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 expect a significant change in ARPANET protocols to be instituted by January 1, 1983. Before then we also expect ARPA to upgrade our Honeywell 516 IMP to a new C-30 IMP and to add a connection between the C-30 and AI&DS, an ARPA contractor. We are also cooperating with the Network Graphics group in the Stanford EE department on the development of an ARPANET gateway that will operate over our local Ethernet.

ETHERNET

A substantial portion of our system effort this past year went into continued development of local Ethernet facilities which link SUMEX resource hardware with other parts of campus (especially to the Computer Science Department building where the Heuristic Programming Project is located). We have also developed a substantial set of software tools and diagnostics as a basis for implementing various kinds of Ethernet servers on both PDP-11 and SUN MC-68000 processors. These have been done in the language C, primarily for portability and 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 Ethernet development include:

1) Leaf server: We implemented the Sequin reliable packet protocol and Leaf byte-level file transfer protocol to enable the Dolphins to access files on the 2020 and KI-10. Code for the VAX/UNIX system was written by a EE student and is available as well. The Leaf protocol is built into the lowest levels of the Dolphin I/O system, and allows any file on a remote file server to be accessed as easily as a disk file in both paged or random access mode. The latest updates to the Sequin transport level have made marked improvements in efficiency. The 2020 now performs Leaf file transfers with a speed approaching that of Xerox's dedicated file server.

2) TENEX/TOPS-20 Ethernet Server: We implemented boot server software to enable Altos, PDP-11's, and SUN terminals to be booted by the KI-10s or the 2020 and also added the address lookup service so that all documented types of miscellaneous services are supported by TENEX and TOPS-20.

3) Ethernet Diagnostic Software: Many difficult situations arose in the past year involving the debugging of locally developed Ethernet interface boards or forcing the gateway link to work over very substandard cable. This necessitated the development of effective diagnostic tools to provide oscilloscope loops, loopback packet checks, local and remote echo checks, and other measures of hardware performance. Such programs were written for the KI-10, PDP-11, and SUN MC-68000 processors

4) Ethernet Gateway: We reported on the initial development of an Ethernet gateway last year. In order to make the gateway between the SUMEX network and other campus networks operate, we had to specially modify the transceivers and interface boards to work with the high-loss, 50 ohm cable currently available. Once working at the hardware level, much work has been done to develop an effective and supportable gateway. This includes a PROM EFTP memory dump and analysis capability to diagnose gateway failures and a miscellaneous services server program to record and analyze hourly event reports on the gateway hardware status. The gateway also now generates periodic Breath-of-Life(*) packets so that we can boot our Dolphins and Alto independently of their disk bootstraps in the event of a hardware failure. The gateway has been up continuously since the middle of December 1981, excepting power failures and outages.

5) Ethernet TIP: The PDP-11 Ethernet operating system and the byte sequential protocol implementation reported last year were combined with serial line handling software and a convenient user interface to create a terminal interface processor (TIP). In order to accommodate buffers for 8 lines in the 32K PDP-11, we wrote an optimizer for the MACRO-11 output from the PDP-11 C compiler. The optimizer provided about a 20% savings in core usage. The TIP code is now stable and the system is in use in Margaret Jacks Hall as the primary means for HPP access to the VAX. Work continues on a similar SUN MC-68000 version of the TIP which should have a substantially greater line capacity.

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

---

(*) These are special non-PUP packets which cannot be generated or handled by the KI-10 or 2020 hosts currently.

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 continued development of prototypical systems for internet mail handling and telnet connections. The mail system (XMAILR) is quite 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. SUMEX has served as the major electronic mail gateway between the Stanford Ethernet community and other ARPANET hosts for the past year. We also brought up a new mail forwarding system (XMAILBOX) that incorporates list expansion, a more convenient syntax for forwarding entries, and allows indirect file references for expanding mail distribution lists.

I.A.3.4 User System 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 made a significant effort to improve the mail reading facilities available under TENEX. With the development of broad internet contacts, the old MSG program was unable to recognize host addresses on networks other than ARPANET. BBN refused to release the sources to MSG for further development, even though they were initially developed outside of BBN. We therefore decided to adapt the MM system developed for TOPS-20. Many changes were necessary to bridge the significant differences that have arisen between TENEX and TOPS-20, pointing up the impediments mentioned earlier to continued import of TOPS-20 community software. We also had to make substantial improvements to the TOPS-20 JSYS emulation software developed earlier. At the present time, MM is running in experimental mode with a subgroup of the SUMEX-AIM community and seems to work well. It provides a full internet capability and is in fact more efficient in terms of machine resources than MSG. We plan to finish adapting documentation for MM shortly and release it to the full community.

We also spent considerable additional effort improving the C compiler at SUMEX to be compatible with other UNIX compilers in terms of the syntax it understands. We also developed a postprocessing object code optimizer for the PDP-11 cross compiler to minimize the memory requirements in the generated code.

We worked on adding functions to Interlisp-10 to make it compatible with Dolphin Interlisp-D. This included adding a LOGIN() function which could be used by the Interlisp-10 FTP and PUPFTP packages, and an EMPRESS function so that Interlisp-10 could produce prettyprinted output with font changes similar to Interlisp-D's prettyprinted output. The LOGIN function has been sent to Xerox PARC. Completion of the EMPRESS function is awaiting some definitions from Xerox.

We also implemented extensions and maintenance updates to many other existing programs including, for example, BBD (the bulletin board system), OMNIGraph (the general graphics package from NIH), MSGFIX (a message file repair program), and FM (a locally developed fast mail reading program). We also participated with other sites at Stanford in a general license for access to the SCRIBE text formatting system from UNILOGIC including versions to run under TENEX, TOPS-20, and UNIX. SCRIBE is rapidly replacing PUB as the preferred tool for text preparation.

I.A.3.5 Documentation and Education

We have spent considerable effort to develop, maintain, and facilitate access to our documentation so as to accurately reflect available software. The HELP and Bulletin Board subsystems have been important in this effort. As subsystems are updated, we generally publish a bulletin or small document describing the changes. As more and more changes occur, it becomes harder and harder for users to track down all of the change pointers. Within manpower limits, we are in a continuous process of reviewing the existing documentation system for compatibility with the programs now on line and to integrate changes into the main documents. This will also be done with a view toward developing better tools for maintaining up-to-date documentation.

I.A.3.6 Software Compatibility and Sharing

At SUMEX-AIM we firmly believe in importing rather than reinventing software where possible. As noted above, a number of the packages we have brought up are from outside groups. Many avenues exist for sharing between the system staff, various user projects, other facilities, and vendors. The advent of fast and convenient communication facilities coupling communities of computer facilities has made possible effective intergroup cooperation and decentralized maintenance of software packages. The TENEX, TOPS-20, and UNIX sites on the ARPANET have been a good model for this kind of exchange based on a functional division of labor and expertise. The other major advantage is that as a by-product of the constant communication about particular software, personal connections between staff members of the various sites develop. These connections serve to pass general information about software tools and to encourage the exchange of ideas among the sites. Certain common problems are now regularly discussed on a multi-site level.

We continue to draw significant amounts of system software from other ARPANET sites, reciprocating with our own local developments. The TENEX community is on the decline, however, and we are becoming more and more isolated in our work for the KI-TENEX system. Still, significant sharing continues with the TOPS-20 and UNIX sites. Interactions have included mutual backup support, experience with various hardware configurations, experience with new types of computers and operating systems, designs for local networks, operating system enhancements, utility or language software, and user project collaborations. We have been able to import many new pieces of software and improvements to existing ones in this way. Examples of imported software include the message manipulation program MM, SAIL, PASCAL, SOS, INTERLISP, the C compiler, VAX Ethernet code, the RECORD program, ARPANET host tables, and many others. Reciprocally, we have exported our contributions such as the crash analysis program, drum page migration system, KI-10 page table efficiency improvements, GTJFN enhancements, PUB macro files, the bulletin board system, MAINSAIL, SPELL, SNDMSG enhancements, our BATCH monitor, and improved SA-10 software.

There are also several important examples of joint development efforts such as the internet mailer program (XMAILR). Because this program incorporates facilities for routing mail through many networks, it is important that the various sections of the program dealing with these specialized protocols be developed by the groups with expertise in the appropriate technology. Network connections have made a joint effort possible involving MIT, Stanford SCORE, and SUMEX. A growing number of network sites are now using this mailer program.

We have spent considerable effort developing a version of a TENEX/TOPS-20 compatibility package. The issue here is that as DEC develops TOPS-20, even though it is TENEX-like, it is not TENEX compatible and vice versa. Thus, a hope was to write a program that would resolve these compatibilities automatically rather than to force special adaptations for the two operating systems. The kinds of incompatibilities that exist include PDP-20 machine instructions that do not exist on earlier machines, new JSYS calls, incompatible changes to old JSYS calls, different syntax and facilities for device/file names, and different handling of error returns (types of return and error codes). It has proven very difficult to effectively handle all of these problems at the user level. Monitor changes are required to implement the widely used error return features (ERJMP/ERCAL) and make handling of other incompatibilities easier. We have not accomplished a complete compatibility package in any sense but have implemented requisite monitor changes and have developed several user packages that help emulate TOPS-20 JSYS calls for programs running on TENEX. We do not foresee being able to completely and effectively solve these problems and expect the differences to increase, especially with user access to extended addressing on the large TOPS-20 systems becoming more frequently used.

Finally, we have also assisted groups that have interacted with SUMEX user projects get access to software available in our community. The DENDRAL project has provided many examples of this kind of sharing -- (see page 96). The most voluminous guest user group is the GENET community started in collaboration with the MOLGEN project (see page 128).

I.A.3.7 Core Research

Over the past year we have supported several core research activities aimed at developing information resources, basic AI research, and tools of general interest to the SUMEX-AIM community. SUMEX is providing only partial support for these projects with complementary funding coming from ARPA, ONR, NLM, and NSF contracts and grants to the Stanford Heuristic Programming Project.

Principal areas of current effort include:

1) AI Handbook -- This is a compendium of knowledge about the field of Artificial Intelligence being compiled by Avron Barr, Paul Cohen, and Professor Feigenbaum. The handbook is broad in scope, covering most of the important ideas, techniques, and systems developed during 26 years of research in AI in a series of articles. Each is short and is a description written for non-AI specialists and students of AI. The handbook is now published in three volumes by William Kaufmann, Inc. The AI Handbook effort is described in more detail starting on page 92 and an outline of the contents of the handbook can be found in Appendix B.

2) AI Tools -- The objective of our core tool building research is to develop new tools and methods for representing knowledge in an attempt to generalize the methods explored in specific applications programs and to respond to problems encountered with existing methods. Several software packages are being developed, each of which has some degree of general applicability. Several have been applied to a variety problems and thus provide experimental data on their strengths and limitations. These packages comprise what we term a Knowledge Engineer's Workbench. They are tools that facilitate the rapid construction, testing, modification, and explanation of knowledge-based expert systems.

   Both EMYCIN and AGE-1 have reached maturity as vehicles of research. Each was completed and documented during the last year. Experiments with each system have also been performed and documented. Both programs are in use on medical and nonmedical problems at several institutions. These experiments, and their analysis, has led to an understanding of the strengths and limitations for each system and will provide valuable information for improved methods. More detailed descriptions of progress on these systems can be found on page 137 and on page 83.

   Even though the Meta-knowledge Representation System (MRS) is still under development, it has been exported to several sites and used experimentally on a variety of problems. A demonstration version of EMYCIN, implemented in MRS, has shown the power of the system's logic-based representation. MRS is implemented in a variety of Lisp's, including MacLisp and InterLisp on the DEC-20, InterLisp on the Xerox Dolphin, and FranzLisp on the VAX 11/780. The system is taught in Stanford's first graduate level AI courses.

In addition to the work on these major systems, improvements were made to the UNITS package, including its transfer to the Dolphin personal workstation.

3) Basic Research -- Building computer programs that reason with and about medical knowledge requires powerful methods for representing that knowledge and for reasoning with it. Basic core research efforts focus on understanding these issues. The three areas that we are addressing in this work are:

   Knowledge Representation -- How can the knowledge necessary for complex problem solving be represented for its most effective use in automatic inference processes? In particular, how can knowledge of causal mechanisms be represented in programs that reason about medical problems?

   Knowledge Acquisition -- How is this knowledge acquired most efficiently from human experts and from natural data gathered by instruments?

   Knowledge Utilization -- By what inference methods can a variety of sources of knowledge of diverse types be made to contribute jointly and efficiently toward solutions?

   Results of recent work are published in technical reports including experiments performed and documented on the power of production rules and of frame-based systems [Memos HPP-80-17, HPP-81-2, HPP-81-17]; and experiments with induction systems and with interactive programs for knowledge acquisition [Memos HPP-81-4, HPP-81-5, HPP-81-13, HPP-81-19]. In addition we have begun work on programs that explore fundamental questions in:

   Causal Reasoning -- How do we represent and use causal mechanisms, e.g., about physiological processes, in knowledge-based computer programs?

   Knowledge Acquisition -- How can we build an intelligent knowledge acquisition system that helps an expert construct a new knowledge base using old ones as models?

I.A.3.8 Resource Operations Statistics

The following data give an overview of various aspects of SUMEX-AIM resource usage. There are five sub-sections containing data respectively for:

1) Overall resource loading data (page 34).

2) Relative system loading by community (page 36).

3) Individual project and community usage (page 40).

4) Network usage data (page 48).

5) System reliability data (page 50).

1. Overall resource loading data

The following plots display several different aspects of system loading over the life of the project. These include total CPU time delivered per month, the peak number of jobs logged in, and the peak load average. The monthly "peak" value of a given variable is the average of the daily peak values for that variable during the month. Thus, these "peak" values are representative of average monthly loading maxima and do not reflect the largest excursions seen on individual days, which are much higher.

These data show well the continued growth of SUMEX use and the self-limiting saturation effect of system load average, especially after installation of our overload controls early in 1978. Since late 1976, when the dual processor capacity became fully used, the peak daily load average has remained between about 5.5 and 6. This is a measure of the user capacity of our current hardware configuration and the mix of AI programs.

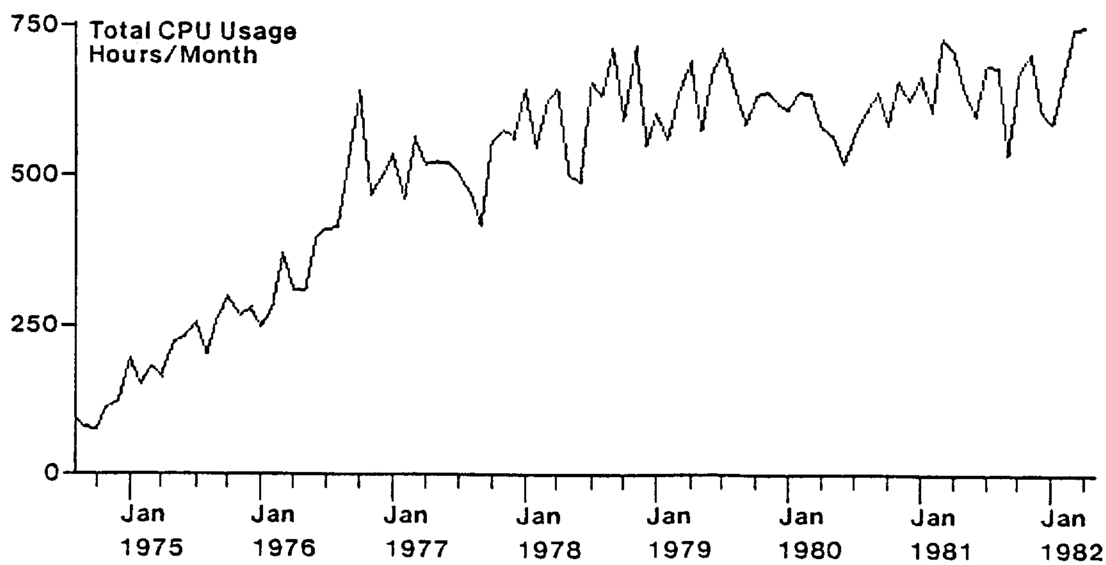

Figure 5. Total CPU Time Consumed by Month

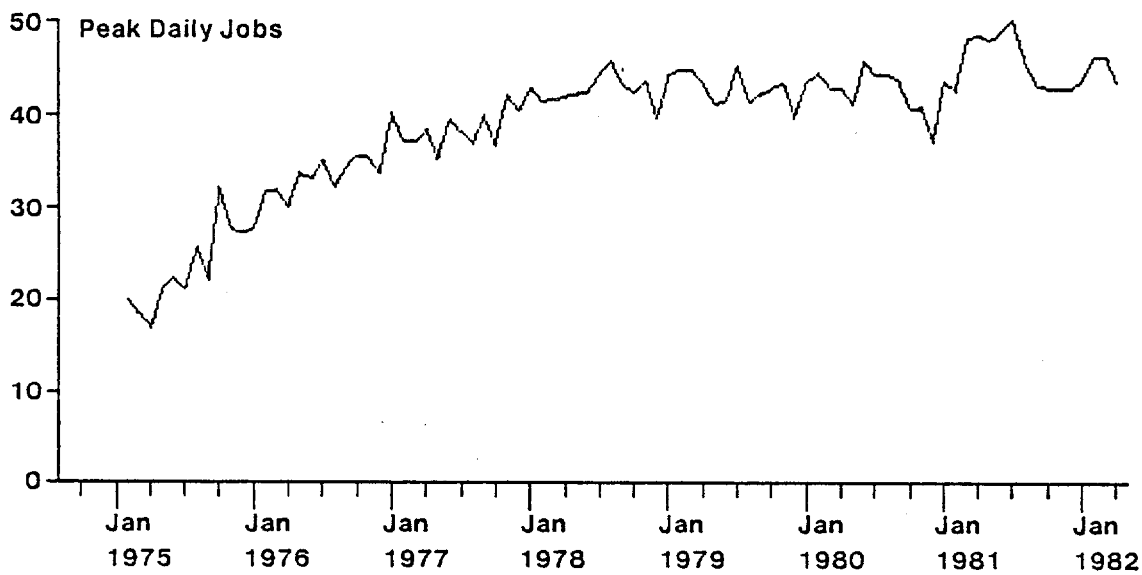

Figure 6. Peak Number of Jobs by Month

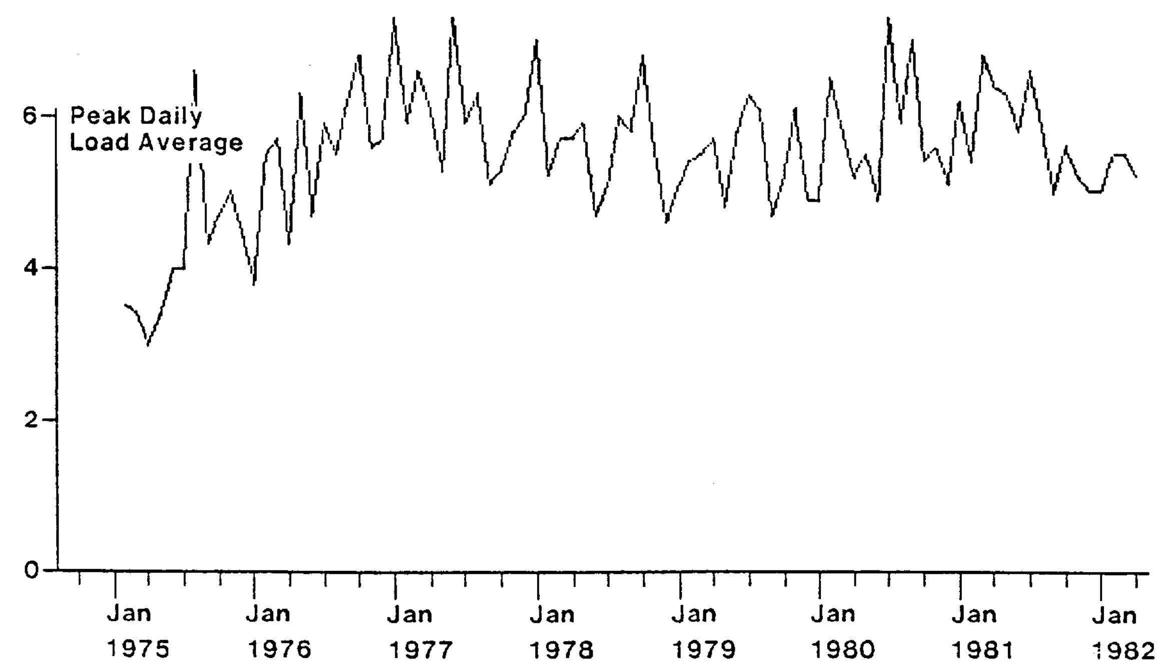

Figure 7. Peak Load Average by Month

2. Relative System Loading by Community

The SUMEX resource is divided, for administrative purposes, into 3 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 BRP at the start of the project, the available system CPU capacity and file space resources are divided between these communities as follows:

Stanford	40%
AIM	40%
Staff	20%

The "available" resources to be divided up in this way 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 and file space resources for each of these three communities relative to their respective aliquots is shown in the plots in Figure 8 and Figure 9. Terminal connect time is shown in Figure 10. In the early years of the SUMEX resource, the Stanford projects held an edge in system usage despite our efforts at resource allocation and the substantial voluntary efforts by the Stanford community to utilize non-prime hours. This reflected the maturity of the Stanford group of projects relative to those getting started on the national side and has correspondingly accounted for much of the progress in AI program development to date. While the Stanford community has remained vigorous, national usage has grown to comparable levels in recent years, in part because of the popular GENET guest facility (see page 40 and page 73).

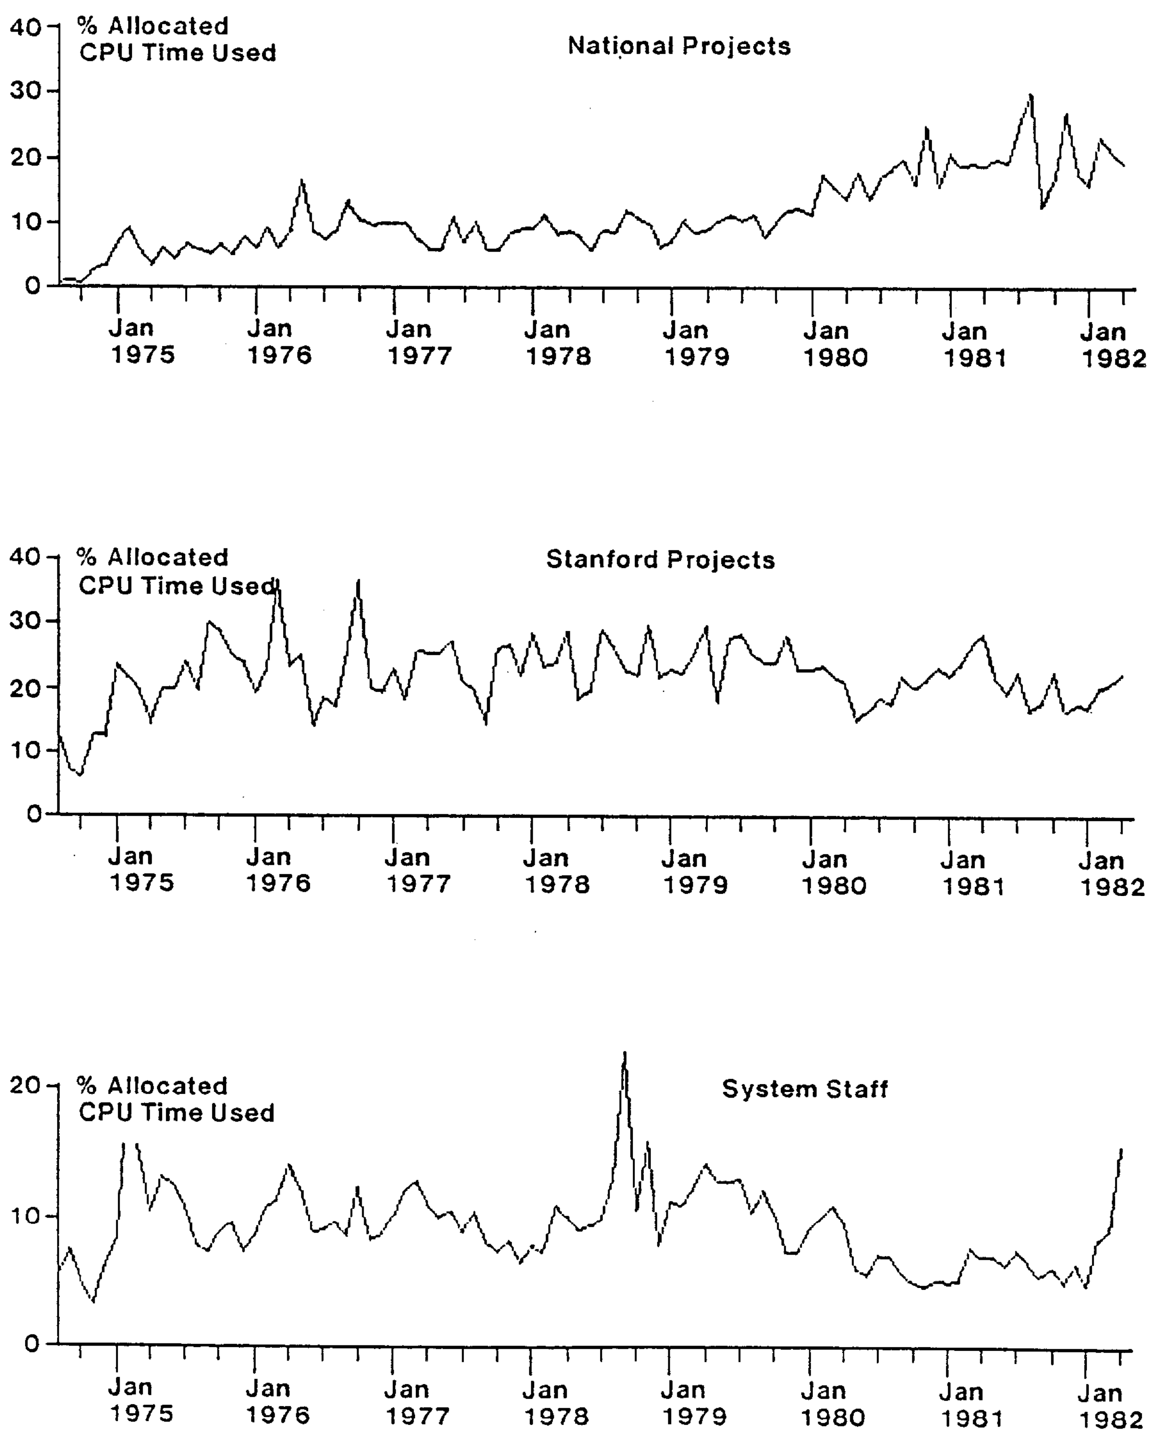

Figure 8. Monthly CPU Usage by Community

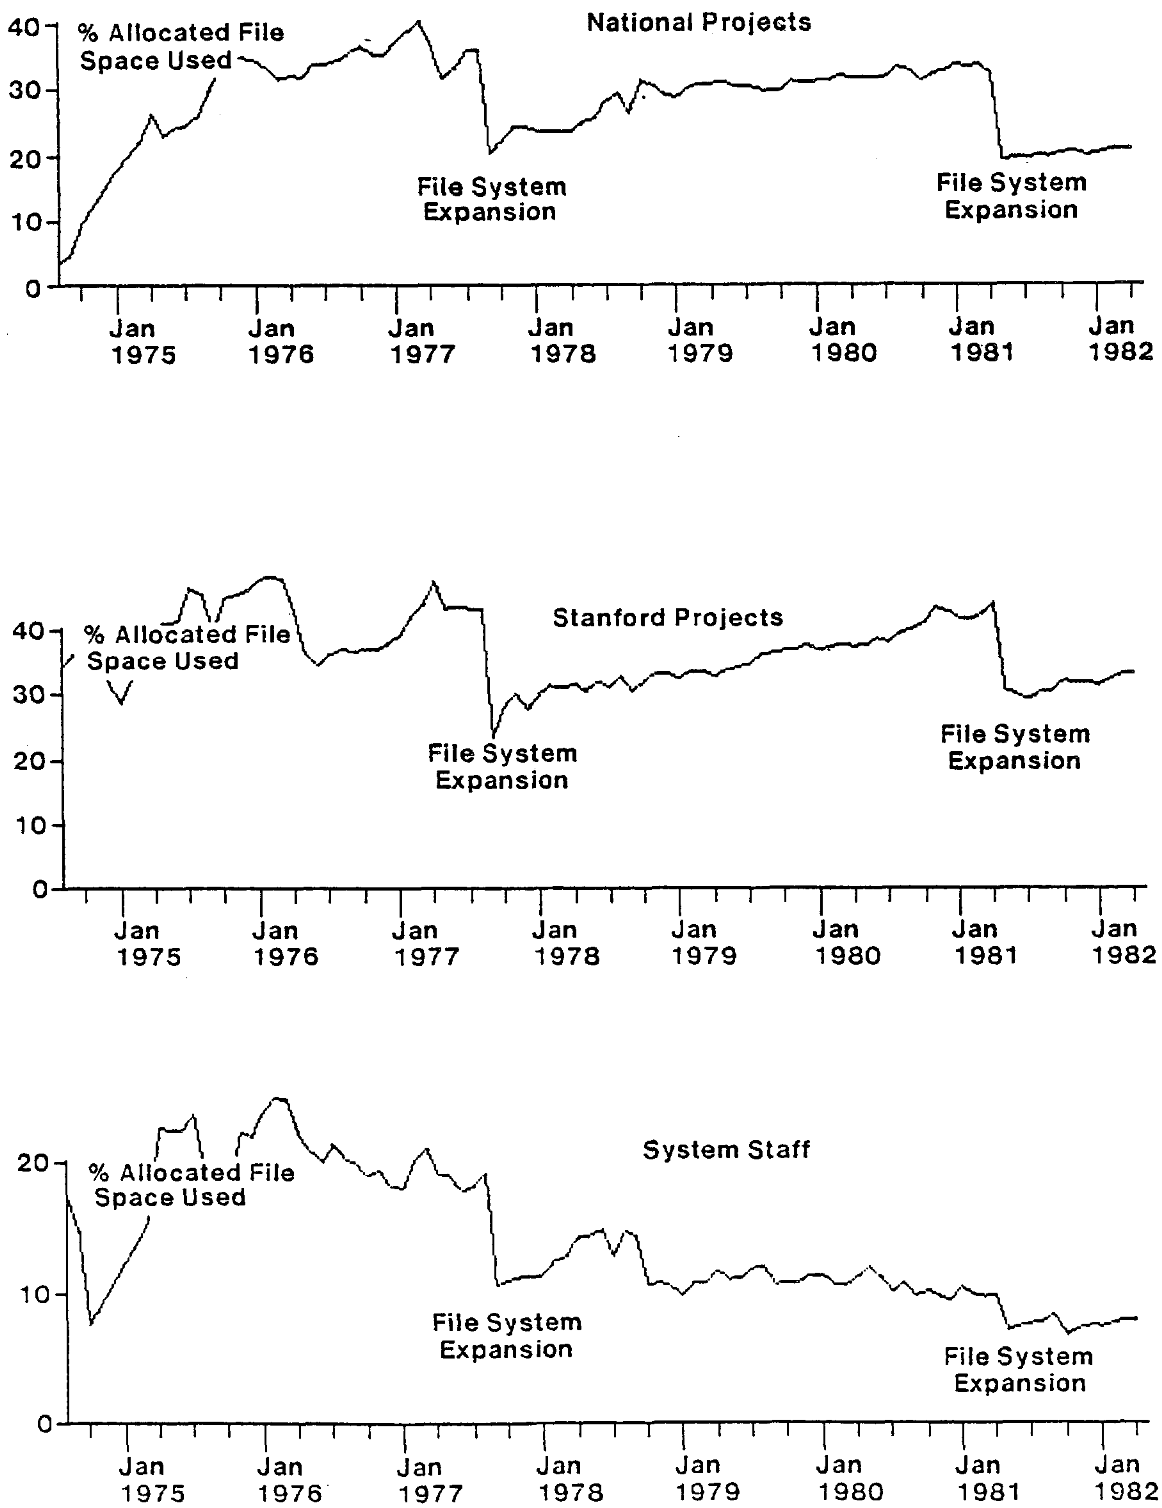

Figure 9. Monthly File Space Usage by Community

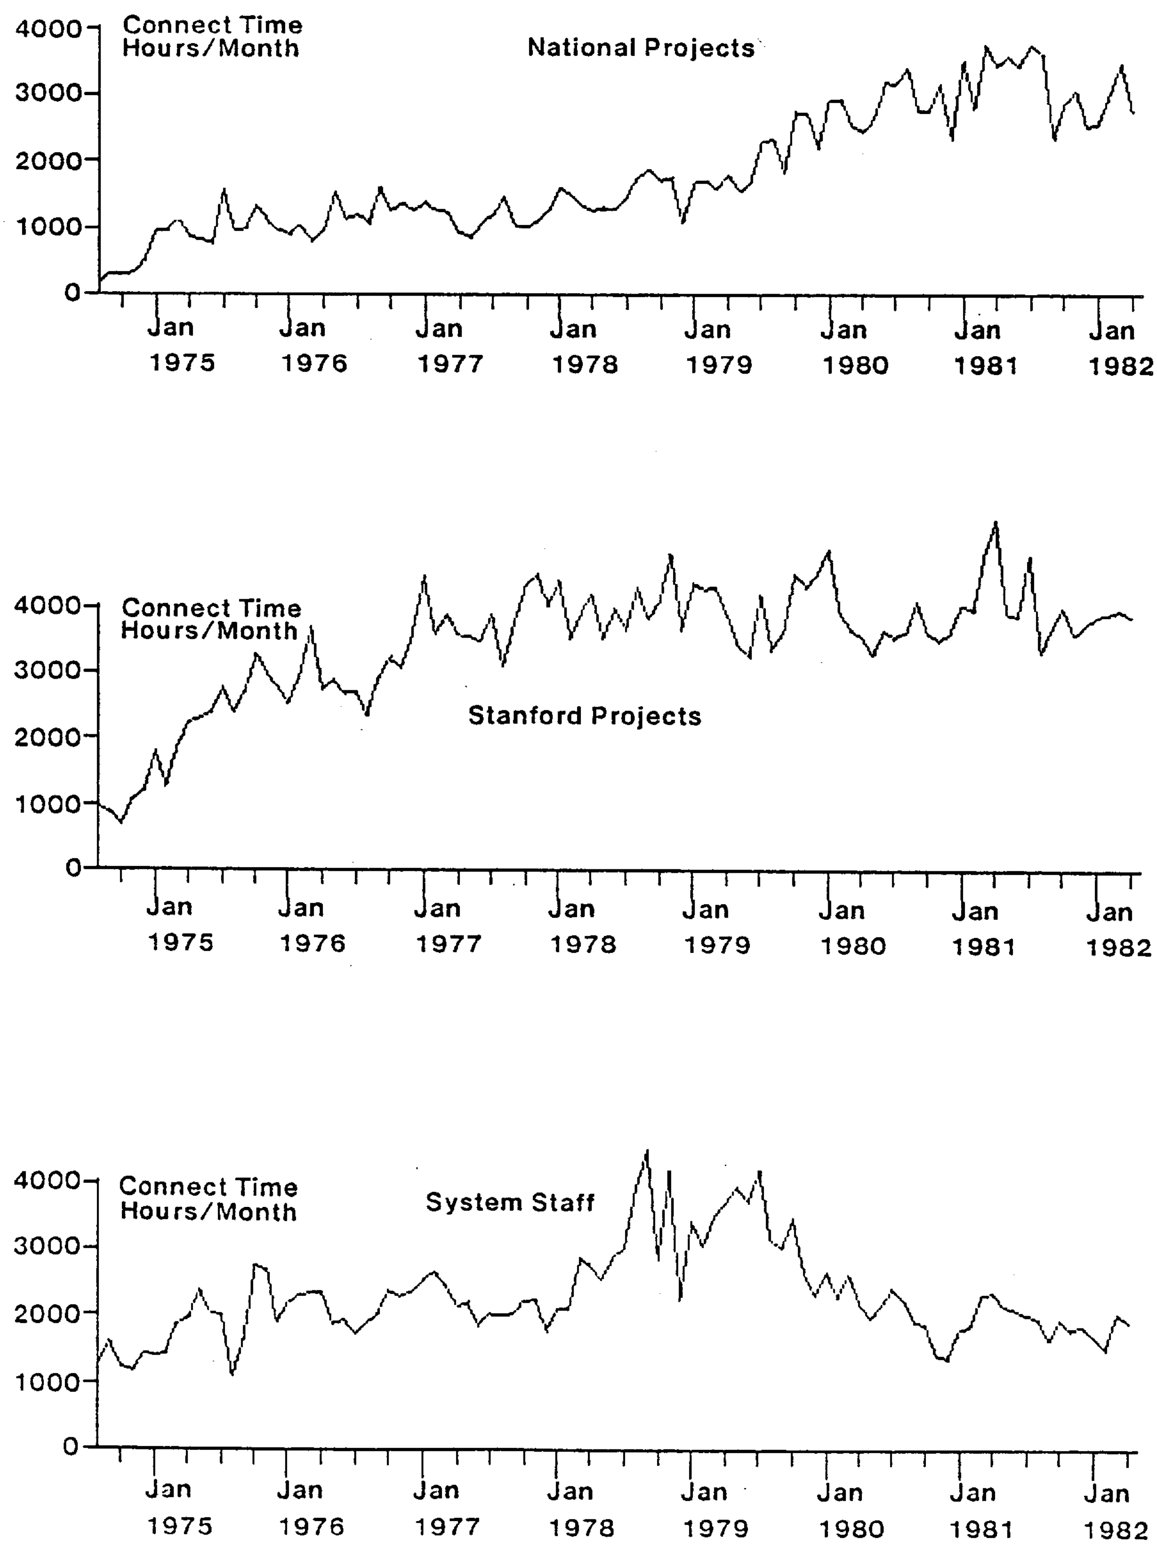

Figure 10. Monthly Terminal Connect Time by Community

3. Individual Project and Community Usage

The following histogram and table shows cumulative resource usage by collaborator project and community during the past grant year. The histogram displays the project distribution of the total CPU time consumed between May 1981 and April 1982. A key point to note is that the total national, Stanford, and Staff resource consumption is approximately in the mandated 40:40:20 ratio. Also note that the large consumption by AIM Users (21.3%) consists mostly of GENET community usage (18.3%), attesting to the great popularity of that facility in the molecular biology community (in spite of the severe access restrictions we have implemented!).

In the table following, entries include a text summary of the funding sources (outside of SUMEX-supplied computing resources) for currently active projects, 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 1981 and April 1982.

It should be noted that the Stanford projects have voluntarily shifted a substantial part of their development work to non-prime time hours which is not explicitly shown in these cumulative data. It should also be noted that a significant part of the DENDRAL, MYCIN, AGE, AI Handbook, and MOLGEN efforts, here charged to the Stanford aliquot, support development efforts dedicated to national community access to these systems. The actual demonstration and use of these programs by extramural users (e.g., the GENET community) is charged to the national community in the "AIM USERS" category, however.

Several of the projects admitted to the National AIM community use the Rutgers-AIM resource as their home base. We do not explicitly list these projects in this annual report covering the Stanford SUMEX-AIM resource. We do record information about the Rutgers resource itself, however, and note its separate resource status with the flag "[Rutgers-AIM]".

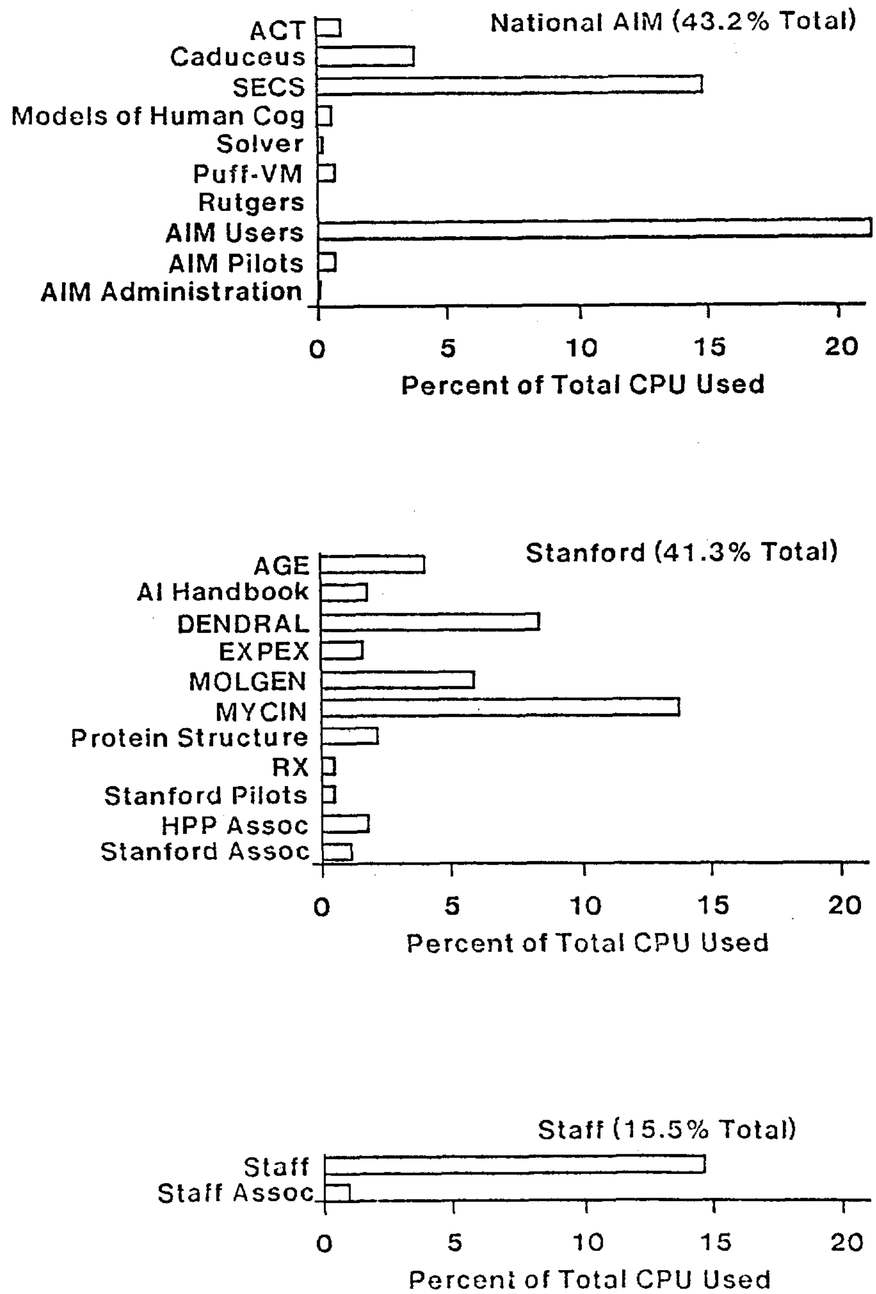

Figure 11. Cumulative CPU Usage Histogram by Project and Community

Resource Use by Individual Project - 5/81 through 4/82

National AIM Community	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) ACT Project "Acquisition of Cognitive Procedures" John R. Anderson, Ph.D. Carnegie-Mellon Univ. NSF IST-80-15357 2/81-2/84 $186,000	51.46	511.92	2702
2) CADUCEUS "Clinical Decision Systems Research Resource" Jack D. Myers, M.D. Harry E. Pople, Jr., Ph.D. University of Pittsburgh NIH RR-01101-05 7/81-6/82 $202,632 NLM LM03710-02 7/81-6/82 $126,746 NLM LM03589-02 Randolph A. Miller, M.D. 7/81-6/82 $29,555 NLM New Invest LM 03889-01 Gordon E. Banks, M.D. 4/82-3/83 $35,725	202.29	5471.02	8487
3) CLIPR Project "Hierarchical Models of Human Cognition" Walter Kintsch, Ph.D. Peter G. Polson, Ph.D. University of Colorado NIE-G-78-0172 9/80-8/81 $46,537 NIMH MH-15872-14-16 7/81-6/82 $37,370 ONR N00014-78-C-0433 Lyle E. Bourne, Jr., Ph.D. 7/81-6/82 $78,125 IBM Michael E. Atwood, Ph.D. 1/82-12/82 $104,000	31.51	725.42	1245

4) PUFF-VM Project "Biomedical Knowledge Engineering in Clinical Medicine" John J. Osborn, M.D. Med. Research Inst., San Francisco Edward H. Shortliffe, M.D.,Ph.D. Stanford University NIH GM-24669 9/78-8/81 $164,000 (*) Johnson & Johnson 1 year $50,000 (*)	34.95	1691.57	2751
5) SECS Project "Simulation & Evaluation of Chemical Synthesis" W. Todd Wipke, Ph.D. U. California, Santa Cruz NIH RR-01059-03S1 7/80-12/81 $36,949 NIEHS ES 02845-01 4/82-3/83 $76,444 NIH GM 31173-01 6/82-5/83 $221,381 (approval pending)	795.25	17123.86	13173
6) SOLVER Project "Problem Solving Expertise" Paul E. Johnson, Ph.D. William B. Thompson, Ph.D. University of Minnesota NICHD T36-HD-17151 (1978-81) NICHD HD-01136 (1978-81) NSF/BNS-77-22075 (1979-82) NSF SE079-13036 (1981-82) NSF SES-8111295 (1981-83)	9.99	306.05	118
7) *** [Rutgers-AIM] ***			
Rutgers Project "Computers in Biomedicine" Saul Amarel, D.Sc. Casimir Kulikowski, Ph.D. Rutgers U., New Brunswick NIH RR-00643 12/81-11/82 $460,944	1.65	49.08	4358
8) AIM Pilot Projects			
AI-COAG (Lindberg)	.56	55.75	86
DATA (Brooks)	36.25	283.01	1027
	-----	-------	----
AIM Pilot Totals	36.81	338.76	1113

9)	AIM Administration	7.83	318.30	3871
10)	AIM Users on Stanford Projects			
	AGE	.09	.95	9
	DENDRAL	105.51	896.34	1253
	HPP	2.25	91.77	204
	MOLGEN	983.58	8288.61	2012
	MYCIN	11.37	428.06	300
	AIM-Associates	25.20	732.48	2408
	Guest (all projects)	14.50	140.06	380
		-------	--------	----
	AIM User Totals	1142.50	10578.27	6566
		------	--------	-----
	Community Totals	2314.24	37114.25	44384

Stanford Community	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) AGE Project (Core) "Attempt to Generalize" Edward A. Feigenbaum, Ph.D. Dept. Computer Science ARPA MDA903-80-C-0107 (**) (partial support)	212.76	3135.87	3333
2) AI Handbook Project (Core) Edward A. Feigenbaum, Ph.D. Dept. Computer Science ARPA MDA903-80-C-0107 (**) (partial support)	96.25	4746.39	2972
3) DENDRAL Project "Resource Related Research: Computers in Chemistry" Carl Djerassi, Ph.D. Dennis H. Smith, Ph.D. Dept. Chemistry NIH RR-00612-13 5/82-4/83 $170,710	448.49	8192.73	19913
4) EXPEX Project "Expert Explanation" Edward H. Shortliffe, M.D.,Ph.D. Michael R. Genesereth, Ph.D. Depts. Medicine/ Computer Science ONR NR 049-479 1/81-12/83 $456,622 NCHSR HS 04422 Randy L. Teach 3/81-8/82 $19,950	84.27	2196.69	1680
5) MOLGEN Project "Computer Sci Application to Molecular Biology" Edward A. Feigenbaum, Ph.D. Bruce G. Buchanan, Ph.D. Laurence H. Kedes, M.D. Douglas L. Brutlag, Ph.D. Peter Friedland, Ph.D. Depts. Computer Science/ Medicine/Biochemistry NSF ECS-8016247 10/81-9/82 $159,785 (*)	316.39	7417.28	8562

6)	MYCIN Projects "Knowledge Engineering for Med. Consultation" Edward H. Shortliffe, M.D.,Ph.D. Bruce G. Buchanan, Ph.D. Depts. Medicine/Computer Science ONR/ARPA N00014-79-C-0302 3/79-3/82 $396,325 NLM LM-03395 (ONCOCIN) Edward A. Feigenbaum, Ph.D. 7/79-6/84 $497,420 NLM LM-00048 7/79-6/84 $196,425	736.59	10591.49	15047
7)	PROTEIN Project "Heuristic Comp. Applied to Prot. Crystallog." Edward A. Feigenbaum, Ph.D. Dept. Computer Science NSF MCS-81-17330 1/82-1/83 $28,976	112.88	1665.96	3990
8)	RX Project "Knowledge from Clinical Databases" Robert L. Blum, M.D. Gio C.M. Wiederhold, Ph.D. Depts. Computer Science/ Electrical Engrg. NLM New Invest LM-03370 7/79-6/82 $90,000 NCHSR 4/79-3/81 $35,000	26.09	1010.03	2299
9)	Stanford Pilot Projects			
	ABARBANEL	2.21	56.78	236
	BRINKLEY	21.95	548.81	763
		------	--------	----
	Stanford Pilot Totals	24.16	605.59	999
10)	Stanford and HPP Assoc.	156.41	6802.63	10146
		-------	--------	-----
	Community Totals	2214.29	46364.66	68941

SUMEX Staff	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) Staff	786.11	21146.04	12560
2) System Associates	47.52	1511.57	4123
3) Misc. Usage	.88	11.01	189
	-------	---------	------
Community Totals	834.51	22668.62	16872

System Operations	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) Operations	2525.35	101687.30	129753
	=======	=========	======
Resource Totals	7888.39	207834.83	259950

(*) Award includes indirect costs.

(**) Supported by a larger ARPA contract MDA-903-80-C-0107 awarded to the Stanford Computer Science Department:

	Current Year (10/81-9/82)	Total Award (10/79-9/82)
Heuristic Programming Project	$ 579,070	$1,613,588
VLSI/CAD Network-based Graphics Development Resource	221,605	685,374
	-----------	----------
Total award	$ 800,675(*)	$2,298,962(*)

4. Network Usage Statistics

The plots in Figure 12 and Figure 13 show the monthly network terminal connect time for TYMNET and ARPANET. This forms the major billing component for SUMEX-AIM TYMNET usage. The terminal connect time does not reflect the time spent in file transfers and mail forwarding.

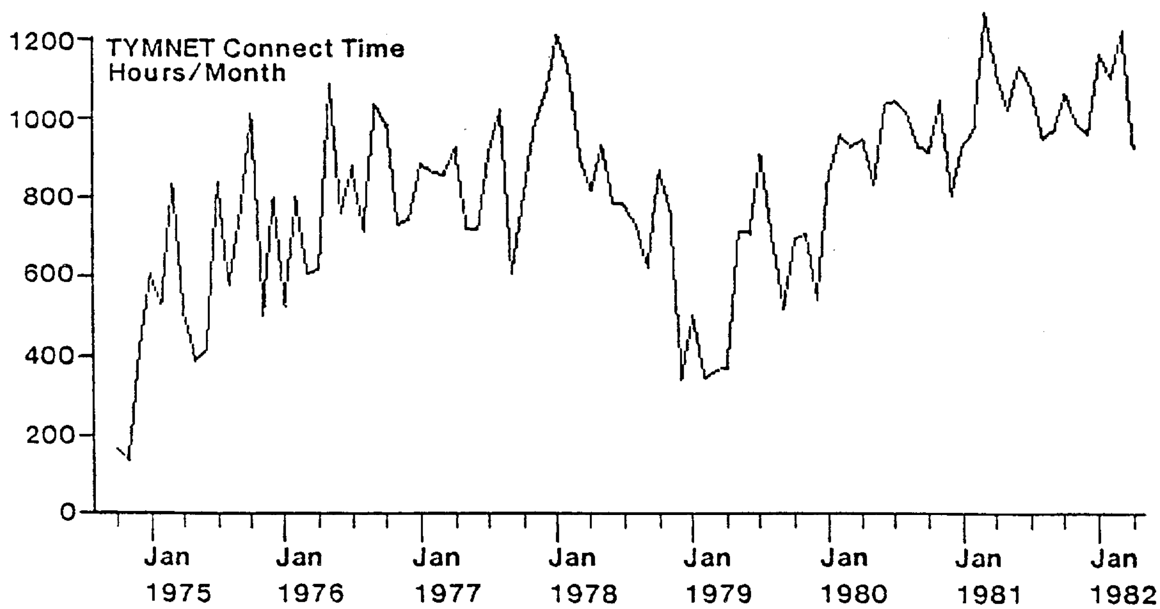

Figure 12. TYMNET Terminal Connect Time

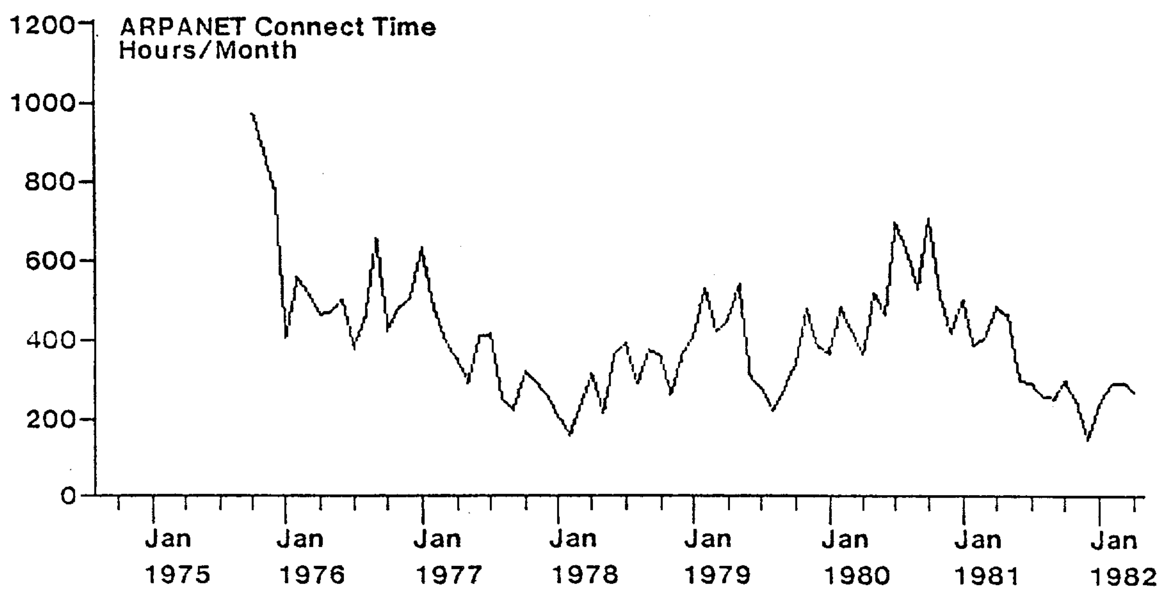

Figure 13. ARPANET Terminal Connect Time

5. System Reliability

System reliability has been very good on average with several periods of particular hardware or software problems. The table below shows monthly downtime for the past year. This includes scheduled maintenance periods, nominally 6 hours per week.

During July and August we experienced a number of serious failures in our Calcomp disk system, induced by poor maintenance practices and parts from Braegen Corporation. In September we switched the maintenance contract for the Calcomp equipment to TRW and have had excellent service and disk reliability since then.

	1981								1982			
	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC	JAN	FEB	MAR	APR
Downtime (Hrs)	44	64	113	100	42	24	13	65	42	32	18	38

TABLE 1. System Downtime by Month