Technology) report from OSTP is titled "A Research and Development Strategy for High Performance Computing. The NSF management overview of NSFNet goals mentioned the "workplace without walls" but the main motivating thrusts were access to supercomputers, central data bases (e.g., GenBank), the MOSIS facility for telemanufacturing, etc. Other uses for general research communication, collaboration, routine resource sharing are mentioned but do not capture primary attention. This balance of emphasis was all reemphasized in discussions with congressional representatives. Current federal funding covers the supercomputer effort and the NSFNet backbone but additional funds (estimated variously at $25M - $100M per year) to support a general National Research Internet (NRI) have yet to be argued for.

13.1.2 - NSFNet Status

The NSFNet consists of a national backbone network coupled to a number of regional networks (currently 11) that link research institutions with each other and with the backbone. The backbone is currently operating with 56 Kbit/sec lines and expects to move to full T1 (1.5 Mbit/sec) links in July this summer. The experimental network is being managed by MERIT, Inc. (a nonprofit consortium of 8 Michigan universities set up long ago in conjunction with the development of the Michigan timesharing system for IBM mainframes). In fact, IBM and MCI are research partners in the NSFNet project. The backbone will link the 6 current NSF supercomputer sites and 7 regional academic computing networks. NSF has put up $14M over 5 years and the State of Michigan has put up $5M from a strategic fund. IBM and MCI are contributing "materials and services".

There are longer term goals of making the configuration of the network more dynamically controllable depending on traffic needs and upgrading lines to DS3 service (45 Mbit/sec). The regional networks are not funded in this grant and many are facing funding crunches later this calendar year. Many regions are interstate and have bigger political barriers to cooperation, as compared to intrastate regional networks like BARRNet or NYSERNet -- although even these are coming under severe funding pressure. Some of the telephone companies (AT&T and MCI) think these services should be paid from from end user fees while many of the university people feel the regional intercampus links should be like the "national highway system" and supported out of general funds. An interesting suggestion was made that an analog to COMSAT might be set up for the NRI, since it will involve a close (nonprofit?) cooperation between federal, state, industrial, and academic interests.

13.1.3 - ARPANET Link

We continue our advantageous connection to the Department of Defense's ARPANET, managed by the Defense Communications Agency (DCA). A recent map of the ARPANET topology is shown in Figure 2. This connection has been possible because of the long-standing basic research effort in AI within the Knowledge Systems Laboratory that is funded by DARPA. Until the advent of NSFNet, ARPANET is the primary link between SUMEX and other university and AIM machine resources, including the large AI computer science community supported by DARPA.

Major changes are underway in the allocation and management of ARPANET resources. DARPA is seeking to control the funding it allocates to network operations and does not see itself as having the mandate of organizing or running an NRI. Effective May 1 of this year, the whole southern ARPANET route (1/2 of the total line miles!) was eliminated, forcing 6 major Texas universities to seek an alternative link with the Texas regional network (SesquiNet) for national

ARPANET Geographic Map, 31 January 1988

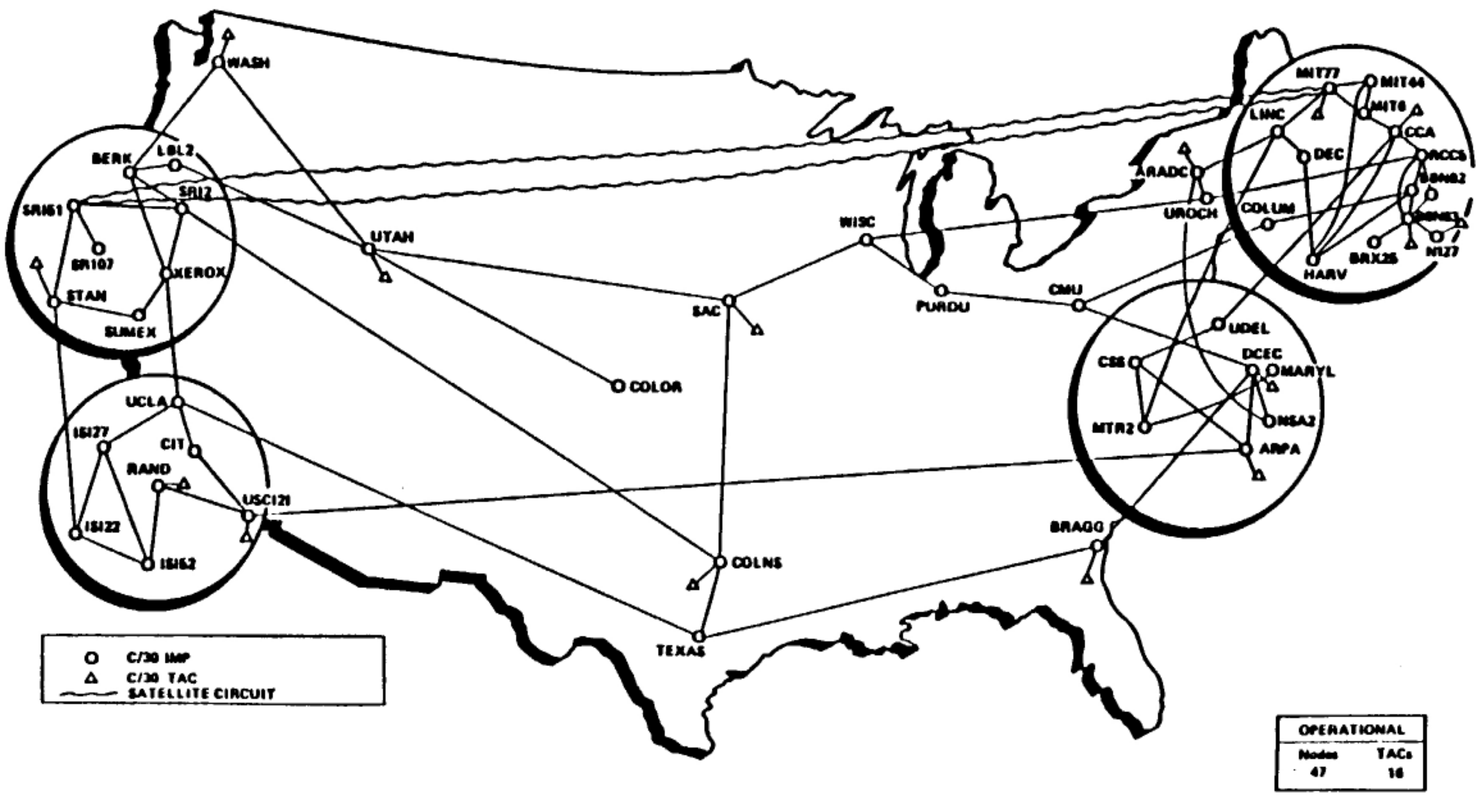

Figure 2: Recent ARPANET Coverage Map

communications through NSFNet. As a result of this cutback, the remaining 56 Kbit/sec ARPANET is becoming even more congested. By this summer, the ARPANET will move to a T1 backbone linking Boston, New York City, Washington DC, Pittsburgh, Chicago, Denver, San Francisco (NASA Ames), Los Angeles, and San Diego -- essentially 1 diagonal line across the US. Locations left off of this route will have to seek connectivity through NSFNet. The reduced coverage and higher speed but cheaper fiber optics lines will significantly reduce medium term network costs. "Multiple network connections" will be eliminated.

In the longer term, DARPA's role will apparently be limited to DoD communications research and support of its contractor community. Thus, the ARPANET will once again become very restrictive in terms of access and use. DARPA expects to construct a Defense Research Internet (DRI) over the next 3 years (in cooperation with DCA and SDI), with the current ARPANET disappearing after that. There is an explicit assumption in what DARPA is doing that the NSFNet or some other NRI will fill in the gaps. Thus, another "lead agency" will be needed. For now, NSF is serving in that role but there is no commitment by them or others that that will continue.

On the operational front, as a member of the ARPANET group, we have an obligation to help with certain network management tasks. For instance, we assisted ARPANET management with installing an additional 56 KB trunk line from our C/30 IMP to a new IMP at the nearby AMES Research Center (NASA). The addition provides an improved network topology in this area. We are working with the ARPANET group to reduce line costs and improve service in other respects also. We still depend heavily on the ARPANET to provide terminal communications service for many of our users as well as file transfer and electronic mail service.

13.1.4 - TELENET Link

The TELENET remains the primary connection mode for those national users without access to ARPANET services. As discussed below, we now have a more general interface to this network that provides an additional password check and outbound as well as inbound services. Because TELENET, like other commercial networks, is oriented to "local-echo" schemes, they do not put a high priority on echo-response across their network. Thus remote users are substantially inconvenienced when using programs requiring close interaction. It appears that our switch from TYMNET to TELENET (made about two years ago) has improved service slightly in this respect.

13.1.5 - X.25/Ethernet Link

The SUMEX-AIM transition to a distributed computing environment and our heavy use of Ethernet motivated our need for a connection between our Ethernet system and the Public Data Networks (specifically TELENET). Because our need involves Ethernet connections to other X.25 based hosts and Packet Assemblers - Disassemblers (PADs), sometimes called Terminal Interface Processors, as opposed to connections between two Ethernets via an X.25 link, we needed a device that provides protocol translation. Such a package would provide SUMEX with both an inbound and outbound capability relative to TELENET and users on SUMEX would be able to access the large variety of hosts and services on the PDNs (such as NLM and Dialog) in a simple and reliable manner. Though the high level protocols for file transfer and mail exchange are developing slowly in the X.25 environment, some progress is being made, so a general purpose interface to these networks is an important asset.

A substantial amount of effort went in to selecting and implementing this equipment. While a large number of vendors claimed to have TCP/IP-to-X.25 interfacing unit, these were generally encapsulating devices intended to connect two TCP/IP networks via an X.25 link. Instead, we needed a device which would provide protocol translation between the X.3 style terminal sessions of the X.25 networks and TCP/IP TELNET support. After a long and frustrating search among potential vendors (e.g., SUN, Bridge, and ACC), we located a suitable system manufactured by Develcon.

The experimental installation of the Develcon Network gateway has provided us with the capability to connect our Ethernet system to a variety of X.25 based networks (for terminal session support). This project has been pursued in conjunction with Stanford's Lane Medical Library and the Stanford University networking group (SUNet).

We have implemented a four-way gateway, connecting our Ethernet system with two X.25 point-to-point lines and the TELENET network (see Figure 8). The result has been a more flexible connection to the TELENET system (at the X.25 level) which gives users the ability to choose a variety of destinations on our Ethernet system when coming into our facility via TELENET. Unlike the simple X.25 PAD which this equipment replaces, the Develcon unit is able to provide both inbound and outbound services.

Our collaboration with Stanford's Lane Medical Library has resulted in this device providing access to the Tandem based Lane On-line Information System (LOIS) via Tandem's X.25 package (no Ethernet package is available), and a Lane sponsored X.25 leased line to the Bibliographic Retrieval Service (BRS) Inc's medical information system (Colleague).

Though this collaboration with the Lane Medical Library has been rather informal, it has proven valuable for both groups. As library science becomes heavily involved with computer based information services the interfacing of their systems with other computer based groups becomes very important. For instance, members of the SUMEX-AIM based Medical Information Sciences (MIS) program are able to make direct transfers of information from BRS to systems on our Ethernet.

The system supplier, Develcon Inc., has recently announced a product based on the work we did with them on this gateway. Further vendor development is needed to support improved accounting for system access and use and for user authentication and access control mechanisms.

13.2 - Microcomputer Networks

We connected our Apple Macintosh computers in 2 buildings with *AppleTalk* and *PhoneNet* network products. More significantly, we integrated them with the rest of our equipment by connecting the microcomputer networks to the campus Ethernet networks using *Kinetics FastPath* gateways, a commercial spin off resulting from the SUMEX work on the SEAGATE gateway.

Software written at Columbia University, Stanford, and elsewhere, makes it possible for a Macintosh to share a VAX file server with the Lisp machines and to access hosts on the ARPA internet as a first-class workstation.

13.3 - Local Area Networks - LAN's

In the past, we have developed local area networking systems to enhance the facilities available to researchers. Much of this work has centered on the effective integration of distributed computing resources in the form of mainframes, workstations, and servers. Network gateways and terminal interface processors (TIP's) were developed and extended to link our environment together. We are purchasing special purpose gateways to interface other equipment as needed too. This year we added several Kinetics Applenet-to-Ethernet gateways to support our Mac II's, LaserWriters, and Micro Explorers, and to link them with the LAN. The original Applenet-to-Ethernet gateways were developed at SUMEX-AIM. Kinetics licensed this software/hardware from Stanford University and used it as a basis for its product. A diagram of our local area network system is shown in Figure 9 and the following summarizes our LAN-related development work.

13.3.1 - Ethernet Gateways

In our heterogeneous network environment, in order to provide workstation access to file servers, mail servers, and other computers within the Stanford local area network, it is necessary to able to route multiple networking protocols through the network gateways. As reported last year, the SUMEX gateways support PUP, Xerox NS, Symbolics/Texas-Instrument CHAOSNET, and the TCP/IP protocols. This support not only provides the routers necessary to move such packets among the subnetworks, but also other miscellaneous services such as time, name/address lookup, host statistics, boot strap support, address resolution, and routing table broadcast and query information. This year a device driver for the Micom-Interlan Ethernet interface was written. This interface has adequate on-board multibus memory to cache up to five 1514 byte input packets (the largest allowed on the network) and more than twenty small packets where "small" is less than 320 bytes. In contrast the older 3COM multibus interface has a fixed two buffer input cache. The larger packet cache minimizes dropped packets at the interface itself, and improves the performance of large file transfers that depend upon the ability to successfully send several back-to-back packets. In fact, we could not run NFS through our gateways with the 3COM interfaces unless we set the maximum number of back-to-back packets that the SUN-3 file server sends to two. With the Micom-Interlan interface, we could change this number to the default which is six, run NFS through the gateways, and achieve maximal performance.

13.3.2 - Terminal Interface Processors

SUMEX-AIM has five TIPs, and in previous years we spent a reasonable amount of time maintaining and augmenting this software. Over the past year, this software has remained stable, and we plan on using vendor supported EtherTIP software for future TIPs. This software originated at SUMEX-AIM and has been commercialized by Cisco Systems, Incorporated, which licensed the EtherTIP/Gateway code from Stanford University.

Currently, one of our TIPs has ten dial-in ports and is used extensively by the local SUMEX-AIM community for dial-in access to the DEC 2060 and our VAX 11/750 servers from home during off hours. The four remaining TIPs are used to access various mainframes during work hours, and one of these runs the Cisco TCP/IP TIP code. As we begin to migrate from the 2060 to our distributed resource, we will in parallel replace our older PUP based TIP software with the above software. We currently run PUP in the majority of our TIPs because the 2060 more efficiently processes PUP, and most of our users login to the 2060 from these TIPs.

14 - Printing Services

Laser printers have become essential components of the work environment of the SUMEX-AIM community with applications ranging from scientific publications to hardcopy graphics output for ONCOCIN chemotherapy protocol patient charts. We have done much systems work to integrate laser printers into the SUMEX network environment so they would be routinely accessible from hosts and workstations alike. This expertise has been widely shared with other user groups in the AIM community and beyond.

SUMEX operates 7 medium-speed (8-20 pages per minute) Imagen laser printers, 2 low-speed (~3 ppm) Apple laser printers, and 1 low-speed (~3 ppm) Xerox laser printer.

Each of the Imagen printers includes an emulator for a line printer, a daisy wheel printer, a Tektronix plotter, and a typesetter (using the *Impress* language). Additionally, the two Imagen *3320* printers implement the *PostScript* typesetter language (also implemented by the Apple LaserWriters) required for printing Macintosh documents. The Xerox printer (an *8046*) interprets the *InterPress* typesetter language.

In total, the laser printers printed about half a million pages of output during the year. Most of the printout was simple text, followed in quantity by formatted text in *Impress* format, Impress-format drawings, and screen dumps. About 16,000 pages each of PostScript-format drawings and formatted text were printed on the Apple LaserWriters and Imagen 3320's (an eightfold increase over last year).

We were one of a small number of sites to beta test Imagen's implementation of the PostScript typesetter language. We submitted a number of bug reports and all the bugs we found were fixed in short order. (Imagen went on to become the first printer manufacturer to ship a non-Adobe implementation of PostScript.) In consideration of our help with the test, Imagen transferred ownership of the test hardware to our lab.

Since the proliferation of Macintosh computers demands higher-speed printing than the Apple LaserWriter can provide, another Imagen 12/300 was upgraded to model 3320 (configured for PostScript) and installed in the other building inhabited by SUMEX and KSL users. Imagen granted a special discount on this upgrade in consideration of our dissatisfaction with the longevity of the 12/300 printers purchased earlier.

15 - General User Software

We have continued to assemble (develop where necessary) and maintain a broad range of user support software. These include such tools as language systems, statistics packages, vendor-supplied programs, text editors, text search programs, file space management programs, graphics support, a batch program execution monitor, text formatting and justification assistance, magnetic tape conversion aids, and user information/help assistance programs.

A particularly important area of user software for our community effort is a set of tools for inter-user communications. We have built up a group of programs to facilitate many aspects of communications including interpersonal electronic mail, a "bulletin board" system for various special interest groups to bridge the gap between private mail and formal system documents, and tools for terminal connections and file transfers between SUMEX and various external hosts. Examples of work on these sorts of programs have already been mentioned in earlier sections, particularly as they relate to extensions for a distributed computing environment.

At SUMEX-AIM we are committed to importing rather than reinventing software where possible. As noted above, a number of the packages we have brought up are from outside groups. Many avenues exist for sharing between the system staff, various user projects, other facilities, and vendors. The availability of fast and convenient communication facilities coupling communities of computer facilities has made possible effective intergroup cooperation and decentralized maintenance of software packages. The many operating system and system software interest groups (e.g., TOPS-20, UNIX, D-Machines, network protocols, etc.) that have grown up by means of the ARPANET have been a good model for this kind of exchange. The other major advantage is that as a by-product of the constant communication about particular software, personal connections between staff members of the various sites develop. These connections serve to pass general information about software tools and to encourage the exchange of ideas among the sites and even vendors as appropriate to our research mission. We continue to import significant amounts of system software from other network sites, reciprocating with our own local developments. Interactions have included mutual backup support, experience with various hardware configurations, experience with new types of computers and operating systems, designs for local networks, operating system enhancements, utility or language software, and user project collaborations. We have assisted groups that have interacted with SUMEX user projects get access to software available in our community (for more details, see the section on Dissemination on page 123).

III.A.2.5. Relevant Core Research Publications

The following is a list of new publications and reports that have come out of our core research and development efforts over the past year. In addition, we include references to earlier reports that are discussed in the text above.

HPP 80-29
H. Penny Nii; **An Introduction to Knowledge Engineering, Blackboard Model and AGE,** March 1980, **45 pages**

KSL 86-19
J.P. Rice; **Poligon, A System for Parallel Problem Solving,** April 1986. To appear in: *DARPA proceedings, Asilomar 1986.* **16 pages**

KSL 86-20
J.R. Delaney; **Multi-System Report Integration Using Blackboards,** March 1986. Submitted for publication to: *1986 American Control Conference.* **12 pages**

KSL 86-38
STAN-CS-87-1147. Barbara Hayes-Roth, M. Vaughan Johnson Jr., Alan Garvey, and Micheal Hewett; **A Modular and Layered Environment for Reasoning about Action,** April 1987. To appear in: *The Journal of Artificial Intelligence in Engineering, Special Issue on Blackboard Systems, October 1986.* **63 pages**

KSL 86-41
H. Penny Nii; **CAGE and POLIGON: Two Frameworks for Blackboard-based Concurrent Problem Solving,** April 1986. To appear in: *DARPA Proceedings, Asilomar 1986* **8 pages** [Superceded by KSL 87-71]

KSL 86-62
STAN-CS-87-1175. David C. Wilkins, William J. Clancey, and Bruce G. Buchanan; **Using and Evaluating Differential Modeling in Intelligent Tutoring and Apprentice Learning Systems,** January 1987. To appear in: *Intelligent Tutoring Systems: Lessons Learned, Lawrence Erlbaum Publishers, 1987.* **37 pages**

KSL 86-63
David C. Wilkins; **Knowledge Base Debugging Using Apprenticeship Learning Techniques,** October 1986. **15 pages**

KSL 86-69
STAN-CS-86-1136. Harold Brown, Eric Schoen, and Bruce Delagi; **An Experiment in Knowledge-Base Signal Understanding Using Parallel Architectures,** October 1986. To appear in: *Parallel Computation and Computers for AI, J.S. Kowalik Editor, Kluwer Publishers.* **39 pages**

KSL 87-27
(Journal Memo) Gregory F. Cooper; **Probabilistic Inference Using Belief Networks Is NP-Hard,** August 1987. **23 pages**

KSL 87-28
(Journal Memo) Eric J. Horvitz; **A Multiattribute Utility Approach to Inference Understandability and Explanation,** September 1987. **34 pages**

KSL 87-34
(Working Paper) Russell T. Nakano; **Experiments with a Knowledge-Based System on a Multiprocessor: Preliminary Airtrac-Lamina Qualitative Results,** June 1987. **72 Pages**

KSL 87-36
Lawrence J. Selig; **An Expert System Using Numerical Simulation and Optimization To Find Particle Beam Line Errors,** May 1987. **39 pages**

KSL 87-38
Naomi S. Rodolitz; **Guidon Manage - Tutoring for Strategic Knowledge,** June 1987. **43 pages**

KSL 87-39
(Journal Memo) Glenn D. Rennels, Edward H. Shortliffe, Frank E. Stockdale, Perry L. Miller; **Updating an expert knowledge base as medical knowledge evolves: Examples from oncology management,** June 1987. To appear in *Proceedings of the AAMSI Congress 87.* **6 pages**

KSL 87-40
Alan Garvey and Barbara Hayes-Roth; **Implementing Diverse Forms of Control Knowledge in Multiple Control Architectures,** June 1987. **32 pages**

KSL 87-41
Thierry Barsalou; **An Object-Based Interface to a Relational Database System,** June 1987. **10 pages**

KSL 87-43
STAN-CS-87-1166. Hiroshi G. Okuno and Anoop Gupta; **Parallel Execution of OPS5 in QLISP,** June 1987. Shorter version to appear in: *Proceedings of the Fourth Conference on Artificial Intelligence Applications (CAIA-88), IEEE, March 1988.* **18 pages**

KSL 87-44
STAN-CS-87-1178. Gregory T. Byrd, Russell Nakano, and Bruce A. Delagi; **A Dynamic, Cut-Through Communications Protocol with Multicast,** August 1987. **23 pages**

KSL 87-45
(Journal Memo) David Heckerman and Holly Jimison; **A Perspective on Confidence and Its Use in Focusing Attention during Knowledge Acquisition,** July 1987. To appear in: *Proceedings of AAAI 87.* **10 pages**

KSL 87-46
(Journal Memo) Michael P. Wellman and David E. Heckerman; **The Role of Calculi in Uncertain Reasoning,** July 1987. To appear in: *Proceedings of AAAI 87.* **12 pages**

KSL 87-48
(Journal Memo) Gregory F. Cooper; **Computer-Based Medical Diagnosis Using Belief Networks and Bounded Probabilities,** February 1988. To appear in *Topics in Medical Artificial Intelligence edited by Perry Miller.* **17 pages**

KSL 87-49
(Journal Memo) Mark E. Frisse and Gio Wiederhold; **Retrieving Information from Hypertext Systems,** August 1987. **14 pages**

KSL 87-50
(Journal Memo) Mark E. Frisse; **Searching for Information in a Hypertext Medical Handbook: The Washington University Dynamic Medical Handbook Project,** August 1987. **19 pages**

KSL 87-51
(Journal Memo) Gregory F. Cooper; **Expert Systems Based on Belief Networks - Current Research Directions,** February 1988. **21 pages**

KSL 87-52
Curtis P. Langlotz; **Advice Generation in an Axiomatically-Based Expert System,** August 1987. To appear in: *Proceedings of the Eleventh Annual Symposium on Computer Applications in Medical Care, Washington, DC, 1987.* **8 pages**

KSL 87-53
(Journal Memo) David Heckerman and Eric J. Horvitz; **On the Expressiveness of Rule-Based Systems for Reasoning with Uncertainty,** August 1987. To appear in: *Proceedings of AAAI, Vol. 1, July, 1987.* **7 pages**

KSL 87-54
Joshua Lederberg; **How DENDRAL Was Conceived and Born,** August 1987. To appear in: *Proceedings of ACM Symposium on the History of Medical Informatics, National Library of Medicince, November 1987.* **19 pages**

KSL 87-57
STAN-CS-87-1184. Hiroshi Okuno, Nobuyasu Osato and Ikuo Takeuchi; **Firmware Approach to Fast Lisp Interpreter,** September 1987. To appear in: *Proceedings of Twentieth Annual Workshop on Microprogramming (MICRO-20), ACM, December 1987.* **25 pages**

KSL 87-58
William J. Clancey; **Knowledge Engineering Methodology: An Annotated Bibliography of NEOMYCIN Research,** September 1987. To appear in: *R. Nossum (ed.), Lecture notes in Computer Science - ACAI'87, Springer-Verlag, 1988.* **10 pages**

KSL 87-59
Janet McLaughlin; **Utility-Directed Presentation of Simulation Results,** December 1987. **57 pages**

KSL 87-60
Richard M. Keller; **Defining Operationality for Explanation-Based Learning,** October 1987. To appear in: *Artificial Intelligence Journal.* **19 pages**

KSL 87-61
STAN-CS-87-1188. Russell Nakano and Masafumi Minami; **Experiments with a Knowledge-Based System on a Multiprocessor,** October 1987. **47 pages**

KSL 87-62
John A. Brugge and Bruce G. Buchanan; **Evolution of a Knowledge-Based System for Determining Structural Components of Proteins,** October 1987. **26 pages**

KSL 87-63
(Journal Memo) Leslie Elaine Perreault; **Automatic Test Case Generation by Modeling Patient States and Physician Actions,** October 1987. **30 pages**

KSL 87-64
(Journal Memo) Eric J. Horvitz; **Problem-Solving Design: Reasoning about Computational Value, Tradeoffs, and Resources,** November 1987. To appear in: *Proceedings of the National Aeronautics And Space Administration Artificial Intelligence Forum, Mountain View, CA.* **19 pages**

KSL 87-65
STAN-CS-87-1189. Bruce A. Delagi, Nakul Saraiya, Sayuri Nishimura, and Greg Byrd; **Instrumented Architectural Simulation,** November 1987. **7 pages**

KSL 87-67
Barbara Hayes-Roth; **Dynamic Control Planning in Adaptive Intelligent Systems,** November 1987. **7 pages**

KSL 87-73
Stephen Barnhouse; **Knowledge Base Tours: Introducing a Knowledge Based System to a Novice User**, December 1987. **15 pages**

KSL 88-01
M. Vaughan Johnson and Barbara Hayes-Roth; **Learning to Solve Problems by Analogy**, March 1988. **15 pages**

KSL 88-02
H. Penny Nii, Nelleke Aiello, and James Rice; **Frameworks for Concurrent Problem Solving: A Report on Cage and Poligon**, March 1988. **28 pages**

KSL 88-04
J. P. Rice; **Problems with Problem-Solving in Parallel: The Poligon System**, January 1988. **21 pages**

KSL 88-06
(Thesis) Mark Alan Musen; **Generation of Model-Based Knowledge-Acquisition Tools for Clinical-Trial Advice Systems**, January 1988. **293 pages** [ONLY ABSTRACT AVAILABLE]

KSL 88-09
(Working Paper) Bruce G. Buchanan and Reid G. Smith; **Fundamentals of Expert Systems**, March 1988. **33 pages**

KSL 88-10
Gregory T. Byrd and Bruce A. Delagi; **A Performance Comparison of Shared Variables vs. Message Passing**, February 1988. To appear in: *May 1988 ISI Supercomputing Conference Proceedings*. **15 pages**

KSL 88-11
Richard M. Keller; **Operationality and Generality in Explanation-Based Learning: Separate dimensions or opposite endpoints?**, February 1988. To appear in: *Proceedings of the AAAI Symposium on Explanation-Based Learning, March 1988, Stanford, CA*. **5 pages**

KSL 88-12
Clifford E. Wulfman, Ellen A. Isaacs, Bonnie Lynn Webber, and Lawrence M. Fagan; **Integration Discontinuity: Interfacing Users and Systems**, February 1988. **12 pages**

KSL 88-13
Eric J. Horvitz, John S. Breese, and Max Henrion; **Decision Theory in Expert Systems and Artificial Intelligence**, February 1988. To appear in: *Journal of Approximate Reasoning, Special Issue on Uncertainty in Artificial Intelligence, July 1988*. **62 pages**

KSL 88-15
Tony Confrey and Fancois Daube; **GS2D: A 2D Geometry Systems**, March 1988. **15 pages**

KSL 88-16
Mark A. Musen; **Conceptual Models of Interactive Knowledge-Acquisition Tools**, March 1988. **9 pages**

KSL 88-19
Barbara Hayes-Roth, Micheal Hewett, M. Vaughan Johnson, and Alan Garvey; **ACCORD: A Framework for a Class of Design Tasks**, March 1988. **12 pages**

KSL 88-20
Barbara Hayes-Roth, Rattikorn Hewett, and Adam Seiver; **Diagnostic Explanation using Generic Models**, March 1988. **10 pages**

KSL 88-21
Alan Garvey and Barbara Hayes-Roth; **An Empirical Analysis of Explicit vs. Implicit Control Architectures**, March 1988. **20 pages**

KSL 88-22
Micheal Hewett and Barbara Hayes-Roth; **Real-Time I/O in Knowledge-Based Systems**, March 1988. **10 pages**

KSL 88-23
Robert Schulman and Barbara Hayes-Roth; **Plan-Based Construction of Strategic Explanations**, March 1988. **13 pages**

KSL 88-24
James F. Brinkley; **The Potential for Intelligent Three-Dimensional Ultrasound**, March 1988. To appear in: *Chervenak, F.A., Isaacson, G., Campbell, S. (eds.), Textbook of Ultrasound in Obstetrics and Gynecology, 1988.* **28 pages**

KSL 88-25
(Working Paper) Greg Byrd; **Modelling a Bus-Based Multiprocessor Using the CARE Simulation System**, March 1988. **11 pages**

KSL 88-26
Mark A. Musen; **Generation of Knowledge-Acquisition Tools from Clinical-Trial Models**, March 1988. To appear in: *Proceedings of Medical Informatics, Europe 1988. Oslo, Norway, August 1988.* **6 pages**

KSL 88-27
(Working Paper) H. J. Suermondt and Gregory F. Cooper; **Updating Probabilities in Multiply Connected Belief Networks**, March 1988. **9 pages**

KSL 88-28
Cooper, G.F.; **A method for using belief networks as influence diagrams**, April 1988.

KSL 88-34
(Thesis) Isabelle de Zegher-Geets; **IDEFIX: Intelligent Summarization of a Time-Oriented Medical Database**, June 1987. **99 pages**

KSL 88-38
Heckerman, D.E.; **An empirical comparison of three scoring schemes**, May 1988.

KSL 88-TBA
Chavez, R.M. and Cooper, G.F.; **KNET: Integrating hypermedia and normative Bayesian modeling.**

KSL 88-TBA
Lehmann, H., **Knowledge acquisition for probability-based expert systems.**

Other Outside Articles:

Cooper, G.F., Expert systems based on belief networks -- Current research directions, *Journal of Applied Statistical Models and Data Analysis*, **4**, 1988.

Cooper, G.F., invited commentary on: Lauritzen, S. and Spiegelhalter, D., Local computations with probabilities on graphical structures and their application to expert systems, *Journal of the Royal Statistical Society*, **B**, **50**, 1988.

Horvitz, E.J., Reasoning under varying and uncertain resource limitations, in: *Proceedings of the National Conference on Artificial Intelligence*, Minneapolis, MN, August, 1988.

Horvitz, E.J., Breese, J.S., and Henrion, M., Decision theory in expert systems and artificial intelligence, *Journal of Approximate Reasoning*, 1988.

III.A.2.6. Resource Equipment

The SUMEX-AIM resource is a complex, integrated facility comprised of mainframes, workstations, networks, and servers illustrated in Figures 3 - 9. A key role of the SUMEX-AIM resource is to continue to evaluate workstations as the technology is changing rapidly. This evaluation includes new hardware and software, 1) to provide superior development and execution platforms for AI research, and 2) to support the ancillary "office environment" (presently carried out on the DEC 2060, which is being phased out). Thus far no single workstation has materialized that provides all the services we would like to see in support of either or both of these missions. This means that for the foreseeable future, we will utilize a multiplicity of machines and software to address the needs of the projects.

Systems based on the Motorola 68020 chip (e.g., SUN Microsystems or Apple Macintosh II workstations), the Intel 80286 and 80387 chips (e.g., IBM PS/1-4 machines), and other newer architectures such as reduced instruction set computer (RISC) chips, have Lisp benchmark data rivaling the performance of existing, specially microcoded Lisp machines (see Appendix B). It is still early to predict how this "race" will ultimately turn out and software environments play an equally important role with raw hardware speed in the decision. For now, the Lisp software environments on the "stock" machines are not nearly so extensively developed as on Lisp machines and conversely, the routine computing environments of Lisp machines (text processing, mail, spreadsheets, etc.) lag the tools available on stock UNIX machines.

In earlier year's we experimentally tried increasing usage of TI and Xerox Lisp machines (purchased as AI research platforms) for text editing, electronic mail, and document formatting with considerable success (although many of these tools were only tested in a prototype form and were not widely distributed). In addition, the use of expensive Lisp machines for routine computing and office applications impacts their availability as research tools.

We had been seeking an integration of both the Lisp machine and stock machine worlds. As discussed extensively in the progress section on Core Systems Research (see Page 39), these two capabilities came together as never before with the Macintosh II and microExplorer coprocessor systems.

1 - Purchases This Past Year

The SUMEX 2060 hardware continues to be stable and the relatively small amount of SUMEX-AIM money for new purchases has been concentrated on experimental workstations and server equipment needed for distributed system development. These purchases are paced carefully with the developments of higher performing, more compact, and lower cost systems. The NIH-funded purchases this past year are summarized below. For the most part, these represent evaluation units used to review the suitability of particular types of equipment preparatory to a larger volume purchase from non-NIH funding sources, as detailed earlier starting on Page 39).

1. Apple Mac-II
2. Rodime Disk Drive 100 Megabyte Disk Drive (for MAC-II)
3. E-Machines Big Picture Display (for MAC-II)
4. U.S. Robotics 9600 Baud Modems (2 ea).
5. Develcon X.25/Ethernet Gateway (shared cost with others)

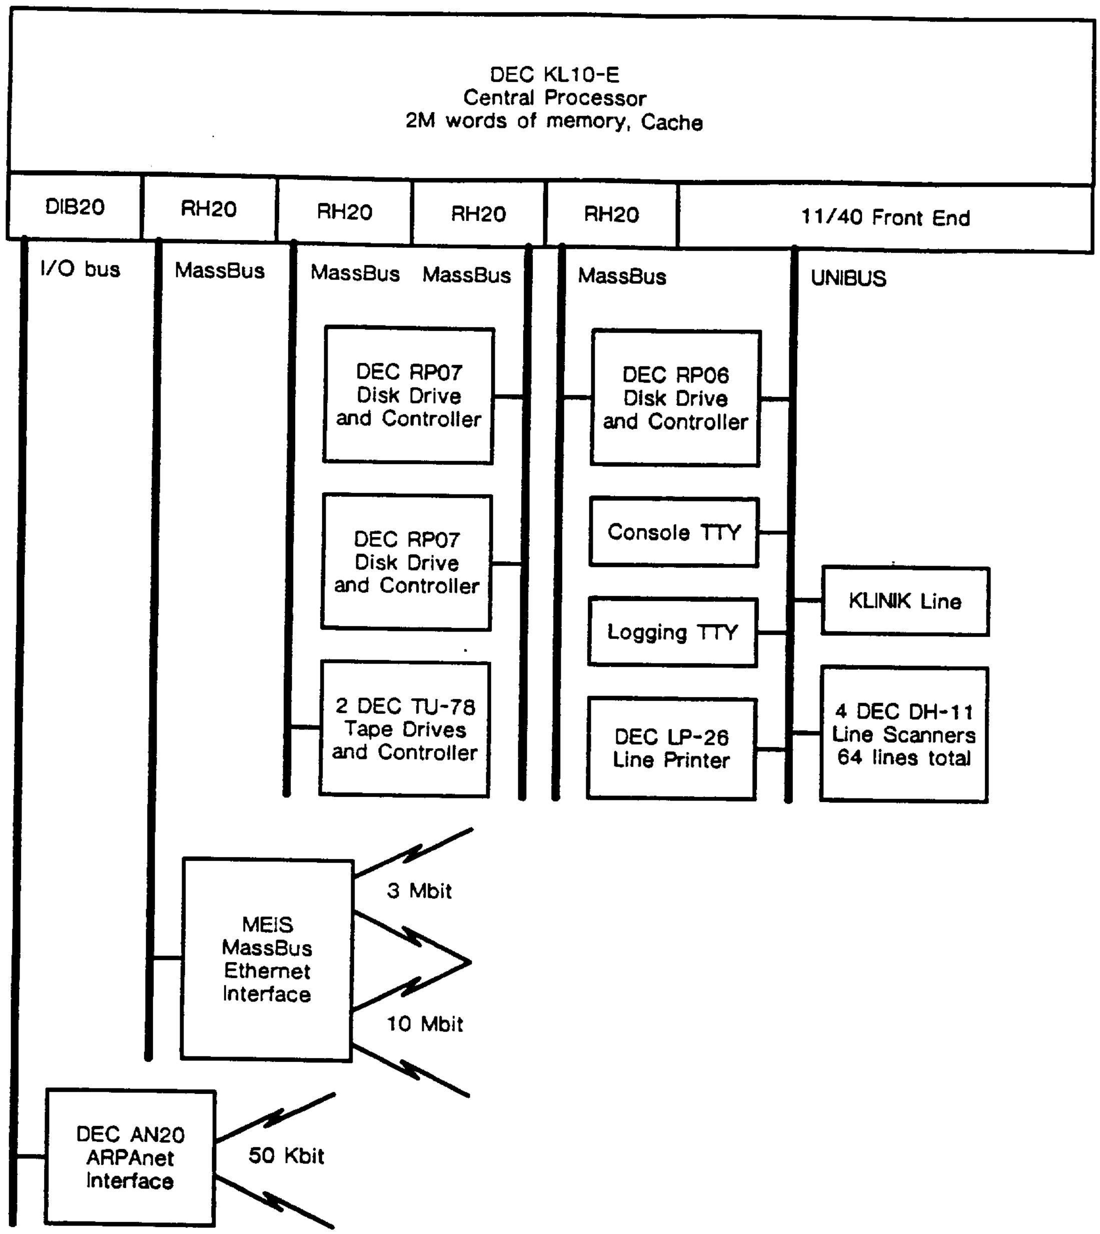

Figure 3: SUMEX-AIM DEC 2060 Configuration

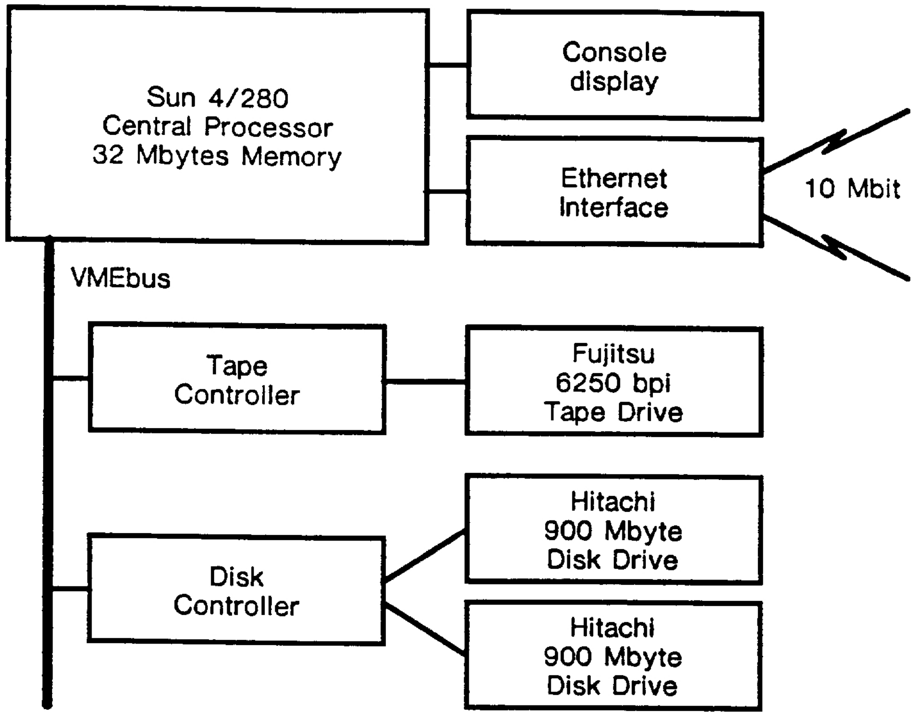

Figure 4: SUMEX-AIM SUN-4 Configuration

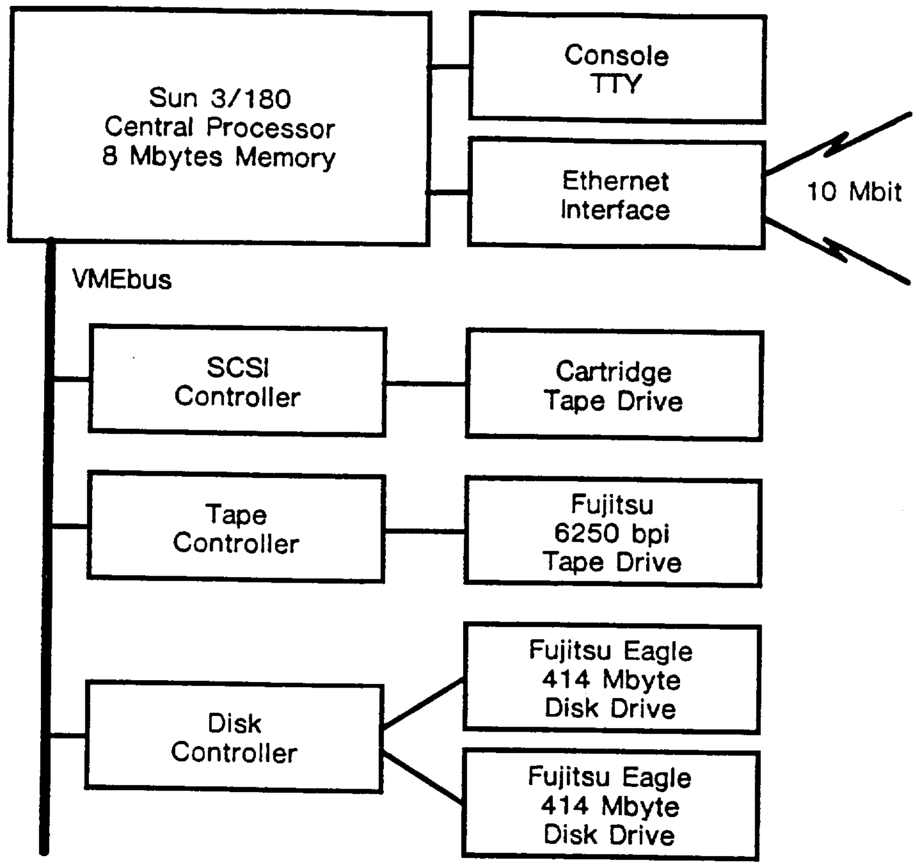

Figure 5: SUMEX-AIM Sun-3 File Server Configuration

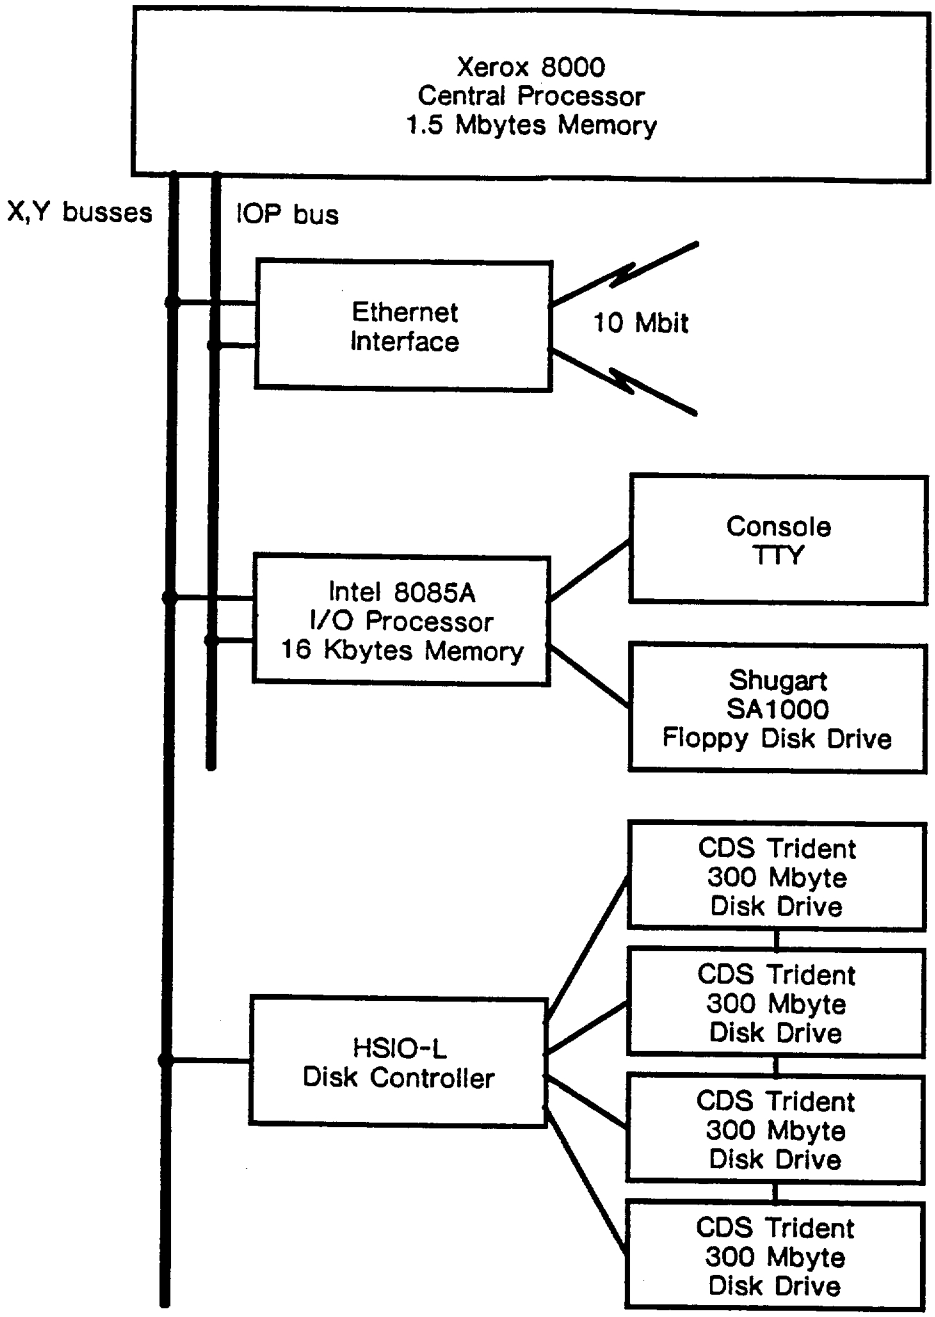

Figure 6: SUMEX-AIM Xerox File Server Configuration

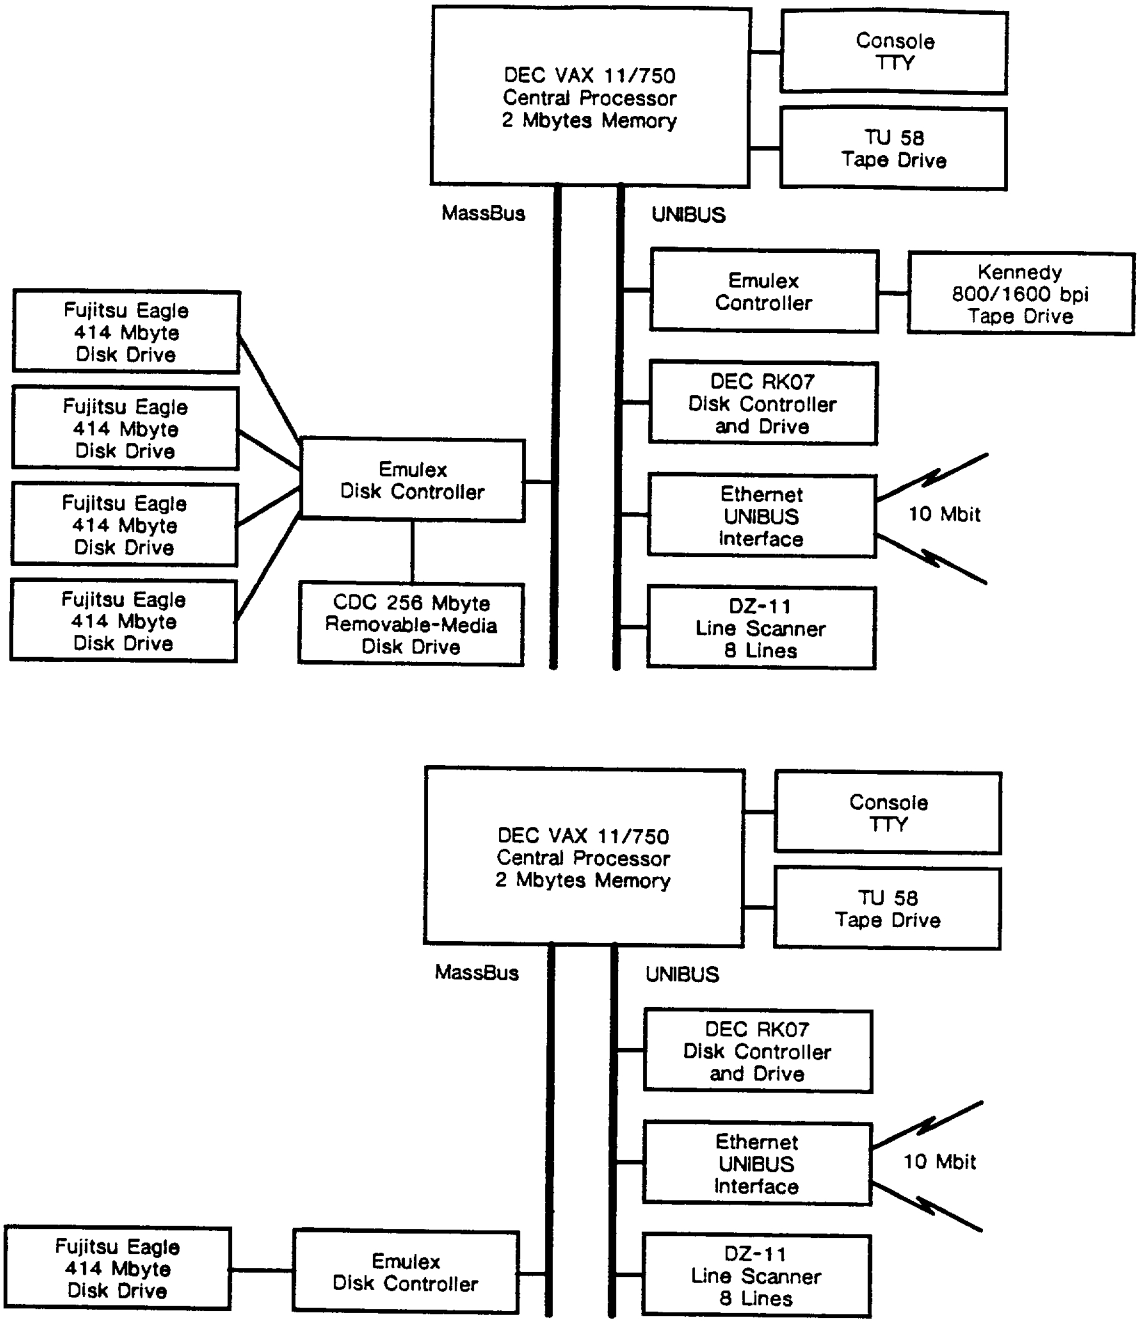

Figure 7: SUMEX-AIM VAX File Server Configuration

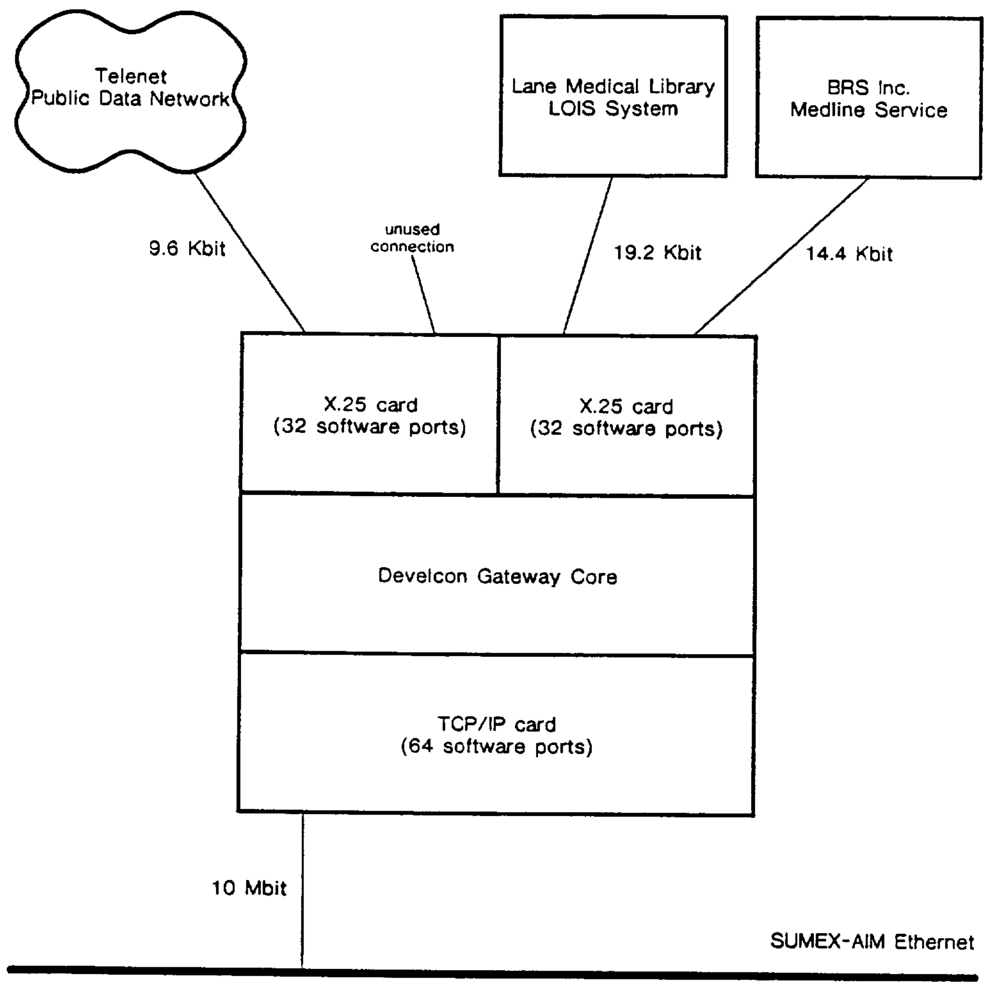

Figure 8: SUMEX-AIM Develcon X.25/TCP-IP Gateway Configuration

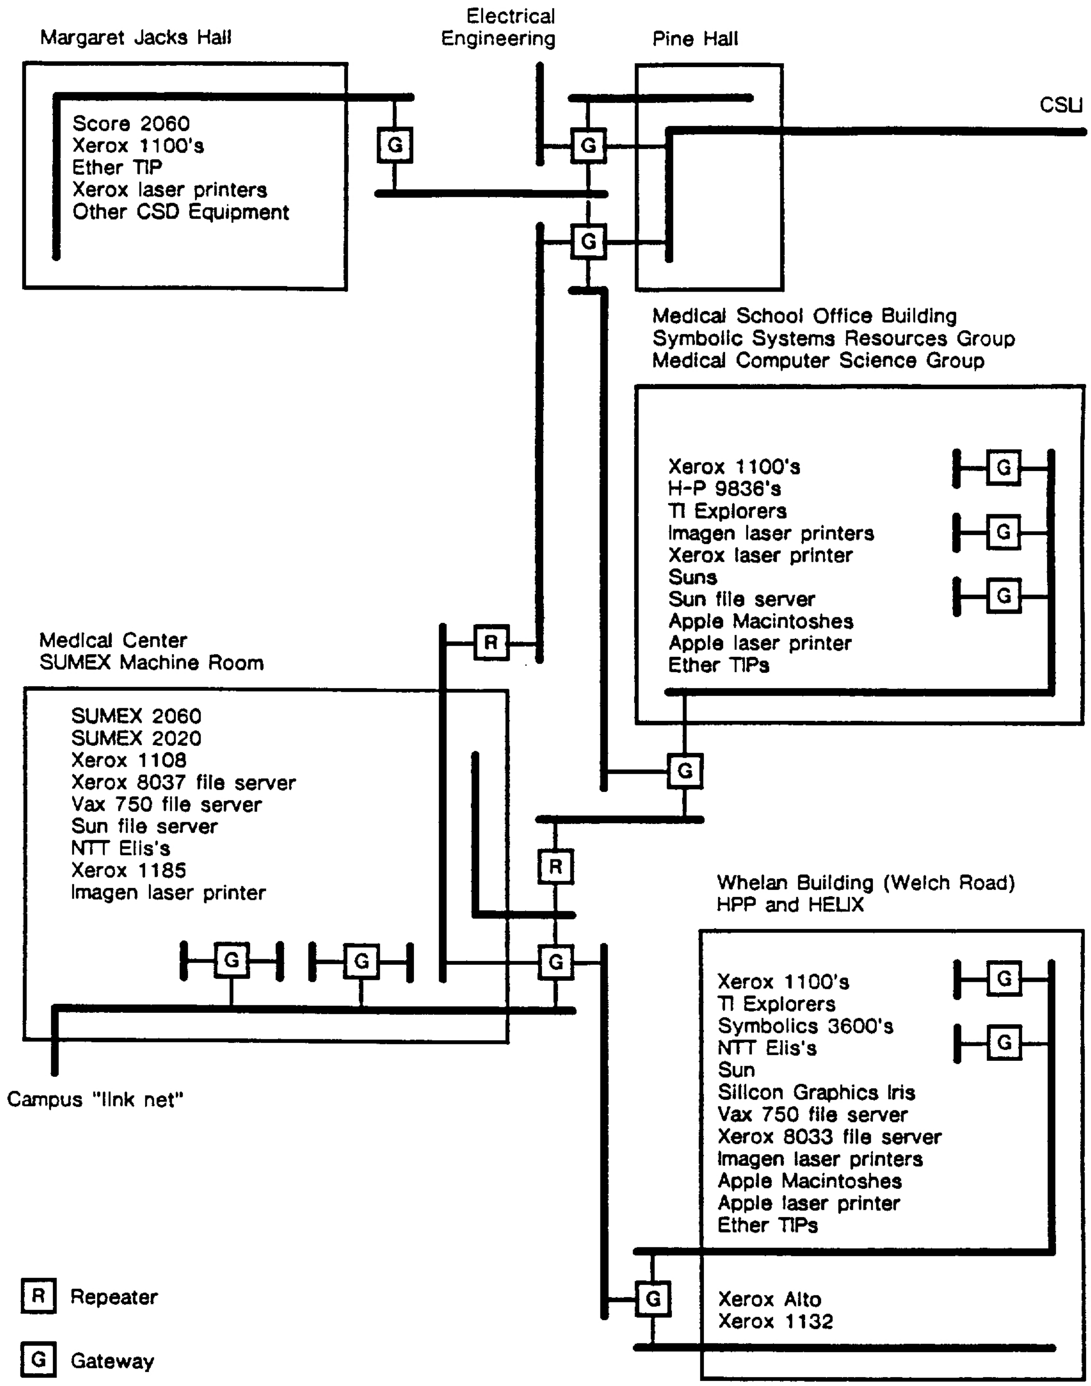

Figure 9: SUMEX-AIM Ethernet Configuration

III.A.2.7. Training Activities

The SUMEX resource exists to facilitate biomedical artificial intelligence applications. This user orientation on the part of the facility and staff has been a unique feature of our resource and is responsible in large part for our success in community building. The resource staff has spent significant effort in assisting users to gain access to the SUMEX-AIM resources at Stanford and use it effectively as well as in assisting AIM projects in designing their own local computing resources based on SUMEX experience. We have also spent substantial effort to develop, maintain, and facilitate access to documentation and interactive help facilities. The HELP and Bulletin Board subsystems have been important in this effort to help users get familiar with the computing environment.

We have regularly accepted a number of scientific visitors for periods of several months to a year, to work with us to learn the techniques of expert system definition and building and to collaborate with us on specific projects. Our ability to accommodate such visitors is severely limited by space, computing, and manpower resources to support such visitors within the demands of our on-going research.

A tutorial was held in January 1988 on the Parallel Computing Architectures Project multiprocessor simulation system (see Page 29). This two-day session was attended by representatives from the DoD, NASA and Boeing, and included descriptions of the CARE/SIMPLE system, as well as the LAMINA programming interface. The attendees received instruction in use of the system for making measurements of the performance of various simulated multiprocessor applications.

Finally, the training of graduate students is an essential part of the research and educational activities of the KSL. Based largely on the SUMEX-AIM community environment, we have initiated two unique, special academic degree programs at Stanford, the Medical Information Science program and the Masters of Science in AI, to increase the number of students we produce for research and industry. A number of students are pursuing interdisciplinary programs and come from the Departments of Engineering, Mathematics, Education, and Medicine.

The *Medical Information Sciences (MIS)* program continues to be one of the most obvious signs of the local academic impact of the SUMEX-AIM resource. The MIS program received recent University approval (in October 1982) as an innovative training program that offers MS and PhD degrees to individuals with a career commitment to applying computers and decision sciences in the field of medicine. In Spring 1987, a University-appointed review group unanimously recommended that the degree program be continued for another five years. The MIS training program is based in the School of Medicine, directed by Dr. Shortliffe, co-directed by Dr. Fagan, and overseen by a group of six University faculty that includes two faculty from the Knowledge Systems Laboratory (Profs. Shortliffe and Buchanan). The specialized curriculum offered by the new program is intended to overcome the limitations of previous training options. It focuses on the development of a new generation of researchers with a commitment to developing new knowledge about optimal methods for developing practical computer-based solutions to biomedical needs.

The program accepted its first class of four trainees in the summer of 1983 and has now reached its steady-state size of approximately twenty-two graduate students. The program encourages applications from any of the following:

- medical students who wish to combine MD training with formal degree work and research experience in MIS;

- physicians who wish to obtain formal MIS training after their MD or their residency, perhaps in conjunction with a clinical fellowship at Stanford Medical Center;

- recent BA or BS graduates who have decided on a career applying computer science in the medical world;

- current Stanford undergraduates who wish to extend their Stanford training an extra year in order to obtain a "co-terminus" MS in the MIS program;

- recent PhD graduates who wish post-doctoral training, perhaps with the formal MS credential, to complement their primary field of training.

In addition, a special one-year MS program is available for established academic medical researchers who may wish to augment their computing and statistical skills during a sabbatical break. As of Spring 1988, 55% of our trainees have previously received MD degrees and another 23% are medical students enrolled in joint degree programs. 27% are candidates for the MS degree, while the rest are doctoral students. The program has seven graduates to date, with several more expecting to complete degrees before the end of 1988.

Except for the special one-year MS mentioned above, all students spend a minimum of two years at Stanford (four years for PhD students) and are expected to undertake significant research projects for either degree. Research opportunities abound, however, and they of course include the several Stanford AIM projects as well as research in psychological and formal statistical approaches to medical decision making, applied instrumentation, large medical databases, and a variety of other applications projects at the medical center and on the main campus. Several students are already contributing in major ways to the AIM projects and core research described elsewhere in this annual report.

We are pleased that the program already has an excellent reputation and is attracting superb candidates for training positions. The program's visibility and reputation is due to a number of factors:

- high quality students, many of whom publish their work in conference proceedings and refereed journals even before receiving their degrees; Stanford MIS students have won first prize in the student paper competition at the Symposium on Computer Applications in Medical Care (SCAMC) in 1985 and 1986, and have also received awards for their work at annual meetings of organizations such as the Society for Medical Decision Making, the American Association for Medical Systems and Informatics (AAMSI), and the American Association for Artificial Intelligence (AAAI);

- a rigorous curriculum that includes newly-developed course offerings that are available to the University's medical students, undergraduates, and computer science students as well as to the program's trainees;

- excellent computing facilities combined with ample and diverse opportunities for medical computer science and medical decision science research;

- the program's great potential for a beneficial impact upon health care delivery in the highly technologic but cost-sensitive era that lies ahead.

The program has been successful in raising financial and equipment support from industry and foundations. It is also recipient of a training grant from the National Library of Medicine. The latter grant was recently renewed for another five years with a study section review that praised both the training and the positive contribution of the SUMEX-AIM environment.

III.A.2.8. Resource Operations and Usage

1 - Operations and Support

The diverse computing environment that SUMEX-AIM provides requires a significant effort at operations and support to keep the resource responsive to community project needs. This includes the planning and management of physical facilities such as machine rooms and communications, system operations routine to backup and retrieve user files in a timely manner, and user support for communications, systems, and software advice.

We spend significant time on new product review and evaluation such as Lisp workstations, terminals, communications equipment, network equipment, microprocessor systems, mainframe developments, and peripheral equipment. We also pay close attention to available video production and projection equipment, which has proved so useful in our dissemination efforts involving video tapes of our work.

We continue to operate the primary elements of our equipment base in a generally unattended manner. Operations costs are kept to a minimum by utilizing a student staff for routine tasks. Senior members of this staff provide improvements to the operations procedures in addition to training and supervising new students. This has provided SUMEX with a cost effective operations scheme, contributed to the education of the students, and assisted students in meeting their obligations in undergraduate financial aid programs.

While most of our equipment is concentrated in three computer equipment rooms, our move towards distributed computing has resulted in a substantial amount of equipment being installed in offices and student carrels. The planning of our project's area during the construction of the Medical School Office Building (described in last year's report) has made this distribution of equipment easier.

2 - Resource Usage Details

The following data give an overview of various aspects of SUMEX-AIM central resource usage. There are 5 subsections containing data respectively for:

1. Overall resource loading data (page 101).
2. Relative system loading by community (page 102).
3. Individual project and community usage (page 105).
4. Network usage data (page 111).
5. System reliability data (page 113).

For the most part, the data used for these plots cover the entire span of the SUMEX-AIM project. This includes data from both the KI-TENEX system and the current DECsystem 2060. At the point where the SUMEX-AIM community switched over to the 2060 (February, 1983), you will notice sharp changes in most of the graphs. This is due to differences in scheduling, accounting, and processor speed calculations between the systems.