

Figure 8: SUMEX-AIM EtherNet Configuration

III.A.3.7. 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 gain access to the central resource system 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.

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). It was Stanford's active on-going research in medical computer science, plus a world-wide reputation for the excellence and rigor of those research efforts, that persuaded the University that the field warranted a new academic degree program in the area. A group of faculty from the medical school and the computer science department argued that research in medical computing has historically been constrained by a lack of talented individuals who have a solid footing in both the medical and computer science fields. 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 graduate students. We do not wish to provide too narrow a definition of what kinds of prior training are pertinent because of the interdisciplinary nature of the field. The program has accordingly encouraged 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 1987, half our trainees have previously received MD degrees and another quarter are medical students enrolled in joint degree programs. One-third are candidates for the MS degree, while the rest are doctoral students. The program has three graduates to date, with several more expecting to complete degrees before the end of 1987.

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.3.8. 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. Of course, the move of our groups to new space in the Medical School Office Building has required major planning and care to ensure minimum downtime for our computing environment and much systems and electronics work to outfit the new space.

Our active participation in the planning of the SUMEX/MCS facility in the MSOB resulted in a coordinated environment for twenty-three staff members and thirty-five student workstations, and included 1000 sq. ft. of computer room space and three conference areas. Provisions were made for easily adding more equipment and networking support. The close interaction with the building designers had the additional effect of increasing the designers' interest and knowledge about planning for computer equipment and networking. We have already seen our insight spread to other building projects on campus and the architectural firms will quite likely spread the insight further. Building design appears to be very much an implementation of standards. We have had a part in moving towards the development of more modern standards; certainly here on the campus and perhaps elsewhere.

We use students for much of our operations and related systems programming work. 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.

SUMEX continues to operate with a generally unattended machine room. Our primary operations staff consists of three part-time student workers. This provides a cost-effective approach and gives these undergraduate students an opportunity to participate in the SUMEX project. The major use of this staff is for moving data files to off-line media and to provide data file backup in case of equipment failure. Though we have had nothing that could be classified as a catastrophic failure in the four years of operating our current 2060 equipment, we have had several failures of drives on the SAFE file server. There have been two cases of "soft" failures of disks on the 2060 system. Though these incidents have consumed substantial staff time to deal with, they have not involved significant time loss to the users.

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 81).
2. Relative system loading by community (page 82).
3. Individual project and community usage (page 85).
4. Network usage data (page 90).

5. System reliability data (page 92).

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.

2.1 - Overall Resource Loading Data

The following plot displays total CPU time delivered per month. This data includes usage of the KI-TENEX system and the current DECsystem 2060.

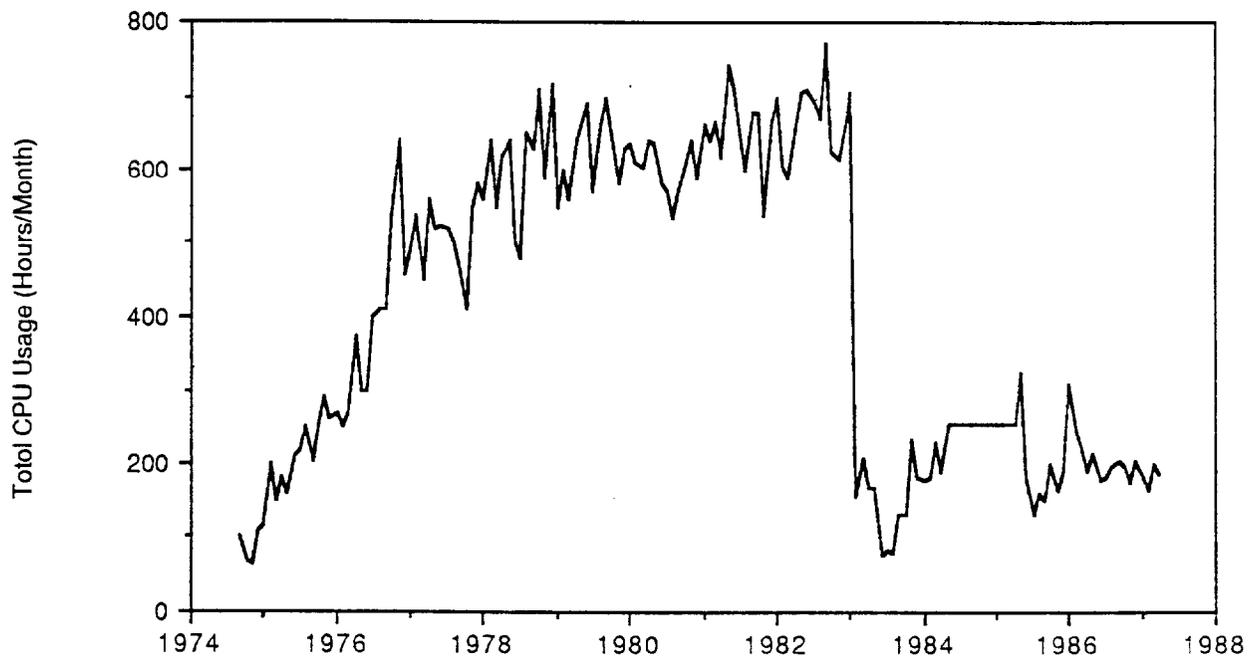


Figure 9: Total CPU Time Consumed by Month

2.2 - Relative System Loading by Community

The SUMEX resource is divided, for administrative purposes, into three major communities: user projects based at the Stanford Medical School (*Stanford Projects*), user projects based outside of Stanford (*National AIM Projects*), and common system development efforts (*System Staff*). As defined in the resource management plan approved by the BRP at the start of the project, the available system CPU capacity and file space resources are nominally divided between these communities as follows:

Stanford	40%
AIM	40%
Staff	20%

The "available" resources to be divided up between these communities are those remaining after various monitor and community-wide functions are accounted for. These include such things as job scheduling, overhead, network service, file space for subsystems, documentation, etc.

The monthly usage of CPU resources and terminal connect time for each of these three communities relative to their respective aliquots is shown in the plots in Figure 10 and Figure 11. As mentioned on page 80, these plots include both KI-10 and 2060 usage data.

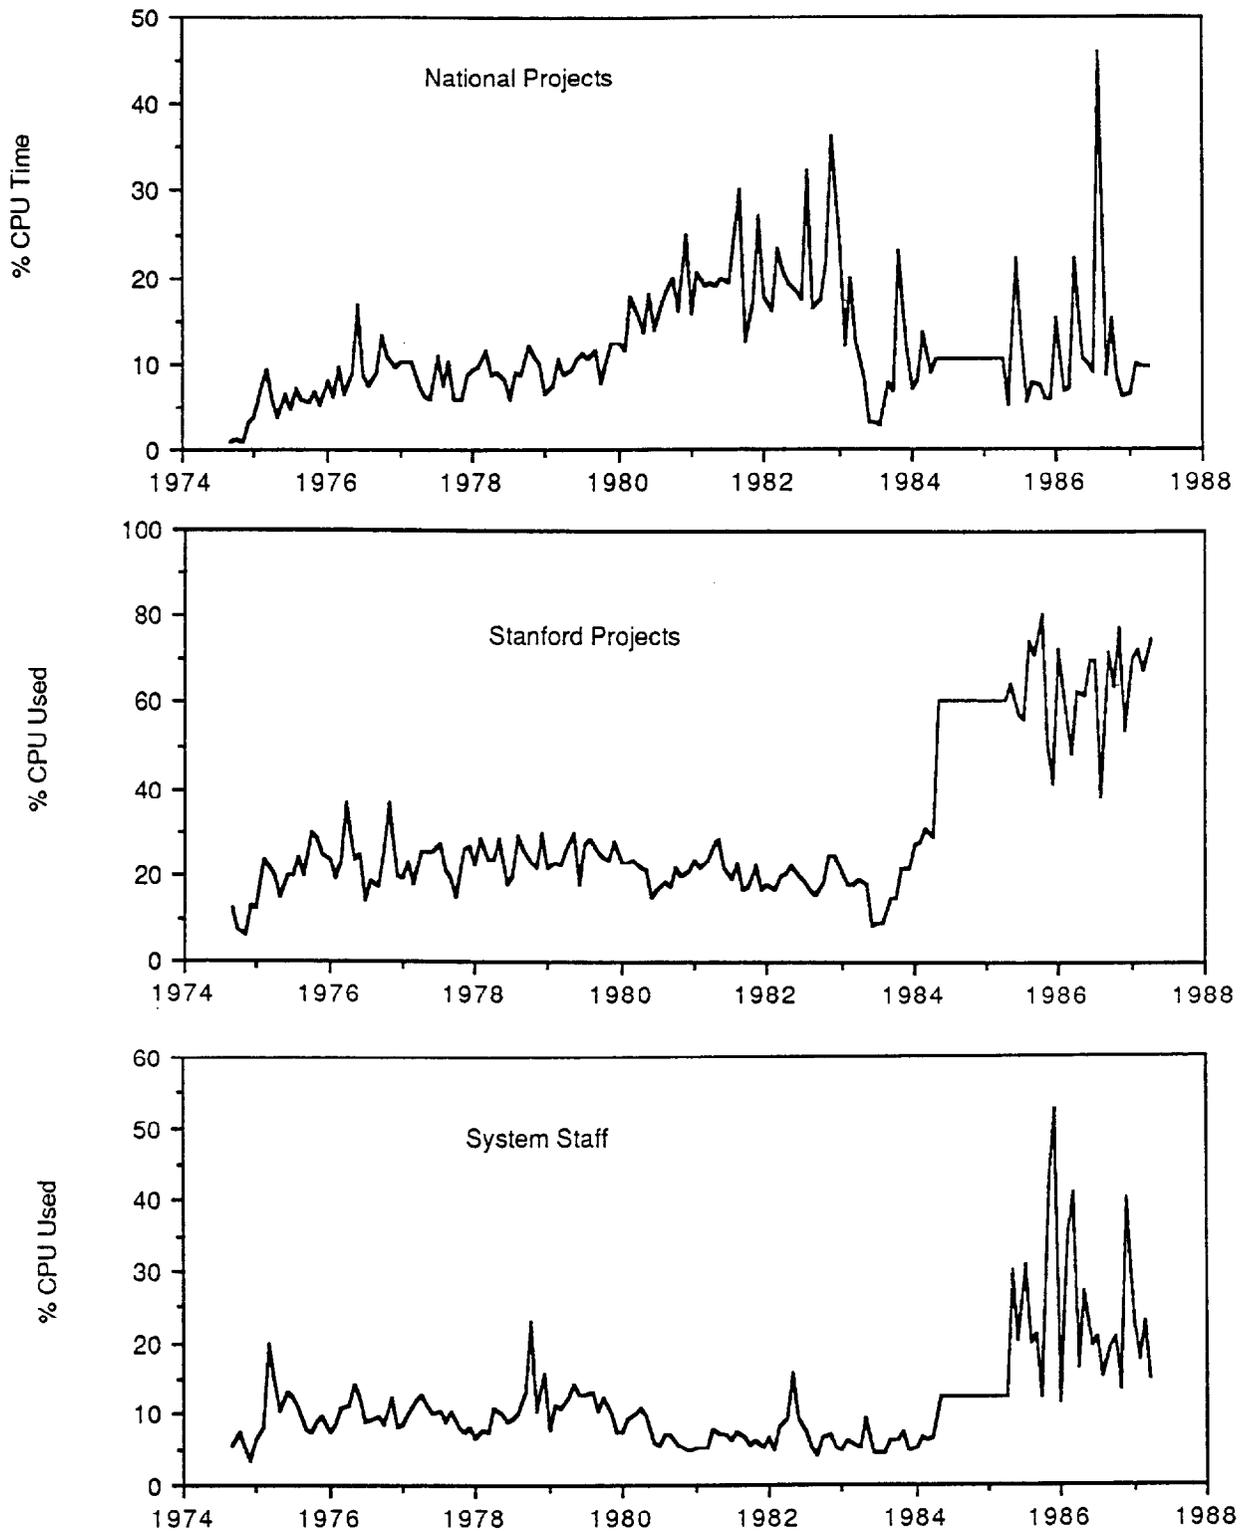


Figure 10: Monthly CPU Usage by Community

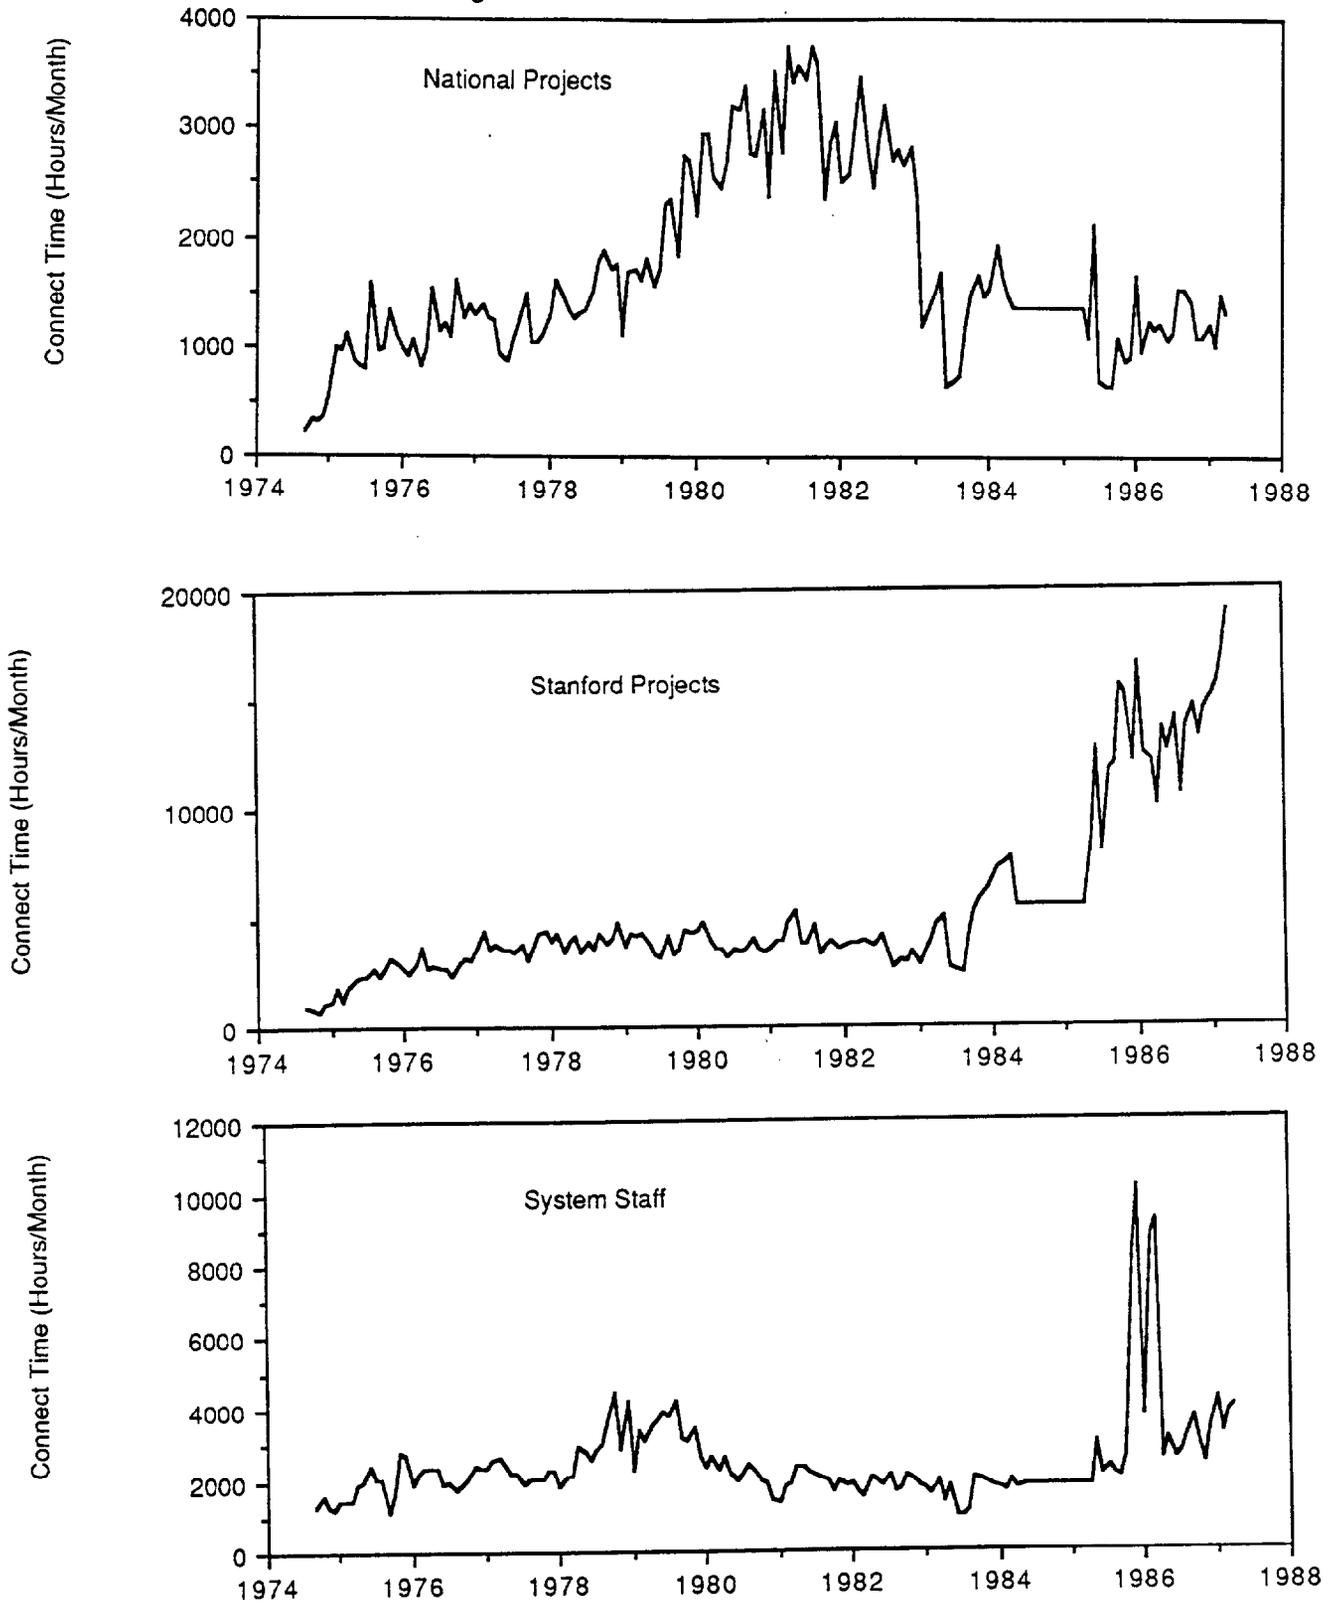


Figure 11: Monthly Terminal Connect Time by Community

2.3 - Individual Project and Community Usage

The following histogram and table show cumulative resource usage by collaborative project and community during the past grant year. The histogram displays the project distribution of the total CPU time consumed between May 1, 1986 and April 30, 1987, on the SUMEX-AIM DECsystem 2060 system.

In the table following, entries include total CPU consumption by project (Hours), total terminal connect time by project (Hours), and average file space in use by project (Pages, 1 page = 512 computer words). These data were accumulated for each project for the months between May 1986 and April 1987.

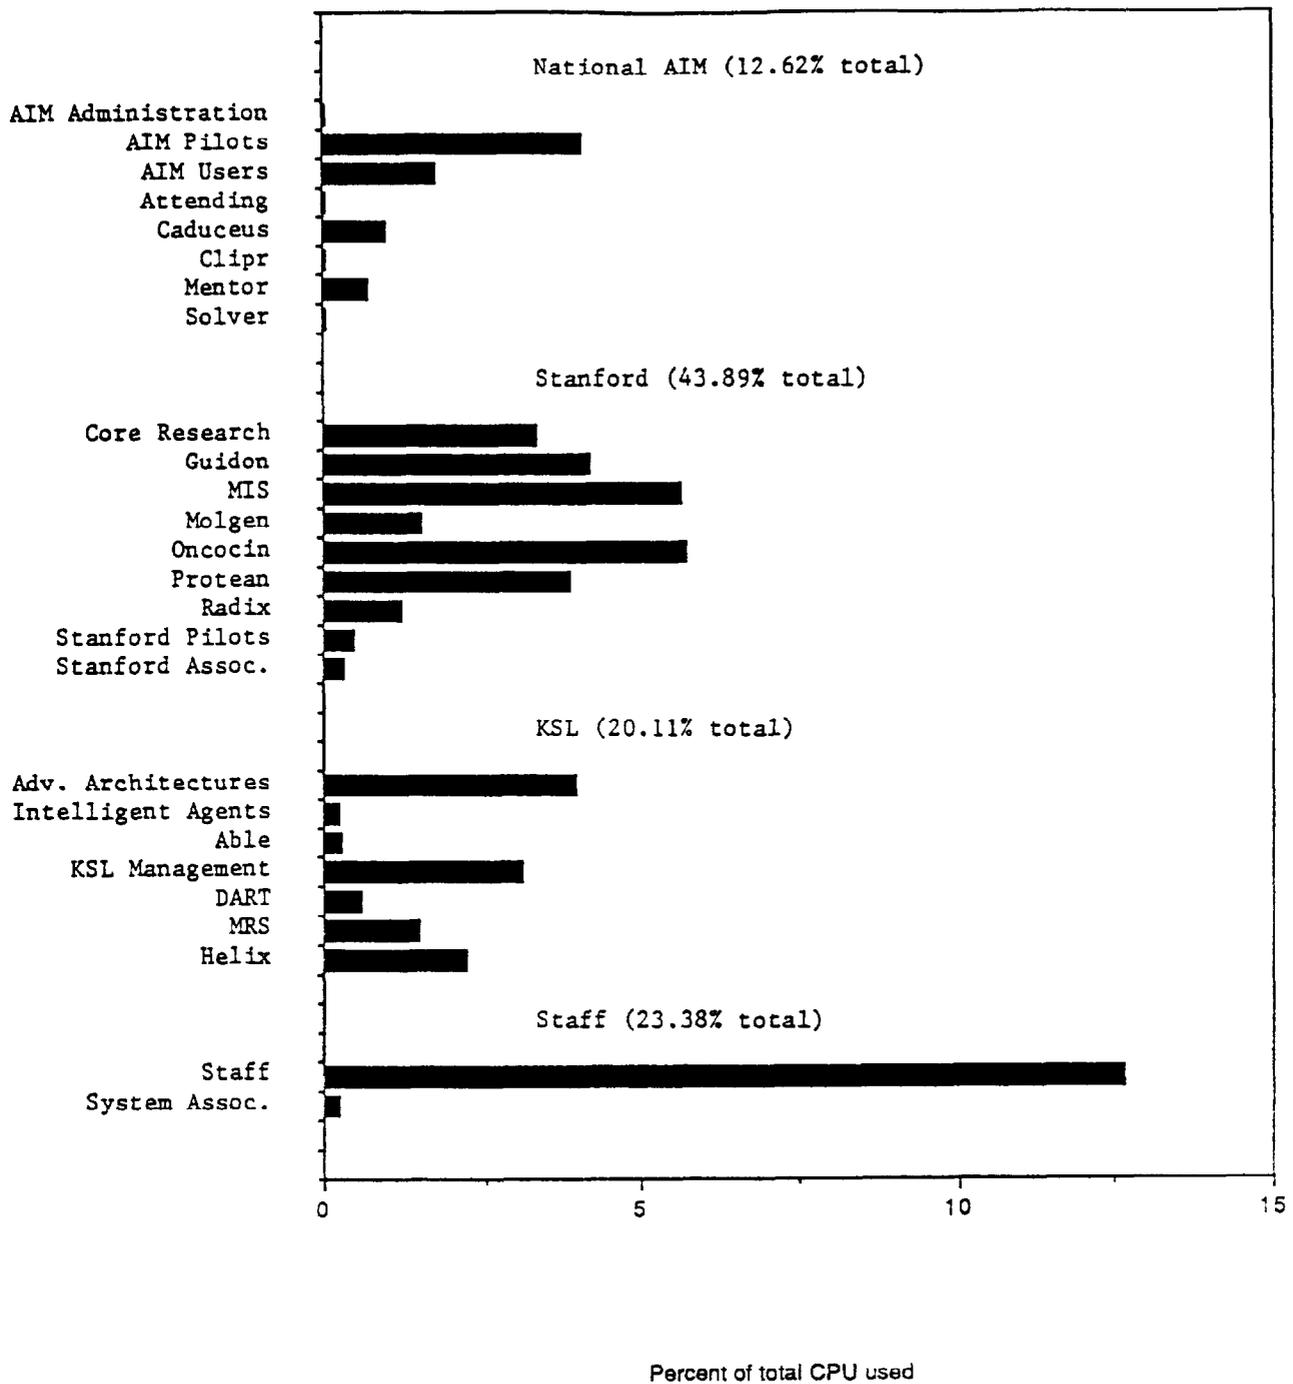


Figure 12: Cumulative CPU Usage Histogram by Project and Community

Resource Use by Individual Project - 5/86 through 4/87

<i>National AIM Community</i>	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) CADUCEUS "Clinical Decision Systems Research Resource" Jack D. Myers, M.D. Harry E. Pople, Jr., Ph.D. Randolph A. Miller, M.D. University of Pittsburgh	21.63	562.85	1066
2) CLIPR Project "Hierarchical Models of Human Cognition" Walter Kintsch, Ph.D. Peter G. Polson, Ph.D. University of Colorado	0.56	144.73	176
3) SOLVER Project "Problem Solving Expertise" Paul E. Johnson, Ph.D. William B. Thompson, Ph.D. University of Minnesota	0.93	133.87	567
4) MENTOR Project "Medical Evaluation of Therapeutic Orders" Stuart M. Speedie, Ph.D. University of Maryland Terrence F. Blaschke, M.D. Stanford University	16.06	6607.36	1044
5) ATTENDING "A Critiquing Approach to Expert Computer Advice" Perry L. Miller, M.D., Ph.D. Yale University School of Medicine	0.16	24.91	3
6) AIM Pilot Projects	92.49	2813.31	836
7) AIM Administration	0.15	23.82	172
8) AIM Users	40.22	4640.05	2308
	-----	-----	-----
Community Totals	172.27	14967.35	6073

<i>Stanford Community</i>	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) GUIDON-NEOMYCIN Project Bruce G. Buchanan, Ph.D. William J. Clancey, Ph.D. Dept. Computer Science	95.67	10725.08	1980
2) MOLGEN Project "Applications of Artificial Intelligence to Molecular Biology: Research in Theory Formation, Testing and Modification" Edward A. Feigenbaum, Ph.D. Peter Friedland, Ph.D. Charles Yanofsky, Ph.D. Depts. Computer Science/ Biology	34.46	7540.02	3109
3) ONCOCIN Project "Knowledge Engineering for Med. Consultation" Edward H. Shortliffe, M.D., Ph.D. Dept. Medicine	131.04	24884.82	2871
4) PROTEAN PROJECT Oleg Jardetzky School of Medicine Bruce Buchanan Computer Science Department	88.64	12876.18	3050
5) RADIX Project Robert L. Blum, M.D. Gio C.M. Wiederhold, Ph.D. Depts. Computer Science/ Medicine	27.01	4356.19	828
6) Stanford Pilot Projects	10.48	1368.01	1690
7) Core AI Research	76.10	16019.67	2732
8) Stanford Associates	6.74	2421.81	179
9) Medical Information Sciences	129.09	16857.16	2060
	-----	-----	-----
Community Totals	599.24	97048.94	18499

	CPU (Hours)	Connect (Hours)	File Space (Pages)
<i>KSL-AI Community</i>			
1) Advanced Architectures	90.41	3643.33	2999
2) FOL	21.80	2275.61	0
3) Intelligent Agent	5.01	824.73	720
4) KSL Administration	70.26	13298.97	2897
5) DART	12.99	3227.10	1577
6) MRS	33.54	9642.75	2205
7) Helix	50.72	13664.49	802
8) ABLE	6.12	3643.33	233
	-----	-----	-----
Community totals	274.49	74378.94	12512
 <i>SUMEX Staff</i>			
	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) Staff	288.63	38743.11	12992
2) System Associates	5.41	681.22	471
	-----	-----	-----
Community Totals	319.20	42391.73	13471
 <i>System Operations</i>			
	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) Operations	918.73	95782.32	24016
	=====	=====	=====
Resource Totals	2283.94	324569.30	74571

2.4 - Network Usage Statistics

The plots in Figure 13 and Figure 14 show the monthly network terminal connect time for the public data networks and the INTERNET usage. The INTERNET is a broader term for what was previously referred to as Arpanet usage. Since many vendors now support the INTERNET protocols (IP/TCP) in addition to the Arpanet, which converted to IP/TCP in January of 1983, it is no longer possible to distinguish between Arpanet usage and Internet usage on our 2060 system.

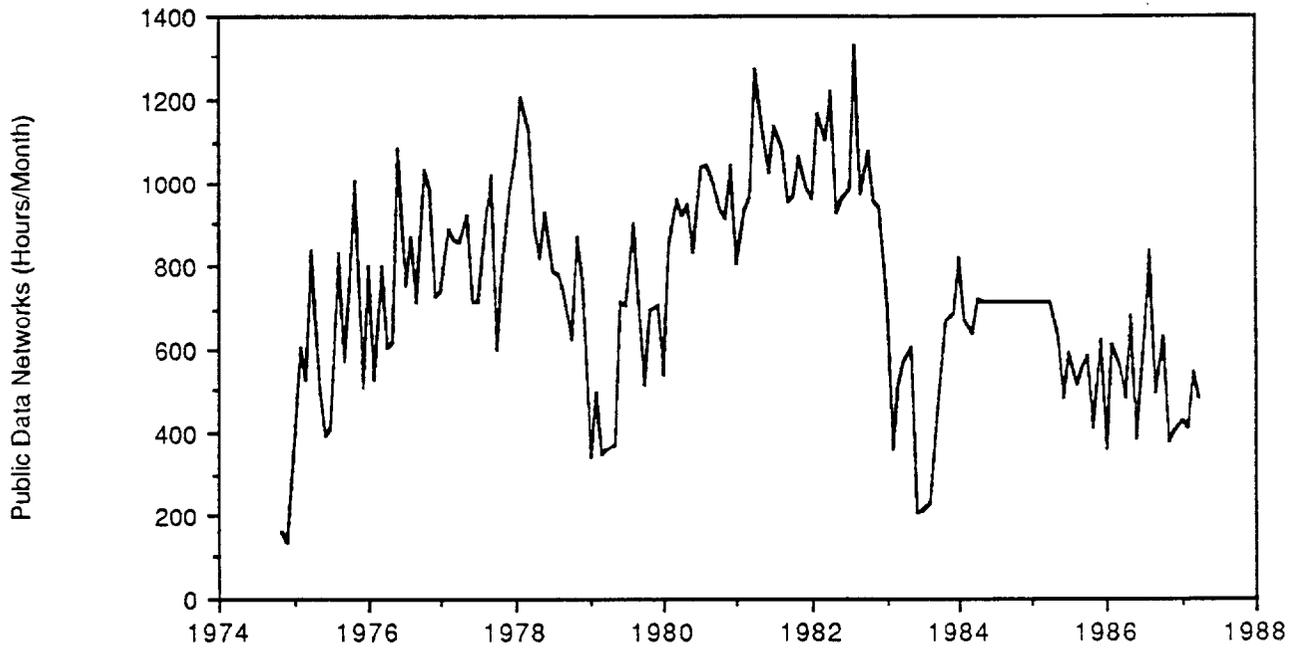


Figure 13: Public Data Network Terminal Connect Time

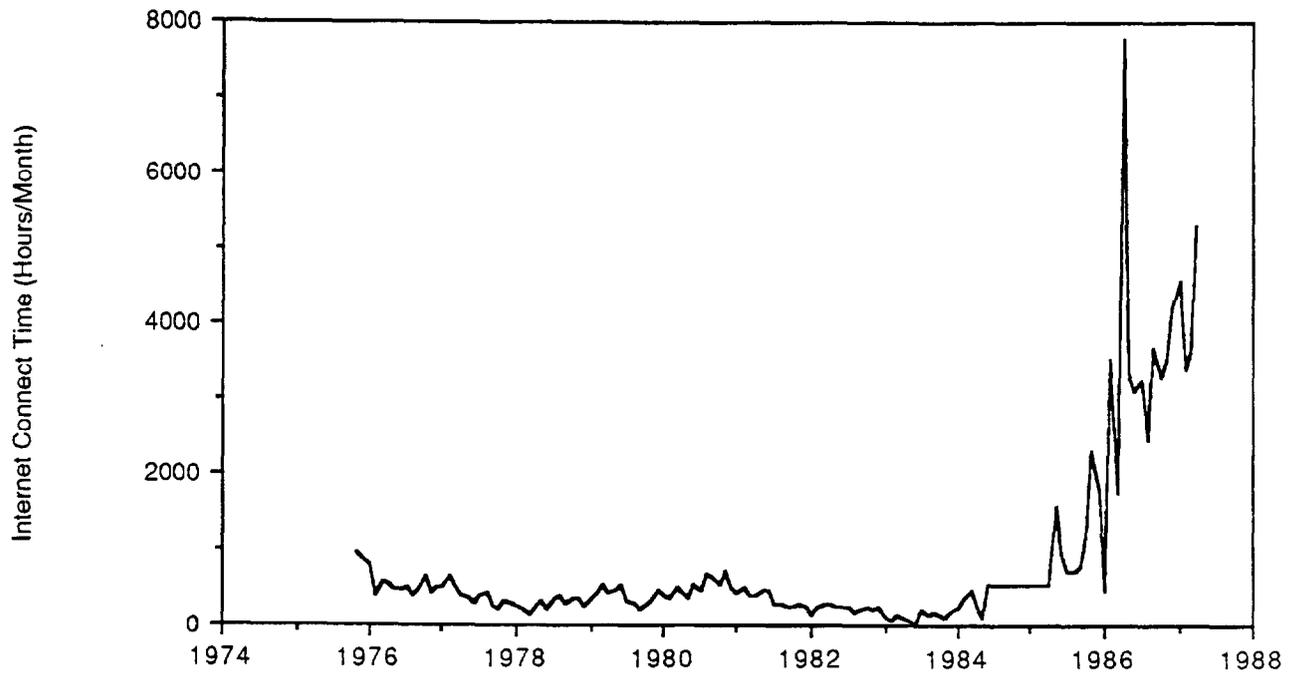


Figure 14: INTERNET Terminal Connect Time

2.5 - System Reliability

System reliability for the DECsystem 2060 has remained quite high in this past period. We have had very few periods of particular hardware or software problems other than while tracking down the internet free space software bug. The data below covers the period of May 1, 1986 to April 30, 1987. The actual downtime was rounded to the nearest hour.

May 1986 - April 1987:

May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr
10	28	13	3	2	2	20	2	2	1	9	11

Figure 15: System Downtime -- Hours per Month

May 1986 - April 1987:

Reporting period	:	365 days, 0 hours, 12 minutes, and 49 seconds
Total Up Time	:	359 days, 23 hours, 8 minutes, and 11 seconds
PM Downtime	:	0 days, 18 hours, 35 minutes, and 5 seconds
Actual Downtime	:	4 days, 6 hours, 29 minutes, and 33 seconds
Total Downtime	:	5 days, 1 hour, 4 minutes, and 38 seconds
Mtbf	:	2 days, 13 hours, 42 minutes, and 29 seconds
Uptime Percentage	:	98.83

Figure 16: Overall System Reliability Summary

III.B. Highlights

In this section we describe several research highlights from the past year's activities. These include notes on existing projects that have passed important milestones, new pilot projects that have shown progress in their initial stages, and other core research and special activities that reflect the progress, impact, and influence the SUMEX-AIM resource has had in the scientific and educational communities.

III.B.1. The MENTOR Project

The MENTOR (Medical EvaluationN of Therapeutic ORders) project, under Drs. Terrence Blaschke at Stanford University, Stuart Speedie at the University of Maryland, and Charles Friedman at the University of North Carolina, seeks to design and develop an expert system for monitoring drug therapy for hospitalized patients. The purpose of the system is to provide appropriate advice to physicians concerning the existence and management of adverse drug reactions.

The computer as a record-keeping device is becoming increasingly common in hospital health care, but much of its potential remains unrealized. Often, information is provided to the physician in the form of raw data. The wealth of such data may effectively hide important information about the patient. This is particularly true with respect to adverse reactions to drugs which can only be detected by simultaneous examinations of several different types of data including drug data, laboratory tests, and clinical signs using sophisticated medical knowledge and problem solving. Expert systems offer the possibility of embedding this expertise in a computer system which would automatically gather the appropriate information and monitor for the prospect or actual occurrence of adverse drug reactions.

The MENTOR project was initiated in December 1983. The project has been funded by the National Center for Health Services Research since January 1, 1985. As of June 1, 1987, a working prototype system has been developed and is undergoing evaluation. The prototype consists of a Patient Data Base, an Inference Engine, an Advisory Module, and a Medical Knowledge Base. The Medical Knowledge Base currently contains information related to Aminoglycoside Therapy, Digoxin therapy, Surgical Prophylaxis, and Microbiology Lab reports. The system is currently implemented on a Xerox workstation. Another version of the Patient Data Base has been developed for a mainframe and is currently being tested. Plans call for the interconnection of the mainframe and the workstation running the inference engine. The mainframe will then be connected to a Hospital Information System for data acquisition.

III.B.2. The GUIDON Project

The GUIDON/NEOMYCIN Project, under Drs. William J. Clancey and Bruce G. Buchanan of Stanford University, is a research program to develop a knowledge-based tutoring system for application to medicine. The primary goal for the GUIDON/NEOMYCIN project is to develop a program that can provide advice similar in quality to that given by human experts, modeling how they structure their knowledge as well as their problem-solving procedures. The consultation program using this knowledge is called NEOMYCIN. The problem-solving procedures are developed by running test cases through NEOMYCIN and comparing them to expert behavior. Also, we use NEOMYCIN as a test bed for the explanation capabilities incorporated in our instructional programs.

Our current emphasis is to construct a knowledge-based tutoring system that teaches diagnostic strategies explicitly. By strategy, we mean plans for establishing a set of possible diagnoses, focusing on and confirming individual diagnoses, gathering data, and processing new data. The tutorial program has capabilities to recognize these plans, as well as to articulate strategies in explanations about how to do diagnosis. The strategies represented in the program, modeling techniques, and explanation techniques are wholly separate from the knowledge base, so that they can be used with many medical (and non-medical) domains.

It has long been felt that medical knowledge, initially codified for the purpose of computer-assisted consultations, may also be used to teach medical students. The technical basis of the system has matured enough that we are now collaborating closely with medical students and physicians to design a useful tutoring program. The system implements a three-step tutorial process in which the student will solve a problem, watch the system solve it, and then explain his solution and seek explanations about the system's solution. In this way, the program will serve as a *model* that the student can study and compare to his own reasoning.

Another tutorial project involves development of a modeling program (ODYSSEUS) aimed at discovering discrepancies between an expert system knowledge base and that of a student or expert problem solver. When ODYSSEUS watches a student, it functions as a student modeling program and when it watches an expert, it functions as a knowledge acquisition program.

The final major effort involves generalizing our expert system tool, HERACLES, so that it can be made available to other research groups wishing to develop knowledge bases that can be used for tutoring.

In our current work, we are focusing on the modeling, explanation, and knowledge acquisition capabilities that will allow the tutor to articulate how a diagnostic solution is flawed and how it can be improved using specific domain knowledge. Thus, we are teaching the students what constraints a good solution must respect and giving them a language for articulating which medical facts are relevant to the case at hand.

Physicians have generally been enthusiastic about the potential of these programs and what they reveal about current approaches to computer-based medical decision making.

III.B.3. The PROTEAN Project

The PROTEAN project, under Professors Oleg Jardetzky and Bruce Buchanan at Stanford University, is concerned with using artificial intelligence methods to aid in the determination of the 3-dimensional structure of proteins in solution (as opposed to crystallized proteins). The molecular structure of proteins is essential for understanding many problems of medicine at the molecular level, such as the mechanisms of drug action. Using NMR data from proteins in solution will allow the study of proteins whose structure cannot be determined with other techniques, and will decrease the time needed for the determination. It is hoped that empirical data from nuclear magnetic resonance (NMR) and other sources may provide enough constraints on structural descriptions to allow protein chemists to bypass the laborious methods of crystallizing a protein and using X-ray crystallography to determine its structure.

During the past year, we have extended our initial prototype program, called PROTEAN, designed using a *blackboard model*. It is implemented in BBl, a framework system for building blackboard systems that control their own problem-solving behavior. The reasoning component of PROTEAN directs the actions of the Geometry System (GS), a set of programs that performs the computationally intensive task of positioning portions of a molecule with respect to each other in three dimensions. The GS runs in the UNIX environment on a Silicon Graphics IRIS 3020 graphics workstation. The reasoning program (in LISP in BBl) is coupled to the GS by a local area computer network developed by SUMEX.

Pictures of the results of GS computations are displayed on the graphics screen of the IRIS workstation, using a locally developed program called DISPLAY to draw the evolving protein structures at several levels of detail. The DISPLAY program can be used to view structures generated by the GS either under the direct control of the user or as directed by the reasoning system running in BBl. MIDAS and MMS are two other molecular modeling and display systems to manipulate protein structures, particularly those obtained from crystallographic techniques as found in the Protein Data Bank. The ability to observe structures in three dimensions is essential to understanding the behavior of the PROTEAN's reasoning and geometry systems and provides essential insights on the problem solving process.

In addition to the Lac-repressor headpiece protein, we have applied PROTEAN to sperm whale myoglobin, T4 lysozyme, and cytochrome B. Each of these latter proteins has a known crystal structure. In each case, we extracted features of the protein and distance constraints to build data sets for PROTEAN. We then applied the PROTEAN system to the resulting data sets to determine the behavior of the system with different kinds of input.

To determine the correctness and capabilities of the PROTEAN method, we applied PROTEAN to sperm whale myoglobin, a molecule whose crystal structure is known. We systematically explored the dependence of the precision and accuracy of the solutions on the quality of the input data available. In all cases, the sets solutions obtained from PROTEAN include the actual structure of the molecule, with the best results coming from data representing many short range constraints.

Work is proceeding on several aspects of the protein structure problem, including assembly of several partial arrangements and integration of these pieces of solution into larger structures, using atomic level volume exclusion of atoms and information on sidechain packing to produce more precise atomic level solutions, and developing more appropriate representations for unstructured coil sections of proteins.

III.B.4. The Medical Information Science Program

The *Medical Information Sciences (MIS)* program continues to be one of the most obvious signs of the academic impact of the SUMEX-AIM resource on Stanford University. 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). It was Stanford's active on-going research in medical computer science, plus a world-wide reputation for the excellence and rigor of those research efforts, that persuaded the University that the field warranted a new academic degree program in the area. A group of faculty from the medical school and the computer science department argued that research in medical computing has historically been constrained by a lack of talented individuals who have a solid footing in both the medical and computer science fields. 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 20 graduate students. We do not wish to provide too narrow a definition of what kinds of prior training are pertinent because of the interdisciplinary nature of the field. The program has accordingly encouraged 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 1987, half our trainees have previously received MD degrees and another quarter are medical students enrolled in joint degree programs. One-third are candidates for the MS degree, while the rest are doctoral students. The program has three graduates to date, with several more expecting to complete degrees before the end of 1987. Research opportunities for students 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.

III.B.5. Remote Virtual Graphics

Lisp workstations of various types have proven extremely powerful, both as development environments for artificial intelligence research and as vehicles for disseminating AI systems into user communities. In addition to the compact, inexpensive computing resources workstations provide, high-quality graphics play a key role in their power. Such graphics systems have become indispensable for understanding the complex data structures involved in developing and debugging large AI systems and are important in facilitating user access to working programs (e.g., for ONCOCIN and PROTEAN).

In the past, members of the SUMEX-AIM community have often watched each others' programs work by linking their CRT terminals to the text output of a running program on the SUMEX 2060. With workstations, though, it is much more difficult to connect to a remote machine and be able to view the complex graphics output of a program. One would like to be able to provide the same powerful graphical tools and programming environment that are available to a user sitting in front of the workstation to the remote user if that user has a low-cost bit-mapped display and mouse.

During this past year, we developed a program called TALK to facilitate interactive, electronic communication between users on independent workstations. Layered on the workstation's native editor, the program allows the full use of all editing capabilities in the process of communication, including deletions, corrections and insertions, font changes, underlining, paragraph formatting, etc. Since the workstation's editor also supports both low- and high-level graphics, the program not only facilitates textual exchanges among users, but also allows the sending of screen images (ONCOCIN flow sheet segments, back traces of program error breaks, code fragments, etc.) as well as structured images (which can be modified on the destination workstation and returned), all interactively.

The TALK program allows the use of different user interfaces, the workstation's document editor being just one possibility. We implemented a simpler terminal mode for compatibility with similar programs on other workstations.

The TALK program has been released gradually to increasing numbers of users in order to get feedback and make changes accordingly. The Medical Computer Science group did an extensive test of the system where for a period they used it in place of their normal electronic and non-electronic communication methods whenever possible. This was both a test of the program and an exploration into what people want in the next generation of electronic communication. The TALK program has been released to the Xerox Lisp workstation community as a whole and researchers at Xerox PARC successfully used the program to hold an interactive, graphic, electronic conversation between users at the PARC facility (in California) and Xerox's EuroPARC facility (in England).

III.C. Administrative Changes

There have been few administrative changes within the project this past year. Professor Shortliffe has been on sabbatical at the University of Pennsylvania as projected last year but has stayed in very close contact with SUMEX and the Medical Computer Science group at Stanford through network connections. During this time, Professor Feigenbaum has acted in the formal role of principal investigator. Professor Shortliffe is expected back at Stanford in mid-July.

The move the Medical Computer Science and SUMEX offices into the newly constructed Stanford Medical School Office Building was completed in June 1986. We now occupy approximately 6500 square feet has almost doubled the space available to us. The design of this space has worked out exceedingly well to improve the interactions within our groups.

We have also designed and implemented a cost recovery system as part of phasing out BRTP subsidy of the DEC 2060 facility. The details of this system are discussed on page 101. In summary, we are successfully recovering the projected 20% of 2060 operations costs this year (\$71,376) from Stanford users, with the continuing component of NIH support used to protect national users from fees for service, including communications. This additional burden on Stanford projects was absorbed almost entirely in existing direct cost budgets since no supplements were forthcoming from other funding agencies in the middle of on-going grant and contract awards for new computing costs. This has affected staffing and student support directly in our labor-intensive research efforts. All of our new support applications are being written with requests for funds to cover computing charges.

This next year we will increase the cost recovery goal to 40% of projected 2060 operations costs as scheduled in our grant application of June 1985.

III.D. Resource Management and Allocation

III.D.1. Overall Management Plan

Early in the design of the SUMEX-AIM resource, an effective management plan was worked out with the Biotechnology Resources Program (now Biomedical Research Technology Program) at NIH to assure fair administration of the resource for both Stanford and national users and to provide a framework for recruitment and development of a scientifically meritorious community of application projects. This structure has been described in some detail in earlier reports and is documented in our recent renewal application. It has continued to function effectively as summarized below.

- The AIM Executive Committee meets regularly by teleconference to advise on new project applications, discuss resource management policies, plan workshop activities, and conduct other community business. The Advisory Group meets together at the annual AIM workshop to discuss general resource business and individual members are contacted much more frequently to review project applications. (See Appendix A on page 217 for a current listing of AIM committee membership).
- We have actively recruited new application projects and disseminated information about the resource. The number of formal projects in the SUMEX-AIM community still runs at the capacity of our computing resources. With the development of more decentralized computing resources within the AIM community outside of Stanford (see below), the center of mass of our community has naturally shifted toward the growing number of Stanford applications and core research projects. We still, however, actively support new applications in the national community where these are not able to gain access to suitable computing resources on their own.
- With the advice of the Executive Committee, we have awarded pilot project status to promising new application projects and investigators and where appropriate, offered guidance for the more effective formulation of research plans and for the establishment of research collaborations between biomedical and computer science investigators. This past year we have admitted projects under Professors Perry Miller at Yale University, Larry Widman at the University of Texas, Ira Kalet at the University of Washington, and Robert Beck at Dartmouth University. The latter two sought access primarily for communication with the AIM community as they have research computing resources of their own.
- We have carefully reviewed on-going projects with our management committees to maintain a high scientific quality and relevance to our biomedical AI goals and to maximize the resources available for newly developing applications projects. Several fully authorized and pilot projects have been encouraged to develop their own computing resources separate from SUMEX or have been phased off of SUMEX as a result and more productive collaborative ties established for others.
- We continue to provide active support for the AIM workshops. The next one will be held at the University of Washington in conjunction with the American Association for Artificial Intelligence meeting in July 1987. It is being organized jointly by Drs. Ira Kalet of Washington and Larry Fagan of Stanford.