons.

## STUDENTS

There are three main types of students enrolled in medical school programs in the United States. First, there are medical students who have not yet earned the M.D.. Second, there are post-doctoral medical residents and interns involved in graduate clinical and non-clinical programs (most often in teaching hospitals). Finally, there are allied-health related students enrolled in a variety of fields in the allied-health professions. The relative numbers of each of these groups in medical-school related programs for the academic year 1986-87 are given in the table below:

---

Table 1. Individuals Enrolled in Medical-School Related Programs during the 1986-87 Academic Year

---

Medical students (all schools)	66142
Graduate medical residents	76815
Dental students	9026
Pharmacy students	4803
Nursing students	12160
Physicians assistants	1863
Undergraduate arts and sciences	13662
Post-doctoral students not in basic sciences	16469
Other health discipline students	1883
Graduate and post-doctoral in basic sciences	<u>21043</u>
Total	<u><u>223866</u></u>

---

Source: Crowley, Etzel and Peterson, JAMA, Vol. 258, August 8, 1987.

---

## CONTINUING MEDICAL EDUCATION (CME)

According to a study commissioned by AMA (Osteen and Gannon, 1987), during the 1986-87 academic year, the faculties of 123 medical schools offered 17,057 continuing education courses. These courses represented a total of 259,623 hours of instruction.

There were 1,031,271 registrants. Felch (1987) argues that there are 10,000 people involved in bringing CME to the physician population, and that CME activities incur \$10 billion in costs each year.. He states that there are more than 6000 hospitals in the United States, and that approximately 50% are of sufficient size to carry out CME programs.

Twenty-seven states and territories currently have legislative requirements for continuing medical education. In these states, physicians are required to obtain a certain number of CME credits in order to re-register to practice medicine. New York currently does not have such a requirement. In addition to legislated requirements, six national medical societies require continuing education credits in order to maintain membership in the society. A few medical specialty boards also require CME credit for recertification to practice their particular medical specialty. Finally, a number of state medical societies require CME credits for membership. The Society in New York has such a requirement.

There are a number of ways in which CME credits can be obtained. Live programs (similar to conferences) in which a number of physicians attend lectures at a conference center, are currently extremely popular. More recently, broadcast video technology has been used to make CME more accessible to a wider audience. One program, sponsored by the American Medical Association, called AMA Video Clinics, is currently being carried over the Lifetime Cable TV Network and the Hospital Satellite Network. Currently there are 42 such clinics in the AMA library. Eight new clinics were produced during 1987.

#### TRENDS IN MEDICAL TECHNOLOGY THAT WILL AFFECT THE DEMAND FOR NETWORKS

There are a number of scientific and technical developments in medicine that will have an effect on the amount and type of networking services demanded by the medical professions. They fall into two different categories: (1) developments which are indirectly related to existing network efforts and (2) developments which are directly related to existing networking efforts. In this section, we outline a number of these developments. We then assess how each will affect the networking requirements of the medical community. For the segments of the medical community involved in education and research, the effect of the developments on demand will

generally be accelerated and the demand for networking to accommodate such developments will occur early relative to that in the segments of the medical community concerned strictly with clinical care.

## INDIRECT DEVELOPMENTS

### Medical Imaging

The numerous developments in medical imaging are the technical innovations which have probably the greatest potential to increase the demand for network bandwidth among the members of the medical community. These include Computed Axial Tomography (CAT), Magnetic Resonance Imaging (MRI), Positron Emission Tomography (PET), and Ultrasound imaging. All of these innovations produce a fine-grained, 2-dimensional matrix of numerical values where each grid point represents a value of some characteristic of biological tissue. The matrices vary in size, with the smallest standard image being 512 x 512 x 8 bits and the largest images having a size up to 2048 x 2048 X 24 bits. Often, a 3-dimensional image is built up out of multiple 2-dimensional images by scanning adjoining "slices" of tissue. These larger images are composed of up to 63 single slices. This places the size of a standard data set at between .25 and 75 Megabytes (Mbyte), with most of the images falling between .25 and 1 Mbyte. These datasets are by no means the largest possible, however. One of our interviewees generated data sets of 1024 x 1024 x 1024 x 8 bits in conjunction with a research project investigating the diagnostic properties of ultrasound. This places the size of the largest dataset we have encountered at approximately 1 billion bytes (1 Gigabyte). Sehr (1988) states that the pediatric radiology department at UCLA Medical School is involved in a project to store images on compact disk. The department produces approximately 209 Mbyte of data per day.

The need to store, manipulate and transfer these images will lead to a tremendous demand for network capacity in the future on the part of both researchers and clinicians. Already, proprietary systems for storing and managing these systems are beginning to be acquired by radiology departments. These picture archiving and communications systems (PACS) are greatly facilitating the process of transmitting images to distant locations. The demand for this type of service will increase as the facilities to produce these images are

acquired by more institutions in the medical community. With the exception of CAT, which is more widely disseminated into the clinical community, these technologies are still concentrated in the research community.

Currently, the amount of information involved and the amount of scarcity of facilities combine to create substantial demands for network bandwidth--ones that often exceed capacity. Research programs often add to the need for connection between medical and clinical facilities. For example, one researcher we interviewed was taking CAT scan data from the Hospital for Special Surgery in New York City and transporting it to Cornell University in Ithaca, NY. The data was fed into a computer program on the Cornell National Supercomputer Facility and that program produced a set of specifications for input to a computer-aided manufacturing (CAM) system. The specifications were sent back to New York City to a CAM machine, where they were further processed and then used to run a computer-controlled machining device. The product of this machining device is a prosthetic implant for use in replacement of one of a patient's joints. The need for networks that can provide accurate and reliable transport, storage and management of information to accomodate more extensive versions of operations like this one is high.

The professor of engineering which is in charge of this operation is extremely interested in gaining network access to Memorial Sloan Kettering Cancer Center in New York City for a similar project. This project uses an almost identical set of procedures, and would make advances in the radiotherapy of tumors. For this project, a cycle of network use would procede as follows: (1)CAT information would be sent from New York to Ithaca, (2)this information would be processed on computers at Cornell and then (3)parameters for computer-numerical control of a radiation treatment machine would be sent back to New York. The datasets for this project would be approximately equal to those discussed previously for the prosthesis project. Currently, there is no available means to transmit the information into and out of the New York City location.

## Laserdisk

A second important area of technical development which will affect the demand for networked services in the medical community is laserdisk technology. A variety of information resources are now being distributed on laserdisk. The products offered via laserdisk exploit its tremendous storage capacity, its inherent potential for fast search and retrieval and its ability to simultaneously link sound, full-motion video, complex graphics and computer programs. Laserdisk products include such options as: (1) traditional physician's desktop reference volumes such as the Physician's Desk Reference, (2) medical textbooks and anatomical reference resources which include such things as nested hierarchies of graphic images displaying both gross anatomy and tissues, (3) expert systems programs which assist physicians in identifying drug interactions, generating diagnoses for sets of symptoms, and outlining pertinent information about various therapeutic protocols and (4) extensive interactive simulations of patient management problems used for computer-aided instruction (CAI) of medical students and continuing medical education (CME) of practicing physicians. The NLM is encouraging the distribution of some of its MEDLINE and other files by disk.

Laserdisk technology has a potential to reduce substantially the amount of traffic carried over networks by providing a low-cost alternative for distribution of large amounts of heavily-used information. The providers of network services can at least partially reduce this negative impact of laserdisk by encouraging the development of products which make cooperative use of both the information resident on laserdisk and the information available via networks. One way to do this is to provide links between the commonly-used information sources that are distributed via laserdisk and information sources that are too expensive, large or rapidly changing to be distributed via disk. An interesting option would be a CAI program which searches information on a disk resident on a user's personal workstation and which automatically connects to and uses a remote network resource when the answer is not found on the videodisk.

Two recent suggestions made by members of Compuserve's Medical Special Interest Group involved just such a use of cooperative remote/local databases to do CAI. In one case, it was suggested that a remote database of full-text journal articles be used as a backup for a more simple help program that would reside on a

local disk. In the other case, it was suggested that blocks of code and information be exchanged with a remote center during a CAI application that simulated a patient management simulation. This scheme would presumably make patient simulation more realistic and challenging.

### Automation of laboratory technology

A third technical trend which will affect the demand for network services is the automation of laboratory facilities. Currently, many research facilities and some clinical facilities have machines that will automatically produce analyses of such things as blood chemistry, tissue chemistry and blood components. Often, the results of these analyses are machine-readable and can be automatically entered into databases within a clinical information system. While automated facilities are most often networked within a particular location (such as in large commercial laboratories and large hospital laboratories) there is a potential for the interconnection of many sites for the purposes of data collection, medical consultation and documentation of pharmacological clinical trials. Many commercial laboratories which serve unusual or highly-specific needs are developing in the medical community because of the economies of scale and expertise that they can provide. The inclusion of biomedical institutions across the state in a networked research project would have a number of similar advantages. It would increase the incidence of recording rare or unusual cases, it would increase the speed with which data could be gathered, it would provide for economies of scale in information storage and it would reduce duplication of effort across different research institutions.

### Computerization of Clinical Information Systems (CIS)

The automation of laboratory facilities and the ease of search for and retrieval of information due to machine-readable storage of results has led to a number of efforts to automate the entire clinical record-keeping process. The efforts are most pronounced at teaching hospitals and medical schools. There are currently many efforts extant to reduce the entire patient record to a machine readable form. Presumably this will reduce the storage volume necessary to maintain patient records over a long period of time and will eliminate the problem of the degeneration of paper records. If the

patient records are reduced to electronic form, the transmission of those records for purposes such as utilization review could probably take place via networks. Also, any economic, administrative or clinical research that used patient records could take place via networks.

One particularly important potential development in the automation of patient records is the introduction of a "smart card." The smart card is a wallet-sized, laser-readable card which has an encrypted version of an individual's entire medical history. Presumably, this card would be carried by all individuals at all times and would be available in an emergency to expedite the process of patient care. It would also eliminate the need for extensive network contact between the variety of medical institutions involved in providing care because there would be little need to locate and remotely access the patient's medical history. Such a plan carries very sensitive issues regarding civil liberties, however, and it may not come to pass in the near future.

#### Consolidation of Hospital Information Systems and Clinical Information Systems

Until recently, the production, storage and processing of information for administrative purposes (Hospital Information Systems) and the corresponding activities for clinical purposes (Clinical Information Systems) have been separate in most health care institutions. With the advent of distributed information systems and the promise of a set of standards for the use of database queries over a network, the possibility of a unified health information system became reasonable. Such a system would rely heavily on high-speed networks to provide contact between specialized workstations located (1) at nursing stations in a clinical facility, (2) in laboratories, (3) in administrative offices, (4) in physicians' offices and (5) in computing facilities. In the more distant future, these network connections could also connect medical research facilities, state accreditation and regulatory agencies, and third-party commercial organizations engaged in insurance and hospital management.

Such a system for local area communication was in the installation phase at the teaching hospital of one of our interviewees. This system involved the complete consolidation of information for the hospital, and required the installation of powerful workstations

at nursing stations within the hospital. It does not stand to have an affect on wide-area communications in the near future, but it's impact on loca-area communications is profound.

#### Computerization of medical examinations

Another development in medicine that should have a substantial effect on the use of networks is a plan by the National Board of Medical Examiners to include a computerized patient management simulation in the national medical board examinations within the next three years. The effect of the announcement on the population of medical students is to increase the amount of interest in and awareness of computing practices, products and hardware. Though the tests will be designed to require minimal computer expertise, some medical students perceive that a thorough knowledge of the technical aspects of computing will increase their chances on the examinations. This should increase the demand for network services as more members of the medical community become sophisticated with computing environments.

#### Development of friendly interfaces

There is a final general technical developments that will have a substantial effect on the demand for networks. This development is the emergence of a widespread concern among producers of software for the medical community to produce user interfaces and applications that are easy to use. Almost all of the producers of small computers, workstations, and software packages are reducing the demands on users to memorize commands, acquire keyboard skills and diagnose error messages. These developments will stimulate use of both computing resources and networks by a wider range of the medical population.

Current developments in the medical community include three new software packages that reduce the amount of knowledge necessary to use on-line medical services. The first is called Paper Chase. It was developed at Beth Israel hospital in Boston for use in conjunction with the MEDLINE database of bibliographic search services at the National Library of Medicine (NLM). The second is the Grateful Med interface, also used in conjunction with services provided by the NLM. Grateful Med acts as an elementary "agent" by automatically logging on to the NLM system, performing the

requested searches and downloading files to the user's local disk. The third interface is a product called Navigator, offered by CompuServe for use in conjunction with the services (including medical services) offered by CompuServe. In the long run, developments in computerized speech recognition, handwriting recognition and speech production, will probably enable interaction to take place independently of a keyboard. This should further stimulate the use of all computerized services, including those over networks.

#### Availability of computerized independent learning programs

Deal (1987) identifies four major ways that the computer has contributed to the modernizing of instruction in medicine. At least two of these are self-paced computerized instruction programs. The first is PLATO, a system developed at the University of Illinois College of Medicine in Chicago. This system is available for both mainframes and microcomputers and contains about 450 separate lessons. The second is a group of clinical simulations developed at Massachusetts General Hospital. These simulations permit individuals to practice their responses to various types of patient-management scenarios by interacting with a computer program. Access to these programs is available over the telephone, over the AMANet network or by purchasing floppy disks. Both of these programs foreshadow the increasing use of computers to increase the richness and availability of scarce knowledge sources for biomedical education. These developments are quite compatible with wide-area networking, and many are already available over wide-area networks.

#### DIRECT NETWORK DEVELOPMENTS

A number of important network services which target medical education and research institutions and individuals in the health professions are already in existence. In this section, we describe the ways in which these services can be classified. The services themselves are described more fully in Appendix B.

Services can be grouped in many ways, but four important divisions are (1) whether or not a service is cooperative, commercial or government-sponsored (2) the primary access method employed by a service (3) how individuals are charged for use and (4) what

populations of biomedical professionals a service is designed to accommodate. We deal with each of these in turn.

## Sponsorship

There are three basic types of sponsorship among network services available for the biomedical community. The first type is a group of bulletin boards and file servers that are largely cooperative. They typically exist under the auspices of individual medical schools or individual teaching hospitals. In almost all cases, they are based on IBM PC's or compatible machines and many of them can be accessed through direct dial-up or through the cooperative personal computer networks like FIDONet or OPUS.

The second type of network services are commercial services. There are three major commercial network service providers for medicine and each one is associated with a large commercial provider of general network services (i.e. CompuServe, Bibliographic Retrieval Services-BRS, and AMANet/GTE-Telenet). There are also a number of other commercial services (e.g. - MacNet, Applelink Personal Network, GENie Medical RoundTable). Finally, there are commercial satellite networks like the Hospital Satellite Network that broadcast relevant programs to member hospitals. The commercial providers of service tend to be focused on providing service to individual practicing physicians and hospital administrators. Many have a number of applications that provide specifically for the needs of individual physicians in their private practices.

The third type of services are those supported either by the national government or by state governments. In almost all cases, these services tend to provide for the needs of the academic and research communities. Many of the most widely used services are available through the National Library of Medicine (NLM). The Canadian government also provides a set of network-accessible services for biomedical research. Finally, the National Science Foundation provides a number of computing centers based at universities throughout the United States. These are also used heavily by individuals involved in biomedical research. There is one major exception to the rule that government-sponsored facilities are organized to serve researchers. The Lister Hill Center, a division of

the NLM, provides on-line access to a wide variety of network-accessible resources for CME. These resources are intended to be made available so the general population of physicians can become familiar with various computer-aided instruction products before they purchase them.

### Access

There are two main types of access for services in the biomedical community. First, users can connect via dial-up facilities. Almost all commercial services and some of the government supported services have this option in some form. Many services are available through commercial carriers like Tymenet or Telenet, which provide local numbers that permit connection to a large number of services for a nominal cost. Most well established services have a dial-in 800-number, for which users are charged an average per-hour surcharge of approximately \$10. The maximum baud rates for almost all dial-up facilities is 9600 baud.

The second widely used method is to provide a connection via established academic networks. Bitnet in particular provides a number of archives, conference facilities and bulletin boards for the biomedical community. The Internet provides some bulletin-board services and some conferencing facilities, but they are more limited.

### Compensation

When individuals use an on-line service outside of the research community, they are almost always charged a fee based on connect time. These fees vary with the service used, the access method used and the time of day in which the connection takes place. Use of an average service during business hours costs approximately \$20-\$30 an hour. With other services (primarily those based in cooperative networks) there is no charge for use of the service. Individual users must pay long-distance charges from their locations to the service, however. In the case of services provided on the BITNET and Internet networks, individuals do not pay for use of the servers, either. Data transport is provided over the leased lines which constitute the networks, and provision for these lines is made by individual universities. Finally, universities pay a fee for

membership in BITNET--this fee is dependent on the size of the individual university members. Because these charges are almost never passed down to individuals, there is no charge at all for use of these services as far as the individual researchers are concerned.

### NEEDS FOR NETWORK SERVICES EXPRESSED BY MEMBERS OF THE BIOMEDICAL COMMUNITY

There are a number of important needs expressed in a variety of ways by the biomedical community. In order to assess these needs, we have tried to obtain information from a variety of sources. We have interviewed individuals in the medical community including medical students, faculty members, members of medical school administrations, researchers, representatives of the press and representatives of government agencies. The names of these individuals are given in Appendix C. We have also surveyed the trade press. We have inquired on the networks, both commercial and cooperative, for contributions from the on-line medical community. Finally, we have consulted archival material from the records of both on-line conferences and on-line publications on the networks. From all these sources we have derived a number of the more important needs of the community. These needs are discussed in the following section.

#### NEED FOR BETTER INTERFACES:

Almost universally, there is concern about the difficulties encountered by the average user when using and particularly when learning to use the networks. Accomplishing any reasonable task using new or unfamiliar network systems is considered an extremely difficult task, even for an experienced user. The situation is considered even worse if it is attempted by a new user. Older users (in chronological age) are even less tolerant and less willing to experiment than the younger users, because of both time and familiarity with the technology. Because of these problems, more and more effort is being devoted to producing easy-to-use interfaces by third party vendors, members of medical school faculty and by students in schools of medicine and computer science departments throughout the country. The major efforts, however, are in designing single applications on stand-alone or PC-based systems. There are some notable exceptions to this, however. Software which provides simpler access to network-based resources is beginning to emerge.

These products include Paper Chase, Grateful Med and the interfaces to some commercial networks.

A proposal for a comprehensive interface to all biotechnical resources is currently under review. This interface, called the Biology Information Network (BIN), will consolidate access to widely-used biotechnical information by providing information about these resources. Patnaik (1988) describes it thus:

Currently, research centers and researchers involved in biotechnology do not have a consistent view of the network. Some researchers have access to ARPAnet, while others are on BITNET and some are fortunate enough to be on more than one network. And there are some researchers and research centers, who do not have access to networks. Unless a university campus (or organization) is participating in computer networks (such as, ARPAnet, NSFnet), connecting a single department to the net is not easy. BIN will provide connection information and encourage centers to join the NSFNet or a regional network.

A major need for the future, then, is for design of an interface that consolidates a large number of network services and which is easy to use. The criterion of ease of use should be judged by what unfamiliar users think, not what design engineers think. The average user in the biomedical population should be considered to be much brighter, but with much less free time to learn a system than the average user.

#### NEED FOR NEW NETWORK SERVICES

A number of network services were suggested by both interviewees and by on-line conferences. Services suggested include:

(1) An on-line financial aid system

Such a system would consolidate the mini database systems currently in use. It would serve financial-aid officers, development officers and students in medical schools throughout the region and/or throughout the country. Major functions would include keyword searching of scholarships, grants, loans and work study programs for students at medical schools. A properly designed

system would enable financial aid officers at medical schools to discover or to register sources of funding for students attending any medical school.

(2) A unified system for matching medical students to residencies

A system that performs this function currently exists. This system, called the National Resident Matching Program, attracted nearly 16,000 graduating medical students nationwide in 1987. It uses IBM-compatible PC machines to query students and administrators of residency programs about their preferences for each other. Software on the PC machines produces floppy disks which are then mailed to a central site, consolidated and evaluated. Recommendations are then made for matching medical students to post-doctoral residencies. The information is transferred using physical transport of the floppy disks. Because this system is currently working, however, there is potentially some resistance to put it on-line. Other individuals, however, would be interested in doing so.

(3) Centralized information system

Because of the value of making an extremely close match between the characteristics of medical school programs and the preferences of potential students, a means by which information about medical schools can be distributed to potential students would also be considered useful. This system could take the form of a database whereby each school provides current information about its policies, plant, programs, student body etc.. It could also take the form of a series of on-line or cooperative electronic mail conferences whereby members of the faculty, administration or student body would answer queries made by potential students. This system would help to minimize drop-outs for academic or personal reasons and would help to make more efficient use of extremely critical societal resources.

(4) Information resources for medical associations

Often there are restrictions in communications between members of various associations of medically related organizations. Teaching hospitals, medical associations, medical schools, and government agencies often have a difficult time reaching one another. Often, members of some organizations do not even know of the existence of other organizations because of the time it takes to locate them. A centralized resource of names, addresses and

computer addresses of organizations would improve communication between these organizations substantially.

#### NEED FOR A CHEAP ALTERNATIVE TO PAY-FOR-USE BULLETIN BOARDS.

For a large number of individuals, particularly medical students and recent graduates, the cost of using commercial bulletin boards is extremely high (\$20-50/hour). Often, those who have the best training to use computerized network services have the least amount of discretionary funds. There is a need, then, for an inexpensive computerized bulletin board system that coordinates activities at medical schools. Areas that would be particularly important would be those pertaining to medical students or faculty at medical schools. There is a movement among some members of The American Association for Medical Systems and Informatics (AAMSI) to create a cooperative bulletin board system connecting medical schools across the country.

#### NEED FOR CONSOLIDATION OF ACCESS TO BIOMEDICAL INFORMATION RESOURCES.

One of the problems often identified by members of the biomedical community, particularly by those doing genetic research, is the lack of a single source for critical research information and communication. Though some commercial sources have consolidated a few of the network services for physicians, there are still a number of critical resources that use different sets of commands, different communication protocols and different conventions for access and operation. Users must learn these multiple systems of commands at the expense of time they could more profitably devote to research. Because the few researchers that are comfortable with and facile with networks are scattered across so many different network systems, a critical mass of communicating, knowledgeable users has not emerged on any of the existing systems. Also, because of the explosion in biomedical information in recent years, it would be impossible for one organization to house and control all information. Therefore, it is necessary to interconnect the various existing systems with a common interface rather than to try to create a new, completely unified system from whole cloth.

Some consolidation of access to specific sets of resources is being currently being attempted by the staff of the National Library

of Medicine's Lister Hill Center. There are currently at least three such projects. The first is a project called Integrated Academic Information Management System (IAIMS) which involves the efforts of 10 universities and which is designed to develop an integrated access to campus-wide integration of information resources like libraries, medical records and administration. The second is a project to consolidate access to all resources important to genetic, biochemical, recombinant DNA and related research. The third is a project called the Unified Medical Language System (UMLS), which will cross-reference information from a variety of bibliographic sources. There is still a role, however, for an organization that will consolidate other information sources, mail systems, and valuable to clinical staff members, research staff members, and medical students.

#### NEED TO CIRCUMVENT THE "CLOCK PROBLEM" IN CONTINUING MEDICAL EDUCATION.

Though standardized video programming for continuing medical education is currently offered through satellite and cable networks, physicians often have a difficult time viewing the programs at a consistent time for several weeks in a row. Emergencies and other events can cause a physician to miss one or more parts of a televised clinic. Most importantly, if a clinic is broadcast nationwide, the differences in time zones often make it difficult for the physicians in some regions to view certain programs. This is remedied in part by multiple broadcasts of the same program. It would be useful, however, if there were facilities for viewing particular programs on demand from particular sites.

#### NEED TO PROVIDE AN EASY-TO-USE METHOD FOR USERS TO CREATE/CUSTOMIZE THEIR OWN ON-LINE SERVICES.

Many researchers, faculty members or students would be willing to create on-line services or conferences, but they are discouraged by enormity of the technical problems involved. These individuals will not have acceptable network services until they can design and create their own. A centralized server or a standardized piece of software which would allow any user to create his/her own archive, bulletin board, on-line conferencing system, on-line database server or remotely-accessible program would be greatly appreciated. Applications should be designed so that servers could be located at a Network Support Center or at any node on a network.

Resources of this type are already available on the PC-based bulletin boards, and they have greatly contributed to their popularity.

### PROBLEMS WITH PROVIDING NETWORK SERVICES TO THE BIOMEDICAL COMMUNITY

There are a number of problems that must be considered by any organization trying to provide network services to the biomedical community. Some of these problems are the same as those experienced by the wider university community, and some are unique to medicine. A few of the most serious of these problems are listed below.

#### LACK OF TIME TO LEARN NETWORK SYSTEMS

As the volume of new information, research and clinical techniques in medicine increases exponentially, the need for physicians to devote time to absorbing it increases accordingly. This situation is made worse by the increasing strictness of malpractice law. Physicians young and old must constantly master new information in order to maintain their standards of practice. This requirement severely limits the amount of time available for learning and experimenting with new network services. Since experimentation is often the best way to learn how to use computing services, they often go unlearned by members of the medical population. It will be necessary to find some way of drastically reducing the time required to both learn and use networks before physicians will be able to use them without hesitation.

#### MANUAL SKILLS

Use of a keyboard to communicate in real time takes a highly-developed set of skills. Many physicians, particularly older physicians, do not have the requisite keyboard skills to use computing resources. This problem is large enough for applications like electronic mail, where there are large amounts of time between communications. For on-line interactive conferencing, where individuals must communicate all responses immediately, lack of accomplished keyboard skills makes participation almost impossible. An easier method of inputting information must be developed and made widely available before many individuals will use networks for these purposes.

## COST

In spite of the fact that physicians are compensated well, there is an issue of minimizing costs associated with providing network services.. Use of a new or complex service on the commercial networks can cost as much as \$50/hour. Annual fees and one-time fees can amount to hundreds of dollars for some services. Medical students in particular cannot afford usage fees like this. Some firms give lower rates to students, but they are still high when compared with the students' ability to pay. In order to teach medical students how to use network services effectively, schools must provide access to the networked services. This can be extremely expensive for a school with more than a few students. Therefore, in order to become widely used, a networking scheme must provide service to the education and research population at a very low cost.

## NETWORK BANDWIDTH:

Many interactive CAI applications require full-screen video interaction in real time. Many medical imaging applications would require bandwidth far in excess of current levels before it becomes practical to use them. Videodisk is an attractive alternative given the speed with which many wide-area network vendors are adding capacity. These must be kept in mind when designing any new network transport services for the medical community.

## SUPERIORITY OF HUMAN CONTACT

In spite of the advantages of many network services, there are parts of the socialization and continued education of many medical personnel simply cannot be performed over networks. Because of the complexity, the need for immediate feedback and the risk involved, human contact will always be necessary for many parts of medical education and research. Providers of network services must keep this in mind when attempting to produce relevant and helpful network services. Anything less will result in products that will be ignored by the physician population.

## CONFIDENTIALITY

There is a dilemma inherent in providing consolidated, accessible and easy-to-use network services to the medical

population. Many records involve extremely sensitive data about individual patients. Because of this, the maintenance of confidentiality of patient records is an extremely important issue. The conflicting requirements of providing quick and easy access to information and providing high levels of security for patient protection makes the provision of network services extremely problematic. Any future network scheme must take both requirements into account.

#### LACK OF STANDARDS

The continuing lack of standards in networking make consolidation of services and communication in the medical field extremely difficult. This lack of standards is made more pronounced by the fact that many research programs at medical schools are in direct competition with one another. Ensuring successful network communications requires achieving cooperation between constituent member organizations. Ensuring this cooperation requires either accommodation or standardization of the hardware and software at all network nodes. Such consolidation will not be an easy task for the medical community.

#### RECOMMENDATIONS

Recommendations concerning NYSERNet's efforts to increase the degree of participation by members of the biomedical community must incorporate a complex set of considerations. In addition to the set of problems and needs described in the previous sections, the recommendations must be consistent with other activities NYSERNet undertakes. The recommendations should also incorporate considerations pertinent to any regional network. For example:

- (1) NYSERNet should attempt to exploit those situations where it has comparative advantage to other existing networks because of (a) its geographic location, (b) its functional characteristics (like transmission speed) or (c) its access to a population of highly trained technical users.
- (2) NYSERNet should serve individuals in its geographic region in at least three ways: (1) as a gateway to network services

located elsewhere in the country, (2) as a provider of network services and (3) as a provider of data transport services.

(3) NYSERNet should also think of itself as a gateway for researchers in the rest of the country to resources in this region. It should create access to those services that are: (a) important to individuals nationwide, (b) located in New York State and (c) not available elsewhere in the country.

(4) NYSERNet should take maximum advantage of any superior knowledge of network hardware, software, resources or procedures that reside in-house.

(5) NYSERNet's Network Information and Support Center is now established. There is strong justification for using this center as a location for regionally-relevant network services.

(6) NYSERNet will benefit by initiating and/or encouraging any kind of collaboration among researchers at its member organizations. This is particularly true for collaborations or initiatives that rely heavily on networks for data transport or information access.

(7) NYSERNet will benefit by services and programs that are developed at its member organizations as well as those developed in-house. It should act as a catalyst where it cannot act as a primary developer.

(8) NYSERNet should continue to have a role in the evolution of standards for communication among members of the education and research community.

(9) NYSERNet must concern itself with survival in the short, intermediate and long term.

The recommendations given below are implied by the principles stated above. They are specific to the biomedical community, and they are ordered roughly by the degree of ease with which NYSERNet could implement them. They represent situations where there are organizations who have expressed interest in cooperating with NYSERNet, where NYSERNet has existing comparative advantage because of geographic, technical, or other

reasons, and where there are large enough numbers of potential users to justify such a course of action.

#### SPECIFIC RECOMMENDATIONS

1. Take steps to provide connection of teaching hospitals and medical schools to NYSERNet.

One reason why a great deal of biomedical researchers do not use TCP/IP networks is that their organizations have no explicit connection to these networks. Often, medical researchers are using time-consuming file transfer methods involving multiple steps or expensive private leased lines to carry data to shared facilities like those at Cornell. Many medical schools and teaching hospitals have no connections to high-speed research networks at all. The inclusion of these teaching hospitals and/or medical schools as nodes on the NYSERNet network would accomplish two things. First, it would increase the base revenues from connection of more institutions. Second, it would foster the development of user created services and cooperation over the network.

2. Ensure availability of BRS Colleague to the medical population

The presence of a service offered by the Bibliographic Retrieval Service (BRS), designed specifically to accommodate biomedical researchers, suggests the wisdom of consciously incorporating it in NYSERNet's BRS interface program. Since NYSERNet is rapidly gaining comparative advantage to other networks in knowledge of BRS, it makes sense that NYSERNet should be the lead organization in doing this. BRS also provides reduced prices to students who use BRS Colleague. The combination of an easy-to-use interface and reduced charges makes it potentially attractive to the medical student population. Implementation of this service explicitly carries out NYSERNet's mandate to provide network services for education.

3. Sponsor a conference of officials in medical community.

At best, this report describes a portion of the needs and requirements of the biomedical community. A more extensive statement of needs might be generated from a group of highly informed officials who bring insights from their separate institutions.

NYSERNet's sponsorship of a short conference which brought together these individuals would do two things. First, it would increase the awareness of NYSERNet on the part of these individuals and hence would increase the probability that they might request connection to the network. Second, it would provide a more accurate, more thorough and more insightful statement of the needs and concerns of members of the biomedical community.

This conference could be inexpensively incorporated into the program of meetings of either national or regional medical societies. Alternatively, it could be organized as an independent event. This would increase the degree of visibility of NYSERNet and the amount of concentration on the previously-outlined problems. The conference could focus either on the problems of the biomedical field in general or on issues associated with a particular specialty of medicine. Researchers involved in the study of medical imaging would probably be the best candidates for a field-specific conference because they would be most sensitive to the technical issues involved.

#### 4. Insure availability of crystallographic information resident at Brookhaven National Laboratories (BNL).

We spoke to at least two of the officials at national laboratories who were very much interested in the on-line availability of crystallographic information that is produced at Brookhaven National labs. These individuals were interested in NYSERNet because of the ease of access which it would provide to Brookhaven. They cited the comparative advantage of electronic mail communication over sending the information on tape. Insuring that the BNL researchers know about this demand and can accommodate it would encourage national awareness of NYSERNet's activities.

#### 5. Explore networking resources in southern Canada

Contact with Dr. John Rodgers at the National Research Council of Canada (NRCC), suggests that there are a number of extremely valuable database, archival and human resources in southern Canada. Because of its geographic proximity, NYSERNet has a strong comparative advantage to other NSF-sponsored networks in providing access to these resources. Currently contact between NRCC

and U.S. research installations takes place through BITNET, which is limited in both connectivity and bandwidth. Discussions with Rodgers indicate that he is extremely interested in working with NYSERNet to gain access to the U.S. research community. He relies heavily on information sent to him on magnetic tape from Brookhaven National Laboratory, a NYSERNet node. Using connection to Rodgers as a justification for providing access to all of southern Canada would open up enormous opportunities on the part of research populations in both the U.S. and Canada.

6. Precipitate development of cooperative imaging research at network nodes.

The presence of a number of extremely well known medical schools and teaching hospitals in NYSERNet's territory suggests that a common research project involving several of them might be possible. Technology is currently available to store, transport and manage body image data over networks. Current offerings include Commview, a product developed jointly by AT&T and Phillips Corporation, and DECRAD, a product offered through a joint venture of Siemens and Digital Equipment Corporation. At least two of the research hospitals in New York City have ordered the Commview product. This suggests that NYSERNet could foster the development of a research project at one of these nodes to provide for network transport of images from around the state. This technology provides for more than just images that are inherently digital (such as CAT or MRI); it also has the capacity to digitize regular flat-film images.

Because of the potential for creation of a database, a network of diagnostic relationships or a network of pedagogical relationships using these technologies in a variety of ways, we feel that NYSERNet should become involved. Though it should not necessarily develop facilities in-house, it should encourage or precipitate creation of such a service among one or more of its constituent members. Following is a list of a few possible places where such facilities might be useful.

- >>Research Database of Body Imaging Information
  - Traditional Body Imaging
    - Digitized Flat Films
  - Computerized Body Imaging
    - Ultrasound Images

- Computerized Tomography
- Magnetic Resonance Imaging
- Positron Emission Tomography
- Holographic Images

>> Database of "Standard" Medical Images

- For Equipment Calibration
- For Equipment Troubleshooting/Diagnosis from a Remote Site
- For Pedagogy
  - Classic Cases/Images
  - Classic Mis-diagnoses
  - Classic Image Artifacts

>>Digitized Images of Objective Symptoms of Various Diseases

>>Research/Educational Database of Tissue Images

- Digitized Images of Gross Anatomy
- Digitized Light-spectrum Micrographs
- Digitized Light-spectrum Ultramicrographs
- Electron Micrographs

>>Pathology Database of Digitized Images of Tissue Samples

>>Histology Database of Digitized Images of Tissue Samples

The creation of databases of the type listed above suggests that there could be other, similar databases that take advantage of the ability of the network to transmit digitized information. These databases could serve persons from other specialties. A representative list of some potential application is given below.

>>Research/Educational Heart-Sounds Database

- Patient Records
- Standard/Classic Problems or Diagnoses
- Research database of individuals with particular ailments

>>Research/Educational Hematology Database

- Blood component data
- Digitized Images of Specimens

7. Work to incorporate network literacy into the daily life of medical students.

One of the most important determinants of whether physicians use computational services is whether or not they learned to use those facilities while they were a student. One of the best ways to cause this to happen is to create a situation where students have access to networks in the same physical locations where they have opportunities to use them. NYSERNet should, as much as possible, try to position network access in a place where medical students will use it. This suggests providing terminals in student lounges, lounges for hospital medical staff and at locations near eating facilities. The increased ease of access might precipitate greater awareness and use of the network. Training sessions three times a year on this equipment would help to stimulate even more use. Finally, the fact that individuals would always be visible when using these facilities would help to generate more interest on the part of other members of the medical community.

A second way of incorporating use of computers into the daily lives of medical students is to push for the incorporation of computer literacy and network literacy into the medical curriculum. This is difficult, because of the already substantial academic load on medical students. If done properly, however, the introduction of computer literacy and network literacy into the population of medical students would result in a net time saving. Gordon Black, a noted advocate of this plan has coined the term "the physicians computer toolbag" to denote the minimum functions that should be covered. These include (Bunn, 1988):

- I. Personal management
  - finances
  - calendars
  - aids to writing (e.g. - producing a *curriculum vitae*)
- II. Word processing
- III. Patient assessment
  - health risk appraisal programs

## IV. Telematics (inter-computer communication)

- medically-oriented bulletin boards
  - Black Bag II
  - MedSig on Compuserve
  - Medline
  - others

## V. Database management

- specialized operating system/database managers
- CD-ROM applications
  - Medline
- hospital medical records

## VI. Clinical Decision Making Support

- simple interpreters of clinical data
  - automatic flagging of abnormal lab results
  - computerized ECG reading
- aiding diagnosis and treatment by
  - automatically scanning charts or records
  - decision aids (available on demand)
- examples
  - Problem Knowledge Couplers
  - Quick Medical Reference
  - DXplain (available on AMANet)

## VII. Practice management

- record keeping
- appointment scheduling

## 8. Support Bottoms-Up Programs

Gordon Black also suggests another program for introducing computer competency into the community of physicians. He instituted this so called "bottoms up" program at Ohio State University (Bunn, 1988). The program exploits the fact that medical students, interns and resident physicians are frequently more adept at using computing and network technology than are older physicians. During their training the junior persons are frequently rotated to different parts of training centers. They are assigned to established physicians during these rotations. At Ohio State, young physicians that are trained to use network and computing resources are then assigned to train the older physicians with whom they

work. This scheme causes the knowledge of how to use the computer and network resources to migrate upward into the established members of the medical community. Also, it takes advantage of the fact that training is best accomplished through person-to-person contact. Such a scheme might be useful in the organizations that NYSERNet serves. Encouraging nodes to adopt such a system will result in the increased use of networks in the medical community.

9. Consider doing sponsored research for an established network service provider.

At least three individuals, (one from AMANet, one from the Lister Hill Center, and one from the National Research Council of Canada) expressed interest in contacting NYSERNet about providing access to their resources. The Lister Hill Center in particular is responsible for distributing funds for projects that improve network access for medical and biological research. The Center is currently interested in providing access to the resources of the National Library of Medicine through the established research networks like BITNET and the Internet. NYSERNet has substantial knowledge of network technology, knows the problems of creating access to resources because of its work with BRS and other institutions and would be considered an acceptable choice for network research funds by the Lister Hill Center or the National Institute of Health.

Also, recent legislation passed by the Congress of the United States provides for \$100 million to be used for, among other things, "a public agency or corporation established for the purpose of developing and operating telecommunications networks to enhance educational opportunities for educational institutions, teacher training centers, *health institutions*, and industry" [emphasis added]. This funding could support a number of the projects suggested in this document.

10. Take steps to encourage the use of electronic mail in the medical community.

John Farrer, a physician who frequently participates in on-line conferences and contributes to on-line forums, makes this suggestion for increasing the use of electronic mail (Farrer, 1987) He argues for placement of supportive statements in the medical journals. These

statements could say that letters and articles may be submitted by electronic mail. This could be augmented by encouraging vendors of things like drugs and medical supplies to provide electronic mail access to their offices or customer service groups. Official sanction of the use of electronic mail for these common information flows can lead to its widespread use in the medical community.

11. Cooperate with operators of PC-based bulletin board systems.

Many of the most relevant and useful services for the medical community are being developed at medical schools using personal computers. These services would be useful to a larger number of individuals if they were available on a nationally-connected network system like the INTERNET or the NSFNET. NYSERNet can play an important role in precipitating this use by acting as a base of operations for these services. Selection of the best solutions and modification of these solutions to run on UNIX-based systems would be extremely valuable in encouraging awareness and use of NYSERNet to the members of the biomedical community.

12. Encourage connection of Macintosh based services.

More and more, applications for the biomedical community are being developed on Macintosh computers. The Macintosh represents the current state of the art in easy of use and friendliness of interface. Already there are network-accessible services for the biomedical community which rely on Macintosh technology. One of the major services, bioMEDICUS, is located in New York State. The availability of technology to link these Macintosh products to TCP/IP networks represents an opportunity to incorporate them into the collection of services that can be accessed via NYSERNet.

## BIBLIOGRAPHY

- Black, Gordon C., 1987, "Computer Literacy Program for medical students and residents," Compuserve, Medical Special Intrest Group, Data Library 4.
- Bunn, Bruce, 1988, "The Physician's Computer Tool Bag" Compuserve, Medical Special Interest Group, Data Library 4.
- Burrows, William R. and Hannah L. Hedrick, 1987, "Allied health education and accreditation," *Journal of the American Medical Association*, Vol. 285, No. 8, pp. 1056-60.
- Crowley, Anne E., Sylvia I. Etzel and Edward S. Peterson, 1987, "Undergraduate medical education," *Journal of the American Medical Association*, Vol. 285, No. 8, pp. 1013-20.
- Crowley, Anne E., Sylvia I. Etzel and Helen A Shaw, 1987, "Graduate medical education in the United States," *Journal of the American Medical Association*, Vol. 285, No. 8, pp. 1031-40..
- Deal, William B., 1987, "Medical education," *Journal of the American Medical Association*, Vol. 285, No. 16, pp. 2296-97.
- Farrer, John, 1987, "Medical telecommunications: A review," Compuserve, Medical Special Intrest Group, Data Library 2.
- Felch, William C., 1987, "Continuing medical education in the U.S.: An enterprise in tradition," *Journal of the American Medical Association*, Vol. 285, No. 10, p. 1355.
- Jolly, Paul, Leon Taksel and Robert Beran, 1987, "U.S. medical school finances," *Journal of the American Medical Association*, Vol. 285, No. 8, pp. 1022-30.
- Journal of the American Medical Association*, 1987, "Medical news," Vol. 286, No. 22, pp. 3216-17.

MEDLARS: The World of Medicine at Your Fingertips, 1986, NIH Publication No. 86-1286, U.S. Department of Health and Human Services, Public Health Service, National Institutes of Health: Bethesda, MD.

National Library of Medicine, 1987, *Long Range Plan for the National Library of Medicine*, U.S. Department of Health and Human Services, Public Health Service, National Institutes of Health: Bethesda, MD.

Newald, Jane, 1987, "Hospitals look to computerization of physician offices as linkage," *Hospitals*, Vol. 61, August 5, pp.92-94.

Osteen, Arthur M. and Michael I. Gannon, 1987, "Continuing medical education," *Journal of the American Medical Association*, Vol. 285, No. 8, pp. 1043-48.

Packer, C. L., 1987 "Use of personal computers increases 40% in '86," *Hospitals*, Vol. 61, February 20, pp.98-99.

Patnaik, Deba, 1988 "Proposal for Biology Information Network," Center for Marine Biotechnology, Baltimore, MD.

Sehr, Barbara, 1988, "Optical disk systems sharpen X-ray image," *Computerworld*, Vol. 22, January 18, p. 71.

Wallace, Cynthia, 1987, "Hospital Satellite Network hopes to post first profit in '87 after start-up losses," *Modern Healthcare*, Vol. 17, January 16, p. 49.

APPENDIX A

MEDICAL SCHOOLS IN AND NEAR THE GEOGRAPHIC AREA  
SERVED BY NYSERNET

---

Table I: *U.S. Medical Schools that fall within the geographic area served by NYSERNet*

---

Albany Medical College of Union University	Albany, NY
Albert Einstein College of Medicine of Yeshiva University	NY, NY
Columbia University College of Physicians and Surgeons	NY, NY
Cornell University Medical College	NY, NY
Mount Sinai School of Medicine of City University of NY	NY, NY
New York Medical College	Valhalla, NY
SUNY at Buffalo	Buffalo, NY
SUNY Health Science Center at Brooklyn	Brooklyn, NY
SUNY at Stony Brook School of Medicine	Stony Brook, NY
SUNY Health Science Center at Syracuse	Syracuse, NY
University of Rochester School of Medicine and Dentistry	Rochester, NY

---

---

Table II.

*Canadian medical schools that fall in geographic areas near those served by NYSERNet*

---

Universite Laval Faculte de Medecine	Quebec, Quebec
Universite de Sherbrooke Faculte de Medecine	Sherbrooke, Quebec
Universite de Montreal Faculte de Medecine	Montreal, Quebec
McGill University Faculty of Medicine	Montreal, Quebec
University of Ottawa Faculty of Medicine	Ottawa, Ontario
Queen's University Faculty of Medicine	Kingston, Ontario
University of Toronto Faculty of Medicine	Toronto, Ontario
McMaster University Faculty of Medicine	Hamilton, Ontario
University of Western Ontario Faculty of Medicine	London, Ontario

---

APPENDIX B

AN OUTLINE OF SERVICES AVAILABLE TO THE BIOMEDICAL  
COMMUNITY

## I. PC BASED COOPERATIVE NETWORKS

### A. Black Bag Networks

These are a collection of bulletin boards and file servers which reside on personal computers typically under the auspices of individual medical schools. The largest and most widely known of these is located at Ohio State University. There are others in operation in Philadelphia and the University of Delaware. Others are being planned at Wayne State University, St. Louis University, The University of Southern Alabama, The University of Michigan and The University of California at Irvine.

### B. Saint Silicon Project

St. Silicon's is a service in operation in Cleveland, Ohio. It is a part of the Cleveland freenet project.

### C. Fidonet

Fidonet is a general cooperative PC Based network that often carries information about PC based software and communications on a variety of subjects. It does not deal explicitly with biomedical issues but sometimes carries information regarding these interest areas.

D. OPUS is a network similar to Fidonet that provides similar services to individuals who own IBM PC and PC-compatible machines.

## II. COMMERCIAL NETWORK SERVICES

### A. CompuServe Medical Special Interest Group (MEDSIG)

This network is one of the largest of the commercial medical network services. It provides a large numbers of on-line services including electronic mail, on-line conferencing, software distribution, expert question and answer, libraries of articles of interest to the medical community, and announcements of events and activities of interest to the medical community. It has a large following.

## B. American Medical Association Network (AMANET)

The AMANET provides several types of service including (1)an electronic communication service, (2)a public information service which gives weekly reports from The Center for Disease Control (CDC), National Library of Medicine, the Surgeon General's office, (3)an easy to use bibliographic research service called Paper Chase which provides access to the services of the National Library of Medicine, (4)a daily news wire service which screens stories from the Associated Press and relays those stories of interest to a particular medical specialty, (5)a library of continuing medical education (CME) interactive courses in a variety of specialities which can be used for credit in CME, (6)a library of support programs for medical decision making which includes DXplain, a diagnostic support program designed at Massachusetts General Hospital and MEDICOM, a database of pharmacological information and (7)a set of general administrative services for the practicing physician.

## C. MacNET

MacNET is a Macintosh based commercial network that contains at least three network services of interest to the biomedical community. These include (1)BioMEDICUS, a general information clearing house for the public and medical professionals which includes software, library services, special interest groups and a drug reference library, (2)medLAW, an on-line information service dealing with legal issues in medicine, and (3)AIDSnet, a centralized information source for researchers investigating various aspects of Acquired Immune Deficiency Syndrome.

## D. Bibliographic Retrieval Service (BRS) Colleague

BRS Colleague is one of the many services provided by BRS. It provides friendly access to databases which pertain to the biological and medical communities. It provides a unified frontend for 150 BRS databases. It allows access to MEDLINE, 70 journals, 25 textbooks, physician referral information and knowledge bases in such areas as oncology and AIDS research.

## E. Hospital Satellite Network

The Hospital Satellite Network (HSN) is a satellite-based network which distributes signals to approximately 550 of the 1100

hospitals who have satellite dishes. . This network provides a variety of programming for continuing education in medicine and hospital administration. It also provides a variety educational and entertainment programming for patients. Programming is continuous and is strictly broadcast. There is currently no provision for two-way video interaction. HSN collaborates with several established organizations in the health-care field including the American Medical Association, The American College of Healthcare Executives and the Healthcare Financial Management Association.

#### F. GENie network

General Electric's GENie network carries a service called the Medical Roundtable. This service provides for on-line conferencing, including such issues as medical ethics, medical practice problems and discussions of medical technology.

#### G. American Hospital Association

The American Hospital Association also sponsors a satellite-based network. It sponsors periodic broadcasts of healthcare conferences from its Media Center. It is concerned primarily with continuing education and important events of interest to individuals in hospital management.

### III. GOVERNMENT SUPPORTED NETWORK SERVICES

#### A. National Library of Medicine (NLM)

The NLM provides probably the most widely used set of Biomedical services in the U.S. today. Through its computerized location and retrieval system (MEDLARS) it provides access to the largest single collection of bibliographic information on biomedical issues in the world. On-line access to the NLM's citation records is provided through its MEDLINE interface, which is used by virtually every institution in the medical community today. Two easy to use interfaces, GRATEFUL MED and Paper Chase, have been developed specifically to make access to MEDLINE more widely available throughout the medical community. In addition to these services, NLM offers a number of other services on-line. These include:

- (1) AVLINE: (AudioVisuals onLINE) provides access to bibliographic citations to audiovisual materials in medicine
- (2) BIOETHICSLINE: provides access to bibliographic citations to material covering ethics in clinical medicine and research.
- (3) CANCERLIT: provides access to bibliographic citations covering topics in oncology
- (4) CANCERPROJ: Provides access to database of summaries of cancer research projects
- (5) CATLINE (CATalog onLINE): provides access to bibliographic records of monograph and book publications in NLM's collections
- (6) CHEMLINE: provides access to an online interactive dictionary of over 650,000 chemical substances
- (7) CLINPROT (CLINical PROTOcols): provides access to summaries of investigations of new anti-cancer drugs and treatment protocols
- (8) DIRLINE (DIRectory of information resources onLINE): provides access to a directory of organizations and research centers with information in specialized medical fields.
- (9) HEALTH (HEALTH planning and administration): provides access to citations to nonclinical aspects of health care
- (10) HISTLINE (HISTory of medicine onLINE): provides access to materials pertaining to the history of medicine.
- (11) MEDLINE: provides access to citations to medical and biomedical journal articles. Nearly 5 million articles from 3200 journals are included in the database.
- (12) MESH VOCABULARY FILE (MEDical Subject Headings): provides access to a comprehensive, cross-indexed dictionary and thesaurus of medical terms

(13) NAME AUTHORITY FILE: provides access to an authoritative listing of names, titles and headings from bibliographic records.

(14) POPLINE (POPulation information onLINE): provides access to citations to information about population issues and family planning

(15) PDQ (Physician Data Query): provides access to summaries of major tumor types and state of the art treatment options, a 10,000 name database for physician referral, listings for 2,000 organizations associated with oncology.

(16) RTECS (Registry of Toxic Effects of Chemical Substances): provides access to a registry of toxicity data for more than 70,000 substances

(17) SERLINE (SERials onLINE): provides access to citations to biomedical serials

(18) TOXLINE (TOXicology information onLINE): Provides access to citations covering biochemical, physiological, pharmacological effects of drugs and other toxic substances.

(19) TOXNET: provides access to an interlocking computerized system of various databanks covering all of aspects of toxic substances

(20) CCRIS (Chemical Carcinogenesis Research Information System): provides information about carcinogenicity, mutagenicity and tumor promotion in approximately 1200 substances

(21) HSDB (Hazardous Substances Databank): provides access to a databank containing toxicological information and research on environment emergency treatment, and regulatory issues for over 4,000 hazardous substances.

## B. Lister Hill Center

The Lister Hill Center is a division of the NLM that is concerned specifically with advancing research and development for the

biomedical community. Among the online services that are offered through the Lister Hill Center are:

(1) A center for continuing medical education that is providing subsidized on-line access to a variety of interactive simulation materials in continuing medical education. Often, software which provides clinical decision support or which is designed for computer-aided instruction is extremely expensive. The Lister Hill Center is committed to providing access to as many of these materials as possible so that individuals can make informed purchase decisions.

(2) A databank of currently available materials for medical education and clinical decision support. This databank also includes information about individuals who are interested in sharing the mastering and production costs to put their applications on compact disk.

#### C. Data Base Services

The U.S. Federal government also supports a variety of database services provided by various agencies within the federal government. The National Institute of Health for example, supports a number of bio-technical databases including: several databases at the National Cancer Institute, several regional information centers of DNA sequence information, a database dealing with Mendellian inheritance information, a protein identification database and a database of crystallographic information.

#### D. Shared Computing Facilities

There are a variety of shared computing facilities sponsored by the federal government, in particular by the National Institute of Health (NIH) and National Science Foundation. The NSF facilities are general purpose computing centers that are often used by biomedical researchers for such research as molecular analysis (including research on molecular folding), medical imaging, and genetic sequence analysis. Because these facilities are shared with research projects from all fields, they tend to be used by only a moderate number of biomedical researchers.

The NIH provides a supercomputing resource specifically for biomedical projects at the campus of the NIH in Frederick, Maryland. This facility is used by biomedical researchers throughout the country.

The NIH also supports a group of regional computing centers at various sites around the U.S.. These include the Molecular Biology Computer Resource in Boston, Massachusetts, Molecular Biology Information Resource at Baylor University in Texas and the BIONET computing resource in California.

#### E. Canadian Biomedical Resources

There are also a number of biomedical resources supported by the Canadian national government. The most significant of these is an extensive database of crystallographic information located in Ottawa, Ontario. This resource is supported by the National Research Council of Canada. Other databases address problems in human gene mapping, oncology, biochemistry, pharmacology and microbiology. These databases are frequently used by researchers throughout the United States via Telenet.

### IV ACADEMIC RESEARCH NETWORKS

There are at least two national networks which provide services to the biomedical community. The first is the Internet, supported by DARPA, the Defense Advanced Research Projects Agency. Besides the usual electronic mail and file transfer facilities, the Internet provides a widely-used collection of bulletin boards called USENET. There are several bulletin boards on USENET of interest to biomedical researchers. These include bulletin boards focusing on general biological and medical issues, biotechnology, medical imaging technology and Acquired Immune Deficiency Syndrome.

BITNET, a cooperative network comprised largely of universities, provides some unique services that fulfill the need of the biomedical community. These include BIOSERVE, a fileserver located at the University of Maryland. BIOSERVE provides archival

information, DNA sequence information for plasmids, and computer conferencing facilities. Access to public domain software and research data are planned in the future. LifeSci is a server running at Technion, the Israeli National Technical University. Its purpose is to increase communication between people in life science related fields. It provides an easy to use interface for creating on-line conferences, archives of messages and distribution of software and data. Finally, BITNET provides a series of mail reflector services where electronic mail addressed to a particular reflector will result in the message being sent to all members on the list. Biomedically related servers include: AIDSNEWS, BIOTECH and NEURON.

APPENDIX C

INTERVIEWEES

The following individuals were interviewed at various times during the course of this project. All of them were helpful, informative and indispensable to the quality of the document. I hope that their statements were successfully incorporated into the document--any errors of omission or commission are strictly due to me.

Michael Ackerman  
Director, The Learning Center for Interactive Technology  
Lister Hill Center  
The National Library of Medicine  
Bethesda, MD 20894

Dennis Benson  
Lister Hill Center  
The National Library of Medicine  
Bethesda, MD 20894

Richard Covert  
Director  
Healthcare Information and Management Systems Society  
American Hospital Association.

Sue Frisch  
Senior Editor  
Computer News for Physicians  
BMI/McGraw-Hill  
New York, NY 10017

Karen Graves  
Director of Reference Services  
American Medical Association  
Chicago, IL 60610

Bob Hall  
System Operator, CD-ROM Bulletin Board System  
Technical Director  
Ellis Enterprises Inc.  
Oklahoma City, OK 73103

Betsy L. Humphreys  
Deputy Associate Director for Library Operations  
The National Library of Medicine  
Bethesda, MD 20894

Fred Keefer  
Code 3  
3M Health Information Systems  
Salt Lake City, Utah 84107

Louis Kettle  
Associate Vice President for Academic Affairs  
Council of Teaching Hospitals  
Association of American Medical Colleges  
Washington, D.C.

Elizabeth Martin  
Director, Communications Department  
Council of Teaching Hospitals  
Association of American Medical Colleges  
Washington, D.C.

Deba Patnaik  
Center for Marine Biotechnology  
600 E. Lombard St.  
Baltimore, MD 21202

Chris Pogorzala  
Program Director  
Department of Medical Records  
School of Allied Health Professions  
Ithaca College  
Ithaca, NY 14850

Thomas C. Rindfleisch  
Director, Knowledge Systems Laboratory  
Stanford University  
Stanford, CA 94305

Timothy P. Rinda  
Representative  
BRS Information Technologies  
St. Davids, PA 19087

John R. Rodgers  
National Research Council of Canada,  
M-55, Room-275,  
Montreal Road  
Ottawa, Ontario K1A 0S2  
Canada

Albert Shar  
Director, Office of Information Technology  
The University of Pennsylvania School of Medicine  
Philadelphia, PA 19104

Dean Taylor  
Associate Professor of Mechanical and Aerospace Engineering  
Cornell University  
Ithaca, NY 14843

William Yasnoff  
Medical Computing Consultant  
Medical Information Network  
The American Medical Association  
Chicago, IL 60610

REPORT III

MARKET AND ECONOMIC  
IMPACT STUDY

VOLUME III: EDUCATION K-12

**NYSERNet**

NIC COLLECTION

**New York State Education  
and Research Network, Inc.**

**1095 Avenue of the Americas  
Room 1702  
New York, NY 10036  
Tel. 212-395-4480**



REPORT III

MARKET AND ECONOMIC  
IMPACT STUDY

VOLUME III: EDUCATION K - 12

Prepared By:

Jim Lombardi  
Cornell University

Tom Vietorisz  
New School For Social Research  
Cornell University

For

MEIS TEAM

### REPORT III: GENERAL INTRODUCTION:

During the period of the last twelve months, the members of the NYSERNet Market and Economic Impact Study (MEIS) Team produced two prior reports that identified many of the needs and problems of individuals using data telecommunications networks for education and research at universities, a paper on the relationship between, "The Economic Benefits from and Public Support of a National Education and Research Network", plus the current report. In the prior reports, the team identified specific groups of individuals who are involved in education and research activities and whose needs required additional study. The several volumes of the current report are devoted to four of those groups:

- (1) Volume I, The Supercomputing Community
- (2) Volume II, The Biomedical Community,
- (3) Volume III, Education: K-12
- (4) Volume IV, Industrial Clients.

Generally speaking, these volumes review the needs of a specific segment of the education and research community, point out current problems associated with providing network service for that segment and outline recommendations for action that can be taken by an organization such as NYSERNet to serve that group most effectively. Taken together, these volumes present a broad spectrum view of the opportunities wide-area networking organizations such as NYSERNet have for making a positive economic impact on their regions--and on the nation as a whole.

Many readers have a primary interest in one or more of the segments covered in the current report, but not all of them. In order to facilitate the distribution of the report to those individuals, we have produced each volume under separate cover.

## VOLUME III: EDUCATION K-12

### I INTRODUCTION

There are a number of efforts underway across the U.S. to introduce telecommunications activities into the elementary and secondary school systems (See Appendix A, Appendix C). These efforts vary in scale and scope, ranging from experiments organized by small groups of individuals to larger implementations involving major players in the telecommunications world.

Despite these efforts, however, telecommunications-based activities remain far-removed from the day-to-day activities of the typical classroom teacher. A number of issues — both technical and social — will need to be addressed before teachers can integrate telecommunications into the classrooms.

There is an imperative to do so. There is a widespread belief that the U.S. is losing ground in terms of economic competitiveness and technological leadership. Telecommunications, computers, and other related technologies are a vital part of the new economic environment, and are rapidly becoming the major tools of the so-called 'knowledge worker' in the 'Information Age'. Our students need to be thoroughly prepared to use these new tools in the workplace, and all of the current Presidential candidates are pushing incentive programs aimed at the educational system to insure that this happens. Telecommunications and education is a timely issue.

In the long-run, telecommunications and computer technology offer the promise of radically transforming our current approaches to education (See for example Lewis J. Perelman's Technology and Transformation of Schools, 10/87). However, in the short-run, telecommunications activities need to be introduced into the existing classroom-based framework.

This section of the report is an attempt to provide NYSERNet with the beginnings of a logical plan to assist in the introduction of telecommunications technology on a broad-based, integrated basis into the New York State elementary and secondary school systems, by:

- identifying some key educationally-related needs of students and teachers which telecommunications technology can help to satisfy.
- discussing some critical issues/needs relevant to the practical implementation of telecommunications applications in the schools.
- making specific recommendations for an action plan for NYSERNet.

## II IDENTIFICATION OF NEEDS

The primary need which telecommunications technology will meet is to free students and teachers from the constraints on education imposed by the most common classroom model — groups of people located in separate physical locations, with unequal, inconsistent, or inconvenient access to resources. Students and teachers also need to be educated about the technology itself.

The needs listed below were derived from: conversations with school personnel (see Appendix A); transcripts of downloaded discussions among interested educators who 'meet' regularly on electronic information services such as CompuServe; and general knowledge of the educational process.

While reviewing these sources of information, the focus was on identifying needs that represent an opportunity to deliver services which allow remote, shared access to scarce resources, and which free people from geographic and time constraints in the communications process; that is, the highest priority was given to identifying needs which telecommunications products are uniquely suited to serve, and which will represent stable market opportunities for telecommunications providers.

### Students

Students need to acquire specific information — they need to know basic facts about the subjects they are studying; they also need to develop a core set of skills which go beyond passive memorization of facts to include active communication and analytic abilities. In an era when the accumulation of knowledge is occurring at an ever-increasing rate, students need to 'learn how to learn'. At a practical level, students also need to prepare to enter the workplace.

This learning will occur best in a context of easy access to a wide variety of information resources; timely, and often interactive, feedback and evaluation; and individual pacing. Traditionally, access to specialized/individualized educational resources has been reserved for students with a variety of handicapping conditions, and/or children considered to be gifted; that is, students at the extreme functional ranges (for example, see Appendix A, Interview

4). The new information technologies, including computers and telecommunications, offer the possibility of greater individual tailoring of the educational process for all students, offering benefits at both the cognitive and affective (ie: motivational) levels.

Some specific student needs which are linked to telecommunications technology include:

- Access to stored sources of data (eg: library and museum holdings, statistical compilations, full-text databases, etc.) for use in writing projects, research reports, etc., and a way to easily search for and retrieve information. The data can include text, graphics, photographic images, sounds, and full-motion video.
- Access to 'experts' in a variety of academic, corporate, and even political settings, who can provide students with current information on topics of interest. Access could eventually be provided in both a time-independent form (eg: e-mail), as well as allowing for real-time interaction (eg: computer-based teleconferencing; and distance learning including audio transmission and full-motion video).
- The ability to distribute work electronically for review by other students, teachers, and 'experts', and to receive feedback on work, interactively and non-interactively .
- An introduction to telecommunications technology as a tool which many students will eventually use in the workplace.
- Access to current information about opportunities and requirements for post-secondary education; and for those who will be seeking work immediately after school, access to current information about careers, job opportunities, etc.

### Teachers

Good teachers are also 'lifelong learners', and to some degree they share many of the same needs as their students. However, the specialized nature of the teaching function within the school system requires that many teachers function in two distinct but related roles

— professional educator, and professional in a given subject area (eg: Math, Chemistry, English, etc.).

Serving in each of these roles, teachers' needs which relate to telecommunications technologies include:

- Access to current information about the potential of telecommunications and other technologies for enhancing the educational process. This need cannot be stressed enough. Many teachers are simply unaware of the new tools which are available.
- The ability to collaborate/consult with other teachers to develop specific curriculum content and standards, share ideas about educational methods, etc., involving the capacity to exchange, solicit, and distribute information, and receive feedback. Again, the specific services could eventually include non-interactive e-mail and file transfer, computer-based teleconferencing, and interactive audio/video transmissions.
- The ability to collaborate/consult with elementary and secondary school teachers, university-based academics, and others working in non-educational settings, who are professionals in the same subject area, to facilitate a teacher's professional development in a given field. Unlike many of their post-secondary counterparts, high school and elementary school teachers can find themselves operating as the only art teacher, or the only physics teacher, etc., in a school or district. Access to other professionals who share their interests will enable teachers to keep current in their field more easily.
- A convenient method for scheduling meetings and conversations, checking on and posting information about upcoming events of interest, etc.
- Access to libraries of developed and tested curriculum materials, including easy-to-use search capabilities which produce relevant materials quickly.
- Access to databases containing information relevant to a teacher's professional field of interest, once again including efficient search capabilities.

- Access to resources for upgrading teaching skills. For example, the ability to communicate easily with a variety of 'master' teachers would be useful. The communications could involve a non-interactive review of lesson plans, two-way interactive audio/video observational capabilities, or the creation of an on-line database of 'master teacher in action' videos which could be downloaded for later study.

### III IMPLEMENTATION ISSUES

In order to meet the needs described above, telecommunications providers must deliver services which help to create the 'virtual' classroom — an electronic classroom which allows students and teachers to reach out beyond the walls of the traditional school. The following basic telecommunications services and applications are or will be necessary, and will form the 'building blocks' out of which specific applications can be built:

- E-mail for quick notes, messages, etc..
- Document transfer between remote sites for more extensive communications, including graphics and text files, with the capability to send fully-formatted files, such as word processing documents with special fonts, boldfacing, graphics, etc.
- Transmission facilities capable of efficiently handling photographic images, sound, and full-motion video.
- Directory services, so that people and resources can be easily located.
- Gateways to remote databases, including flexible search and retrieval capabilities. 'Database' is defined broadly to mean commercial services such as CompuServe, private bulletin boards, libraries, museums, etc.
- Teleconferencing services, including on-line computer-based forums, and more sophisticated services such as two-way interactive audio/video capabilities, and simultaneous voice/video/data communications.

Before these services become fully integrated into the school environment, however, a number of key issues must be addressed:

- Methods must be devised for systematically educating teachers about the potential for telecommunications to enrich the process of education, and their own professional development.

A variety of media and forums, including telecommunications networks themselves, should be utilized.

- Teachers do not exist in a vacuum. Strategies for convincing school district administrators of the benefits of telecommunications services will also need to be developed.
- Teachers must be taught how to use the technologies, and ongoing technical support must be provided. Training and support materials need to be developed, including interaction with and access to 'live' experts who can guide teachers through rough spots in learning to use the new tools.
- Resistance on the part of individual teachers and/or organized groups of teachers will almost certainly be encountered. This resistance will have many roots, including: fear of the unknown and the new; anxiety about loss of jobs; concerns about the real educational value of specific telecommunications services; fear that the education process is being 'mechanized and dehumanized'; and a perception on the part of already overburdened teachers that they are being forced to accept yet another responsibility. All of these factors will need to be dealt with.
- Systematic methods for teaching students to use the technologies must be devised. These methods will need to take into account factors such as grade level of students, individual abilities and motivational factors, and special circumstances such as handicapping conditions. Teacher positions for individuals who specialize in information technologies may need to be created. Materials appropriate for different groups of students will need to be developed.
- The technologies themselves must be made easy-to-use, and must fit into the school environment in a seamless way. Teachers do not think about how to use their blackboards. They must not be distracted from their instructional role by the technology.
- The telecommunications infrastructure in schools will need to be assessed, and upgraded where necessary. At a minimum, the existence of phone lines in every classroom may be necessary.

- In the short-term, existing equipment must be supported. In practical terms, that means telecommunications support for Apple II and IBM PC personal computers is a necessity, as these are currently the most widespread devices in use.
- A variety of input and output devices will eventually need to be accommodated, including a variety of computers, specialized terminals, digitizing scanners for graphic reproduction and text recognition, fax machines, optical disks, etc.
- The role of home-based access to services for students and teachers needs to be investigated in more detail.

As one example, some students already have access to computers and telecommunications services from their homes. Equalizing the distribution of resources between technological 'have and have-nots' in the home raises a number of issues involving economics, family background and lifestyles, etc. One solution might be the development of a low-cost student workstation which could be widely sold or distributed.

'Home delivery' of actual instruction by means of telecommunications services may be questioned strongly by teacher unions. There have already been examples of distance learning projects opposed by teacher groups [get reference].

- Procedures for logons, passwords, distribution of computer accounts, etc., which reflect the educational setting, need to be developed. For example, if a class connects to a teleconferencing session, does each member of the class need to logon separately, or can some type of group logon be devised?
- There are a number of sources of educationally relevant material which are already accessible on-line (See Appendix C for a partial listing). Should educators and telecommunications providers work actively with existing sources to add to and refine their existing services, or should resources be utilized in creating duplicate services? For example, should NYSErNet create its own educational E-mail service, or should it provide a link to an existing service?

- Active relationships need to be developed between the K-12 sector and universities, colleges, corporations and other sources of 'live' experts who can contribute to the resource pool available over educational networks.
- The services must be cost-effective (ie: cheap). Lump-sum payments for services, rather than pay-by-the drink billing methods, may need to be adopted. Teachers are not given open-ended invoices by school districts.

#### IV ACTION RECOMMENDATION FOR NYSERNET — THE STAR SCHOOL LEGISLATION

As stated in the introduction, NYSERNet has an opportunity to assist in the introduction of telecommunications technology on a broad-based, integrated basis into the New York State elementary and secondary school systems. The need exists, and the time is right.

The scope of current educational needs and implementation issues indicates that no single entity can be expected to manage the undertaking. A group effort will be required. The 'glue' which could cement a joint partnering is the so-called 'Star Schools Program Act' — a bill which has recently been approved by the U.S. U.S. Congress. This bill proposes:

"To authorize a star schools program under which grants are made to educational telecommunications partnerships to develop, construct, and acquire telecommunications facilities and equipment in order to improve the instruction of mathematics, science, and foreign languages, and for other purposes."

This bill is particularly relevant for NYSERNet, given the way that a "telecommunications partnership" is defined:

- 1) "a public agency or corporation established for the purpose of developing and operating telecommunications networks to enhance educational opportunities for educational institutions, teacher training centers, health institutions and industry; or"
- 2) "a partnership which includes three or more of the following which will provide a telecommunications network:
  - (A) a local educational agency
  - (B) a State educational agency
  - (C) an institution of higher education
  - (D) a teacher training center, or
  - (E) a public agency or private nonprofit organization with experience or expertise in

the operation of a telecommunications network."

Eligible partnerships can qualify for a demonstration grant of up to 20 million dollars in a single fiscal year. NYSERNet should consider the formation of a partnership to apply for a Star School grant. In combination with the other players discussed below, NYSERNet is in a good position to apply, since the bill favors a telecommunications partnership which "...will serve the broadest range of institutions, including public and private elementary and secondary schools, institutions of higher education, teacher training centers, research institutes, and private industry". To round out the partnership, NYSERNet should also cultivate a relationship with professional organizations, such as the NEA. As noted in the introduction, the new information technologies have the potential to radically change the educational process, and teachers need to be involved actively in planning as well as implementation efforts.

Interested 'players' have been identified; they want to explore a working relationship with NYSERNet; and NYSERNet can fill a role in helping to tie the pieces together.

The New York State Education Department is taking an active role in telecommunications through the Technology Network Ties project (TNT, see Appendix A Interviews 2 & 3) — an effort designed to develop an "...integrated, statewide network linking BOCES, school districts, and other educational agencies with the State Education Department". TNT staff are potentially interested in obtaining 'bandwidth' from NYSERNet.

The TNT network is organized in three tiers — a statewide backbone, regional networks, and local networks. The basic functionality envisioned for the network includes : e-mail; file upload/download capabilities which including merging of host-based data with workstation programs like Lotus 123 or dBASE III Plus; computer conferencing; and bridge/gateway services to other networks and e-mail systems. Migration to ISDN services is planned for the future. The interface to the network is also an integral part of the overall plan. TNT staff intend to provide end-users with a transparent interface to all network services, and have already made progress in this area.

The provision of data relevant to student and financial management is currently a high priority of the TNT project. However, TNT involvement with the instructional process is also planned. The three-year old Model Schools Program (see Appendix A, Interview 3), is "...doing research to evaluate the extent to which computer technology can improve and enhance instruction in elementary and secondary schools". Teachers from 17 school districts have received training in a variety of telecommunications and computer technologies. The experience gained in running this program will help to more effectively integrate the new information technologies into the school system.

At the local level, staff at the Binghamton City School District and the Broome-Delaware-Tioga BOCES (see Appendix A, Interview 2) have expressed interest in exploring the possibility of joint projects with NYSERNet. A 'wish list' of services which NYSERNet might realistically provide in the near-term was prepared by Binghamton staff (see Appendix B), and includes access to libraries, museums, university staff, supercomputer facilities, and information retrieval services.

NYSERNet can offer a valuable service to the K-12 educational sector by providing a relatively high-speed, reliable 'pipe', and by delivering additional connectivity through links to college, university, and corporate resources, such as databases, computing power, and 'live' experts. NYSERNet should bring this level of connectivity and its networking expertise to other efforts which are underway around the State, creating just the type of partnership called for in the bill.

---

Interviewee: Dan Lake, Systems Consultant — Computer Assisted Instruction

Institution: Onondaga/Cortland/Madison BOCES

Date of Interview: 4/12/88

Interviewed by: W. Dougan, J. Lombardi

---

Dan Lake has been experimenting with word processing in the classroom since 1982, and in the process became interested in telecommunications. His perspective is that of a former writing teacher, and he views the addition of telecommunications to the writing process as providing a motivating factor for the development of good writing skills. In his own words:

"... I have left the classroom with a mission: to show teachers how to use new methods and tools to channel student ideas to real and widespread audiences. If you as ... teachers can create these channels, you can more effectively train your children to reach out with the written word into the world beyond the classroom. These channels of communication can be of benefit to both yourself and your students."

DL has been very active in the 'on-line community' on services such as CompuServe. As a result, he was contacted by representatives of the AT & T Consumer Labs in Basking Ridge, NJ, for consulting assistance on projects in the area of telecommunications in the classroom, including:

- The AT & T Long Distance Learning network — which is a Bulletin Board system linking a number of schools and classrooms, allowing text to be freely passed between participants for joint writing projects. The system uses an 800 number for access.
- Educational networking trials which involve international links.
- A curriculum project in which several teachers were commissioned to write lesson plans which involve telecommunications as an integral part of the activities.

Although primarily interested in the writing process, DL identified several generally applicable educational needs which telecommunications can help to meet:

For Teachers:

- Teachers need to be made more aware of the new computer and communication technologies, and to be educated about the potential which they have for improving the quality of instruction.
- Teachers in secondary and pre-secondary settings need access to colleagues who teach in the same content areas, in order to develop

professionally. Unlike many of their post-secondary counterparts, high school and elementary school teachers can find themselves operating as the only art teacher, or the only physics teacher, etc., in a school or district. Telecommunications can help by providing a means of communication which avoids scheduling conflicts (which are often a big concern) and which allows for a relatively rapid way of providing mutual feedback.

For Students:

- Students need access to sources of current information for research of topics and writing, including databases and 'live experts'.
- Students also need an efficient way to distribute their work and/or receive timely feedback from multiple individuals, including other students, teachers, and 'experts'. DL refers to this as developing a greater 'sense of audience'.

For Both Teachers and Students:

- Underlying the above statements, teachers and students need connectivity which includes many individuals and classrooms, and which spans diverse geographic locations, in order for telecommunications to have maximum effect. The possibility for sharing of information among teachers and students with diverse points of view and experiences will enrich the educational process for all.

DL listed several factors which he felt were necessary for the practical implementation of telecommunications activities in the classroom:

1. Cost-effectiveness is critical.
2. A lump-sum arrangement for telecommunications services is essential. Teachers are not given open invoices by school districts to spend money.
3. The software which drives telecommunications activities must be easy to use.
4. Any system implemented must use existing equipment, which in many cases means Apple II and/or IBM PC compatibility.
5. The system must be instructionally based, and the software should be customized for the classroom environment.
6. Teachers are busy people. Any system which is implemented should not detract from teaching time by causing teachers to pay attention to irrelevant levels of detail concerning the technology.

7. Active support of teachers as they adopt the new technology will be necessary. There is fear and misunderstanding to overcome, and well as general lack of information.

DL also briefly discussed the need for home delivery of telecommunications services in an educational context. DL noted that any hint that services implemented would replace or bypass teachers in the classroom would meet with strong resistance from teachers' unions at this time. He felt that there might be a market for educational services for 'homebound' students, but it was currently minimal. He mentioned several uses for home access to educationally related Bulletin Board Systems such as: facilitating PTA interactions; inviting communication from the community around school curricula, with different boards organized by content area (eg: history, science, etc.); and student access for submission of assignments, extension of classroom-based activities, etc.

#### Other Topics of Interest:

- DL has a counterpart in the Rochester—area BOCES, who is also actively involved in telecommunications in education.
- DL discussed the TNT network, a proposal to connect the 13 BOCES Regional Computation Centers. Although the proposed network will focus initially on administrative applications, instructional services will almost certainly be added.
- DL noted with interest that Apple Computer's extension of Appletalk-based networking to the Apple II series of computers will allow LANs to be created within and between schools. Gateways to other telecommunications services can be accessed from such LANs.
- DL discussed the attractiveness of a combination of CD-ROM storage of data locally, with automatic and transparent searching of remote databases if information targeted by a search is not found in the local CD-ROM database.
- DL also mentioned the possible application of telecommunications for facilitating inter-library loans, and foreign-language boards for the teaching of language skills.
- DL noted that the value of distance-learning projects which consisted of nothing more than students viewing 'talking heads' would be of minimal value at best. True interactive capabilities are necessary, but DL felt that the cost of such services might be prohibitive at this time.

- DL informed us that schools in New York State were mandated to provide 7th and 8th graders with an introduction to technologies that they would encounter in the workplace — telecommunications being one such technology. However, experimentation is already occurring at the younger grade levels. Systematic introduction at the lower grade levels is needed.

---

Interviewee 1: Doug Green, Director of Computer Services  
Institution 1: Binghamton City School District

Interviewee 2: Jim Moody, Assistant Superintendent for Management and  
Computer Services  
Institution 2: Broome-Delaware-Tioga BOCES

Date of Interview: 4/12/88  
Interviewed by: W. Dougan, J. Lombardi

---

The discussion at this meeting was focused on the Technical Network Ties Project (TNT) — a proposed educational network that will be statewide, eventually accommodating services such as e-mail, computer conferencing, data and voice communications, and full-motion video.

The TNT plan is for every building in every school district to be able to access network services. The backbone of the network will consist of connections between the 13 BOCES Regional Computer Centers and their IBM mainframes. 400 of the State's 700 school districts already have connections to a BOCES Computer Center. Within 6 months to one year, these districts will have the opportunity to connect to e-mail services.

At the present time, the services which are planned are focused on administrative support. The intent is to provide instructional services in the future. For example, one priority will be to create a mechanism for teachers to review planned syllabi, and jointly develop curricula.

Moody noted that there is an effort underway involving 5 BOCES regions serving 77 school districts to "aggressively pursue the complete vision" which is being projected for TNT. An initial \$12,000 in planning money was used to cement good relationships among the 5 BOCES — a key first step. Now, the intent is to demonstrate to the 77 school districts that useful projects can, in fact, be implemented. As an example, there is a plan to have a 'limited interactive TV' trial before the end of the year, in cooperation with WSKG TV. The interactive capability will be provided by a dial-up feature provided in conjunction with the programming. Initial topic areas include teen-age stress and sports medicine.

According to Jim Moody, when he first became aware of NYSERNet, he did not see it as being particularly relevant to TNT. His perception was that NYSERNet was strictly focused on that segment of the academic community which required access to supercomputers. However, given the possibilities which exist for NYSERNet to become an active player in the K — 12 market, he now feels that cooperation between TNT and NYSERNet would be beneficial, and should be explored in more detail. Strong interest in a pilot project involving the Broome-Delaware-Tioga BOCES, the Binghamton School District, NYSERNet, TNT, and telecommunications vendors was expressed. Green noted that such an experiment in this region might be particularly effective, given the relatively high level of computer expertise in the area.

---

Interviewees: Denis Martin, Mike Radlich  
Institution: New York State Education Department (SED) — Technology  
Network Ties Project (TNT)

Date of Interview: 5/12/88

Interviewed by: W. Dougan, J. Lombardi

---

The TNT Project is overseen by staff of the Office of Elementary/Secondary/Continuing Education, Planning and Support Services. The goal of the project is to develop an "... integrated, statewide network linking BOCES, school districts, and other educational agencies with the State Education Department", providing both management and instructional services. There are 41 BOCES districts, with 13 of the BOCES designated as computing centers. Currently, 86% of New York State's 735 school districts receive some form of computing services from BOCES.

The impetus for the TNT Project came from a Coopers and Lybrand audit and review of computing in the New York State K-12 sector, conducted 5 years ago. The primary recommendation of the report was that SED should implement a strategic plan for computing statewide. This was particularly important in order to improve the timeliness of data used for planning and administrative purposes at the SED.

The overall architecture of the TNT network (as described in documents obtained from TNT staff) consists of three tiers:

- a statewide SNA-based backbone connecting BOCES Regional Information Centers, the New York City Board of Education, and the State Education Department.
- a series of X.25 regional networks, which provide gateways to the statewide backbone by accessing the Regional Information Centers, and which provide connections among BOCES Computing Centers, other BOCES centers, and school districts within a region.
- a variety of local area networks, which will provide connectivity at the classroom, office, building, campus, and district level (although Radlich noted that the telecommunications infrastructure at the classroom level needs to be built).

The basic functionality envisioned for the network includes : e-mail; file upload/download capabilities which including merging of host-based data with workstation programs like Lotus 123 or dBASE III Plus; computer conferencing; and bridge/gateway services to other networks and e-mail systems. Migration to ISDN services is envisioned for the future. The provision of data relevant to student and financial management is currently a high priority.

The interface to the network is also an integral part of the overall plan. TNT staff intend to provide end-users with a transparent interface to all network services, including IBM PROFS. During our discussion, we were given

a demonstration of a graphics-based interface which used the metaphor of 'doorways' as access points to network services. Denis Martin, Project Manager for TNT, noted that he and his staff are attempting to specify a uniform 'window' onto the network in terms of functionality, while recognizing that the actual form of the interface may be system-dependent at the user-end.

Martin also briefly described the three-year old Model Schools Program, which is "...doing research to evaluate the extent to which computer technology can improve and enhance instruction in elementary and secondary schools". Teachers from 17 school districts have received training in:

- 'tool software' applications, such as word processing, spreadsheets, databases, graphics, etc.
- access to mainframe databases
- communications and networking
- curriculum development
- computer-assisted design
- interactive video
- desktop publishing

Part of the discussion with Radlich and Martin involved potential links with NYSERNet. Several issues were raised. Access to reliable bandwidth at a low cost is the primary concern, along with retention of network management capabilities by TNT personnel. The need for a continuing dialogue with NYSERNet focusing on the specifics of any joint project was stressed.

There is also some interest in NYSERNet because of the access to academic and corporate resources which NYSERNet might provide. However, Radlich noted that work needs to be done to create greater interest in the academic community for cooperation with elementary and secondary schools.

Finally, Radlich and Martin stressed the need to move beyond the pilot project stage to begin implementation of services

---

Interviewee: Jay Innes — math teacher, also teaches computer skills

Institution: Kendall Demonstration Elementary School and Model  
Secondary School for the Deaf, Gallaudet University

Date of Interview: 4/7/88

Interviewed by: Thomas Victorisz

---

### General information\*

KendallNET is a computer network linking schools for the deaf throughout the US and Canada. It is supported by the Kendall Demonstration Elementary School on the campus of Gallaudet University in Washington, DC. Gallaudet is financed in large part by direct appropriations from Congress, and is dedicated to advanced education of the deaf, using American Sign Language as its language of instruction.

KendallNET links users at Kendall and elsewhere on the Gallaudet campus; at 9 satellite centers located as far away as Florida, California, Oregon, and Ontario; and at remote sites. Through KendallNET, users can share information related to the CMES (Computer Managed Education System), other computer uses, and issues of importance to the education of the deaf. The network provides facilities for E-mail and supports 16 public bulletin boards which are updated at Kendall and at each satellite center every 24 hours. Users can call in to Kendall Central via Telenet, or to one of the regional satellite centers via public phone lines.

### Some bulletin boards currently maintained by KendallNET

- Satellite center news: contact persons and consultants.
- Software notes.
- Hardware notes.
- General forum: discussion of issues relevant to deafness.
- Kid notes: students from any school/program for the hearing impaired.
- Parent notes: parents and educators of hearing impaired children.
- Professional development notes: deaf education professionals to learn of professional activities.
- Teaching strategies: teachers of the deaf using computer technology.
- Learning Tools notes: Users of Learning Tools, Inc., software.

### Professional tools developed at Gallaudet

The two key professional tools developed at Kendall that teachers of the deaf can share via KendallNET are:

- ENFI: English through Natural Form Instruction
- CMES: Computer Managed Education System.

---

\* Information obtained during the interview is supplemented by information taken from KendallNET Update, June 16, 1986; and Kendall's CMES (Computer Managed Information System) descriptive folder.

The basic idea of ENFI is that if kids use English to communicate, they learn the language at a high level of motivation. Pre-lingually deaf students (who were born deaf or who lost their hearing prior to learning to talk) have special difficulty in learning the correct use of English. Beside providing a high level of motivation, the flow of communications in the form of messages over the network can be monitored and controlled for mistakes, and thus provides the raw material for effective teaching.

CMES is an educational package which is tailored to the specialized needs of education of the deaf, and which takes into account the generally limited financial means of schools of the deaf. It has three main software components:

- A curriculum management system.
- A planning system for individual students.
- An administrative planning system.

In addition, a PC-based E-mail and bulletin board system, using a hard disk, has been designed for installation in schools that serve as satellite network sites.

#### Network use

The following are some current KendallNET uses:

- Students communicating with each other and with their teacher.
- Students at one site (e.g., Kendall) communicating with students at a second site (e.g., The Sir James Whitney School in Belleville, Ontario).
- A mainstreamed student (a deaf student placed in a hearing school) communicating with an itinerant teacher.
- A student doing a student interest poll.
- Satellite center representatives sharing information.
- Family Education Coordinator communicating with other sites to plan or discuss parent education programs at their sites.
- Educational professionals seeking and sharing information on computer hardware and software.

At Kendall, many students contact the local bulletin boards; but this is not a formal program; they do it on their own. Mr. Innes is teaching networking to students in his math class. In the Fall of 1988, a new computer literacy course will have a unit on networking.

The Model Secondary School for the Deaf and Kendall work together closely and share many members of their teaching staffs. MSSD serves as a clearinghouse for deaf education. The CMES system defines the curriculum, but the courses themselves are not available on line. The school has a large collection of software programs available for loan (a printed catalog, "Software to Go," has been published in 1988) but these programs are not yet on the network, either.

On the other hand, the Gallaudet University library catalog is now accessible via modem through the central Vax of the system, both from inside and outside the university campus.

The capabilities of the existing system, both from the technical and the educational point of view, were demonstrated two years ago when a private foundation provided a grant for a pilot transmission of data to England, to a hearing high school. MSSD students exchanged data with English high school students on water quality tests following the Chernobyl nuclear power plant accident. At the end of the interchange program, a still picture of the participating students was transmitted; the kids were thrilled. Regularizing such exchanges is, for now, a matter of financial resources.

#### The potential of networking

Mr. Innes sees the following potential for expanded future capabilities and use of the network:

- For students: expanded information search capability; course data banks for e.g., math, social studies, English. Information, review, and practice materials for any course.
- For teachers: more intensive networking within a school, within a district, or with a local university, for the purpose of quickly getting information; also, expanded use of CMES software for keeping attendance records, student records, etc.
- For schools: pooling of resources and information; pooling data banks.

The next steps at the Model Secondary School for the Deaf will be:

- To integrate fully the computer into the classroom.
- To make sure that every teacher has a computer.

Mr. Innes also commented that he has not yet seen the computer fully integrated into day-to-day classroom instruction; it is treated as an extra, rather than as an organic part of education. One reason is that the computer is expensive. Especially in deaf schools, financial pressures are always great since most of them are private institutions. When he visits deaf schools, he finds it hard to press for the purchase of computers while the school is having difficulties purchasing textbooks.

Mr. Innes sees the role of the computer as giving basic information. The computer, in his view, can only supplement but not displace the efforts of the teacher, especially in three critical functions:

- To motivate.
- To discuss.
- To integrate.

What the computer is likely to displace is not the teacher. "The computer will displace the overhead projector."

Wish List\*Near Term:

- Libraries and Museums
  - search holdings
  - download descriptions and sources
  - arrange for visitations
  - arrange for on-site use or viewing
- Distance learning
  - search for resources for distance learning projects
  - download lectures and papers
  - obtain information on college course offerings and prerequisites
  - set up projects involving students and staff at participating institutions. Projects may involve one-to-one or one-to-many relationships. University staff could:
    - provide instruction plans and resources for small groups or classes
    - respond to questions using E-mail facility
    - advise a school staff member's research project
    - conduct their own research in cooperation with the school
    - supervise student interns
    - advise student projects
  - obtain supercomputer access for exemplary student projects
- Access various information retrieval systems
- Communicate with NYS Education Department
  - download syllabi and other curriculum material
  - get answers to numerous questions via E-mail system

---

\* Adapted from a list prepared by Doug Green, Director of Computer Services, Binghamton City School District

Long Term:

- Libraries and Museums
  - downloading print materials in digital format
  - downloading still or full-motion video including:
    - interactive guided tours
    - images of holdings
- Distance learning
  - interactive lectures (audio and video)
  - downloading of still and full-motion video including:
    - film, tape, and resources
    - campus tours
    - field-trip style videos of on-campus and off-campus sites

---

---

RESOURCE GUIDE FOR INTRODUCING ONLINE RESEARCH AND COMPUTER-BASED  
NETWORKS IN A CLASSROOM SETTING\*

---

---

SECTION I: ARTICLES

This bibliography covers a wide range of articles that address the concerns of educators who are exploring telecommunications through online research projects and computer networking. National, commercial services as well as local databases and BBSs are covered.

Aversa, E.

Teaching online searching: a review of recent research and some recommendations for school media specialists. *SCHOOL MEDIA QUARTERLY*, Summer 1985. Provides useful recommendations for curriculum planning and identifies areas that require further investigation in school settings.

Bigley, T. and Tapanila, G.

PC communication packages. *INFOWORLD*, Dec 7, 1987. Comprehensive feature article comparing telecommunication software programs for the IBM PC. Covers shareware as well as commercial products; lists sources of availability.

Brawer, J.

Teachers' toolbox: a guide to on-line communications services for the classroom. *A+ MAGAZINE*, Mar 1988. Describes several of the newer online services, including: MIX, from McGraw-Hill; Einstein, from Addison-Wesley; and the Learning Link, operated by WNET public television in New York.

Butler, G.

Keyboard connections. *THE COMPUTING TEACHER*, Apr 1987. Describes the world-wide network of FrEdWriter users who exchange disks using the mail for a variety of curriculum-related activities. Butler is based in Australia; a related article by Holly Jobe provides an American perspective on Keyboard Connections.

Craver, K., & Quanian, L.

An introduction to online bibliographic searching for high school students: a successful approach. *EDUCATIONAL TECHNOLOGY*, June 1984. Describes how high school students were taught to develop search strategies. The searches were executed by trained professionals.

Dodge, B. and Dodge, J.

Readiness activities for telecommunications. *THE COMPUTING TEACHER*, Apr 1987. Describes a wide range of specific activities that teachers can use to help prepare students for online telecommunications; discusses the benefits of off-line activities and describes available resources.

---

\* Prepared by Chuck Lynd, Director, Information Services, LINC Resources, Inc. Copyright 1988, LINC Resources, Inc. All rights reserved.

Ellenwood, W.

Commodore users: on line. *THE COMPUTING TEACHER*, Apr 1987. Describes a cost-effective telecommunications program using Commodore computers in an elementary school; helpful suggestions cover set-up operations through implementation.

Fischer, M.

Selecting software: a survey of telecommunications software for the Apple II family. *A+ MAGAZINE*, Apr 1987. Comprehensive survey of telecommunications software for the Apple II family; includes helpful descriptions of each program and indicates which packages are suitable for the IIs.

Kleiner, A.

"Telecommunicating," in *WHOLE EARTH SOFTWARE CATALOG*. New York: Doubleday, 1986. This section contains an excellent (albeit somewhat dated) overview of telecommunication services, databases, and networks; also offers recommendations for modems, telecommunications software, and BBS programs.

Knapp, L.

Teleconferencing: a new way of communicating for teachers. *CLASSROOM COMPUTER LEARNING*, Mar 1987. Explains and describes how both teachers and students can use computers to communicate in "real time" conferences; covers both local and national conferencing systems and discusses the pros and cons of each.

Lake, D.

Beyond word processing. *CLASSROOM COMPUTER LEARNING*, NOV/DEC 1985. Describes a variety of activities that enable English teachers to develop writing projects that utilize electronic mail and the exchange of text between classrooms across town or across the country.

Levinson, M., & Walcott, A.

Online databases--a school project. *MEDIA AND METHODS*. Sep/Oct 1985. A project for junior and senior high school students in a prep school in New Jersey is described.

Lodish, E. and Caputo, A.

Researching on line: schools and electronic databases. *THE COMPUTING TEACHER*, Apr 1987. Describes an innovative program for high school students in Montgomery County in Maryland; emphasizes Dialog's Classmate program for training students how to conduct online research using a variety of databases.

Mendrinis, R.

Computers as curriculum tools: exceeding expectations. *MEDIA & METHODS*, Jan/Feb, 1987. An excellent overview of some of the more advanced high-tech applications in schools; special emphasis is placed upon telecommunications and database projects.

Morabito, M.

Teaching telecomputing in schools. *LINK-UP*, Sep 1986. Provides an overview of telecomputing applications and options that can be integrated into a school's computer literacy program.

Riedl, R.

CompuServe in the classroom. *THE COMPUTING TEACHER*, Mar 1986. A project in Fairbanks, Alaska used telecomputing resources to support a magazine for teens. Provides a good analysis of online costs and benefits.

Schack, M.

The electronic link: computer bulletin boards. *MEDIA & METHODS*, Jan/Feb 1987. Provides an overview of activities that schools are doing with local BBS software; includes tips on managing a BBS and lists several educational BBSs that schools may contact.

Scrogan, L.

Telecommunications: how to overcome the roadblocks. *CLASSROOM COMPUTER LEARNING*, Feb 1987. Addresses the commonly reported problems of costs, availability of equipment, and teacher training; suggests a variety of practical solutions, including a listing of off-line programs that simulate telecommunications.

Scrogan, L.

The online underworld. *CLASSROOM COMPUTER LEARNING*, Feb 1988. Identifies ethical problems that arise when teachers introduce telecommunications in the classroom; describes steps teachers can take to deal with piracy, illegal use of phone numbers, and potential damage or invasion of privacy on others computer systems.

Scrogan, L.

What's new in online information services. *CLASSROOM COMPUTER LEARNING*, Apr 1987. Provides an overview of how schools can tap into the commercial online information networks; covers Dialog, BRS, Dow Jones, and the new online services for schools from Addison-Wesley and McGraw-Hill.

Slatta, R.

Setting up your own bulletin board. *LINK-UP*, Dec 1985. Practical considerations in planning for the operation of a local BBS. Brief, but includes useful references.

Solomon, G.

Electronic research. *ELECTRONIC LEARNING*, Mar 1986. An overview of how online databases can support student research goals. Includes suggested activities.

Tamashiro, R. and Hoagland, C.

Telecomputing a chain story. *THE COMPUTING TEACHER*, Apr 1987. Describes the popular language arts activity of writing a chain story using an electronic bulletin board; specific benefits for both students and teachers are outlined.

Teague, M., Teague, J., & Marchionini, G.

The high tech road to research. *THE COMPUTING TEACHER*, Nov. 1986. A media specialist and two researchers describe a project that introduced the online use of Grolier's encyclopedia on CompuServe to elementary students in Maryland. Recommendations are offered based upon the analysis of their experience.

Tennis, Jean.

Telecomputing Diary. *ELECTRONIC LEARNING*, Sep 1985 through May 1986. A series of articles about a classroom teacher who introduced telecomputing (CompuServe) to a middle school class.

Tenopir, C.

Student online data base searching--part I. *THE COMPUTING TEACHER*, Apr 1986. Describes the benefits of online searching in school settings and the equipment needed to get started.

Tenopir, C.

Student online data base searching--part II. *THE COMPUTING TEACHER*, Apr 1986. Provides information on the major online search systems that offer special alternatives to schools.

Wheeler, F.

The new ready-made databases: what they offer your classroom. *CLASSROOM COMPUTER LEARNING*, Mar 1987. Provides an overview of pre-packaged databases and outlines questions teachers need to ask before using one in the classroom. The author works in the Media Group at the Bank Street College of Education.

**SECTION II: PRE-PACKAGED DATABASE PROGRAMS.**

All of the sources listed publish databases for students ranging from middle school to senior high level. Some databases are designed to operate on a stand-alone basis, while others require the use of an existing database program such as Bank Street Filer, pfs:file, etc. These requirements are noted, and compatible computer systems are also indicated. Write to the companies for descriptive literature and current price information.

**ACTIVE LEARNING SYSTEMS**

5365 Azenida Encinas Suite J  
Carlsbad, CA 92008  
800-423-0818, 619-931-7784

Publishes USA IN PROFILE and ONE WORLD COUNTRIES DATABASE; Apple II, Commodore 64/128, IBM

**GROLIER ELECTRONIC PUBLISHING, INC.**

Sherman Turnpike  
Danbury, CT 06816  
800-858-8858, 203-797-3530

Publishes U.S. & WORLD FACTS, SCIENCE AND NATURE FACTS, TRIVIA HIT PARADE; requires Grolier's Friendly Filer program; Apple II, Commodore 64/128, IBM

**D.C. HEATH AND COMPANY**

125 Spring Street  
Lexington, MA 02173  
800-235-3565 (Information) 800-428-8071 (Ordering) 617-862-6650

Publishes HEATH SOCIAL STUDIES and HEATH AMERICAN READERS; requires the DCH Notebook Filer program; Apple II

**HRM SOFTWARE**

175 Tompkins Ave.  
Pleasantville, NY 10570  
800-431-2050, 914-769-7496

Publishes SOCIAL STUDIES FACT FINDER; requires pfs:file; Apple II

**MECC**

3490 Lexington Ave.  
North St. Paul, MN 55126  
800-228-3504, 612-481-3500

Publishes MECC DATAQUEST series, with modules called THE 50 STATES, THE PRESIDENTS, THE WORLD COMMUNITY; Apple II family

## SCHOLASTIC SOFTWARE

730 Broadway

New York, NY 10003

212-505-3000 Ordering/sales information: 800-325-6149; 800-392-2179 (in MO)

Publishes three databases: WORLD GEOGRAPHY, CULTURES, ECONOMICS; requires pfs:file; Apple II series

## SUNBURST

39 Washington Ave.

Pleasantville, NY 10570

800-431-1934, 914-769-5030

Publishes six databases collectively called the NORTH AMERICA DATABASE; requires Bank Street Filer; Apple II, Commodore 64, 128

SECTION III: OFF-LINE SIMULATION PROGRAMS

This is a source list for programs that can be used to prepare students before they log on to a database research service or telecommunications network.

ERIC/INFORMATION RESOURCES

Attn: Publications 036 Huntington Hall  
Syracuse University Syracuse, NY  
13244-2340

ERIC MICROsearch is an inexpensive database program that can be used to guide students through the principles of online searching; uses ERIC database subsets; Apple II or IBM.

EXSYM

7016 Dellwood Avenue, N.E.  
Albuquerque, NM 87110  
505-881-3670

Publishes several useful programs authored by Len Scrogan: THE ELECTRONIC VILLAGE, a communications primer and simulation of an electronic bulletin board system; THE ELECTRONIC MAILBAG, which simulates electronic mail and includes a curriculum materials; and WINDOWS ON TELECOMMUNICATIONS, which simulates the processes underlying the vocabulary of telecommunications; Apple II and IBM. A new program, ETHICS ONLINE, addresses the problems associated with the ethics of using telecommunication networks.

GROLIER ELECTRONIC PUBLISHING, INC.

Sherman Turnpike  
Danbury, CT 06816  
800-858-8858, 203-797-3530

Publishes THE INFORMATION CONNECTION, a communications primer and simulation of an online computer network (similar to CompuServe); also contains a fully operational telecommunications program for use when students are ready to go online; Apple II, Commodore 64/128, and IBM.

RANDOM HOUSE MEDIA

400 Hahn Road  
Westminster, MD 21157  
800-638-6460, 800-492-0782 (in Maryland)

Publishes INFORMATION FOCUS: A RESEARCH SKILLS PROGRAM; simulates electronic mail, computer conferencing, and database research; Apple II or IBM.

REGIONAL CONSORTIUM FOR EDUCATION AND TECHNOLOGY

Attn: The Electronic Mailroom  
10601 Clayton Road  
St. Louis, MO 63131

Publishes a BBS program that does not require modems or phone lines; Apple II or IBM.

## SOFTSWAP

P.O. Box 271704  
Concord, CA 94527  
415-685-7265

Softswap publishes several public domain programs for the Apple II that simulate online activities or function as a modemless network: SIMUCOMM simulates telecommunications software, electronic mail, bulletin boards, and an online database service; MODEMLESS CMS, designed to simulate a local BBS, allows polling and helps train young "sysops" or BBS system operators; and KIDMAIL, an electronic mail simulator that encourages writing skills, is the key to a national exchange of disks between KIDMAIL users. (To participate in the exchange of KIDMAIL disks, contact Wayne Ayers, Culver City High School, 4401 Elenda St., Culver City, CA 90230.)

SECTION IV: COMMERCIAL DATABASE VENDORS

This is a brief guide to companies that provide online research services of interest to educators.

**BRS INFORMATION TECHNOLOGIES**

1200 Rt. 7  
Latham, NY 12110  
800-345-4277

Educators may contact:  
Paula Swope, Education Services Coordinator  
BRS Information Technologies  
555 E. Lancaster Avenue St.  
Davids, PA 19807  
215-254-0236

BRS offers a strong collection of databases of interest to teachers and other education professionals, and they offer special discount rates and menu-driven search services designed for educators. Ask about their BRS/EDUCATOR and BRS/INSTRUCTOR packages. They also publish a three-part BRS VIDEO TRAINING COURSE, including student workbooks; available in 3/4 inch VHS or Beta format.

**DIALOG INFORMATION SERVICES, INC.**

Marketing Department  
3460 Hillview Avenue  
Palo Alto, CA 94304  
800-3-DIALOG, 415-858-3785

Dialog is the largest database vendor and, although they offer fewer professional level databases for educators than BRS, they provide access to many resources of interest to students. Dialog offers special discount rates for educators, especially through their Classmate Instruction Program (CIP) that is designed to train high school students in online research skills. A wide range of supplementary training resources are available, including a demonstration disk for their KNOWLEDGE INDEX databases.

**DOW JONES NEWS/RETRIEVAL**

P.O. Box 300  
Princeton, NJ 08540  
609-452-2000

Educators may contact:  
Margaret Bakes 609-520-4069

Dow Jones is a business-oriented database and news service. Educational accounts are available, however, and include an Educator's Guide in addition to the normal User's Guide. A tutorial disk with sample sessions and specific suggestions for teachers is also available for Apple II and Commodore 64/128 computers.

## ADDISON-WESLEY INFORMATION SERVICES

2725 Sand Hill Road  
Menlo Park, CA 94025  
415-854-0300

Educators may contact:  
Mark Battey, Manager, Information Services

A-W offers EINSTEIN: THE INFORMATION ACCESS TOOL, a service that provides access to dozens of databases selected because of their relevance to teachers and high school students. Einstein uses a "gateway" service that automates the online search process through the use of specially designed menus that prompt the user or student for search terms that describe the subject of interest; Einstein then performs the actual search using the appropriate databases and database vendors. Excellent training materials are available, including a newsletter with searching tips and stories about how the service is being used in classrooms. Einstein is also accessible via the CompuServe Information Service; see Section V.

SECTION V: COMMERCIAL TELECOMMUNICATIONS NETWORKS

This brief guide includes companies that offer specific services for educators. Contact the companies for current information about subscription fees, online connect charges, and availability of local network access numbers in your area. Most services will offer free demonstration accounts for a limited time period and for a specific purpose if requested on school stationery.

COMPUSERVE INFORMATION SERVICE

5000 Arlington Centre Boulevard  
P.O. Box 20212  
Columbus, OH 43220  
800-848-8199 614-457-8650

Educators may contact:

Cindy Benis, Education Products Manager  
614-457-8600

CompuServe is the largest telecommunications network and offers access to electronic mail, bulletin boards, computer conferencing, and databases of interest to both students and professional educators. Of note is the wide variety of educational forums organized for special interest groups such as students, science and math teachers, foreign language educators, Logo users, computer using educators, etc. Several menu-driven databases are available, including EPIE's The Educational Software Selector, resources for handicapped users, and an encyclopedia. Einstein, the gateway service to educational databases, is also available through CompuServe. Public domain software and technical support for all computer brands are readily available. Educators may purchase a business account that permits schools to administer multiple passwords with a single billing option.

LEARNING LINK WNET/THIRTEEN

356 West 58th Street  
New York, NY 10019  
212-560-6610

Educators may contact:

Shirley Gillette, Director of Educational Outreach

The Learning Link is a unique computer-based educational network designed to help teachers use television more effectively. In addition, the service includes topical bulletin boards for teachers and a "gateway" service that provides access to more than 100 databases. Although based in New York City, plans for national expansion are underway.

**MIX/MCGRAW-HILL**

9855 West 78th Street  
Eden Prairie, MN 55344  
800-622-6310, 612-829-8200

Contact: Steve Laliberte

The McGraw-Hill Information Exchange for Teachers, or MIX, is designed primarily as an idea-exchange for teachers. Ongoing message exchanges are organized around topics of interest to teachers. A print newsletter is available for MIX users. Special billing options are arranged to help schools manage online costs, and online "connect time" is sold in blocks of allotted time.

**THE SOURCE**

1616 Anderson Road  
McLean, VA 22102  
800-336-3366, 703-734-7500

The Source is a full-featured information service that offers electronic mail, bulletin boards, conferences, and databases. Though smaller than CompuServe, The Source offers specific services for educators, including EDLINE, an electronic news service sponsored by the National School Public Relations Association. A new Special Interest Group for educators has also been formed recently. Public Domain software and support for specific computer brands are also available. Support materials and a demonstration disk are available for training.

**SPECIALNET**

National Association of State Directors of Special Education  
2021 K Street, N.W., Suite 315  
Washington, DC 20006  
202-296-1800

SpecialNet is a computer-based telecommunications service designed specifically for special educators. SpecialNet features electronic mail, more than 50 national bulletin boards, as well as private bulletin boards and related information management services for education agencies. As the service has expanded new bulletin boards of interest to regular educators have been added, including the Education Daily newsletter. The service is aimed primarily at state and local education agencies and other educational institutions as opposed to individual subscribers.

REPORT III

**MARKET AND ECONOMIC  
IMPACT STUDY**

VOLUME IV: INDUSTRIAL CLIENTS

**NYSERNet**

NIC COLLECTION



**New York State Education  
and Research Network, Inc.**

**1095 Avenue of the Americas  
Room 1702  
New York, NY 10036  
Tel. 212-395-4480**

REPORT III

MARKET AND ECONOMIC  
IMPACT STUDY

VOLUME IV: INDUSTRIAL CLIENTS

Prepared By:

Tom Vietorisz  
New School For Social Research  
Cornell University

For

MEIS TEAM

### REPORT III: GENERAL INTRODUCTION:

During the period of the last twelve months, the members of the NYSERNet Market and Economic Impact Study (MEIS) Team produced two prior reports that identified many of the needs and problems of individuals using data telecommunications networks for education and research at universities, a paper on the relationship between, "The Economic Benefits from and Public Support of a National Education and Research Network", plus the current report. In the prior reports, the team identified specific groups of individuals who are involved in education and research activities and whose needs required additional study. The several volumes of the current report are devoted to four of those groups:

- (1) Volume I, The Supercomputing Community
- (2) Volume II, The Biomedical Community,
- (3) Volume III, Education: K-12
- (4) Volume IV, Industrial Clients.

Generally speaking, these volumes review the needs of a specific segment of the education and research community, point out current problems associated with providing network service for that segment and outline recommendations for action that can be taken by an organization such as NYSERNet to serve that group most effectively. Taken together, these volumes present a broad spectrum view of the opportunities wide-area networking organizations such as NYSERNet have for making a positive economic impact on their regions--and on the nation as a whole.

Many readers have a primary interest in one or more of the segments covered in the current report, but not all of them. In order to facilitate the distribution of the report to those individuals, we have produced each volume under separate cover.

## VOLUME IV: INDUSTRIAL CLIENTS

### PRINCIPLES OF AN INDUSTRIAL MARKETING STRATEGY

We propose that NYSErNet's marketing strategy in relation to industrial users be based on three interrelated principles:

\*TO OFFER NEW ADDED VALUE TO INDUSTRIAL USERS, in order to provide an incentive for industries to seek connection to the network--mostly as Category II users who also generate Category III connections. This means:

- New on-line services;
- Accessible in a transparent fashion.

\*TO ENSURE THE EDUCATIONAL AND RESEARCH CHARACTER OF NEW ADDED-VALUE SERVICES FOR INDUSTRY, maintaining the clearcut separation between "new" information that can be carried over the network and "current" (commercial) information that cannot:

- This is essential for complying with NYSErNet's charter, and coincides with the preferred strategy of soliciting government support (for "new" information only);
- This also satisfies the needs of some key clients (e.g., I.B.M.: see Appendix) who wish to restrict their exposure to litigation by avoiding any ambiguity in regard to their compliance with regulatory provisions that restrict NYSErNet to educational and research traffic.

\*TO OFFER NEW ADDED VALUE WITH HELP/SERVICES RELATED TO SECURITY in handling company-confidential information. Large companies can afford to handle their own access controls or encrypting, but such help or services might make a real difference to smaller industrial users in applications connected with education/training and R&D.

NYSERNet can also provide security assistance to all sizes of firms that have administrative SNA networks. That is, for security purposes network incompatibility between TCP/IP networks and SNA networks--generally the bane of the telecommunications managers--may be an asset for helping to assure security (and legality).

IBM's four internal networks (see Appendix) are being unified under IBM's proprietary, hierarchial networking scheme, SNA. SNA networks are currently incompatible with NYSERNet since NYSERNet is a TCP/IP network. Paradoxically this could make NYSERNet attractive as a partner to IBM in its networking operations. Since NYSERNet is intended for research and education traffic ("new" information) and not for traffic associated with normal business operations ("current" information), this incompatibility can greatly facilitate the separation of these two types of networking. "New" information can be exchanged over NYSERNet without interfering with or providing unauthorized access to "current" information that is exchanged over the SNA networks.

### THREE APPROACHES TO ADDED-VALUE SERVICES

The question at this point must be: precisely what new added-value services to generate? We wish to approach this question from three different directions, each of which aims to explore the needs of industry that can be reasonably addressed at different levels of affordable bandwidth. Each of the three approaches leads to specific implications for action. We will translate these implications into recommendations for concrete marketing strategies and pilot projects.

The three approaches to generating value-added services are:

(1) What we might call the "standard" approach in the context of NYSERNet--aiming to build customer demand progressively from the base of E-mail to instantaneous file transfer of KB, then MB, and eventually of GB size; and finally to two-way live packet-switched TV image transfer.

This approach assumes that bandwidth increases will follow upon massive increases in demand.

- (2) An offshoot of the standard approach aims to offer affordable service for the transfer of increasingly large datasets under realistic bandwidth constraints.

This approach (inspired by I.B.M.'s BDN and VNET services: see Appendix) branches off from instantaneous file transfer at some critical file size where a store-and-forward service option becomes attractive, with delay (as a function of file size) in the range of a few hours to a couple of days.

- (3) Another offshoot of the standard approach addresses not increasingly massive dataset transfers but increasingly intensive user interactions--such as are likely to build up in industrial education and training services.

A new technology of "connected education" is now under rapid development that organizes asynchronous electronic seminars and even "electronic campuses." This technology permits offering industry an attractive alternative to high-cost electronic classrooms based on dedicated two-way TV connections.

#### THE "STANDARD" APPROACH

The basis of this approach is E-mail; from the user's point of view file transfers of increasing size typically develop from and build upon this fundamental usage. The first-time network user must overcome a variety of resistances before he/she will be comfortable with E-mail--and some potential industrial users never will be. One of our project monitors from industry--be it said with all love and charity--to this day refuses to go on line with us by E-mail.

This is by no means exceptional. NYSEG (New York State Electric and Gas), the upstate New York utility, has installed literally thousands of E-mail connections for its managerial employees, but the system, up to now, has apparently had disappointingly little use. At Westvaco, (see Appendix) one manager in charge of a number of paper plants habitually receives field reports overnight by E-mail, but handles his answers, as well as interchanges with other executives, by voicemail (Voicecom and Dial).

Nevertheless, not everybody slights E-mail. At IBM, 350,000 of 400,000 employees worldwide are now connected and the intensity of network usage (including file transfers) is such that, according to

Dr. David Heggie, Director of Data Distribution Networks, it, "has become part of the fiber of our business and as more applications are justified to run on the network, the value continues to grow."

We believe that in regard to network usage, IBM offers a crystal ball that permits us to gaze into the future. IBM employees, for obvious professional reasons, cannot afford cavalier attitudes toward on-line communications; so no matter what their personal resistances, they have to grit their teeth and get on with it. Of course, after these initial resistances are overcome, the technology becomes hugely attractive. Our prediction is that once networking gains momentum within industry, no company will be able to ignore it, and usage will snowball at the 30-40 percent annual rates that have characterized growth at IBM. It must be recognized that these growth rates at I.B.M. are in relation to services that are "free" to the user. There is no explicit incremental charge for use of the networks.

We have addressed many of the implications of the "standard" approach to generating network usage in our previous reports. In brief:

\*IN GENERATING ENTRY-LEVEL USAGE, i.e., in the early stages of learning, the essentials are transparency of network access and provision for some handholding in coping with personal resistances.

Making a major effort to expand entry-level (E-mail) network usage is important as a stepping stone to generating more massive file-transfer traffic. Its importance as a marketing tool is analogous to Detroit's traditional efforts to capture the entry-level new-car buyer who can later be channeled toward more upscale company products.

\*IN GENERATING MORE UPSCALE USAGE, two basic types of connected on-line services play a decisive role:

- Data-access: E-mail and other directories, data banks, current information.
- Processing-power-access: supercomputers; other specialized computers (e.g., parallel processors); access to on-line desktop publishing and similar facilities.

\*ANTICIPATED FUTURE DEVELOPMENT is assumed to carry, by progressive stages, toward packet-switched live two-way TV image transfer over communications channels offering massive bandwidth.

### MASSIVE FILE TRANSFER: STORE-AND-FORWARD SERVICE

Under NYSERNet's current bandwidth availability, transmission of gigabyte-sized files, or even much smaller ones, is highly problematical. With exclusive access to a T-1 trunk, a gigabyte requires about an hour and a half to transmit. This time increases some twentyfive-fold for 56K access lines, and of course stretches out a lot more with shared access. Besides, large files congest the network.

Yet, many concrete research problems routinely involve datasets of such size, especially those that generate images requiring computerized handling. We have encountered an instance at the Cornell Medical School involving ultrasonic imaging of the eye which creates 1 gigabyte of data per run (and could easily produce a run each day).

The traditional store-and-forward technology of many telephone-based networks offers an approach to this problem which remains attractive for large files despite NYSERNet's much greater bandwidth. We regard it as a marketing opportunity for NYSERNet to offer a special service for large-file transfer, using off-peak network capacity for overnight transmission.

Following IBM's practice for internal company traffic (see Appendix), priority might be given within such a service to smaller files which would thus arrive with shorter delays. Establishing the service requires that the routing computers should have attached storage capacity for holding files until off-peak periods. As the bandwidth of the network increases, the definition of a "large" file can be expected to shift progressively upward and delay times can be expected to decrease.

## MASSIVE INTERACTION BY USERS: CONNECTED EDUCATION

The concept underlying connected education has been presented elsewhere in this report (K-12 Education Section). In the present context, a noteworthy observation is that industrial training applications of connected education are as important as, if not more important than, academic applications.

The role of education and training as a lifetime concern has been emphasized over and over again in discussions of how industry needs to adapt to the information age. Most recently, the federal government's Office of Technology Assessment has issued its report of a four-year study (TECHNOLOGY AND THE AMERICAN ECONOMIC TRANSITION, 1988; reviewed in SCIENCE, May 20, 1988). The report urges a major commitment both by the federal government and industry to research on understanding how people learn and to the development of new educational technologies.

Connected education is such a new educational technology which is exceptionally adapted to the needs of industry. While placing "an unprecedented demand on the intellectual skills and knowledge of American workers," industry is also facing an unprecedented need for both the training and the periodic educational updating of its workers. Connected education, operating in an asynchronous mode by an extension of E-mail and file-transfer technologies, is particularly well adapted to meeting this emerging need.

Connected education differs from ordinary E-mail and file transfer by the massive interactions it brings about between (1) teachers and students; (2) students and students; and (3) students and on-line teaching/training materials. The latter materials are themselves characterized by rapid turnover.

The spread of this technology will create an increasing need for on-line service facilities--that may be referred to collectively as an "electronic campus"--to mediate these intensive interactions. It is thus a natural marketing opportunity for NYSERNet to take a leading role in both stimulating the spread of connected-education technologies and in serving the needs which these technologies create.

NYSERNet may wish to do so either entirely on its own, or in association with pioneering institutions such as Connected Education of New York City or EIES (Electronic Information Exchange System) of the New Jersey Institute of Technology.

### ACTION RECOMMENDATIONS

#### "STANDARD" APPROACH:

##### \*Entry-level usage

- Generate an industrial outreach program aimed at
  - (1) connecting individual industries as Category II users
  - (2) connecting the employees of such industries as Category III users.
- Provide "intelligent node" and other service support for new industrial clients and their employees, in order to overcome resistances to the intensive use of E-mail.
- Create a "connected education" program for novice users of E-mail, with formal enrollment, instructors, teaching materials, and an "electronic campus," as a regular course with a diploma at the end, financed by the companies in question.

##### \*More upscale usage

- Implement earlier recommendations for creating on-line directories and data access facilities. The joint arrangement with BRS is one of several possible models to follow.
- Implement earlier recommendations for creating on-line special-purpose server facilities in:

computing (parallel machines);  
typesetting and printing;  
desk-top publication;  
hypercard teaching material preparation.

MASSIVE FILE TRANSFER: Store-and-Forward Service

\*Create storage, forwarding, and traffic management facilities, with user-friendly and transparent access procedures and with the provision of anticipated time delay estimates.

\*Create special access routines and transport protocols for frozen-frame TV, graph, CAD/CAM, and other image transfer.

CONNECTED EDUCATION:

\*Following Office of Technology Assessment recommendations, formulate a proposal for federal government support of a pilot project in connected education oriented to the training and educational updating needs of industries:

- How does an "electronic campus" oriented to industries have to be different from one that is academically oriented?

- What kinds of connected-education services and facilities are in most demand by potential industrial users?

- What is the proper division of functions between the federal government and the states in supporting industry-oriented connected education and its national networking infrastructure?

\*Following up an earlier meeting with DARPA, formulate a proposal for a pilot project that would bring together the many contractors on a particular project around an asynchronous "electronic roundtable."

- The method of approach here would draw on existing experience with connected education, both in the industrial and academic context.

- NYSERNet would take specific responsibility for the exploration of security considerations having to do with data transfer and access.

\*Explore the best way for NYSERNet to take a role in the organization of an industry-oriented connected-education consortium of the research universities now on line with NYSERNet plus other institutions engaged in connected education (e.g., the New School).

- The experience of the National Technological University might be a useful model to draw on. NTU is a consortium of 26 engineering schools, based in Fort Collins, Colorado, whose subscribers, mainly industrial companies, are distributed all over the United States.

### CONCLUSION

The focus of NYSERNet's thinking about its potential market among industrial companies has so far been largely on the R&D area. This section indicates that equally promising prospects exist in the area of industrial training and education.

## INDUSTRIAL INTERVIEW SERIES

Date of interview: 6/2/88, 6/3/88

Company: GENERAL ELECTRIC

Persons Interviewed: James Lommel, 518-387-6162  
Manager of the Information Center for the  
Research & Development Center, GE Schenectady, NY

Dr. Rolland Redington, 518-387-6402  
Director of NMR Research and Medical Imaging  
Research and Development Center  
Schenectady, NY

Lloyd Smith, 518-387-6319  
Director of Networking for the Research and  
Development Center  
Schenectady, NY

Interviewed by: Alan Mc Adams

A GE View of the Status of Medical Imaging

Dr. Redington is very pessimistic about the potential for broad bandwidth transmission of medical images. He makes two major points on this subject:

(1) A radiologist reads x-rays at the rate of 30 seconds per film. Within this period a judgement must be made as to whether the film represents normal or abnormal conditions and whether later films represent change from earlier films. Radiologists are remarkably quick and accurate with these judgements. They must have a precise memory for the normal and be adept at seeing discrepancies from the normal. They must also be adept at recognizing changes over time (e.g. for the degree of healing of a broken bone). It is not unusual for a radiologist to view sequentially 12 to 15 images for a single patient. It is just economically not feasible to make comparisons of this kind electronically within a similar time span given today's technology and costs. Further, with electronic imaging one loses the advantage of having a technician load the "light wall" well in advance of viewing of the images. With a single monitor on a workstation, images must be viewed sequentially in real time.

(2) Until large x-rays can be captured directly and stored in digitized form, it will remain uneconomic to transmit x-rays in any volume over networks. Chest x-rays cannot yet be captured initially and stored in digitized form. Rather, they must be captured first on film and then digitized--a procedure that is prohibitively expensive. The physical storage of x-ray films is expensive. Chest x-rays are large (14"x17"). These dimensions determine the requirements for storage facilities. Until it is possible to create a digital chest x-ray directly, it will remain economically infeasible to change the storage medium for x-rays from

film to electronic. Two forms of capture--film then digital--and two forms of storage--physical and electronic--make the approach infeasible.

Dr. Redington did identify areas where digital transmission could be economic. The most important of these is for manipulation of the image by surgeons or researchers as compared to radiologists. The surgeon is accustomed to providing individual attention and significant blocks of time to a patient; not so the radiologist. The surgeon needs to manipulate and study the patient's x-ray images to facilitate the use of his or her expert skills. Equally important, the surgeon is compensated for the skills devoted to the treatment of the individual patient. The radiologist is compensated for the rapid, accurate screening of large numbers of films.

The stark contrast in the modus operandi of the radiologist vs the surgeon also provides the basis for the next conclusion reached by Redington. Electronic transmission of the medical image as the basis for a surgical referral is likely to become widespread soon, while transmission for a diagnostic referral reading by a radiologist is likely to remain much more limited, perhaps essentially experimental.

The same analysis also establishes the most likely new product in the medical field: the surgical workstation. Images from CT's, NMR's and office x-rays are likely to be transmitted over LAN's to the surgical workstation as a support facility for the surgeon during an operation.

One can infer from the above discussion that the introduction of a chest x-ray machine that can capture a digital image is the key to the next level of breakthrough in this field. Once images can be directly captured in digital form, it will become economic to store them on optical disks. In turn, optical disk storage can make possible the changeover from the prior physical storage medium.

Dr. Redington states that today there is a great deal of talk, but very little action in the transmission of medical images over networks. Experimentation is taking place, however. For example, at Duke University Medical School, local area nets have been established connecting equipment of multiple types--CT's, NMR'S, SUN workstations, spectrometers, PC's, etc.--and transferring images throughout the network as required. Similarly, at the University of Pennsylvania Medical School internal networking is well advanced with workstations for data distributed throughout the hospital. DECNET protocols are used on this system which has been installed with the cooperation and partial financial support of Digital Equipment Corporation. To date, however, this latter installation does not provide for the transmission of medical images. Dr. Redington estimates that the costs of the workstation to support imaging today would be in the range of approximately \$50,000, not \$10,000. This is because of a combination of huge memory requirements and powerful computational capabilities necessary to create 3-D images with these workstations.

Dr. Redington commented about those medical personnel who currently do use digitized medical images i.e., those who currently make use of CT's and NMR's. He states that typically these doctors, "don't want anyone to lose any bits". This implies that with current technology they will not permit

compression of images unless the original resolution can be fully restored. CT's and NMR's require 12 bits due to the "dynamic range required". The current video storage devices are generally 8 bit digital devices. As Redington puts it, "8 bits won't hack it".

Dr. Redington and the interviewer agreed that hypermedia medical records represent a much more likely next step in the medical field. Voice, data and video transmission capability will be required to support such hypermedia records. The bandwidth that this implies is much broader than is currently used, but falls well short of that required for the high-tech medical images (12x512x512 bits).

### GE's Internal Networks

General Electric has a single internal Ethernet backbone at its Schenectady facility with subnets isolated to serve affiliated groups. The network supports 30-35 VAX's, a IBM 3081 host and approximately 250 UNIX workstations, mainly SUN's. The workstations are subnetted together for project groups and groups that are geographically contiguous. Where appropriate, such subnets are served by print servers, file servers and, currently on a trial basis, by compute servers.

The local area net backbone is a 10 megabit Ethernet. There are 17 subnets that are "bridged off" from the backbone. This has proven to be necessary especially with the subnets served by SUNservers for diskless SUN workstations. Nonetheless, GE has experienced considerable congestion on its Ethernets even after they are bridged off. This problem was encountered relatively recently. When an Ethernet is providing connectivity for terminal servers, even hundreds of terminal servers, the network is not stressed. However, the presence of a very small number of diskless workstations with SUNservers can, "bring the net to its knees". Problems are also encountered by GE with the DECNET/SNA gateways which cannot handle broadcast traffic. The solution that has been implemented is to hide these gateways behind a bridge which filters out broadcast traffic. This has not slowed transmissions significantly.

GE has experienced traffic over its Ethernets up to 50-70% of capacity. Traffic is currently running at 15-20% of capacity. These numbers contrast with those for Ethernets connecting terminal services; they run at 1-3% of capacity. Other organizations such as Michelein are experiencing similar Ethernet problems. Mr. Smith is a member of the GE Corporation's Telecommunications Council. He reports that similar problems are beginning to crop up elsewhere in GE with similar solutions being implemented. GE's use of E-mail is discussed below. The T/C council is currently investigating creating a uniform corporate wide E-mail system.

GE has several local area APPLE TALK networks at 230K. These tie together between 100 and 200 Macintosh desktop systems. The philosophy at GE is to have all its computer systems interconnected and to have no freestanding systems or networks. Macintosh users can use PHONENET to access TCP/IP networks. They connect to TELNET and use FTP. The local nets use Kinetics gateways and TOPS software. Most PC's are free standing, however.

GE is investigating the installation of Wellfleet equipment with a 40 megabyte backplane to replace the current backbone. The backbone currently is served by DECLAN 100 bridges that "are expensive and offer no migration path to higher bandwidth through FDDI or otherwise". The Wellfleet bridges are expected to offer FDDI as their vehicle to higher bandwidth, and a price approximately half that of the DECLAN 100's. SUN is reported also to be investigating FDDI directly for use with its systems.

#### GE's Remote Computing/Telecommunications

GE makes use of local and remote computational facilities. It achieves access to supercomputers at three locations, the University of Minnesota, Carnegie-Mellon University, and its internal facility at Cincinnati (Evendale), Ohio. These remote facilities are accessed through leased lines. The latter facility is accessed by two leased lines, the first at 56K multiplexed down to 19.2K using DECNET protocols; and the second, a 56K TCP/IP network providing access for the UNIX SUN workstations. GE's supercomputer is a two processor Cray X-MP with 8 megawords of main memory. Jobs requiring huge real memory are undertaken on the Cray II at the University of Minnesota. The leased line to Minnesota is a 56K TCP/IP dedicated line.

GE uses E-mail extensively. Of the 1800 employees of the R&D Center, approx 70% have E-mail addresses. There are two in-house corporate E-mail systems: A VMS system for the VAX's which is transmitted over the GE DECNET network and PROFS software which includes an E-mail component for the IBM systems. There are 2100 nodes on the DECNET network throughout the country; there are 12 locations connected through the IBM PROFS system, mainly on the East coast.

Researchers at the R&D Center do a considerable amount of work from their homes. Recently the range of usage has been between 25-35,000 minutes of dial-up use per week through 2400 Baud modems. These modems work through northern telecom PBX's that provide 2500 phone connections and approximately 2200 data connections. There are two such PBX's, one at each of two locations. GE has banks of modems which can be selected through the PBX system, but require an authorization code. As compared to the period before authorization codes were required, usage dropped by 15% when authorization was required (presumably since the user of the modem can now be explicitly identified).

GE has internal telephone directories which provide a great deal of information about all of the GE employees. This directory is no longer printed in hardcopy. It is updated weekly for changes in phone numbers and twice weekly for changes in E-mail addresses. Users can use a laser printer to create a hardcopy of this phonebook if they wish. Both E-mail and voice messaging systems are growing rapidly within GE. (The R&D Center does not have voice messaging, however, and therefore relies on E-mail.) Mr. Smith states that more of his E-mail messages are answered, and answers come more promptly for E-mail than for phone messages. He states further that he answers all of his E-mail messages, but only some of his phone messages. Why this is so, we don't know.

GE is concerned about security of its networks. For its MILNET work with the government it is necessary for authorized users to sign on to a front-end VAX system. For classified operations, networks are not used. For the existing networks at GE only information which is releasable to the public is put over the network. Most of the research information is far enough removed from "product" that it can be treated in this way.

GE does feel that attachment to the INTERNET through TCP/IP may provide further security exposure, possibly permitting access from the INTERNET to GE's internal networks. The dimensions of the security exposure have not yet been identified. A solution to this problem has not yet been found.

Most of the work done in conjunction with Carnegie-Mellon is related to joint research on computer system architecture and system software development. Access to Carnegie is done at 9.6K which GE sometimes runs at 14.2. A number of the projects are for the joint development of software. Few if any large files are exchanged over these networks. The project with MIT is also related to computer system architecture. Connection to MIT is by dial-up.

#### GE and NYSErNet

General Electric has joined NYSErNet on a limited basis. It is installing its own Proteon box so as to provide flexibility for future expansion of bandwidth. Currently it is achieving NYSErNet access through a leased line at 9.6K to the RPI node. Mr. Lommel took pains to clarify that GE is not a "node" on NYSErNet, nor does it intend to be in the future. It uses its interconnection to NYSErNet to achieve access to the vast INTERNET networking facilities, but it does not "...want to accept the responsibility of retransmitting NYSErNet traffic to other locations. The universities are doing a good job of that". Rather, GE wishes to be a "dead-end node".

GE has joined the Cornell Theory Center as an Associate. The motivation for this came from their computational fluid dynamics group which was seeking access to Cornell's Hypercube. GE's approach to remote computing is to optimize to the existing network bandwidth. In the current situation this implies that only commands and small files will be transmitted over the 9.6K line. Mr. Lommel indicated that GE's research personnel are, "clever at using what's available". This means that they structure their problems to achieve significant preprocessing and postprocessing in order to minimize telecommunications transport. There appear to be rules of thumb used by managers within GE in relation to communications costs. An example given by Mr. Lommel to illustrate: if given \$100,000 for computing, \$10,000 for communications would be OK; however, \$30,000 would not.

GE's approach to its connection to NYSErNet is initially to provide the 9.6K line, to expand that to 14.2K when the traffic warrants, and then to further increase the bandwidth when demand builds up. Mr. Lommel stated that it would greatly facilitate GE's use of NYSErNet if it were possible to charge users directly for their use of the network. He stated that it was part of the GE corporate culture to have specific charges for all activities and to minimize corporate-wide "utilities". The interviewer noted that this does not preclude establishing "local utilities". For example, networking managers at

GE found that users were dialing up Carnegie-Mellon and "costing us a lot of money". It was at that point that the leased line was put in place. This is a good illustration of the way in which GE information managers function. It is likely that in the not too distant future the connection to NYSERNet can be used in place of this leased line, thus generating further internal savings to GE.

The advantages that GE sees for tying to NYSERNet focus on achieving broader access to the field of university research and are not limited to NY State. GE seeks access to a national research and education network through its connection to NYSERNet. It delayed joining NYSERNet until it was satisfied that nationwide connectivity would be achieved. They hope to remove the "geographic barrier" to interaction with universities throughout the country..

When the interviewer provided a brief outline of IBM's experience with its four networks and pointed out that all of these are "free" to the user, Mr. Lommel suggested that such a system would not be possible at GE. He noted with envy the rapid growth of usage in IBM's networks and volunteered similar envy for the worldwide corporate networking systems of Texas Instruments, Inc.

## APPENDIX

## INDUSTRIAL INTERVIEW SERIES

Date of interview: 5/12/88

Company: IBM

Persons interviewed: Jordan Becker  
Manager of Laboratory Automation  
Research Division  
T.J. Watson Research Center  
Yorktown Heights, NY  
David Heggie,  
Director, Data Distribution Networks  
White Plains, NY

Interviewed by: Thomas Vietorisz

#### IBM R&D Activities and Need for High-Speed Packet Switching Networks

Data transmission from the Brookhaven synchrotron to Yorktown at 9.6 KB is now unsatisfactory. For materials research, IBM has available 2 beam lines on Brookhaven synchrotron (beam lines function like ports on an accelerator). Data are generated at the rate of 1 MB per run. Now 1-2 runs are undertaken per day; with higher bandwidth, over NYSERNet, this could increase to 5-6 runs per day.

Moreover, with higher bandwidth other things would also become possible. Data analysis and graphics on IBM's large systems could be handled remote from Brookhaven. Now, Jordan says, "we just transmit to Yorktown."

NYSERNet is expected to make an important contribution to IBM's interaction with universities. For example, a cooperative project is being undertaken with the Cornell physics department, involving electronic structure calculations for various materials: semiconductors, materials for integrated circuits, silicon, GaAs, superconductors.

IBM also has an extensive need for simulation and modelling of mechanical devices: disk heads, disk surfaces, devices under design. This kind of development work is often done in cooperation with universities, involving faculty and an exchange of services.

In terms of networking, these activities require E-mail, file transfers, and remote logon to computer facilities.

#### R&D versus Commercial Network Traffic

NYSERNet has a need to separate IBM's R&D traffic, which can be legitimately carried over the network (e.g., between an IBM research lab and a

university) from commercial traffic which NYSERNet under its charter cannot carry (e.g., between two IBM locations).

IBM has equally good reasons to make sure that a clean separation of the two kinds of traffic is accomplished, because the company is at a high risk of litigation whenever restrictions on the use of some asset exist and questions can be raised about compliance. IBM has a company policy for such cases involving not networks as such but all intellectual property; the policy is to pay whatever it costs for IBM to have free access.

Thus in the case of a packet switching network such as NYSERNet-- where individual packets are not opened at nodes for inspection and accounting--IBM would prefer to see a monitoring or allocation system that would enable the company to pay for a certain bandwidth that it might then use for any purpose.

Mr. Jordan illustrated his point by analogy to an allocation system that IBM uses for large computers. Allocations are stated in terms of the number of MIP-minutes allowed per 3-month period, and usage is debited against these allocations whenever there is no spare capacity (when there is spare capacity, it can be used freely by anyone). If a user goes over his allocation, his bandwidth will be restricted. This takes the responsibility for capacity utilization off the back of the network manager and puts it on the user; it is then up to the user to contract for a large enough allocation.

#### Other Networking Activities and Resources

IBM supports connections to about 50 inter-enterprise networks for E-mail and database access. The company is currently interested in BRS (B... Reference System ??) and looks to the NYSERNet connection for BRS access.

IBM has taken a key role in a recent NSF project to integrate NSF's regional high-speed packet switching networks in the Northeast. Meritt (?) at the University of Michigan is the main contractor, with IBM and MCI acting as subcontractors.

IBM also has an internal voice-and-video communications network that is used, for example, for transmitting lectures, other training materials, and conferences; and a large internal packet switching data network system called VNET, managed by Mr. David Heggie. The rest of the IBM interview material was obtained from Mr. Heggie.

#### IBM's Internal Data Network (VNET): Overview

VNET (from VM or "virtual machine" network) is one of the largest, if not the largest, data transmission network in existence. It had its origins in 1972 when employees at the Scientific Center in Cambridge, Mass., wanted to reduce delays in getting their paychecks from Poughkeepsie. The two sites devised a computer link for sending messages back and forth. More sites joined the hookup, and the corporate job network was born.

I.B.M.'s current data communications system as a whole is often referred to as "VNET," even though that name is also used more narrowly for one of the

system's major components. There are four such major components, each serving different functions. The components had originally been designed in partially incompatible ways and are only now being gradually moved toward a common SNA protocol. They are:

(1) The Bulk Data Network (BDN), one of I.B.M.'s two Data Distribution Networks. BDN carries very large files (over 1 MB) in a store-and-forward mode, with an average time delay of 24-48 hours. BDN has 250 stations around the world.

(2) VNET proper, the second of IBM's two Data Distribution Networks, with 3000 stations in 235 cities and 55 countries. It is used for the transmission of files of a few hundred bytes to some 100KB, but in no case to exceed 1 MB. This network also operates in the store-and-forward mode, with transmission times generally in the range of minutes to hours. Priority is given to smaller files, so if a file approaching the upper limit is sent, it may sit around until the weekend. Large files can be sent more expeditiously if broken up for transmission by the user and reassembled at the destination.

(3) PVM or "pass-through VM". Mr. Heggie calls this a "pseudo-interactive" mode of operation which works through the same machines as the store-and-forward VNET network; 2000 of the 3000 VNET nodes are now equipped to handle PVM. The user "dials up" his destination, i.e., he establishes a path, link by link, following a menu on his PC screen. Response time, depending on the path, is 2-20 seconds. PVM is used heavily by persons on trips who wish to dial in to their home hosts for mailbox access or the remote use of other facilities.

(4) Consolidated Corporate Data Network (CCDN), the U.S. field transaction network, with a 2-second response time. This is IBM's "flagship" network, built with the latest company products and designed to display IBM capabilities. Terminals give logical access to over one hundred applications ranging from order entry to time card processing. The network works with a leased T-1 backbone, 56K trunks, and 9.6K access connections, some of which are still analog; it covers the entire U.S., with gateways to Europe and the Far East.

Of I.B.M.'s 400,000 employees worldwide, 350,000 are now connected with a network address. The increase in the number of connections over the last year and a half has been 150,000, or almost 50% per year; this obviously represents the closing phase of connectivity expansion, because at this rate every employee worldwide would be connected up in four months.

The number of network nodes has over the last five years been increasing some 30 percent per year, a rate barely beginning to slacken off. The weekly traffic over the main trunk line now runs at 30 gigabytes for file transfer (2.7 million files of 11K average length) plus 6.8 gigabytes of PVM traffic (180 thousand sessions of 38K average length). Mr. Heggie says that over a 5-year horizon from 1986 to 1991, planning for the two Data Distribution Networks assumes a three to fivefold expansion of traffic, corresponding to a yearly growth rate of 25 to 38 percent.

## VNET: Technical Detail

(From documentation provided by Mr. Heggie.)

As a file makes its way through the network to its destination, the store and forward services for RSCS traffic come into play. RSCS is the "Remote Spooling Communications Subsystem" which serves traffic that is not using the pass-through (PVM) feature of the same facility (PVM: see point 3, above). The store-and-forward services allow files to be sent to intermediate systems that assume ownership pending receiving site availability. The transport facility provides routing, maintaining network directories in support of connectivity requirements.

Expanded connectivity is provided as the backbone, originally a bisync network, migrates to SNA. An SNA implementation establishes a common transport protocol for the two Data Distribution Networks, VNET and BDN (BDN: see point 1, above). This permits them to utilize, at a minimum, the same backbone links; the VNET backbone has, accordingly, been re-named "DTN" (Data Transport Network). With SNA, interconnection to other networks also becomes feasible. Still another advantage is the installation of common software across networks, for consistency in network operations and management.

The move to SNA is a migration during which both bisync and SNA traffic will be supported. Beside increased connectivity, another incentive to convert to SNA is the option of bypassing the store-and-forward process, by establishing "Thru sessions" for heavy traffic patterns. While store-and-forward has its pluses, it adds to transit time and storage requirements. SNA thus provides some extra flexibility in resource management.

The processors at VNET host sites range in size from S/36 or S/38 to a 3090, with a 43X1 used most frequently. The 43X1 is the processor used at backbone sites. This minimizes the space required for backbone operation, since these systems are air cooled. The operating system used in approximately 70% of the sites is VM. Operating systems at the remaining sites range from S/36 with POWER to MVS with JES2 or JES3. The backbones are all VM, with an MVS system and MVS backup system for central network management control, measurement collection, and report generation.

There is a management system in place for the backbone network, ensuring that service levels are met across the backbone. The backbone sites are remotely managed and operated. There is standard hardware and software at the backbones. As the management system is being extended beyond the backbone, greater conformity is being promoted and more is expected in terms of service agreements by those who connect to the backbone. For backup purposes these sites are expected to support connections to two backbone sites.

Backbone links and links to the backbone will be predominantly 56KB. These links will be encrypted. The direction is to continue to improve network service performance, with the appropriate security, for end-to-end user satisfaction.

## VNET: Applications and Benefits

The list of business processes using VNET embraces practically the entire gamut of IBM corporate functions, from financial systems and business planning to development laboratory work, engineering design, manufacturing automation, or international distribution.

A basic function of the network is to make existing information more accessible. This can be achieved in part by formal means, such as IBM's Technical Information Retrieval Center (ITIRC), and partly by informal means, such as intelligent networking via E-mail.

On the formal side, ITIRC--which is accessible via VNET--maintains massive files of technical articles and reports, many of them proprietary. According to IBM technical information and communications staff quoted in an internal publication made available by Mr. Heggie, "The challenge is to use what has already been done to make people more productive"; and "There's less need, much less need to reinvent the wheel now. Quite often, someone has tackled the same problem before." (Think magazine, November 1986, article on VNET entitled "Very Special Delivery," p.32.)

Beyond formal search of internal company databases, informal networking via E-mail is a highly effective way of identifying existing knowledge, especially in a huge organization such as IBM. The article in Think magazine cited above gives an example of an IBM employee in Australia who was trying to win a contract for an IBM mainframe that had to operate in a DEC environment at the university of Queensland. Needing help, he sent out a plea over VNET. Responses started coming in almost immediately and kept coming for several months, eventually adding up to a 500-page collection known at IBM as the "U.Q. Notebook." Not only was the contract won, but a valuable reference source was created in the process for anyone faced with a similar marketing problem.

Mr. Heggie offered an example of VNET application from another area of IBM business. Within the company, "any plant, any lab" has become the engineering to manufacturing systems strategy. Development labs and manufacturing plants are located worldwide; therefore it is necessary to get the designs to manufacturing locations which may be remote from the locations where products are designed. Rather than mail documents around the world, the design engineers transmit the designs, which have been created with computer assisted design tools, across the network. Data communications thus allow the entire company to function as though it had one plant and one lab.

Productivity gains from effective data communications are achieved in faster product development, shorter planning cycles, better response times to customers, or earlier entry into new markets, and a series of other areas of qualitative and quantitative improvement. Annual benefits to the company have been minimally estimated at 150M dollars (for 1986) on the very conservative assumption that only a 1% productivity gain is achieved by VNET. (Article in Think magazine, cited above, quoting George Kettell who at the time headed the VNET project office in White Plains).

Actual benefits are likely to be much higher. Yet, by now the use of VNET has become so integral to the way IBM is doing business that separating the benefit attributable to the network, from the benefit attributable to other aspects of the company's functioning, is considered difficult if not impossible. "VNET has become part of the fiber of our business and as more applications are justified to run on the network, the value continues to grow."

## INDUSTRIAL INTERVIEW SERIES

Date of interview: 11/30/87

Company: WESTVACO CORPORATION  
Park Avenue, New York, NY 10171

Introduced by: Harold Sutphin, Vice President

Persons interviewed: Girard L. Calehuff, Manager  
Facilities Planning  
Fine Papers Division

G. Ronald Brown, Ph.D., Technical Assistant to Vice  
President  
Corporate Research Director

Interviewed by: Thomas Vietorisz, MEIS team  
Gerd-Ulf Krueger, The Conference Board

## Accent on Quality

\*"The paper machine is god. Quality information is more key than anything; production rate is secondary."

This has become the key operational principle during the last ten years. Before then, quality was secondary. The increased strength of the dollar at that time brought much stronger international competition based on quality.

\*Quality of product is determined by:

- Pulp (wood species)
- Weight
- Shade
- Brightness
- Finish

\*Increasing attention is now given to quality by species of wood. A mix of species is generally used: hardwood pulp for uniformity, softwood pulp for strength. A new mill in Canada is using pulp from as many as four kinds of hardwood. Some of the low-density hardwood such as cottonwood "is not best." Eucalyptus pulp is now very highly regarded.

Westvaco has excess pulp capacity and sells market pulp; but it also buys quite a bit of pulp, depending on kind of pulp (by species) needed.

\*Westvaco reviews all quality complaints by customers. There are three problem areas in printing applications of fine papers by customers:

- Paper
- Ink

- Printing machine

Sometimes, for a large customer, considerable time and money may be spent on trouble shooting.

Labor and Computerization

\*Computerization is not a threat to employment; its aim is not so much reduction of employment as the upgrading of quality.

\*When Westvaco introduced computerization, it also reorganized the work structure. Previously there had been 7 skill levels; this was reduced to 3. There is now a much wider span of responsibility for an individual worker. Yet, while the number of skill levels was reduced, the number of skilled workers has not contracted.

\*Reorganization was made easier by a special feature of the work process on paper machines. Whenever a paper break occurred, even under the old system everybody was used to working together cleaning it up, including salaried workers. During the clean-up process, assignment to detailed individual work stations was irrelevant. Thus, run without detailed work station assignments all the time.

\*Mr. Calehuff expressed his feeling that operator skills were generally underestimated. "They learn the ins and outs of the process fast; they buy their own personal computers to understand the process better." (We double checked to make sure that he was referring to ordinary production workers!) Westvaco nurtures these positive attitudes by employee involvement programs which have been set up at several mills and at one of the research labs.

\*Westvaco workers have at least a high school education; there is no hiring below this level. Workers get a very involved training program before they start. The company also developed a fairly good dynamic simulation for a 90 million dollar boiler; workers learned to "fly the boiler" on the simulator before startup.

\*Three quarters of the operators do well. The jobs are complex; people who have filtered up to the top operating jobs are good. Sometimes, a worker does not take a promotion that requires a higher level of control, or even steps down!

\*The level of computerization is such that the plant will run itself unattended for a long time, but over an extended period, e.g., more than 24 hours, the plant conditions will begin to drift.

A skilled operator looks for and controls trends. An anecdote was recounted to us about a chief operator who noted a drift in the ash content of the paper and, failing to find an explanation after checking all possible sources of change on the paper machine itself, walked over to the pulp mill to find out what was going on. Actually, a small change had been made on the previous day in the foamer used in the pulp mill; gradually, this worked its way downstream into the characteristics of the paper produced.

There are still too many open loops in the control logic which allow such drifts to occur unchecked. It is not possible to run blind; a skilled operator expects the unexpected. There are also paper breaks which must be cleaned up manually; and there are changeovers between different products. For efficient production, one hopes for no more than one change a day, but there can be as many as 3 or 4 changes when business is bad and runs are short.

## Networks

\*Computer networks play an important role in the operations of the company. Mr. Cahuff estimates that as many as a couple of hundred low-level networks are installed in mills and offices, coordinated by higher-level networks.

\*Mainframes serve individual plant locations and corporate functions; they are interconnected into a network and mutually back up each other.

\*Separate networks are used for:

- Production control and internal communications. The internal system is all-IBM, utilizing a number of leased lines.

- Inventory control and customer contact. The latter system runs on its own computer and is completely separated from other company communications for security reasons, so that hackers will not be able to break into the company's confidential records or tamper with production and quality control systems. GEISCO (General Electric) network facilities are used for reaching customers who are linked to the system by Hayes 2400-Baud modems.

\*Three aspects of networking at Westvaco will be further discussed below:

- (1) Networks for process control
- (2) Networks for internal communications
- (3) Networks for customer contact

### (1) Networks for process control

\*Computers handling process control at each plant location generate large sets of production and quality data which are needed for analysis. This information is dumped into a main frame operating at the plant location.

\*Within each mill, there are one or two (higher-order) local networks providing for the needed connectivity. All process control information is handled at the same geographical location where the process is undertaken.

\*Plants are located at:

- Luke, MD
- Wycliff, KY
- Covington, VA
- Charleston, SC
- Tyrone, PA: a "dry" plant, no pulp mill; small, only 2 machines; but very advanced in terms of networks.

(2) Networks for internal communications

\*All local mainframes at manufacturing plants are tied into a corporate mainframe (with backup at Springfield, MA). For most mills, the local mainframe is backed up by another main-frame elsewhere.

\*Westvaco R&D labs and corporate divisions are likewise electronically linked through the corporate mainframe.

\*The system supports E-mail and file transfers. Mr. Calehuff receives daily production and status reports from local mills overnight which he checks on an IBM PC first thing every morning.

\*Yet Mr. Calehuff, like most executives at corporate headquarters, does not himself use the E-mail system for communications with other executives and with managers in the field. For this purpose, Westvaco maintains for Mr. Calehuff's division a telephone operated voice-message system, "Voicecom," and a similar system, "Dial," for some of the other divisions.

\*CAD is used by the design division of the corporate engineering department which is located at corporate headquarters. They design a capital expansion plan and specific projects and transfer drawings from headquarters to individual plants. They also interact with outside engineering firms.

Designs can be transferred from one location to another by the computer network, but it is so time consuming that CAD tapes are mailed; in outside interactions there is the additional problem that the CAD standards used are not always the same.

(3) Networks for customer contact

\*Shipping warehouses guarantee customers a shipping time of 5 shipping days. It is traditional in the industry to work with a sequence of order schedules every 30 days.

\*Westvaco has a proprietary system, WesTrak, for customer interactions. The system runs on an IBM 4381 whose operations are isolated from data flows over the process control and internal communications networks of the company. WesTrak is used internally only in so far as divisions make purchases (e.g., of pulp) from each other.

\*Customers need only establish contact with a sales office via telephone and modem in order to connect to WesTrak. The system has the following capabilities:

- It allows the customer to generate and electronically transmit to Westvaco purchase orders. The WesTrak purchase order module is a software program furnished to the customer on an IBM-compatible floppy disk.

- The system generates and electronically transmits to the customer a shipping manifest as the shipment leaves the Westvaco warehouse. Thus, the customer is informed ahead of time what/when will arrive at his shipping dock.

- The system generates bar-coded shipping labels attached to outgoing merchandise. When the shipment arrives, it is scanned with a portable bar code reader at the customer's plant, and the information is electronically transferred to the customer's computer which compares it with the manifest already logged in.

\*The WesTrak network also updates inventories in Westvaco's warehouses.

- An addition is under preparation that will permit customers to inquire about specifically what kind of paper is available at any time at each warehouse.

- Electronic invoicing is further down the road.

\*There is a move in the industry to give the customer quality control data. No indication was received in the interview as to how such data transfers might be integrated with WesTrak.

#### Research and development

\*Westvaco's R&D labs are linked to each other and to corporate divisions and manufacturing plants through the corporate mainframe. Technical memos are being sent back and forth over this network using IBM protocols.

\*R&D activities are directed at (1) product improvement; and (2) process improvement. The first type of activity is now the more important one, undertaken in support of the operating divisions. It aims to develop specialty products for customers.

\*The issue arises: Is this a mature industry from the R&D point of view?

Mr. Calehuff considers in response that "we tend to lurch"--meaning that periods of sudden change alternate with periods characterized by more static conditions. For example, the main effort in process control used to be to keep the product uniform lengthwise along the roll of paper as it came off the paper machine, but there was very little in terms of sensors and automated controls to maintain uniformity sideways across the paper sheet as it was

being formed on the machine. Then cross-machine control became an important issue 10-15 years ago.

\*A related issue has also been addressed: Is there increasing market segmentation in this industry, and thus a need for R&D to support it?

Mr. Calehuff suggests that the nature of the market is undergoing a change, but not in the direction of increasing segmentation. Printing presses used to run at 1000 ft/min; this has now been raised to 1500-2000 ft/min, together with an increase in desired printing quality. This has put a premium on the uniformity of the product as an absolutely essential characteristic. R&D efforts must be directed toward improvements in this respect.

\*No specific instances of R&D interactions by Westvaco with academic institutions have surfaced during the interview; a need in this respect may, however, exist. We raise this issue because Mr. Calehuff noted that some of the very highest quality papers that are being manufactured in Germany cannot be replicated here; academic cooperation in R&D might well offer a way to break through this limitation.

In this regard, Westvaco's rigorous separation of the internal and external data transfer networks is interesting. For security reasons, it might in general be advantageous for business firms to channel academic R&D cooperation activities through separate computers dedicated to external interactions.

Such a policy might also be advantageous for NYSERNet or to the proposed national network, because it would make it easier to ensure that the educational and research network is not abused for branch-to-branch operational data transmissions of a business.

## INDUSTRIAL INTERVIEW SERIES

Date of interview: 1/6/88

Company: LOUIS BERGER GROUP OF COMPANIES  
100 Halstead Street  
East Orange, NJ 07019

Persons interviewed: Derish Wolfe, President  
Nicholas J. Masucci, Vice President,  
Louis Berger International, Inc.

Carter Brandon, Director  
Development Economics Group  
Louis Berger International, Inc.

William A. Rosenberg, Archaeologist and  
System Analyst, Cultural Resource  
Group, Louis Berger and Associates Inc.

Interviewed by: Thomas Victorisz

The Louis Berger Group, headquartered in East Orange, NJ, is one of the top two international firms in engineering construction, design, and economic consulting, with operations all over the United States as well as 60 countries abroad, doing several hundred million dollars' worth of business per year. Telecommunications functions are basic to the firm's day-to-day operations; thus it is of great interest for the MEIS project to find out how these functions are organized and undertaken.

The key interview on telecommunications was conducted with William A. Rosenberg whose background is in archaeology and who has gone intensively into computer work over the past five years.

Federal legislation concerning the preservation of archaeological and historical materials (NEPA, 1966, as amended) and state legislation of a similar character passed by a number of large urbanized states, mandates impact statements in the case of many major construction projects that might affect artifacts, generally more than 50 years old, of archaeological or historical importance. Databases of cultural resources are being set up in a number of states, and a federal resource database is now being organized. This whole field gives rise to projects with significant funding in many of which Louis Berger & Associates have been and are actively involved.

The firm has no centralized computer/information division; individual areas of operation have a high degree of autonomy, but Mr. Rosenberg sees an increasing need for coordination. For telecommunications, the firm now uses Telex, Telefax, and to a lesser extent MCI Mail.

Telex works with dedicated terminals and pay-by-the-message. Telefax is now at the point where anyone can buy a Group 3 (9600 Baud) Telefax

machine which sends compressed bit-mapped images. The compression routine can be done by a special board set in a personal computer; this enables one fax setup to send images over the phone to another. There is a choice of image resolution; in standard mode it takes 15-30 seconds to transmit one 8-1/2 x 11 page. At higher resolution, depending on resolution and image detail, it takes longer.

An internal network for the company has recently been set up by Mr. Rosenberg, in the form of a message and bulletin board system, initially for the 8-9 U.S. offices of Berger's Cultural Resources Group but available to other divisions as well. The system uses a PC-XT host and Fido software. Mr. Rosenberg is preparing to expand the system for general company use all over the U.S. and world-wide; on a trial basis he has already exchanged messages and files on the system between headquarters and operations in Seoul, South Korea. He expects substantial savings over previous communications media.

The Fido system was originally put together for hobbyists by a California programmer named Tom Jennings. It used to be in shareware; the latest version is a two-part system combining mail and bulletin board. Though still made available for free to amateur groups, it is now commercially sold for about \$150--a very low price indeed.

According to Mr. Rosenberg, for businesses Fido represents a major, totally under-utilized resource. Installed on a personal-computer host (IBM compatible, DEC, Rainbow; but not Macintosh), it runs unattended under control of a batch file. The system operator generally does a few minutes to an hour of daily maintenance (file cleanups, etc.) Mr. Rosenberg activates the system only at night and on weekends; during office hours he uses the computer for other work.

The Fido user interface is simple, with good help facilities; after a ten-minute introduction, I was able to sign on, read and write messages, and download and upload files, even prior to reading the excellent ten-page documentation. Fido software underlies the amateur public network, Fidonet, maintained by IFNA (International Fido Net Association). The organization of Fidonet--with regional nodes that communicate among each other at the lowest night-time rates--has served as the model for Berger's company network. Use of the system by businesses for internal networking is spreading--but slowly, as there is no distribution organization, no advertising, no fancy packaging, etc. (Business use of the Fido system is documented in InfoWorld, 4/21/86, in two short articles.)

#### DISCUSSION OF THREE KEY ISSUES FOR MEIS

##### (1) The user interface of the Fido internal business network

In our first report, we have emphasized that the transparency and ease of use achieved by a Macintosh-like graphic user interface is critical to expanding applications of high-speed networks, since the typical user is interested in getting to the utility he wishes to connect to and cannot be bothered with complicated communications procedures.

The question arises: how successful can a Fido-based internal company network be if it does not have a transparent graphics interface? Berger's internal company network will furnish a good test case on this question as the network is implemented worldwide. Will training costs be reasonable or as high as advocates of graphic interfaces claim?

As a test case, Berger's network is particularly valuable, because it shows an excellent adaptation to Berger's requirements and company philosophy (see Key Findings section). Thus, it will be interesting to see how successful it is in attracting use by key management, engineering, and professional personnel in the U.S. and abroad.

In particular, it will be interesting to observe if the IBM OS-2 based graphic "Presentation Manager" (expected for the end of 1988) will come into use with Fido; and how fast it will be adopted as a feature of the Berger network if and when the software allows for it. It is noteworthy that since the OS-2 system will require AT-compatible PC's or better, Mr. Rosenberg is already encouraging a company-wide upgrade from XT-compatibles; in fact, he is laying the groundwork for leapfrogging the AT and moving directly into PC's with the 32-bit 386 chip, such as the IBM clone Compaq, now selling in the \$ 5000 range.

(2) The low speed of Berger's internal Fido network, and complementarity between low-speed and high-speed networks from the user's point of view.

The issue is this: assume Berger's low-speed, low-cost worldwide internal network is a success, and that the business use of this type of network spreads rapidly over the next 1-3 years. Where does this leave NYSERNet and other high-speed networks?

Spreading use of low-speed networks may (1) whet the appetite of users for greater bandwidth; or (2) create entrenched habits standing in the way of change to higher-speed options. Mr. Rosenberg and I agree that the former is likely to be the case.

If so, we can regard the spread of Fido systems in business as paving the way toward the use of high-speed networks such as NYSERNet.

(3) Complementarity between low-speed and high-speed networks from the geographical point of view.

A discussion of Berger's global traffic patterns with Mr. Rosenberg indicates that the really massive data transfers (data-bases, large sets of engineering drawings, etc.) are likely to occur within rather than between major world regions. The reason for this is that the generation and implementation of major projects is largely decentralized to such regions. Much of the data transfer in connection with such projects is likely to take place between regional headquarters and decentralized locations.

An example of such massive datasets is the large number of engineering drawings often needed for implementing construction projects. (The construction of 250 miles of railroad in Nigeria required 1500

engineering drawings of culverts and other track detail.) When such drawings are electronically transferred, it is possible to restrict the file transfer to the parameters from which the drawings are generated by CAD equipment; in fact, the Fax card which is installed in PC's for compressing pixel-by-pixel information for Telefax also has the capability of handling files of drawings using the AUTOCAD standard.

Despite the compression, thousands of drawings still represent massive datasets measured in at least the tens of megabytes. So do large databases collected for engineering and economic project evaluations. The need for transferring these between regional headquarters and decentralized project locations is likely to generate increasing pressure for the use of high-speed communication networks within the regions.

Between headquarters and the major regions--or between any two of these--data transfers are much more likely to consist of relatively short messages and files of moderate size. Berger's evolving Fido-based low-speed network is thus likely to remain adequate in both the short and the intermediate term for serving the needs of these worldwide data transfers between major regional centers.

We conclude that for organizations operating in Berger's decentralized pattern--clusters of local activities exchanging very large datasets with regional headquarters in major world regions, while regional and world headquarters exchange only moderate-sized datasets among themselves--complementarity exists between high-speed and low-speed networks. This pattern of activity calls for high-speed networks within rather than between major world regions.

## INDUSTRIAL INTERVIEWS SERIES

Date of interview: 5/20/88

Company: XEROX

Persons interviewed: Larry Pace  
Director of Human Resources Planning

Interviewed by: Alan Mc Adams

When Xerox thinks of NYSERNet, it thinks of it only in the context of basic, publishable research that requires a supercomputer. It was for this reason that the liaison person for interactions with NYSERNet was the Director of Research Computing. Xerox has a very limited budget for basic research totalling between \$100 and \$200,000. Currently Xerox scientists do some work at CNSF. They use dial-up for their network connection. This limits them to 9.6K baud, a bandwidth that does appear to be sufficient to their current needs. NYSERNet, thus does not appear to them to be cost effective.

From our discussions with the Human Resources personnel at Xerox, however, we come to the conclusion there is a significant future opportunity for interaction between Xerox and Cornell through NYSERNet. Even under current conditions there are a number of interactions outside of research that are capable of being catalyzed. These would require a significant marketing effort, however.

In the immediate future, Xerox anticipates activating an internal fiber optic network to replace and augment its multiple Ethernets that use coaxial broadband cable. As our contact person stated, "Once we get our fiber network in and people begin to make greater use of graphics, there will be a much greater opportunity to work with NYSERNet."

There are a number of impediments to an effective interaction between Xerox and Cornell. The most important of these is the question of security. This comes up in a number of contexts. Xerox has an internal TCP/IP network between its East coast and West coast facilities, a network that to date has been successfully insulated from external access. Xerox is worried that once they interconnect to the NYSERNet/NSFNET networks that their internal network would become vulnerable to security penetration. Further, there is some question in the minds of Xerox people about what would happen to information they make available to an open research/education facility such as CNSF. The presence of many industrial firms on site as affiliates of the CNSF, appears to be giving them pause. In addition, they recognize that TCP/IP networks have never been known as secure networks in any case. This later point is especially bothersome when Xerox people contemplate use of a network such as NYSERNet for marketing, sales or other proprietary information in any context.

### The Future

Perhaps fortunately for themselves, Xerox, sees the future as one in which businesses will continue to require hard-copy information. Yet they

also recognize the revolution in electronic processing of data and information. In the not too distant future, Xerox anticipates that they and their competitors will be offering multi-function devices: for digital scanning, thermofaxing, electronic manipulation and editing as well as copying. As such equipment becomes available and its cost falls to a reasonable range, we can anticipate a need for high telecommunications bandwidth for rapid exchange of information in increasingly sophisticated form. It is likely that Xerox as an organization that is an early adopter, will make use of this equipment internally. Having once experienced the quick, easy transfer of graphics and text in its internal operations, Xerox will be likely to adopt the same approach for all of its interactions with the outside world. This is the opportunity for NYSERNet.

#### Xerox's Product Development Process (PDP)

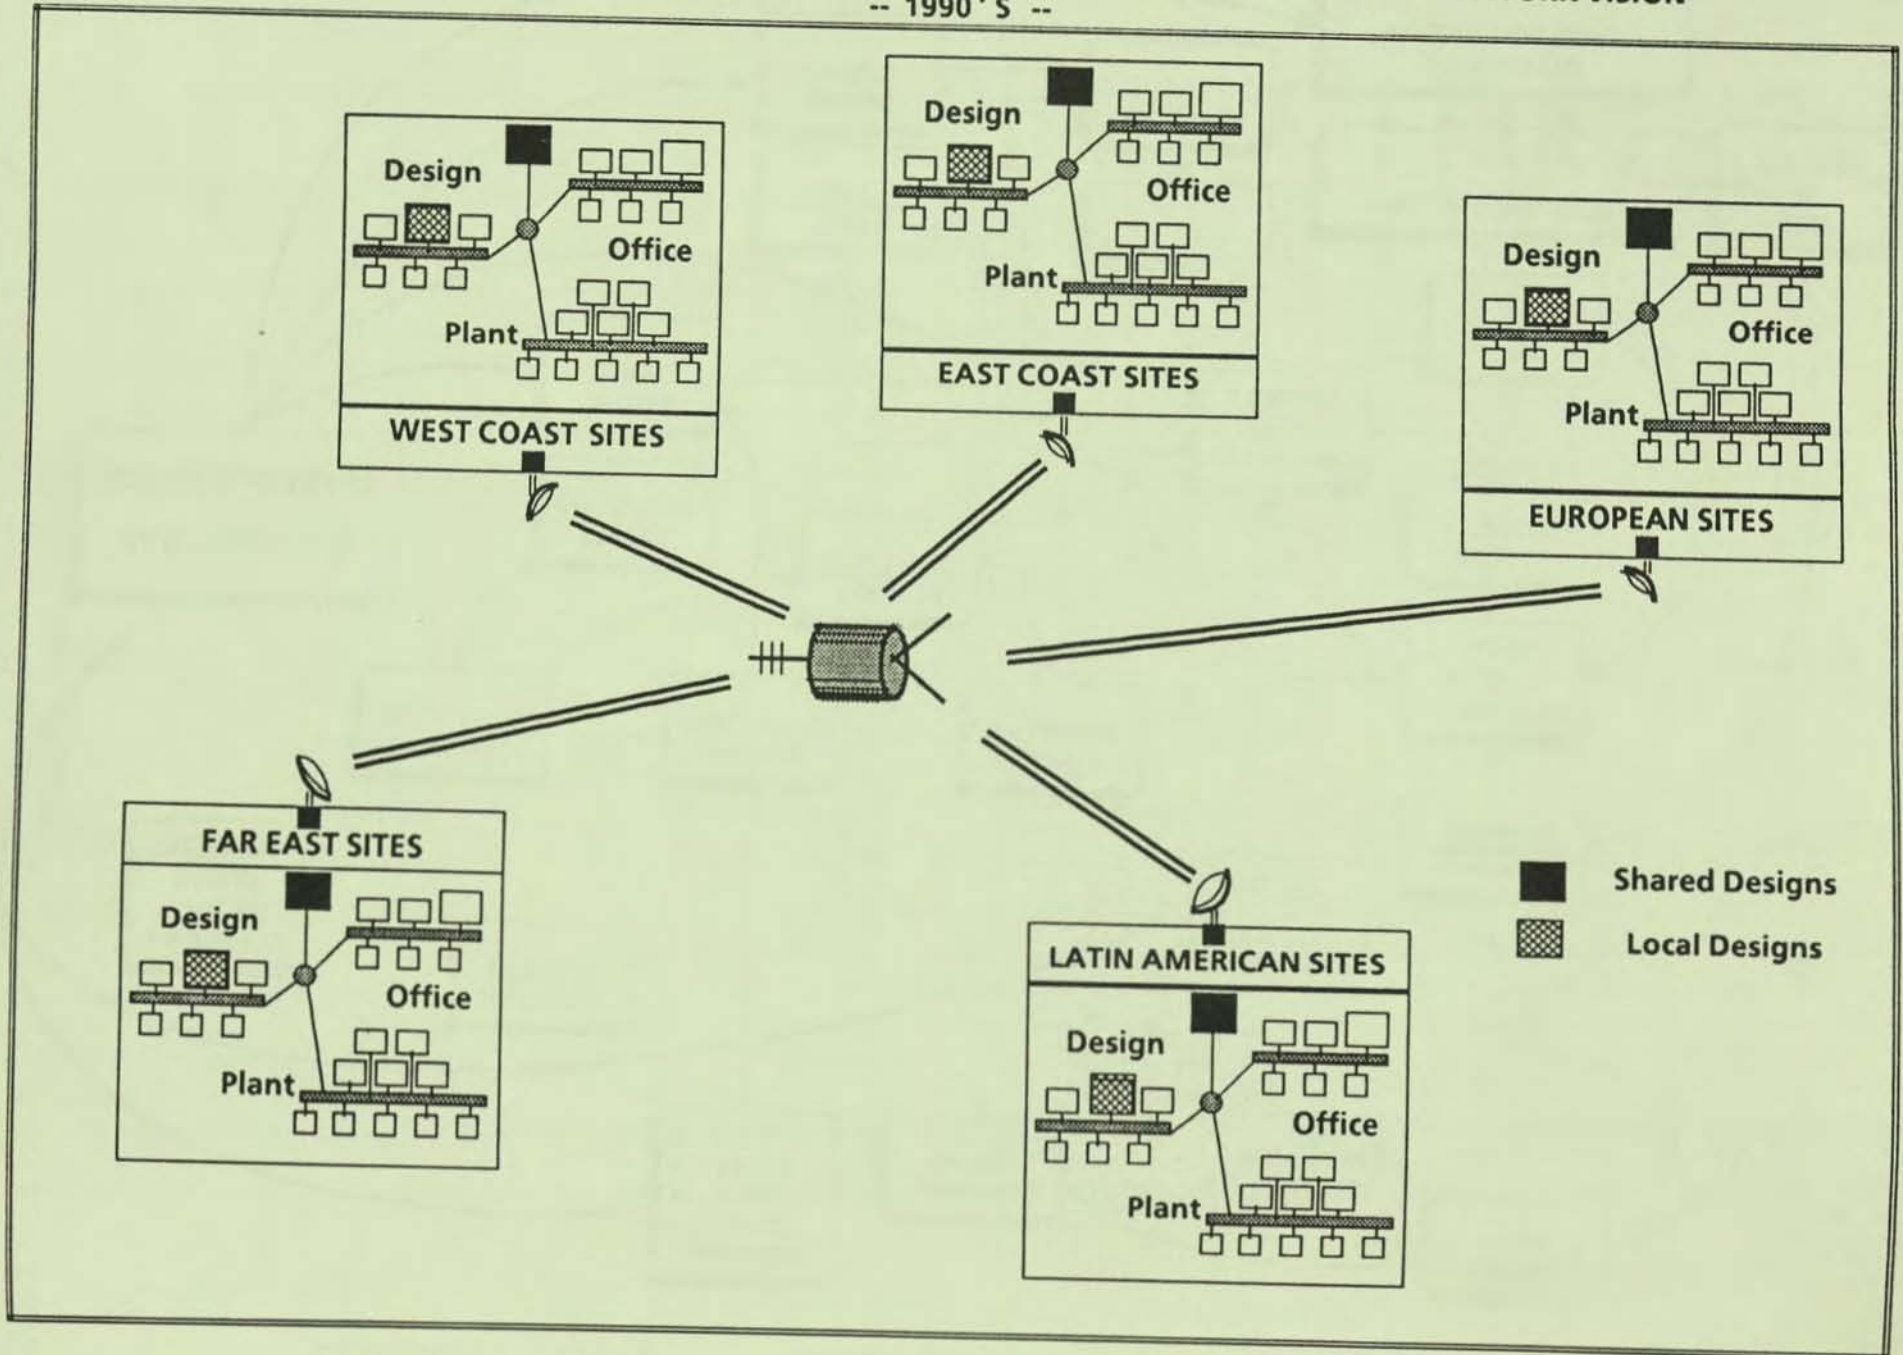
Xerox has one of the most sophisticated product development processes in the country today. The process is fully "staffed" with both personnel and resources. For each product under development a "Chief Engineer" (C.E.) is assigned along with a core group to shepherd the product from conception, through first customer-ship, to product "upgrade" (end of product life). At each stage, teams with overlapping competence are pooled together. Formal reviews are made at each benchmark point to insure that development to that stage is complete before moving to the next stage. An objective of the process is to be able to freeze the design prior to the beginning of manufacturing. Engineering changes are to be concentrated during the development stages. Engineering resources are committed disproportionately during these stages to make possible the achievement of these twin goals: optimal design and uninterrupted efficient ease of manufacture.

Exhibit one represents the current state of internal networking at Xerox. Exhibit two outlines the desired state of internal networking envisioned for the early 1990's.



# Desired State

XEROX COMPUTER INTEGRATED ENGINEERING & MANUFACTURING INFORMATION NETWORK VISION  
-- 1990'S --



## SUPERCOMPUTER INTERVIEW SERIES

Date of Interview: 4/25-26/88

Company: MICROELECTRONICS CENTER OF NORTH  
CAROLINA (MCNC)Persons Interviewed: Alan Blatecky,  
Director of University Relations, MCNC  
Research Triangle Park  
North CarolinaPat Richardson  
Communications System Data Network Supervisor,  
MCNC  
Research Triangle Park  
North CarolinaPeggy Montgomery  
Video Operations Supervisor, MCNC  
Research Triangle Park  
North CarolinaSamuel Averitt  
Manager, Communications and Hardware  
Computing Center, NCSU  
Raleigh, North CarolinaDr. Jim M. Fikry  
VBEE Administration  
Industrial Extension Service  
School of Engineering, NCSU  
Raleigh, North Carolina

Interviewed by: Alan Mc Adams

There are two important networks at the Microelectronics Center of North Carolina (MCNC). The first is the high bandwidth of two-way, full motion, video network. The second is the 45 megabit high-speed data network. Each of these networks is intended to connect six major educational institutions in North Carolina. The six are :

University of North Carolina at Chapel Hill  
University of North Carolina at Charlotte  
Duke University  
North Carolina Agricultural and Technical College  
North Carolina State University  
Research Triangle Institute

The networks do perform their intended function of providing a sufficient critical mass for establishing significant microelectronics design capability in North Carolina.

On the educational side, the television network permits the offering of a broad range of graduate courses in microelectronics. The objective eventually is to integrate the course offerings of the various institutions to permit graduate students the richest possible educational experience. To date, the offerings have been approximately 12 per semester. However, with the exception of the originating institution, the numbers of students who make use of courses through the network is quite limited. For example, at Duke, the usual enrollment is 1-3 students. The modal enrollment at any remote site is 3.

There are a number of advantages to all of these institutions from the educational network interaction. Probably the most significant is that it clearly differentiates the universities in this region from peer institutions. The presence of courses from multiple institutions of higher learning is a powerful selling point for prospective students and faculty. Second, in situations such as occurred in the semester just completed, it is often possible for a given institution to substitute for a course offered on the network for a local offering that was or is difficult to staff. This occurred in the case at Duke in that the computer science department at Duke accepted a course on Construction of Compilers offered by Chapel Hill to fulfill the requirement for such a course at Duke. Since the course was advertised to graduate students at Duke as fulfilling this requirement in the department of computer science, 12 graduate students took the course for credit through the network. Since there was a large number of students taking the course at Duke, the department assigned a course T.A. for the Duke students. The T.A. worked intensively with the students at the local site to the great satisfaction of all concerned. This is an example of what can be achieved through the two-way television educational process.

From the point of view of the computer science department at Duke, the compiler course was a great success both as an educational experience for the students and as a resource conservation step for the department. The faculty member who ordinarily would have offered the course locally at Duke was on sabbatic leave and did not need to be replaced.

This instance illustrated the factors which optimized the effectiveness of the real time television educational opportunity. 1) The receiving school had a real need for the course through the network. 2) The transmitting school had highly qualified well accepted instructional capability for the course being offered. 3) The students were provided by the receiving institution with the appropriate "signal" that the course was of a quality appropriate to their educational experience, i.e., Duke indicated that the course would be accepted as meeting distributional requirements for its computer science department. This instance contrasts with what to date has been the usual situation. Usually students are made only generally aware of the availability of courses on the network by the institutions that could receive

them. Courses are generally not included on the menu from which courses meeting distributional requirements must be selected. For both students and faculty, the network course is looked upon as a possibly interesting elective. Given that the total graduate student enrollment in computer science at Duke is 60, it is not surprising that the number of students selecting such an elective is small.

The above contrast in response rate by students to the network offering provides a powerful lesson for those who seek to optimize the contribution of the network to the educational program in microelectronics. The three key elements required to achieve this are now clearly understood. When they are present, the contribution of the network can be optimized. When they are absent, small numbers of students at receiving institutions should come as no surprise.

### The Quality of the Educational Experience of Courses Through the Network

A course offered throughout the network is different than a course offered through a standard classroom situation. It appears, however, that the educational experience can be quite comparable. "Something" is lost with the loss of direct human contact. There is a significant difference in the way in which written material is presented during the classroom portion of the course. Nonetheless, the educational experience is comparable.

What are these differences? First, there is no blackboard. It has proven virtually impossible to transmit information to remote site through standard approaches as use of a blackboard. For both the local and the remote sites, an overhead camera is used to replace the blackboard. That is, the instructor will present written material on 8.5 by 11 inch paper, photographed from directly overhead and displayed on the TV monitors. Note that this occurs both at the local and the remote site. What students get at both locations is identical. This requires some adjustment by both the instructor and the students. Once again, the medium is different but the results are comparable. What is lost is the record on the blackboard which is continually available to the students (until it is erased). What is gained is the ability to present written material of much greater sophistication and precision. One instructor commented that the freedom to select complex diagrams and equations from multiple text sources and to have them immediately available through the overhead camera proved to be a great educational boon. A second major advantage identified by this instructor was that material presented live by creating notes on 8.5 by 11 paper constituted a permanent record of what had been transmitted. (It is not erased at the end of the period--or before) This instructor is so pleased with his permanent record that he tends to make it into a formal textbook over the summer for one course and into a technical note for the other.

There are additional advantages/disadvantages from the differences between the network course and the course in the usual classroom. From end-of-course questionnaires, the reaction of a number of students to the network

course is that it is much better organized and much more formally presented. The instructors appear to be better prepared. Students report a downside to all this however. They state that "things sometimes go to fast" and that instructors appear to jam a significantly greater amount of material into a one semester course.

Students at remote sites did not appear to be inhibited from asking questions during a classroom session any more than were those at the originating site. Since the network provides two-way video communication and requires that the receiving site "sponsor" the course locally, including providing local camera operators. The questioner at the receiving site can be seen by both the instructor and the students at the sending site. Figure 1 shows a diagram of a typical TV classroom with the instructor at the front of the room, flanked by two 42 inch projection monitors. Three small monitors are available to the instructor on the podium; one to show what is going out over the network, one to show what is available through the overhead camera, and one to show students at the remote site(s). Any of the screens can be split four ways. Also facing the instructor is the main front view camera plus a large monitor, generally showing students at remote sites. From our discussions with a number of instructors, we learned that each has "tricks" for involving the students at remote sites. Often the remote sites will be polled for questions, with polling taking place alphabetically. As in normal classrooms, it is also possible for all sites to have persons speaking simultaneously, and this sometimes happens.

From the point of view of the instructor there are also additional advantages to teaching on the network. First is the feeling that one is "doing one's part", by originating a course from one's own institution in exchange for those that are received. ( There is no financial incentive provided.) Second, this is a new, challenging and exciting experience. Third, this may be the only way that a critical mass of students can be assembled to make it economic to offer the given course. Fourth, on another dimension, the instructor is provided with greater flexibility. If a student has missed a class, the instructor is not called upon to provide "private instruction"; the student can be directed to a video tape. Similarly, if the course had previously been offered on the network and the instructor finds himself/herself unable to meet a scheduled class, it may be possible to substitute a video tape from a prior period offerings. ( In fact this was done by a young instructor who had assisted one of the giants of computer science, Fred Brooks, in offering a course the prior year. The young instructor was able to substitute a video tape of a Fred Brooks class for one of his own and Professor Brooks agreed to answer questions in person for the fifteen minutes at the end of the session.) Fifth, it may be possible for colleagues to provide assistance even in real time. For example, a young instructor found himself in difficulty with a complex concept during a network transmitted class. A colleague who happened to have his video monitor tuned to the network, observed the problem, descended to the stairs, joined the class, and resolved the problem to the great satisfaction of all. In fact, a similar incident occurred at a later time in the same course. (We were fortunate enough to have heard about these incidents both from the instructor in question and from others who were aware of them. All parties spoke of these instances in very positive terms.)

There are also disadvantages. First, it is a great deal of work. There is no doubt that offering a course through the network requires significantly greater preparation. Second, there is some risk since this is a new medium. Even experienced instructors are not assured of success in presenting a course this way. Third, the excitement wears off relatively quickly. Fourth, you really never know who actually is watching the network and where; there is the potential for "big brother" to be watching. Fifth, the problems of coordinating the delivery of materials to multiple sites in a timely fashion are significant. (MCNC is assisting the local sites financially to permit them to have a network educational coordinator to see to the myriad of problems and details that arise.)

### Video Entrepreneurship at NC State University

While the other sites on the network deal generally with only the real time broadcast of course, NC State has an extension program that is highly dependent on video tape distribution. The School of Engineering has its own "industrial extension" program that deals largely with video tape. This group attempts to achieve multiple goals through its television activities. It participates in four different programs:

- 1) MCNC
- 2) The National Technical University (NTU)
- 3) Instructional Television Fixed Service (ITFS)
- 4) Video Based Courses

The above discussion applies generally to MCNC type course whether originated at State or at another site. NTU offers graduate education in engineering nationwide, generally, to businesses on their premises. This is done through satellite transmission. ITFS instruction is essentially a subset of the NTU activity. It is generally limited to small, medium size companies in North Carolina and thus is not directly competitive with NTU offerings which are generally made under contract to national organizations. State also offers credit courses and short courses to organizations of all sizes, directly through its own efforts.

In developing new courses for credit, State tries hard to make these offerings available through more than one of the above channels. For example, it will select approximately 24 courses to be offered for credit through video tape, 12 in each of 2 semesters. From these offerings, it will identify 5 or 6 each semester to be offered to NTU to be used on the national satellite network. In turn, NTU selects from the offerings of its various participating institutions, those which it decides to broadcast. For the courses it selects, NTU will provide an honorarium of approximately \$800 to the instructor and then beyond a certain minimum, will provide an additional fee per student enrolled. In addition, NTU makes similar payments directly to State, also with additional fees per student beyond a given minimum.

State will not offer its short courses to NTU for a very good reason. These short courses are a gold mine. They support the the for credit offerings by generating revenues in excess of their costs. The short courses are developed in the following way. First, a topic of interest is identified. Then a

person competent to offer such a short course is also identified (not necessary from NC State itself). Once the instructor has agreed to offer the course, then State will provide financial and technical support for its development. The usual initial stipend for development of a short course is \$2000. The instructor must develop a complete study guide for the course and make that available to the producer of the TV presentation. Jointly, they will select those elements which are most appropriate for the video portion of the course, leaving others for inclusion only in the study guide. The production team is also prepared to provide "field shots" as appropriate. If this is a first time video experience for the instructor, the production staff.....The instructor is required also to prepared and present a 2-3 minute promotional "plug" for the course. This serves several purposes, not the least of which is to permit the instructor to see himself/herself on tape. Usually this approach is all that is required to acquaint a first time video presenter with enough information to overcome any minor problems of presentation. The instructor is compensated further by 35% of the revenue net of production costs that are achieved from the short course. The more successful of the short courses generate up to \$10,000 to the instructor. NC State retains the other 65% and recoups its full production fees. The video section generates approximately \$600,000 per year to NC State. These funds are generally plowed back into new production efforts. The value of the raw tape in inventory is also approximately \$600,000 without consideration of any value added.

NC State uses professionals in the preparation of its videos. Now there is a role for students but only under the supervision of full time professionals. The staff that supports these operations numbers 10, 4 of whom are exempt personnel. With the exception of approximately half of the total personnel and the physical space provided, the operations are self-funding. This includes the highly sophisticated production equipment that is required for the operation. The physical facilities involve two video classrooms, plus a video conference room. Each of the classrooms is configured generally as described above. One classroom seats approximately 40, the other approximately 60. There are occasions when each is filled to capacity. However, given the small, average number of students who generally enroll in networked courses at receiving sites, State has found that the small conference room is greatly to be preferred for use by these students. There are a number of plans for additional small video classrooms to accommodate up to 12 or 15 persons. Funding is available for a number of these rooms which are referred to as "satellite" facilities.

#### Live Broadcast vs. Videotape

Depending upon the site of the discussions and the persons involved, we heard different perceptions of the relative merits of live network courses as compared to course work provided through videotape. After attempting to sort out the conflicting appraisals, we offer the following hypothesis. Each is more effective in particular circumstances. The case for videotape is perhaps more easily made. For the types of students who are the clientele of the engineering industrial extension service, videotape is greatly to be preferred. This clientele is made up of full time employees who have professional responsibilities competing with their educational efforts. For such people, to arrange to be present at a fixed time and place to participate in a live broadcast of a classroom activity can be extremely confining and may often be

impossible. Such clients greatly prefer the flexibility provided by the videotape. For these persons the substance of the materials on the tape is what is important and not the presentation. They will find a time to view the tape and absorb its message.

On the other hand, full time students have as their highest priority the pursuit of their course work. They are accustomed to being at a particular place at a particular time for their courses. They also have become acclimated to interaction with instructors and their classmates. This latter point is very important. Often this can be the most rewarding portion of the classroom experience--and its extension to the preparation and follow-up to class. For these people the videotape can be perceived as barren while the ability interact in real time with an instructor and with classmates provides a significant added dimension of stimulation. For such people, there is no substitute for the live classroom experience.

### The Data Networks

The data networks tie together the major research campuses identified above. The development of this network has been gradual leading to the 45 megabit pipe size among the main campuses. The timing of the development is as follows:

January, 1985	12 megabit
June, 1986	25 megabit
March, 1988	45 megabit

As the bandwidth was increased, equipment required for the lesser level of bandwidth was shifted outward in the network. The current 45 megabit pipe size provides multiple T-1 lines. To date, only a modest portion of the available capacity is in use, and some T-1 lines are used exclusively for development work.

MCNC has been serving in a catalytic role. It has had an enormous advantage for this function in that it has had significant funding available for it to permit it to provide resources directly to the universities. This has been a mixed blessing. The universities actual resent the fact that MCNC has resources available in a manner that is totally outside of their control. On the other hand, MCNC has found itself powerless to "require" the universities to take any given action. It has consistently functioned through persuasion and example. The MCNC network is a star network. The center has been able to provide the capability in-being to each of the universities only to find the universities were not in a position to interconnect. One university questioned the need for the network at all. The situation has changed very rapidly, however. The university just mentioned now has the largest number of LAN's in place as Ethernet.

The objective of the data network is to permit joint research to go forward among these universities. This is only beginning to happen. However, each of the universities has excellent communication now with MCNC itself. Faculty members of each university commute to and from MCNC and are now able to function effectively from either end of the network connection with their home institution.

We visited three campuses and were able to discuss their networking activities in some detail at each campus. Our impressions necessarily were fragmented and new information provided new perspectives throughout the visit.

Perhaps the most impressive single installation is that of the computer science/computation center building at UNC Chapel Hill. Norman Vogel of the computer science department was our host. He is a truly dynamic and impressive person. We joined a discussion by Dr. Vogel with a group from the SRC, the semiconductor research corporation. Vogel was providing a detailed demonstration of the communications scheme for the building. Copies of his transparencies are included as Appendix Y. In a nutshell, the building is "wired" to a higher degree than virtually any other facility with which we are familiar.

Vogel is attempting to make use of Ethernets to a degree and for purposes beyond which they were designed. As a result, his installation is experiencing some unusual and transient difficulties. The resolution of these difficulties is expected to advance the state-of-the-art.

From later discussion with technical persons elsewhere in the system, we learned that the computer science building is "an oasis" on the North Carolina campus. So far it appears that there is little coordination of communications function by persons in various authority positions.

This situation contrasts rather strongly with that which we discovered at NC State University. We had an extended discussion with Sam Averitt who is manager of communications and hardware at the Computing Center. Averitt is a no-nonsense guy with an excellent reputation among his peers. From his position he has been able to bring about some order on the campus of NC State. Virtually the entire campus is now wired through broad-band cable. This was installed only after an experimental facility had been established and tested within the Computing Center. Most of the current needs of the campus, including the production origination, reception and up-link of televised courses can be met through the broad-band. Averitt is currently designing an initial experimental fiber-net again to be tested over a significant period prior to any final decision on the introduction of fiber elsewhere on the campus.

(The state of North Carolina has made available a site adjacent to the current campus with acreage greater than that of the current campus for the development of the next stages of NC State.)