

I Narrative Description

This is an annual report for the Stanford University Medical Experimental computer resource for applications of Artificial Intelligence in Medicine (SUMEX-AIM). It covers the period between May 1, 1980 and April 30, 1981.

We are about to begin a 5-year renewal of the SUMEX resource grant which will launch an important and exciting new phase for SUMEX-AIM community research. Recent successes in developing expert systems, many of them stemming from projects in the SUMEX-AIM community, have stimulated increasing interest in AI research from many fronts. At the same time, the on-going revolution in computational tools, made possible by larger and larger scale microelectronic integration, is making routine applications of AI systems more practical and effective. Our approved renewal goals focus principally on a merging of state-of-the-art community research in biomedical AI applications with these new computing tools and on the challenges they will bring to the SUMEX-AIM community and resource. We expect that the integration and exploitation of these emerging computer technologies that will have a profound effect on the development and export of practical biomedical AI programs.

This report on the last year in our current 3-year grant is thus, in a sense, a culmination of the early phase of the SUMEX resource. This phase has been characterized by the building of a national community of biomedical AI collaborators around a central resource located at Stanford University. Beginning with 5 projects in 1973, the AIM community grew to 11 major projects at our renewal in 1978 and currently numbers 16 fully authorized projects plus a group of 7 pilot efforts. Many of the computer programs under development by these groups are maturing into tools increasingly useful to the respective research communities. The demand for production-level use of these programs has surpassed the capacity of the present SUMEX facility and has raised important issues of how such software systems can be optimized for production environments, exported, and maintained.

To be sure, we will continue to seek interesting new AI applications in an expanding community of biomedical and computer scientists interacting through electronic media. However, we expect the SUMEX-AIM community to develop a somewhat different character in the coming years. It will become more decentralized in terms of computing resources, more diverse in scope, and even more heavily dependent on network communication facilities for interactions, collaborations, and sharing.

The following sections report on the activities of the SUMEX-AIM resource this past year including brief summaries of the objectives of SUMEX-AIM, a characterization of biomedical AI research, resource organization and operating procedures, recent core progress in system development and basic AI research, and progress in the collaborative projects.

I.A Summary of Research Progress

I.A.1 Overview of Objectives and Rationale

SUMEX-AIM ("SUMEX") is a national computer resource with a dual mission: a) promoting applications of computer science research in artificial intelligence (AI) to biological and medical problems and b) demonstrating computer resource sharing within a national community of health research projects. The central SUMEX-AIM facility is located physically in the Stanford University Medical School and serves as a nucleus for a community of medical AI projects at universities around the country. SUMEX provides computing facilities tuned to the needs of AI research and communication tools to facilitate remote access, inter- and intra-group contacts, and the demonstration of developing computer programs to biomedical research collaborators.

I.A.1.1 What is Artificial Intelligence

Artificial Intelligence research is that part of Computer Science concerned with symbol manipulation processes that produce intelligent action [1 - 7]. By "intelligent action" is meant an act or decision that is goal-oriented, is arrived at by an understandable chain of symbolic analysis and reasoning steps, and utilizes knowledge of the world to inform and guide the reasoning.

Placing AI in Computer Science

A simplified view relates AI research with the rest of computer science. The ways in which people use computers to accomplish tasks can be "one-dimensionalized" into a spectrum representing the nature of the instructions that must be given the computer to do its job; call it the What-to-How spectrum. At the How extreme of the spectrum, the user supplies his intelligence to instruct the machine precisely how to do his job, step-by-step. Progress in computer science may be seen as steps away from that extreme How point on the spectrum: the familiar panoply of assembly languages, subroutine libraries, compilers, extensible languages, etc. illustrate this trend.

At the other extreme of the spectrum, the user describes What he wishes the computer to do for him to solve a problem. He wants to communicate what is to be done without having to lay out in detail all necessary subgoals for adequate performance. Still, he demands a reasonable assurance that he is addressing an intelligent agent that is using knowledge of his world to understand his intent, complain or fill in his vagueness, make specific his abstractions, correct his errors, discover appropriate subgoals, and ultimately translate What he wants done into detailed processing steps that define How it shall be done by a real computer. The user wants to provide this specification of What to do in a language that is comfortable to him and the problem domain (perhaps

English) and via communication modes that are convenient for him (including perhaps speech or pictures).

The research activity aimed at creating computer programs that act as "intelligent agents" near the What end of the What-to-How spectrum can be viewed as a long-range goal of AI research.

Expert Systems and Applications

The national SUMEX-AIM resource is an outgrowth of a long, interdisciplinary line of artificial intelligence research at Stanford concerned with the development of concepts and techniques for building "expert systems" [1]. An "expert system" is an intelligent computer program that uses knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solution. For some fields of work, the knowledge necessary to perform at such a level, plus the inference procedures used, can be thought of as a model of the expertise of the expert practitioners of that field.

The knowledge of an expert system consists of facts and heuristics. The "facts" constitute a body of information that is widely shared, publicly available, and generally agreed upon by experts in a field. The "heuristics" are the mostly-private, little-discussed rules of good judgment (rules of plausible reasoning, rules of good guessing) that characterize expert-level decision making in the field. The performance level of an expert system is primarily a function of the size and quality of the knowledge base that it possesses.

Currently authorized projects in the SUMEX community are concerned in some way with the application of AI to biomedical research (*). The tangible objective of this approach is the development of computer programs that will be more general and effective consultative tools for the clinician and medical scientist. There have already been promising results in areas such as chemical structure elucidation and synthesis, diagnostic consultation, and modeling of psychological processes.

Needless to say, much is yet to be learned in the process of fashioning a coherent scientific discipline out of the assemblage of personal intuitions, mathematical procedures, and emerging theoretical structure comprising artificial intelligence research. State-of-the-art programs are far more narrowly specialized and inflexible than the corresponding aspects of human intelligence they emulate; however, in special domains they may be of comparable or greater power, e.g., in the solution of formal problems in organic chemistry.

(*) Brief abstracts of the various projects can be found in Appendix A on page 278 and more detailed progress summaries in Section II on page 89.

I.A.1.2 Resource Sharing

Besides the biomedical AI research theme of SUMEX-AIM, another central goal is an exploration of the use of computer-based communications as a means for interactions and sharing between geographically remote research groups engaged in biomedical computer science research. This facet of scientific interaction is becoming increasingly important with the explosion of complex information sources and the regional specialization of groups and facilities that might be shared by remote researchers [8]. We expect an even greater decentralization of computing resources in the coming years with the emerging VLSI (*) technology in microelectronics and a correspondingly greater role for digital communications.

Our community building effort is based upon the current state of computer communications technology. While far from perfected, these developing capabilities offer highly desirable latitude for collaborative linkages, both within a given research project and among them. A number of the active projects on SUMEX are based upon the collaboration of computer and medical scientists at geographically separate institutions; separate both from each other and from the computer resource. The network experiment also enables diverse projects to interact more directly and to facilitate selective demonstrations of available programs to physicians, scientists, and students.

We have actively encouraged the development of additional affiliated computing resources within the AIM community and expect such decentralization to become the "way of the 80's". Since 1977, the facility at Rutgers University has allocated a portion of its capacity for national AIM projects and our network connections to Rutgers and common facilities for user terminals have been indispensable for effective interchanges between community members, workshop coordinations, and software sharing. In addition, the "Caduceus" project (**) (page 187) is expecting delivery of their own machine momentarily, the "Simulation of Cognitive Processes" project (page 226) already is doing most of their work on their own VAX computer, and several more projects have proposed machines dedicated to their own use.

The proliferation of distributed machines will serve to increase the importance of electronic communications to facilitate interactions and sharing. Even in their current developing state, communication facilities enable effective access to the SUMEX community resources from a great many areas of the United States and to a more limited extent from Canada, Europe, Japan, Australia, and other international locations.

(*) Very Large Scale Integration

(**) Previously called "Internist".

I.A.1.3 Impact of AI in Biomedicine

Artificial Intelligence is the computer science of symbolic representations of knowledge and symbolic inference. There is a certain inevitability to this branch of computer science and its applications, in particular, to medicine and biosciences. The cost of computers will continue to fall drastically during the coming two decades. As it does, many more of the practitioners of the world's professions will be persuaded to turn to economical automatic information processing for assistance in managing the increasing complexity of their daily tasks. They will find, from most of computer science, help only for those of their problems that have a mathematical or statistical core, or are of a routine data-processing nature. But such problems will be relatively rare, except in engineering and physical science. In medicine, biology, management -- indeed in most of the world's work -- the daily tasks are those requiring symbolic reasoning with detailed professional knowledge. The computers that will act as "intelligent assistants" for these professionals must be endowed with symbolic reasoning capabilities and knowledge.

The growth in medical knowledge has far surpassed the ability of a single practitioner to master it all, and the computer's superior information processing capacity thereby offers a natural appeal. Furthermore, the reasoning processes of medical experts are poorly understood; attempts to model expert decision making necessarily require a degree of introspection and a structured experimentation that may in turn improve the quality of the physician's own clinical decisions, making them more reproducible and defensible. New insights that result may also allow us more adequately to teach medical students and house staff the techniques for reaching good decisions, rather than merely to offer a collection of facts which they must independently learn to utilize coherently.

The knowledge that must be used is a combination of factual knowledge and heuristic knowledge. The latter is especially hard to obtain and represent since the experts providing it are mostly unaware of the heuristic knowledge they are using. Medical and scientific communities currently face many widely recognized problems relating to the rapid cumulation of knowledge, for example:

- codification of theoretical and heuristic knowledge
- effective use of the wealth of information implicitly available in textbooks, journal articles and from practitioners
- dissemination of that knowledge beyond the intellectual centers where it is collected
- customizing the presentation of that knowledge to individual practitioners as well as customizing the application of the information to individual cases

We believe that computers are the most hopeful technology to help overcome these problems. While recognizing the value of mathematical modeling, statistical classification, decision theory and other techniques, we believe that effective use of such methods depends on using them in conjunction with less formal knowledge, including contextual and strategic knowledge.

Artificial intelligence offers advantages for representing information and using it that will allow physicians and scientists to use computers as intelligent assistants. In this way we envision a significant extension to the decision making powers of individual practitioners without reducing the significance of the individuals.

Knowledge is power, in the profession and in the intelligent agent. As we proceed to model expertise in medicine and its related sciences, we find that the power of our programs derives mainly from the knowledge that we are able to obtain from our collaborating practitioners, not from the sophistication of the inference processes we observe them using. Crucially, the knowledge that gives power is not merely the knowledge of the textbook, the lecture and the journal but the knowledge of "good practice" -- the experiential knowledge of "good judgment" and "good guessing", the knowledge of the practitioner's art that is often used in lieu of facts and rigor. This heuristic knowledge is mostly private, even in the very public practice of science. It is almost never taught explicitly; almost never discussed and critiqued among peers; and most often is not even in the moment-by-moment awareness of the practitioner.

Perhaps the the most expansive view of the significance of the work of the SUMEX-AIM community is that a methodology is emerging therefrom for the systematic explication, testing, dissemination, and teaching of the heuristic knowledge of medical practice and scientific performance. Perhaps it is less important that computer programs can be organized to use this knowledge than that the knowledge itself can be organized for the use of the human practitioners of today and tomorrow.

The researchers of the SUMEX-AIM community currently constitute a large fraction of all the computer scientists whose work is aimed at the development of symbolic computational methods and tools. SUMEX-AIM is laying the scientific base so that medicine will be able to take advantage of these technological opportunities for inexpensive computer power. Medical diagnostic aids and tools for the medical scientist that operate in a environment of a network of "professional workstation" computers have the practical possibility of large-scale and low-cost use because of anticipated near-term developments in the computing industry.

I.A.2 Synopsis of Recent Progress

As we complete year 08, we can report substantial further progress in the overall mission of the SUMEX-AIM resource. We have continued the refinement of an effective set of hardware and software tools to support the development of large, complex AI programs for medical research and to facilitate communications and interactions between user groups. We have worked to maintain high scientific standards and AI relevance for projects using the SUMEX-AIM resource and have actively sought new applications areas and projects for the community. Many projects are built around the communications network facilities we have assembled; bringing together medical and computer science collaborators from remote institutions and making their research programs available to still other remote users. As discussed in the sections describing the individual projects, a number of the computer programs under development by these groups have matured into tools increasingly useful to the respective research communities. The demand for production-level use of these programs has surpassed the capacity of the present SUMEX facility and in preparation for our renewal goals, we have been investigating the general issues of how such software systems can be moved from SUMEX and supported in production environments.

A number of significant events and accomplishments affecting the SUMEX-AIM resource occurred during the past year:

- 1) In August 1980, under the chairmanship of Prof. Ted Shortliffe and with the assistance of Drs. L. Fagan and R. Blum, Stanford hosted the sixth AIM workshop. This workshop was innovative in that the presentations were fully "demo-based" using a live video projection of program typescripts and actual running sessions. The purpose of this approach was to allow participants to see more deeply into the inner workings of the various systems under development.
- 2) In conjunction with the 1980 workshop, Drs. Clancey and Shortliffe organized a continuing education tutorial for practicing physicians. The tutorial session was attended by over 135 doctors and included an introduction to computing, background information on decision theory and database applications in medicine, and presentations on a number of AI systems by 15 members and affiliates of the SUMEX-AIM community.
- 3) In November 1980, we defended our pending renewal application before a peer review site visit team. The SUMEX-AIM community was represented by several members of the AIM Executive Committee. A strong endorsement for future SUMEX goals and a recommendation for a 5-year renewal period resulted. These were confirmed by study section and council action. The technical substance of our future goals are outlined beginning on page 47.
- 4) The SUMEX-AIM collaborator project community has continued vigorous development of their respective programs. Details are reported by the individual investigators in Section II. The VM and ONCOCIN projects have begun preliminary clinical testing/evaluation this past year using SUMEX network and computing resources. The CADUCEUS

(INTERNIST) and SIMULATION OF COGNITIVE PROCESSES projects have been funded for and are setting up their own local VAX computing resources which should help reduce the load on SUMEX for newer pilot efforts. We have continued to work hard to meet the needs of collaborating projects and are grateful for their expressed appreciation.

- 5) We supported a highly successful, experimental dissemination of the MOLGEN programs into the molecular biology community. "Advertised" through presentations and demonstrations by MOLGEN investigators at several professional conferences, over 200 molecular biologists have used the system and most have found it easy to learn and highly effective as a research tool for their investigations.
- 6) We have continued development of the SUMEX facility hardware, software, and network systems to enhance throughput and to assist user access to existing and planned resources. A good range of internetwork software is available now including telnet, file transfer, and mail handling. Following the council recommendation for approval of our renewal application, our request to augment the AMPEX memory was funded by BRP. We have installed the new memory and are in the process of tuning the monitor to optimize use of the increased user memory.
- 7) We have actively explored options for professional workstation and VAX LISP systems in preparation for our renewal research. The current state of available systems is encouraging. However, delays in an operational version of Interlisp-VAX and an earlier than expected availability of Interlisp-Dolphin workstations has led us to recommend beginning the workstation phase of our research first.

I.A.3 Details of Technical Progress

The following material covers SUMEX-AIM resource activities over the past year in greater detail. These sections outline accomplishments in the context of the resource staff and the resource management. Details of the progress and plans for our external collaborator projects are presented in Section II beginning on page 89.

I.A.3.1 Facility Hardware

Over the past year, the SUMEX facility hardware configuration, including the main KI-10 machine (Figure 1), the 2020 satellite machine (Figure 2), and system network interconnections (Figure 3), have continued to develop according to plan and to operate effectively within capacity limitations. The primary facility hardware development efforts this year have been directed at:

- 1) Augmentation of the 256K word AMPEX memory to 512K words.
- 2) Implementation of Ethernet interface equipment for the KI-10 and other network server facilities.
- 3) Investigation and planning of hardware alternatives for the system development goals of our renewal grant.
- 4) Support of local project hardware needs.

Memory Augmentation

The SUMEX-AIM facility has been operating at capacity in terms of prime-time computing load for the past several years as documented in our previous reports. In spite of implementing a number of strategic facility augmentations over the years, we have not been able to satisfy the computing demands of our community. This condition has constrained the growth of the AIM community and our ability to bring AI programs nearing operational status in contact with potential external user communities while continuing to support on-going program development efforts. We have taken active steps to transfer prime time interactive loading to evening and night hours as much as possible including shifting personnel schedules (particularly for Stanford-based projects). We have implemented tools to control the fair allocation of CPU resources between various user communities and projects and have encouraged jobs not requiring intimate user interaction to run during off hours using batch job facilities. And we have acquired a 2020 system to offload program demonstrations and evaluations from the main research machine. Despite these efforts, our prime time loading has remained at saturation. Perhaps the most significant effect of the resulting poor response time is the deterrence of interactions with medical and other professional collaborators experimenting with available AI programs, whose schedules cannot be adjusted to meet computer loading patterns.

From the SUMEX viewpoint, we have attempted to do everything feasible and economically justified within available budgets to maximize the use of the existing hardware for productive work. One remaining step has been the expansion of our AMPEX memory from its current 256K word complement to its full 512K word capacity. The effect of this upgrade is to make more physical memory available to user programs thereby reducing swapping overhead (page faults and interrupt handling) and smoothing out system responsiveness under heavy load by keeping more working sets in core.

We requested approval for this expansion in May 1980. Following council approval of our renewal grant application, we received funding for the upgrade. The added memory was received and installed May 14, 1981, checked out during the following week, and a new 786K monitor brought up on May 21. This addition has increased user memory by about 60%. It is still too early to draw detailed conclusions about the effect of this enhancement and further tuning of monitor parameters controlling process scheduling and working set management needs to be done. We will report detailed results next year.

Local Network Interfaces

The initial design of the SUMEX system was that of a "star" topology centered on the KI-10 processors. In this configuration, all peripheral equipment and terminal ports were connected directly to the KI-10 busses. With the addition of new satellite machines, a unique focus no longer exists and some pieces of equipment need to be able to "connect" to more than one host. For example, a user coming into SUMEX over TYMNET will want to be able to make a selection of which machine he connects to. Another TYMNET user may want to make another choice of machine and so the TYMNET interface needs to be able to connect to any of the hosts. This could be accomplished by creating separate interfaces for each of the hosts to the TYMNET, each with a different address. Besides being expensive to duplicate such interface connections, it would be inconvenient for a user to reconnect his terminal from one host to another.

Over the past year and a half we have been developing a local, high-speed Ethernet to provide a flexible basis for our planned facility developments. The KI-10's and the 2020 were connected in time to support the AIM workshop last summer. Our development of Ethernet facilities has been guided by the goals of providing the most effective range of services for SUMEX community needs while remaining compatible with and able to contribute to and draw upon network developments by other groups. Since the early 3 Mbit/sec Ethernet was given to Stanford and several other universities by Xerox, an agreement has been reached between DEC, INTEL, and Xerox on the standards for an even higher performance network [13]. The new network runs at 10 Mbits/sec and supports a significantly larger packet address space. Xerox has started to market products for the new network but debugged interfaces, software, etc. for general use are not routinely available yet. Furthermore, even though three companies have agreed on a set of low level protocols and interface conventions, the rest of the world may not go along. There is already an alternative (but closely related) IEEE specification in preparation. Even among the three parties in the Ethernet specification, there is no agreement on higher level protocols.

All of this suggests that it is not time to jump to the newer and faster networks yet. We feel the 3 Mbit/sec network is adequate for our bandwidth needs in the near future and there is already a significant investment in 3 Mbit/sec network equipment at Stanford related to SUMEX community interests. In the longer term, we will want to upgrade to whatever hardware and protocol standard is broadly adopted.

In the meantime we are continuing to develop our 3 Mbit/sec PUP network services. This places a heavier burden on us to develop and maintain our own equipment for Ethernet support. We have tried to minimize the "home-brew" nature of this work by sharing common hardware and software designs with other groups in the same situation

The initial KI-10 interface was made via a PDP-11 connected to the I/O bus which is inefficient under heavy traffic. In anticipation of increased Ethernet demands on the KI-10's for high-speed terminals, file transfers, and other server functions, we have been designing and implementing a more efficient direct memory access interface. This interface uses a phase decoder (design borrowed from the SUN terminal project at Stanford) to detect the incoming serial Ethernet signal, an internal packet buffer to prevent overruns to and from the TENEX time-sharing system, and a memory bus interface to transfer data. The KI-10 DMA interface is partially debugged while highest priority work is proceeding on a gateway to the computer science building across campus.

In our initial connections of the KI-10 and 2020, we used a UNIBUS interface board designed by E. Markowski at Xerox. Because of the limited availability of these boards for our future work (an immediate need being for a gateway between various campus Ethernets), we began work on a PDP-11 interface board. This design is similar to that of the KI-10 interface and shares the serial phase decoder network front end. It provides several features not available on the Xerox board including more explicit error information and a more sophisticated filter on source addresses for incoming packets.

Planning for the Renewal Period

Over the past year we have spent considerable effort evaluating strategies and alternatives for planned system development in our renewal grant. Pending funding, council has approved our plan to acquire two VAX machines, five professional workstations, and a file server for the SUMEX resource starting in August 1981. We have debated at length the appropriate timing for purchases of this equipment within budget constraints. The Initial Review Group and Council enthusiastically endorsed the importance of optimizing the timing of our planned hardware acquisitions to coincide with the availability of desired technological developments and community needs. They recommended in their report that we be allowed considerable flexibility as to phasing of equipment purchases within the 5-year renewal period.

The rapidly changing technical and commercial situation vis a vis the research computing equipment we plan to buy if funded, indicates that there would be significant advantage to the SUMEX-AIM community in exercising

this flexibility by delaying the purchase of our first VAX until the second renewal year (grant year 10) and advancing the purchase of the Professional Workstations to the first year (grant year 09). The rationale for this switch is as follows:

- 1) The INTERLISP language has been the basis for most SUMEX-AIM community AI research. Development of the VAX INTERLISP system at USC-ISI is substantially behind schedule. The most current estimate for completing a usable system is mid-1982 and no viable alternative version of LISP, with a fully developed programming support environment, will be available any sooner than that. Thus, if we purchased a VAX in year 09, we could not offer effective VAX LISP services before year 10. Strong pressure does exist within the ARPANET community to get VAX INTERLISP completed as expeditiously as possible so we believe that VAX will be a good machine choice by year 10 once INTERLISP is running. We are undertaking a separate study of this situation to assess the likelihood of VAX/INTERLISP being completed in a timely fashion and to estimate its performance characteristics on the VAX 11/780.
- 2) If the interchange in timing of the purchases of the first VAX and the Professional Workstations is approved, we have agreed that SUMEX-AIM will have shared access next year to the VAX 11/780 funded by ARPA to support Stanford Heuristic Programming Project research. This will minimize any delays in SUMEX-AIM work involving VAX that is not dependent on INTERLISP and will enable necessary systems development work and preliminary experimentation by SUMEX-AIM users to proceed without having to commit NIH grant funds. Because of long term commitments for the ARPA VAX and expected growth in SUMEX community needs, however, it can only substitute on a temporary basis during the first renewal grant year. After that a VAX dedicated solely to SUMEX-AIM community use will be needed.
- 3) The DEC VAX product line is continuously changing and there are some indications that new products may be offered on the 1982-1983 time scale that would be advantageous to SUMEX-AIM research. These may include features that enhance technical performance and/or cost effectiveness for our purposes. By year 10, we should be able to make a more judicious choice of the best configuration for our needs.
- 4) While VAX/INTERLISP is delayed, a suitable model of the professional workstation we need for our experimentation is available earlier than expected. The Xerox Dolphin is a system that has been in use as a research machine within Xerox for some time. It meets our technological needs including a high-bandwidth bit-mapped display terminal, full TENEX INTERLISP software compatibility, increased address space over the PDP-10 (but not as large as will be available on the VAX), acceptable capacity (roughly twice the single-user KA-10 speed), and existing Ethernet hardware/software support. Dolphins will be produced shortly in limited quantity by Xerox EOS, primarily for the ARPANET computer science community, and will be available for delivery beginning in August 1981. Their cost is currently higher than that expected for comparable systems several years from now. However, the

immediate purchase of the limited quantity planned will be cost-effective in allowing research to proceed in the SUMEX-AIM community on software that will be needed to exploit these later systems.

- 5) If we purchase the five INTERLISP-Dolphins in year 09, a significant increase in LISP processing capacity will be added to the SUMEX-AIM resource earlier than would be possible with VAX/INTERLISP. Even though these are intended primarily for stand-alone use, they nevertheless will afford badly needed relief for the overloaded central machines since the people using the Dolphins will not be running INTERLISP simultaneously on the KI-10's or 2020. In quantitative terms, taking a Dolphin to be about equal in speed to two KA-10's, the five Dolphins will roughly treble our current dual KI-10/2020 computing capacity.

Based on this rationale, it seems clearly to be in the best interest of the SUMEX-AIM community to delay the acquisition of the first VAX system and to accelerate the purchase of the five Professional Workstations.

Other Hardware Development

We have undertaken other hardware efforts as appropriate during the past year. Most significant of these was the development of a controller for a printer in the Stanford Oncology Clinic to support the ONCOCIN evaluation getting underway. This printer is part of an existing internal information system in use by the clinic. In order to integrate the printout from ONCOCIN sessions, we needed to provide a flexible connection to the SUMEX facility spoolers. We built a Z80-based microprocessor controller that senses status of the printer and performs buffering, flow control, and data rate conversion so it can act as a remote printer to the SUMEX machines when needed for ONCOCIN sessions.

In addition we have provided broad support to users for terminal and communications connections and repairs.

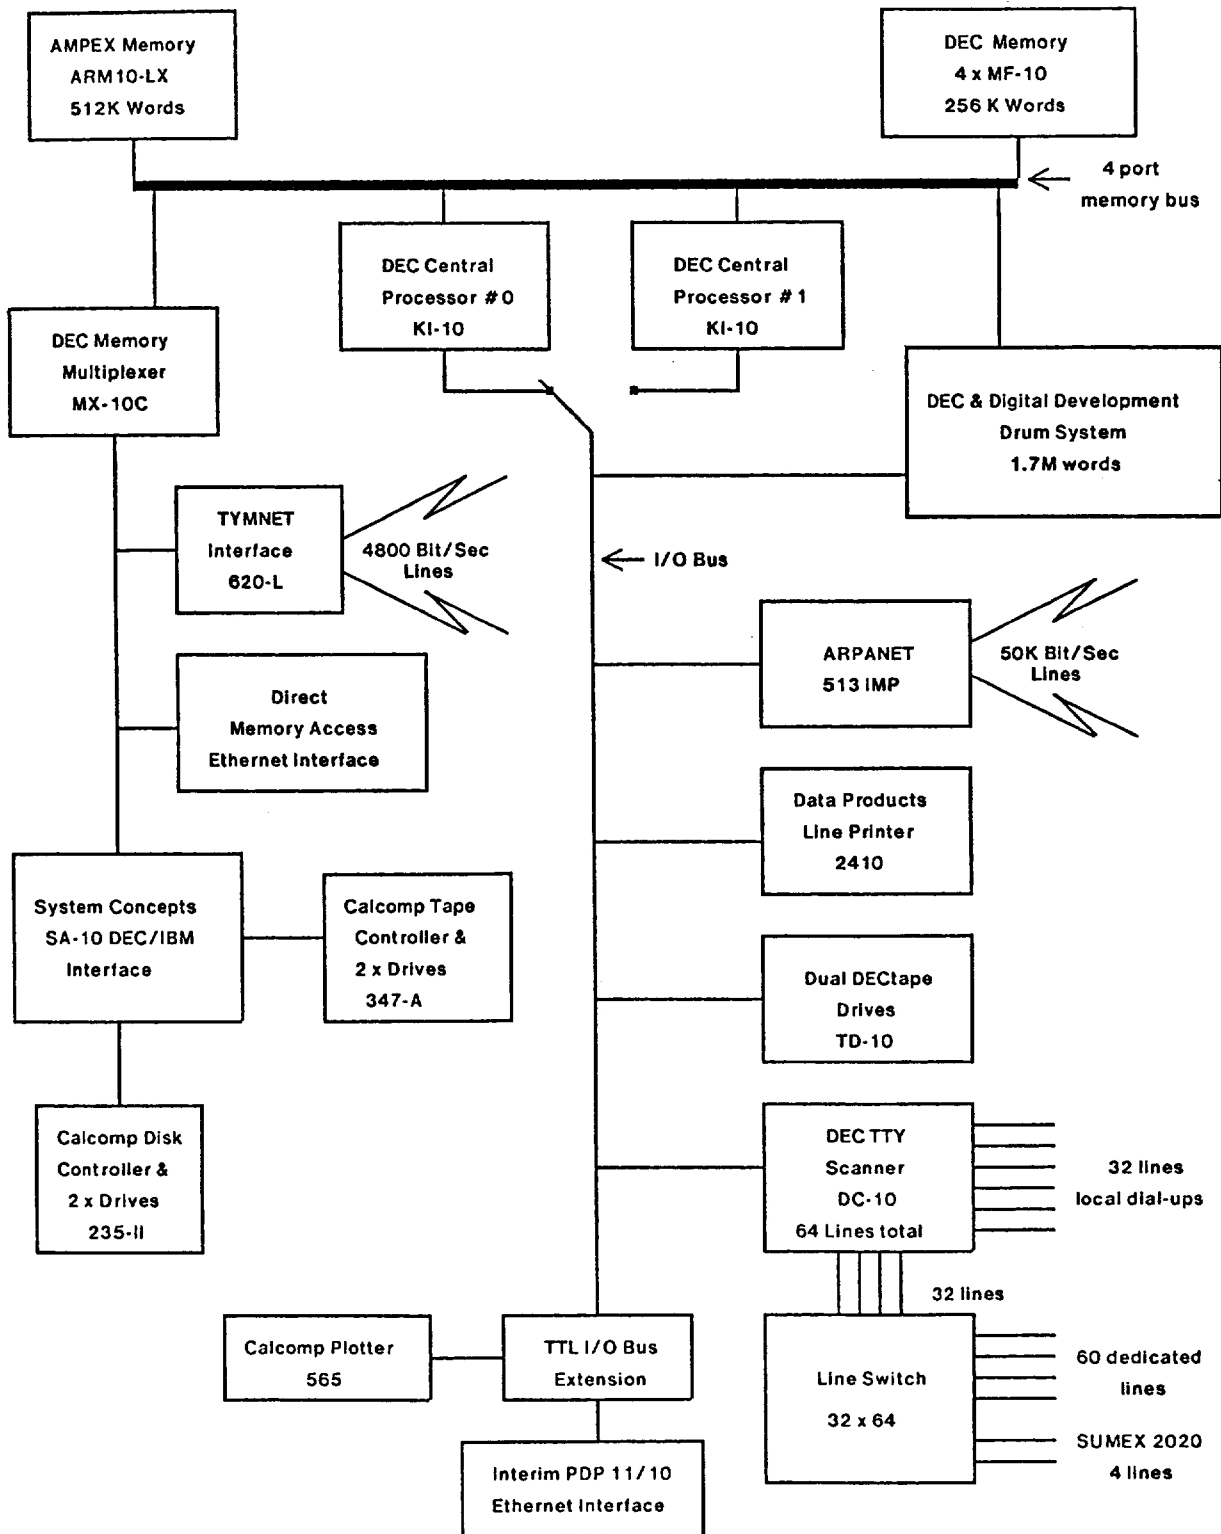


Figure 1. Current SUMEX-AIM KI-10 Computer Configuration

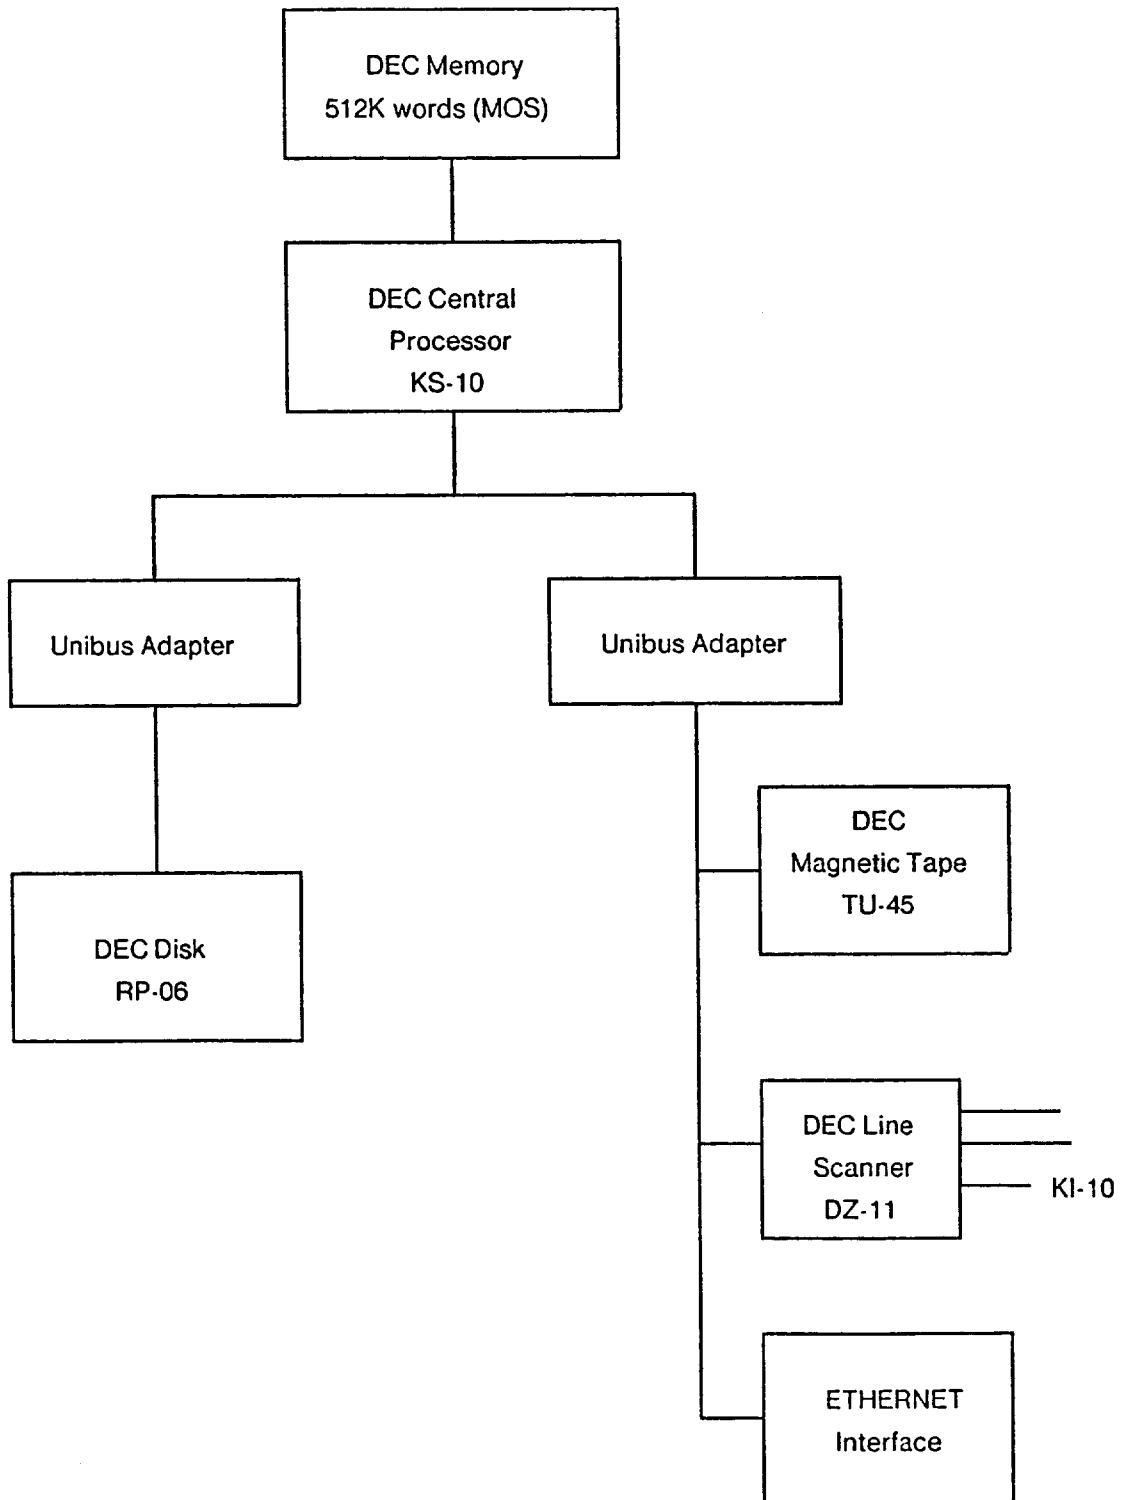


Figure 2. Current SUMEX-AIM 2020 Computer Configuration

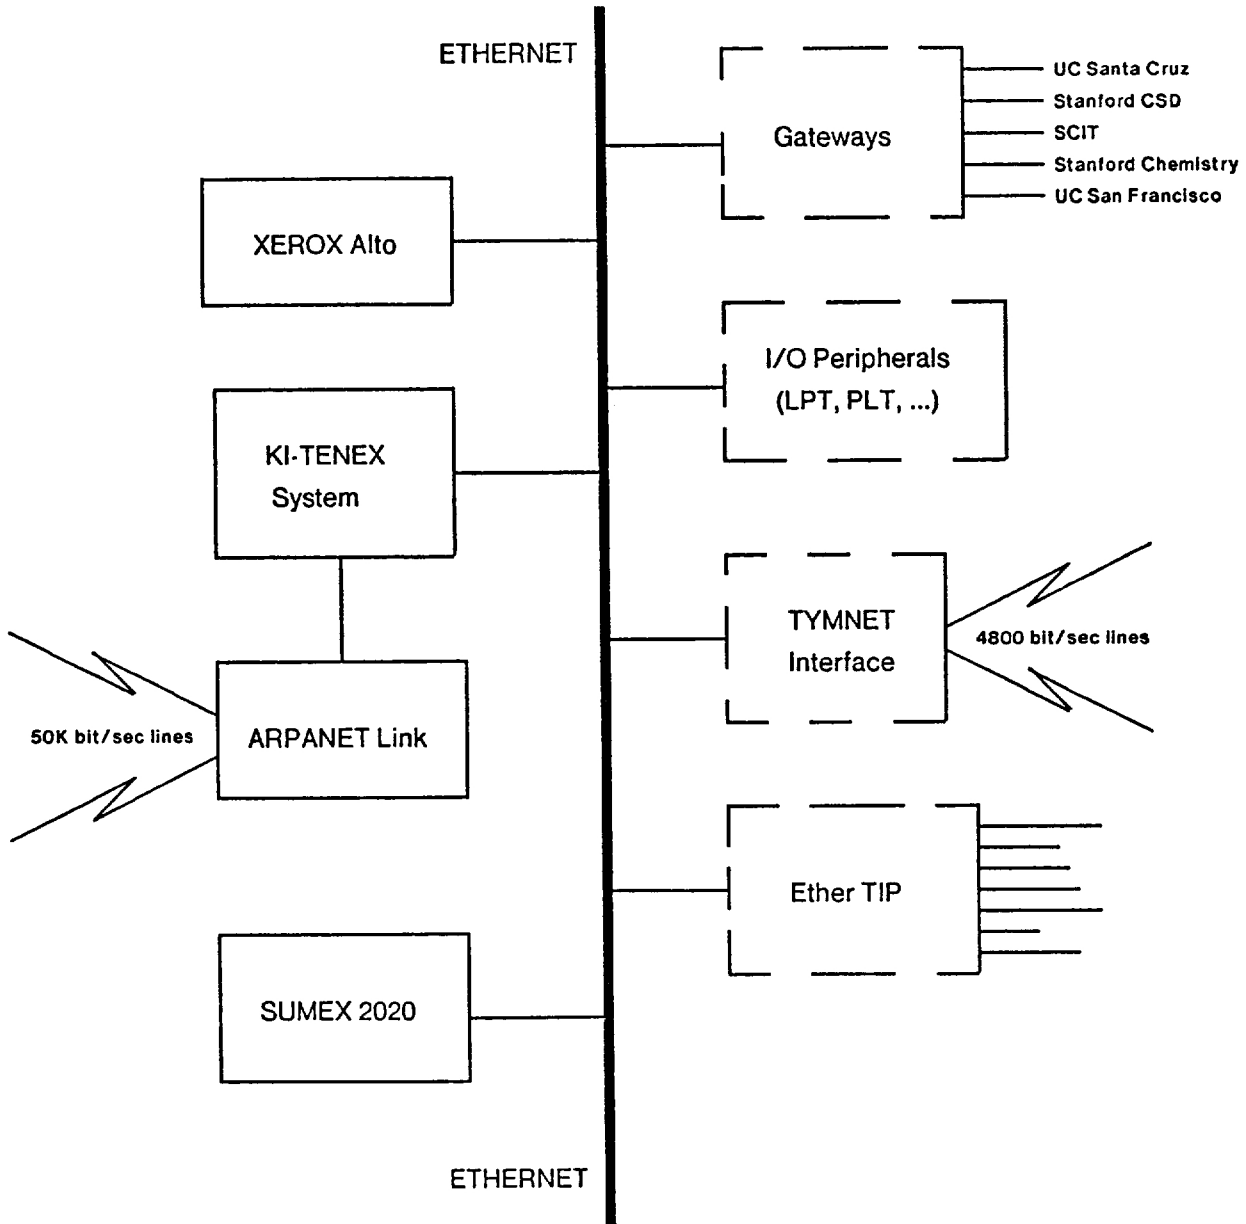


Figure 3. Intermachine Connections via ETHERNET

I.A.3.2 System Software

Our monitor software work this past year has concentrated on several areas including changes to support hardware development projects, upgrading and enhancing network interface service, correcting encountered system bugs, and implementing new features for better user community support. In addition we have invested substantial effort in becoming familiar with the VAX/UNIX system which will play a key role in our future research efforts.

Hardware Implementation

In parallel with our principal hardware efforts this past year to extend and improve local Ethernet connections, the necessary monitor changes to support the new hardware are being made. The largest effort, for which debugging is still on-going, has been the direct memory access Ethernet interface for the KI-10 duplex. This interface has been partially completed, including developing new interrupt service routines and facilities for doing hardware debugging during time-sharing so as not to disrupt availability of the system. Completion of this work has been delayed by placing higher priority on building a gateway connection to the Department of Computer Science building across campus. Since the KI-10's have a working, albeit inefficient, PDP-11 interface already, we decided our limited development resources were better used in establishing this badly needed new capability.

Additional work has gone into upgrading software support for the terminal hardline switch (SLM) developed last year. These improvements were to correct several problems in assigning terminals to available lines and to improve user feedback on system status while negotiating for a connection.

Network Interface Service

Effective January 1, 1981, the ARPANET formally changed the standard for packet "leaders" to allow addressing more hosts on an Interface Message Processor (IMP) and more IMP's on the network. This change required substantial upgrades to the monitor ARPANET service routines including the internal handling of data packets and two new JSYS's that communicate with user programs about network information. We imported much of the new code from ARPANET sites working on the development of network software (especially USC-ISI and SRI) but considerable work was required to adapt it to our dual-processor monitor and operating environment. The changeover went extremely smoothly with most users unaware that a change was taking place. We expect further changes to be required by early 1983 when the higher level communication protocols will move from NCP (network control protocol) to TCP (transmission control protocol).

This past year we have also continued to develop the Ethernet PUP software including improved hardware interfaces discussed above and numerous bug fixes. Many of the bug fixes relate to interactions between the PUP management software and other parts of the monitor such as the teletype handler. The Ethernet software is running very reliably now.

Monitor Bug Fixes and Improvements

We have continued to repair important bugs in our TENEX monitor. In general the system runs extremely reliably with most problems coming from explicit hardware malfunctions or periods of instability following significant monitor changes.

We found an additional number of subtle bugs in the system this past year that had been causing various problems. By now, all of the "obvious" bugs have been located and so those remaining are much more elusive, occurring infrequently or only after a long chain of rare events that is difficult to reconstruct. Examples of fixes include:

- 1) After an extended period of uptime, TYMNET users found all the ports to SUMEX in apparent use. This only happened after about 6.5 days of continuous uptime, itself a rather rare event. After a long search, we found an invalid index into one of the TYMNET connection database arrays which instead of testing the appropriate state bit, was testing a high order bit of a timer field. Thus, when that bit came on after being up for a long period, the test erroneously detected the port in use.
- 2) With the installation of the operational Ethernet, the overall timing of system functions changed, including the management of the drum service. Commands for page transfers were sent to the drum controller asynchronously as the requests were placed on the queue. This was done in the drum interrupt service and timing was such that new transfers could be posted "on the fly". As the Ethernet became operational, the timing of interrupt handling changed so that attempts to post these new transfer requests came at the wrong time for the drum controller and caused command sequence errors. A temporary fix was made to avoid this conflict but we still want to rework the drum management software to optimize performance.
- 3) Finally, users are invariably able to design system call arguments that present special cases to the monitor routines which don't work. We have repaired a number of such problems in various string handling JSYS's and in the floating point output JSYS.

System Loading Controls

We previously reported on the system load controls we have implemented on the KI-10 duplex to allocate available system capacity effectively among projects and users according to Executive Committee guidelines. These continue to operate effectively and we have not made any substantial changes in this area. All communities (National, Stanford, and Staff) are under load controls now. We have adjusted relative priorities for projects in the national community in accordance with Executive Committee reviews of the community in August 1980.

We have instituted a mechanism for reserving the 2020 for demonstrations and developmental testing of various expert systems (e.g., DENDRAL, ONCOCIN, etc.). Because of the unpredictability of usage during

these reserved times, we feel that too much of the 2020 capacity is lost by simply dedicating the machine to such users. We are now reevaluating the reservation system, probably in favor of a "pie-slice" system that will guarantee dedicated users a large fraction of the machine but which allows other useful work to go on when their demand is low.

Executive Program

We have made several changes and improvements in the SUMEX EXECutive program this past year:

- 1) Many of the features of the EXEC that enhance its "friendliness" require access to auxiliary files. When we come up after a crash and there is file system damage to repair, these files may be compromised and in general extraneous file access at such times is undesirable. Thus we carefully reviewed the internals of the EXEC and made changes so that in debugging mode it operates "bare bones". All unneeded file accesses and interactions are eliminated.
- 2) Because of the difficulty in collecting definitive data about user experiences with network connections, we implemented a log of involuntary disconnects from network terminals. This log allows us to better correlate disconnects, looking for instances when all TYMNET users are disconnected or all users from a given node are disconnected as opposed to drops by individual users which may be caused by hanging up the telephone. We have now collected a database of these disconnects covering several months and indications are that TYMNET users are being dropped occasionally through some sort of network glitch. We are developing programs to better analyze these data so we can distinguish problems at the SUMEX end from those in the network so appropriate solutions can be worked out.
- 3) We implemented several layers of access constraints for GENET users (see page 84) including a limit to the number of simultaneous login's and a requirement for a user password to restrict access for commercial users. These developments have in fact limited the growth of the GENET community as recommended by AIM Executive Committee policy.

I.A.3.3 Network Communications

A highly important aspect of the SUMEX system is effective communication with remote users and between the growing number of machines available within the SUMEX resource. In addition to the economic arguments for terminal access, networking offers other advantages for shared computing. These include improved inter-user communications, more effective software sharing, uniform user access to multiple machines and special purpose resources, convenient file transfers, more effective backup, and co-processing between remote machines.

We continue to base our remote communication services on two networks - TYMNET and ARPANET for reasons detailed in previous annual reports. Users asked to accept a remote computer as if it were next door will use a local telephone call to the computer as a standard of comparison. Current network terminal facilities do not quite accomplish the illusion of a local call. Data loss is not a problem in most network communications - in fact with the more extensive error checking schemes, data integrity is higher than for a long distance phone link. On the other hand, networking relies upon shared community use of communication lines to procure widespread geographical coverage at substantially reduced cost. However, unless enough total line capacity is provided to meet peak loads, substantial queueing and traffic jams result in the loss of terminal responsiveness. Limited responsiveness for character-oriented TENEX interactions continues to be a problem for network users.

TYMNET

TYMNET provides broad geographic coverage for terminal access to SUMEX, spanning the country and also increasingly accessible from foreign countries. Technical aspects of our connection to TYMNET have remained unchanged this past year and have continued to operate reasonably reliably. As noted earlier, however, users have complained periodically about having their connections dropped and we have implemented a data collection facility in the EXEC program to help document and classify these failures. There are definitely episodes in which all connections are lost and the jobs are detached. These occur about once every few days but we are still analyzing these data to try to separate out local from network causes.

TYMNET has made few technical changes to their network that affect us other than to broaden geographical coverage. The previous network delay problems are still apparent although better cross-country trunks into New York and New England are available improving service there. TYMNET is still primarily a terminal network designed to route users to an appropriate host and more general services such as outbound connections originated from a host or interhost connections are only done on an experimental basis. This presumably reflects the lack of current economic justification for these services among the predominantly commercial users of the network. Whereas TYMNET is developing interfaces meeting X.25 protocol standards, the internal workings of the network will likely remain the same, namely, constructing fixed logical circuits for the duration of a connection and multiplexing characters in packets over each link between network nodes from any users sharing that link as part of their logical circuit.

We have continued to purchase TYMNET services through the NLM contract with TYMNET, Inc. Because of current tariff provisions, there is no longer an economic advantage to this based on usage volume. SUMEX charges are computed on its usage volume alone and not the aggregate volume with NLM's contribution to achieve a lower rate. We have implemented the "dedicated port" charging system for SUMEX use and have realized a substantial reduction in monthly usage costs. We will continue to work closely with NIH-BRP and NLM to achieve the most cost-effective purchase of these services.

ARPANET

We continue our advantageous connection to the Department of Defense's ARPANET, now managed by the Defense Communications Agency (DCA). Current ARPANET geographical and logical maps are shown in Figure 4 and Figure 5 on page 23. This connection has facilitated close collaboration with the Rutgers-AIM facility which is also on the net. Consistent with our long-standing agreements with ARPA and DCA we are enforcing a policy that restricts the use of ARPANET to users who have affiliations with DoD-supported contractors and system/software interchange with cooperating network sites. We are somewhat unique in this policy among other network sites since NIH has not become a member of the "sponsor's group" for the network. We would strongly encourage this step so that biomedical users could have more uniform access to the superior facilities of the ARPANET. This will become increasingly important as more NIH-sponsored sites desire access to the net and each other.

We have maintained good working relationships with other sites on the ARPANET for system backup and software interchange. Such day-to-day working interactions with remote facilities would not be possible without the integrated file transfer, communication, and terminal handling capabilities unique to the ARPANET. The ARPANET is also key to maintaining on-going intellectual contacts between SUMEX projects such as the Stanford Heuristic Programming Project authorized to use the net and other active AI research groups in the ARPANET community.

As indicated in the discussion of monitor software development, we implemented a significant change in ARPANET software support this past January 1, 1981. This change added support for the extended (96-bit) leader for packets that allow more Interface Message Processors (IMP) on the network and more ports per IMP. Substantial changes to the monitor network control program were necessary as well as to various user-level programs (TELNET, FTP, NETSER, RSSER, NETSTAT, etc.). The changeover went extremely smoothly with most users unaware of any effect.

ETHERNET

A substantial portion of our system effort this past year went into continued development of local network facilities to refine the connection between the KI-10 duplex and the 2020, to extend our network ties to other parts of campus (especially to the Computer Science Department building where the Heuristic Programming Project sits), and to prepare for the addition of new hardware in the renewal grant. As indicated in the earlier sections on monitor software and hardware, much has been done to implement more effective and efficient low-level system network connection facilities for our host systems.

We have also developed a number of software tools as a basis for implementing various kinds of Ethernet servers. These have been done in the language C, primarily because it is the language on which UNIX is based, has an active support community, and is being used for other network software that may be useful for our work. Specific areas of development include:

- 1) Server operating system: We have developed a simple operating system for use in servers that provides low-level interface to the Ethernet, hardware dependent interrupt service, process scheduling capabilities, and a series of defined monitor calls for invoking communication functions. This system is written initially for the PDP-11 but will also be portable to the MC-68000.
- 2) Higher Level Protocols: We have written routines that provide datagram, rendezvous/termination, and byte sequential protocol facilities on which other services such as EFTP, TELNET, etc. can be based.
- 3) We have written software for an Ethernet-to-Ethernet gateway that will establish connections between the SUMEX machine room and the Computer Science Department across campus. This system runs currently on a PDP-11/10 and supports dynamic assimilation of a routing table, periodic broadcast of this information to other hosts on connected networks, routing of addressed packets between connected networks, forwarding of key broadcast packets to allow distribution of network directories, and recording of gateway event status reports.
- 4) We have developed a wide range of diagnostic programs to assist in Ethernet software development including hardware diagnostics and downloading and debugging software.
- 5) We are actively working on the design of an Ethernet TIP to provide more terminal ports for the SUMEX system.

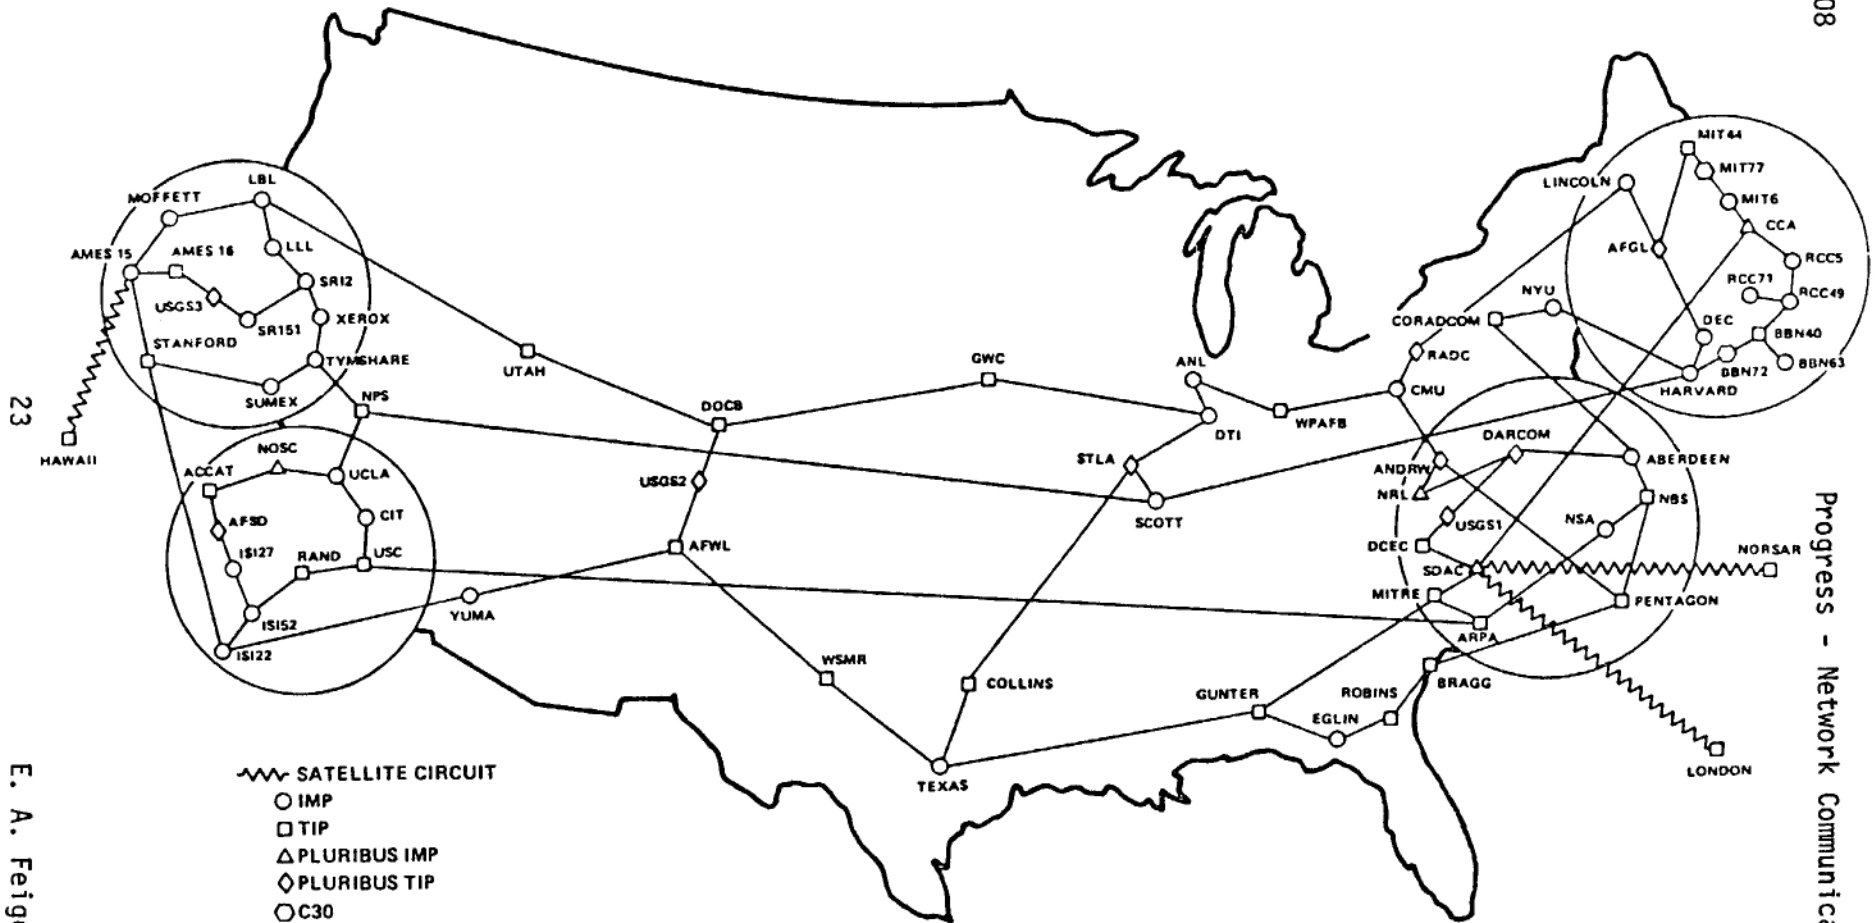
INTERNET SOFTWARE

One of the issues confronting the development of complex network-based systems, interconnected by gateway machines, is the support of internet communication of various kinds. For example, when a user at one of the Stanford Ethernet hosts wants to send a message to someone at MIT on a Chaosnet host, the mail handling programs have to know how to do the routing and the mail server programs have to be prepared to receive such mail for forwarding. Similarly, when establishing terminal telnet connections between such sites, the path of the link should be established automatically with the intervening sites merely acting as relay stations. In conjunction with groups at MIT and Stanford CSD, we have been developing prototypical systems for internet mail handling and telnet connections. The mail system is most highly developed and currently knows how to route messages between hosts on the Stanford Ethernet, the ARPANET, the MIT Chaosnet, the MIT LCSnet, and the Dialnet. This system has been operating since February.

We are also running a version of TELNET developed by Mark Crispin at SU-SCORE that allows a user to establish a connection across network boundaries without having to log into each intervening gateway and telnet further to the next station of a path to the desired destination host.

Figure 4. ARPANET Geographical Network Map

ARPANET GEOGRAPHIC MAP, MARCH 1981



(NOTE: THIS MAP DOES NOT SHOW ARPA'S EXPERIMENTAL SATELLITE CONNECTIONS)
 NAMES SHOWN ARE IMP NAMES, NOT (NECESSARILY) HOST NAMES

ARPANET LOGICAL MAP, MARCH 1981

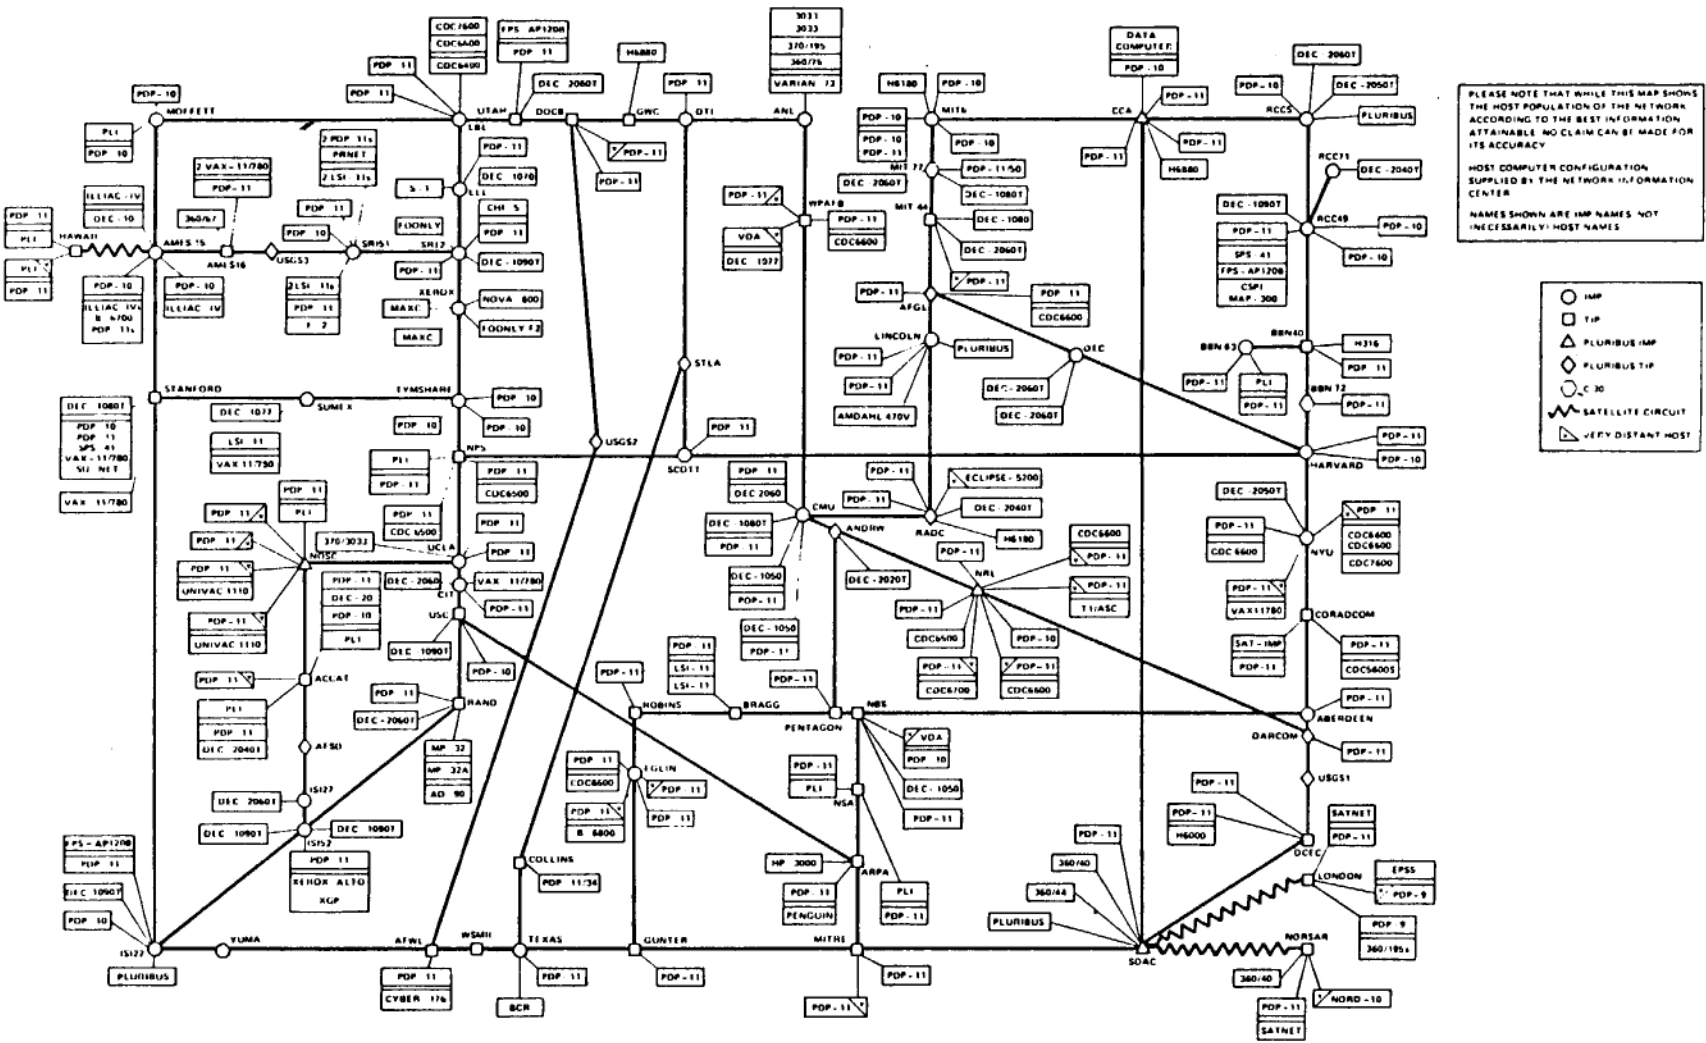


Figure 5. ARPANET Logical Network Map

I.A.3.4 User Software

We have continued to assemble and maintain a broad range of utilities and user support software. These include operational aids, statistics packages, DEC-supplied programs, improvements to the TOPS-10 emulator, 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 many more. Over the past year we have undertaken several significant development efforts to provide needed new programs to the SUMEX-AIM community. These include:

- 1) TTYFTP - Many groups have had the need to move files between computers and do not have the sophisticated facilities of ARPANET or other local networks to help. These include for example the transfer of data between the PUFF project at Pacific Medical Center in San Francisco and SUMEX and the movement of instrument data in support of the Ultrasound Imaging (Ob-Gyn) project. We reported last year on the development of a file transfer program usable over any teletype line (hardline, dial-up, TYMNET, etc.) which incorporates appropriate control protocols and error checking. The design is based on the DIALNET protocols developed by Crispin at the Stanford AI Laboratory and extended by our group to achieve machine and data source independence. This past year we have had a number of requests from outside groups in similar situations for copies of this software. We have distributed copies to Rutgers, Stanford Research Institute, and the University of Texas.
- 2) C Compiler - We spent considerable effort bringing up a workable C compiler at SUMEX that would generate code for our KI-TENEX system and also cross compile to generate code for PDP-11's and other machines. We imported an early version of a TOPS-20 C compiler from MIT and adapted it for our system. The linker, code generator, and runtime package for this system were suitably modified to work under TENEX and code generators for other machines developed. The PDP-10 version still generates quite inefficient code in that bytes occupy full 36-bit words rather than being packed. This is satisfactory for debugging purposes but would have to be fixed if C were to become a system programming language for future TENEX work (we do not anticipate this).
- 3) TV Editor and Display Terminal Support - Much work was done to extend the TV editor which is widely used at SUMEX. This was done by importing the work done by Hedrick at Rutgers, adapting it to our needs, and extending it for additional features. Important improvements include multiple string searches, string replacement facilities, large block text relocation, and support for additional display terminals (Infoton, Zenith H-19, Concept-100, ADDS Regent 60, and Hewlett-Packard 2600 series terminals). We have also agreed to unify the sources for TV so that closer compatibility with other groups will exist (Rutgers, SUMEX, USC-ISI, SRI, and Stanford CONTEXT).