disease profile. We are experimenting with different scales for scoring each feature-value pair, and several methods for combining the scores to form a differential diagnosis. A disease-independent import is also assigned to each feature-value but only a two-valued scale is used. This is because, in PATHFINDER, imports are only used to make boolean or yes/no decisions (see below). In addition to import, PATHFINDER utilizes the concept of *classic* features for a disease -- within each disease frame, the pathologist marks those feature-value pairs which are considered to be part of the classic pattern of the disease.

The PATHFINDER knowledge base contains information about obvious association between features. This information is of the form: "Don't ask about feature x unless feature y has certain values." For example, it wouldn't make sense to ask about the degree or range of follicularity if there are no follicles in the tissue section. The feature links also serve to identify interdependencies among features. Feature interdependence is a problem because it can lead to inaccuracies in scoring hypotheses.

The prototype knowledge base was constructed by Dr. Nathwani. During the beginning part of 1984, we organized two meetings of the entire team including the pathology experts to define the selection of diseases to be included in the system, and the choice of features to be used in the scoring process.

*D. Publications Since January 1984*

Horvitz, E.J., Heckerman, D.E., Nathwani, B.N. and Fagan, L.M.: *Diagnostic Strategies in the Hypothesis-directed PATHFINDER System, Node Pathology*. HPP Memo 84-13. Proceedings of the First Conference on Artificial Intelligence Applications, Denver, Colorado, Dec., 1984.

*E. Funding Support*

Research Grant submitted to National Institutes of Health, March, 1984.
Grant Title:"Computer-aided Diagnosis of Malignant Lymph Node Diseases"
Principal Investigator: Bharat Nathwani

Professional Staff Association, Los Angeles County Hospital, $10,000.

University of Southern California, Comprehensive Cancer Center, $30,000.

Project Socrates, Univ. of Southern Calif., Gift from IBM of IBM PC/XT.

## II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

*A. Medical Collaborations and Program Dissemination via SUMEX*

Because our team of experts are in different parts of the country and the computer scientists are not located at the USC, we envision a tremendous use of SUMEX for communication, demonstration of programs, and remote modification of the knowledge base. The proposal mentioned above was developed using the communication facilities of SUMEX.

*B. Sharing and Interaction with Other SUMEX-AIM Projects*

Our project depends heavily on the techniques developed by the INTERNIST/CADUCEUS project. We have been in electronic contact and have met with members of the INTERNIST/CADUCEUS project, as well as, been able to utilize information and experience with the INTERNIST program gathered over the years through the AIM conferences and on-line interaction. Our experience with the extensive development of the pathology knowledge base utilizing multiple experts should provide for intense and helpful discussions between our two projects.

The SUMEX pilot project, RXDX, designed to assist in the diagnosis of psychiatric disorders is currently using a version of the PATHFINDER program on the DEC-20 for the development of early prototypes of future systems.

*C. Critique of Resource Management*

The SUMEX resource has provided an excellent basis for the development of a pilot project. The availability of a pre-existing facility with appropriate computer languages, communication facilities (especially the TYMNET network), and document preparation facilities allowed us to make good progress in a short period of time. The management has been very useful in assisting with our needs during the start of this project.

## III. RESEARCH PLANS

*A. Project Goals and Plans*

**Collection and refinement of knowledge about lymph node pathology**

The knowledge base of the program is about to undergo revision by the expert, and then will be extensively tested. A logical next step would be to extend the program to clinical settings, as well as possible extensions of the knowledge base.

Other possible extensions include: developing techniques for simplifying the acquisition and verification of knowledge from experts, creating mapping schemes that will facilitate the understanding of the many classifications of non-Hodgkin's lymphomas. We will also attempt to represent knowledge about special diagnostic entities, such as multiple discordant histologies and atypical proliferations, which do not fit into the classification methods we have utilized.

**Representation Research**

We hope to enhance the INTERNIST-1 model by structuring features so that overlapping features are not incorrectly weighted in the decision making process, implementing new methods for scoring hypotheses, and creating appropriate explanation capabilities.

*B. Requirements for Continued SUMEX Use*

We are currently dependent on the SUMEX computer for the use of the program by remote users, and for project coordination. We have transferred the program over to Portable Standard Lisp which is used by several users on the SUMEX system. While the switch to workstations has lessened our requirements for computer time for the development of the algorithms, we will continue to need the SUMEX facility for the interaction with each of the research locations specified in our NIH proposal. The HP equipment is currently unable to allow remote access, and thus the program will have to be maintained on the 2060 for use by all non-Stanford users.

*C. Requirements for Additional Computing Resources*

Most of our computing resources will be met by the 2060 plus the use of the HP9836 workstation. We will need additional file space on the 2060 as we quadruple the size of our knowledge base. We will continue to require access to the 2060 for communication purposes, access to other programs, and for file storage and archiving.

*D. Recommendations for Future Community and Resource Development*

We encourage the continued exploration by SUMEX of the interconnection of workstations within the mainframe computer setting. We will need to be able to quickly move a program from workstation to workstation, or from workstation back and fcrth to the mainframe. Software tools that would help the transfer of programs from one type of workstation to another would also be quite useful. Until the type of workstations that we are using in this research becomes inexpensive ($5000 or less), we will continue to need a machine like SUMEX to provide others with a chance to experiment with our software.

# 6.4.2. RXDX Project

RXDX Project

Robert Lindsay, Ph.D.
Michael Feinberg, M.D., Ph.D.
Manfred Kochen, Ph.D.
University of Michigan
Ann Arbor, Michigan

## I. SUMMARY OF RESEARCH PROGRAM

### *A. Project Rationale*

We are developing a prototype expert system that could act as a consultant in the diagnosis and management of depression. Health professionals will interact with the program as they might with a human consultant, describing the patient, receiving advice, and asking the consultant about the rationale for each recommendation. The program uses a knowledge base constructed by encoding the clinical expertise of a skilled psychiatrist in a set of rules and other knowledge structures. It will use this knowledge base to decide on the most likely diagnosis (endogenous or nonendogenous depression), assess the need for hospitalization, and recommend specific somatic treatments when this is indicated (e.g., tricyclic antidepressants). The treatment recommendation will take into account the patient's diagnosis, age, concurrent illnesses, and concurrent treatments (drug interactions).

### *B. Medical Relevance and Collaboration*

There has been a growing emphasis in American psychiatry on careful diagnosis using clearly defined clinical criteria (Feighner, et al., 1972; Spitzer, et al., 1975, 1980; Feinberg and Carroll, 1982, 1983). These efforts have led to several sets of criteria for the diagnosis of psychiatric disorders. The "St. Louis" criteria (Feighner, et al., 1972) were succeeded by the Research Diagnostic Criteria (RDC), formulated by researchers from St. Louis and New York (Spitzer, et al., 1975). The RDC led directly to the criteria that are now quasi-official in American psychiatry, DSM-III (Spitzer, et al., 1980). All of these criteria lists were based on a combination of clinical opinion and literature review, and use a decision-tree approach to making a diagnosis. These diagnostic systems have been shown to be acceptably reliable, but their validity remains untested. Other groups have used a multivariate statistical approach to diagnosis. Roth and his colleagues (Carney, et al., 1965) published a discriminant index for distinguishing "endogenous" from "neurotic" depressed patients. This work was repeated by Kiloh, et al. (1972) with much the same results, confirming the findings of Carney, et al. (1965).

We have done similar work, deriving two discriminant indices for separating endogenous depressed patients (unipolar or bipolar) from nonendogenous (neurotic) patients. We cross-validated these indices in separate groups of patients, and also validated them against an external standard, the dexamethasone suppression test (Feinberg and Carroll, 1982, 1983). At the same time, we and others have been further developing this and other biological measures that may differentiate between patients with endogenous and nonendogenous depression. These include neuroendocrine tests such as the dexamethasone suppression test (DST) and quantitative studies of sleep using EEG. Carroll, et al. (1981) have shown that the DST is abnormal in about 67%

of patients with endogenous depression (melancholia) and only 5-10% with nonendogenous (neurotic) depression. Kupfer, et al. (1978) and Feinberg, et al. (1982) have similar results with EEG studies of sleep. These biological markers may be useful for routine clinical use, and can certainly be used as external validating criteria to test the performance of different clinical diagnostic methods, including those mentioned above. Furthermore, we have developed biological criteria for "definitely endogenous" depression and "definitely nonendogenous" depression based on DST and sleep EEG. (Carroll, et al., 1980). Our goal is to use these criteria as an external validating criterion for assessing the performance of various new or different diagnostic schemes, in particular an expert system of the sort we are developing.

*C. Highlights of Research Progress*

We examined two other SUMEX-based psychiatry projects, the BLUEBOX project of Mulsant and Servan-Schreiber (1984), and the HEADMED project of Heiser and Brooks (1978, 1980). Mulsant and Servan-Schreiber visited us at Michigan and discussed the rationale and progress of their project. Heiser also visited with us and agreed to collaborate with our project as a consultant.

At Michigan, we encoded the Hamilton Rating Scale (Hamilton, 1967) into EMYCIN rules. This is the standard scale (in English) for rating the severity of depression, and many of the items in it are relevant to our consultant program. We moved our work to the AGE system, breaking the Hamilton scale into its component subscales and adding other components to determine patient demographic information, personal and family psychiatric history, and other rating scale information. We then introduced other knowledge sources to construct a differential diagnosis list for psychiatric illnesses based on our expert's taxonomy and methods. We are now focussing on rules that discriminate endogenous from non-endogenous depression. Concurrently we are developing a treatment knowledge base on a LISP workstation. Thus far, the treatment knowledge base contains information about drug therapies, including types, dosages, activities, interactions, and side effects.

We have conducted interviews with patients recently admitted to the University of Michigan Adult Psychiatric Hospital. They are interviewed by Feinberg and the interviews are observed by Lindsay plus a group of psychiatric residents, psychiatrists and psychologists. After the interview, Feinberg is debriefed by Lindsay, and then the others discuss the case. These data are the initial source of the expert knowledge base for our consultant.

*D. List of Relevant Publications*

This project has not yet produced any publications. The following list contains the references cited above, including our previous publications relevant to the RxDx Project.


1. Carney, M. W. P., Roth, M. and Garside, R. F.:*The diagnosis of depressive syndromes and the prediction of ECT response*, Brit. J. Psychiatry, 111, 659-674, 1965.

2. Carroll, B. J., Feinberg, M., Greden, J. F., Haskett, R. F., James, N. McI., Steiner, M., and Tarika, J.: *Diagnosis of endogenous depression: Comparison of clinical, research, and neuroendocrine criteria*, J. Affect Dis., 2, 177-194, 1980.

3. Carroll, B. J., Feinberg, M., Greden, J. F., Tarika, J., Albala, A. A., Haskett, R. F., James, N. McI., Kronfol, Z., Lohr, N., Steiner, M., de Vigne, J-P, and Young, E.:*A specific laboratory test for the diagnosis of melancholia*, Standardization, validation, and clinical utility. Arch. Gen. Psychiatry, 38, 15-22, 1981.

4. Feighner, J. P., Robins, E., Guze, S. B., Woodruff, R. A., Winokur, G., and Munoz, R.: *Diagnostic criteria for use in psychiatric research*, Arch. Gen. Psychiatry, 26, 57-63, 1972.

5. Feinberg, M. and Lindsay, R. K.: *Expert systems*. Proceedings of the NCDEU Annual Meeting, Key Biscayne, Florida, May 1985.

6. Feinberg, M. and Carroll, B. J.: *Separation of subtypes of depression using discriminant analysis: I. Separation of unipolar endogenous depression from non-endogenous depression*, Brit. J. Psychiatry, 140, 384-391, 1982.

7. Feinberg, M. and Carroll, B. J.:*Separation of subtypes of depression using discriminant analysis. II. Separation of bipolar endogenous depression from nonendogenous ("neurotic") depression*, J. Affective Disorders, 5, 129-139, 1983.

8. Feinberg, M. and Carroll, B.J.: *Biological markers for endogenous depression in series and parallel*, Biological Psychiatry 19:3-11, 1984.

9. Feinberg, M. and Carroll, B.J.: *Biological and nonbiological depression*, Presented at Annual Meeting of the Society of Biological Psychiatry, Los Angeles, May, 1984, Abstract #81.

10. Feinberg, M., Gillin, J. C., Carroll, B. J., Greden, J. F., and Zis, A. P.:*EEG studies of sleep in the diagnosis of depression*, Biological Psychiatry, 17, 305-316, 1982.

11. Heiser, J. F. and Brooks, R. E.:*Design considerations for a clinical psychopharmacology advisor*, Proc. Second Annual Symp. on Computer Applications in Medical Care. New York: IEEE, 1978, 278-285.

12. Heiser, J. F. and Brooks, R. E.:*Some experience with transferring the MYCIN system to a new domain*, IEEE Trans. on Pattern Analysis and Machine Intelligence, PAMI-2, No. 5, 477-478, 1980.

13. Kiloh, L. G., Andrews, G., and Neilson, M.:*The relationship of the syndromes called endogenous and neurotic depression*, Brit. J. Psychiatry, 121, 183-196, 1972.

14. Kupfer, D. J., Foster, F. G., Coble, P., McPartland, R. J., and Ulrich, R. F.:*The application of EEG sleep for the differential diagnosis of affective disorders*, Am. J. Psychiatry, 135, 69-74, 1978.

15. Mulsant, B. and Servan-Schreiber, D.:*Knowledge engineering: A daily activity on a hospital ward*, Computers in Biomedical Research, 1984.

16. Spitzer, R. L., Endicott, J. and Robins, E.: *Research diagnostic criteria*, (2d ed.) New York State Department of Mental Hygiene, New York Psychiatric Institute, Biometrics Research Division, 1975.

17. Spitzer, R. L.: (Ed.)*Diagnostic and statistical manual of mental disorders*, (3d ed.). Washington, D. C.: American Psychiatric Association, 1980.

18. Van Melle, W.:*The EMYCIN Manual*, Computer Science Department, Stanford University, Report HPP-81-16, 1981.

*E. Funding Support*

We have received support from the Vice-President for Research at the University of Michigan, and from the NIH "Small Grants" Program (Grant Number ro3MH40239-01; Total Direct Costs = $13,850). These funds have enabled us to gather the pilot data for a grant application to be submitted to NIH on July 1, 1985.

## II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

### *A. Medical Collaboration and Program Dissemination via SUMEX*

We have established via SUMEX a community of researchers who are interested in AI applications in psychiatry. We also have used the message system to communicate with other AI scientists at SUMEX and elsewhere.

### *B. Sharing and Collaboration with other SUMEX-AIM Projects*

Our use of EMYCIN and AGE has been of major importance. In addition, we have worked with Dr. Larry Fagan to learn about his Pathfinder program. We used that program, on SUMEX, to obtain some information for the RxDx project by applying it to data we previously collected on depression symptom frequencies.

### *C. Critique of Resource Management*

We have been using EMYCIN and AGE in our work, and have found these programs very valuable, saving us many hours of programming in LISP. There are some problems with them, many of which center around discrepancies between the versions described in the manuals and the versions actually running on SUMEX. We would suggest that software be more strongly supported than is now the case, if it and SUMEX are to be even more useful to beginners in AI in Medicine.

SUMEX itself has been invaluable. We don't have ready access to any other machine of equal computing power which also has a strongly supported LISP available. Specifically, the LISP compiler available on the Amdahl 5860 here differs from those used at major AI centers such as Stanford and MIT. We have also made good use of the ARPANET connections that SUMEX offers. Feinberg spent a month of his sabbatical working with Prof. Peter Szolovits at MIT, learning about AI in Medicine. This visit was arranged using computer mail through SUMEX. Lindsay and Feinberg were able to continue their collaborative work while the latter was in Cambridge, using the same medium. The alternative would have been days lost in the mails and many dollars spent on phone calls. We have also been able to get help with problems that arise with EMYCIN and AGE using computer mail.

Most of the limitations of SUMEX, and they are often severe, derive from the necessity to access it via TYMNET. Response time is often impossibly slow, and even at its best the delays are annoying and frustrating, even for editing and debugging. For example, editing is limited to a primitive line editor, since EMACS interacts with the network XON/XOFF handshaking in a disastrous way. The staff has not been helpful in solving these network related problems, probably because they do not have to live with them in their own interactions with the system. In any case, many of the problems are beyond the reach of the Sumex staff. The future of long-haul network collaborations depends critically on increased bandwidth and faster response times.

It would have been helpful to us to obtain the AGE system that runs on a Xerox 1108. However, the $530 price, though perhaps modest in comparison to its development costs, was beyond the reach of our budget. It would be helpful if distribution costs for software could be held under $100.

## III. RESEARCH PLAN

*A. Project Goals and Plans*

Our immediate objective is to develop an expert system that can differentiate patients with the various subtypes of depressive disorder, and prescribe appropriate treatment. This system should perform at about the level of a board-certified psychiatrist, i.e. better than an average resident but not as well as a human expert in depression. Eventually, we plan to enlarge the knowledge base so that the expert system can diagnose and prescribe for a wider range of psychiatric patients, particularly those with illnesses that are likely to respond to psychopharmacological agents. We will design the system so that it could be used by non-medical clinicians or by non-psychiatrist MD's as an adjunct to consultation with a human expert. We plan also to focus on problems of the user interface and the integration of this system with other databases.

*B. Justification and Requirements for continued SUMEX use*

The access to SUMEX resources is essentially our sole means of maintaining contact with the community of researchers working on applications of AI in medicine. Although we plan to move our system to local workstations as soon as we are able, the communications capability of SUMEX will continue to be important.

We anticipate that our requirements for computing time and file space will continue at about the same level for the next year.

*C. Needs and Plans for Other Computing Resources*

As our project evolves and we run into the limitations of the time-shared SUMEX facility, we anticipate employing different expert systems software. At this time, we are not at a stage to say exactly what that will be, but our project is not sufficiently large that we will be able to mount such a software development project ourselves, so we will depend on development and support elsewhere. Ultimately, when our consultant is made available for field trials and clinical use, it will need to be transported to a personal computer that is large enough to support the system yet inexpensive enough to be widely available. A LISP machine is an obvious candidate. While current prices of the necessary hardware are too high, computer prices are continuing to drop. Our design strategy is to avoid limiting ourselves and our aspirations to that which is affordable today; instead we will attempt to project the growth of our project and the price-performance curve of computing such that they meet at some reasonable point in the future.

*D. Recommendations for Future Community and Resource Development*

Valuable as the present SUMEX facilities are to us, they are in many ways limited and awkward to use. The major limitation we feel is the difficulty and sometimes the impossibility of making contact with everyone who could be of value to us. We hope that greater emphasis will be put on internetwork gateways. It is important not only to establish more of these, but to develop consistent and convenient standards for electronic mail, electronic file transfers, graphic information transfer, national archives and data bases, and personal filing and retrieval (categorization) systems. The present state of the art feels quite limiting, now that the basic concepts of computer networking have become available and have proved their potential.

We expect that the role of the SUMEX-AIM resource will continue to evolve in the direction of increased importance of communication, including graphical information, electronic dissemination of preprints, and database and program access. The need for computer cycles on a large mainframe will diminish. We hope to have continued access to the system for communication, but do not anticipate continued use of it as a LISP computation server beyond the next year or eighteen months.

If fees for using SUMEX resources were imposed, this would have a drastically limiting effect on the value of the system to us. Even if we had a budget to purchase such services, the inhibiting effect of having a meter running would cause us to make less use of it that we should. We have been conscious of the costs of the system and feel that we have not used it imprudently, even though we have not directly borne its costs.

# Appendix A

# Stanford Knowledge Systems Laboratory

## ARTIFICIAL INTELLIGENCE RESEARCH IN THE KNOWLEDGE SYSTEMS LABORATORY (Incorporating the Heuristic Programming Project)

Stanford University
Department of Computer Science/Department of Medicine
April 1985

The Knowledge Systems Laboratory (KSL) is an artificial intelligence research laboratory of about 90 people -- faculty, staff, and students -- within the Departments of Computer Science and Medicine at Stanford University. KSL is the new name for the interdisciplinary AI research community that has evolved over the past two decades. Begun as the DENDRAL Project in 1965 and known as the Heuristic Programming Project from 1972 to 1984, the new organization reflects the increasing complexity and diversity of the research now under way. The KSL is a modular laboratory, consisting of five collaborating yet distinct groups with different research themes:

- **The Heuristic Programming Project (HPP)**, Professor Edward A. Feigenbaum, scientific director -- blackboard systems, concurrent system architectures for AI, and the modeling of discovery processes. Executive director: Robert Engelmore. Research scientists: Harold Brown, Byron Davies, Bruce Delagi, Peter Friedland, Barbara Hayes-Roth, and H. Penny Nii. Consulting professor: Richard Gabriel.
- **The HELIX Group**, Professor Bruce G. Buchanan, scientific director -- machine learning, transfer of expertise, and problem solving. Faculty: Paul S. Rosenbloom (joint appointment, Computer Science and Psychology). Research scientists: James Brinkley, William J. Clancey, Barbara Hayes-Roth.
- **The Medical Computer Science (MCS) Group**, Professor Edward H. Shortliffe, scientific director (Department of Medicine with courtesy appointment in Computer Science) -- research on and advanced application of AI to medical problems; includes the Medical Information Sciences (MIS) program. Research scientist: Lawrence M. Fagan.
- **The Logic Group**, Professor Michael R. Genesereth, scientific director -- formal reasoning and introspective systems. Research scientist: Matthew L. Ginsberg.
- **The Symbolic Systems Resources Group (SSRG)**, Thomas C. Rindfleisch, scientific director (joint appointment, Computer Science and Medicine) -- research on and operation of computing resources for AI research, including the SUMEX facility. Assistant director: William J. Yeager.

Tom Rindfleisch serves as KSL project director.

This brochure summarizes the goals and methodology of the KSL, its research and academic programs, its achievements, and the research environment of the laboratory.

## Basic Research Goals and Methodology

Throughout a 20-year history, the KSL and its predecessors, DENDRAL and HPP, have concentrated on research in expert systems -- that is, systems using symbolic reasoning and problem-solving processes that are based on extensive domain-specific knowledge. The KSL's approach has been to focus on applications that are themselves significant real-world problems, in domains such as science, medicine, engineering, and education, and that also expose key, underlying AI research issues. For the KSL, AI is largely an empirical science. Research problems are explored, not by examining strictly theoretical questions, but by designing, building, and experimenting with programs that serve to test underlying theories.

The basic research issues at the core of the KSL's interdisciplinary approach center on the computer representation and use of large amounts of domain-specific knowledge, both factual and heuristic (or judgmental). These questions have guided our work since the 1960s and are now of central importance in all of AI research:

1. **Knowledge representation.** How can the knowledge necessary for complex problem solving be represented for its most effective use in automatic inference processes? Often, the knowledge obtained from experts is heuristic knowledge, gained from many years of experience. How can this knowledge, with its inherent vagueness and uncertainty, be represented and applied?
2. **Knowledge acquisition.** How is knowledge acquired most efficiently -- whether from human experts, from observed data, from experience, or by discovery? How can a program discover inconsistency and incompleteness in its knowledge base? How can knowledge be added without perturbing the established knowledge base?
3. **Use of knowledge.** By what inference methods can many sources of knowledge of diverse types be made to contribute jointly and efficiently toward solutions? How can knowledge be used intelligently, especially in systems with large knowledge bases, so that it is applied in an appropriate manner at the appropriate time?
4. **Explanation and tutoring.** How can the knowledge base and the line of reasoning used in solving a particular problem be explained to users? What constitutes a sufficient or an acceptable explanation for different classes of users? How can problem-solving systems be combined with pedagogical and user knowledge to implement intelligent tutoring systems?
5. **System tools and architectures.** What kinds of software tools and system architectures can be constructed to make it easier to implement expert programs with greater complexity and higher performance? What kinds of systems can serve as vehicles for the cumulation of knowledge of the field for the researchers?

## Research and Academic Programs

### CURRENT RESEARCH PROJECTS

The following list of projects now under way within the five KSL research groups gives a brief summary of the major goals of each project and lists the personnel (staff and Ph.D. candidates) directly involved. More complete information on individual projects can be obtained from the person indicated as the project contact. Inquiries should be addressed in care of:

Knowledge Systems Laboratory
Department of Computer Science
Stanford University
701 Welch Road, Building C
Palo Alto, CA 94304
415-497-3444

**The Heuristic Programming Project**

- **Advanced Architectures Project** -- Design a new generation of computer architectures to exploit concurrency in blackboard-based signal understanding systems.
*Personnel*: Edward A. Feigenbaum (contact), Harold Brown, Byron Davies (TI), Bruce Delagi (DEC), Richard Gabriel, Penny Nii, Sayuri Nishimura, Jim Rice, Eric Schoen, Jerry Yan.
- **Knowledge-Based VLSI Design Project** -- Study the hierarchical design process involved in the development of complex very large scale integrated circuits.
*Personnel*: Harold Brown (contact), Jerry Yan.
- **Blackboard Architecture Project** -- Integrate current knowledge about blackboard framework problem-solving systems and develop a domain-independent model that includes knowledge-based control processes.
*Personnel*: Barbara Hayes-Roth (contact).
- **MOLGEN** -- Study the processes of scientific theory formation and modification, using recently developed models of genetic regulation as an example.
*Personnel*: Peter Friedland (contact), Charles Yanofsky (Biological Science), Peter Karp.

**The HELIX Group**

- **PROTEAN** -- Study complex symbolic constraint-satisfaction problems in the blackboard framework with application to protein structure determination from nuclear magnetic resonance data.
*Personnel*: Bruce Buchanan (contact), Oleg Jardetzky (Nuclear Magnetic Laboratory), Jim Brinkley, Barbara Hayes-Roth, Russ Altman, Olivier Lichtarge.
- **NEOMYCIN/GUIDON2** -- Develop knowledge representation and explanation capabilities for the computer-aided teaching of diagnostic reasoning.
*Personnel*: Bill Clancey (contact), Stephen Barnhouse, Diane Hasling, David C. Wilkins.
- **SOAR** -- Develop a general production-system-based problem-solving architecture that integrates reasoning, domain expertise, learning, and planning of problem-solving strategies.
*Personnel*: Paul Rosenbloom (contact), Andrew Golding, Amy Unruh.
- **Knowledge Acquisition Studies** -- Study the processes for transferring knowledge into a computer program, including learning by induction, analogy, watching, chunking, reading, and discovery.
*Personnel*: Bruce Buchanan (contact), Li-Min Fu, Russell Greiner, Ramsey Haddad, David C. Wilkins.

**The Medical Computer Science Group**

- **ONCOCIN** -- Develop knowledge-based systems for the administration of complex medical treatment protocols such as those encountered in cancer chemotherapy. *Personnel*: Ted Shortliffe (contact), Charlotte Jacobs (Oncology), Larry Fagan, David Combs, Gregory Cooper, Jay Ferguson, Christopher Lane, Janice Rohn, Homer Chin, Holly Jimison, Curt Langlotz, Mark Musen, Glenn Rennels.

- PATHFINDER -- Develop a knowledge-based system for diagnosis of lymph node pathology.
  *Personnel*: Ted Shortliffe, Bharat Nathwani (USC), Larry Fagan (contact), David Heckerman, Eric Horvitz.

**The Logic Group**

- **Metalevel Representation System (MRS)** -- Study logic-based introspective programs that can reason about and control their own problem-solving activities.
  *Personnel*: Mike Genesereth (contact), Matt Ginsberg, Russ Greiner, Ben Grosof, Yung-Jen Hsu, David E. Smith, Devika Subramanian, Richard Treitel.
- **The DART/HELIOS Project** -- Study an integrated design environment that includes capabilities for design specification, refinement, and validation; fabrication engineering; and failure diagnosis and testing.
  *Personnel*: Mike Genesereth (contact), Glenn Kramer (Fairchild), Narinder Singh.
- **Intelligent Agent Project** -- Study planning and problem-solving activities for an intelligent interface between human users and complex computing environments.
  *Personnel*: Mike Genesereth (contact), Matt Ginsberg, Jeff Finger, Jeff Rosenschein, Jock Mackinlay, Vineet Singh.
- **Intelligent Task Automation** -- Build a program that can use the description of a manufacturing task to develop a plan by which a robot can carry out the task.
  *Personnel*: Mike Genesereth (contact), Matt Ginsberg, Jeff Finger, David E. Smith, Richard Treitel.

**The Symbolic Systems Resources Group (SSRG)**

- **SUMEX-AIM Resource** -- Develop and operate a national computing resource for biomedical applications of artificial intelligence in medicine and for basic research in AI at KSL.
  *Personnel*: Tom Rindfleisch (contact), Bill Croft, Frank Gilmurray, Christopher Schmidt, Andrew Sweer, Israel Torres, Bob Tucker, Nicholas Veizades, Bill Yeager.
- **Financial Resource Management** -- Develop an expert system for financial resource planning.
  *Personnel*: Tom Rindfleisch (contact), Bruce Buchanan.

**Other Projects**

The KSL also has close ties to collaborative projects. These include PIXIE, developing an intelligent tutoring system, under Derek Sleeman in the School of Education, and RADIX, studying discovery of knowledge from databases, under Bob Blum in Computer Science.

## STUDENTS AND SPECIAL DEGREE PROGRAMS

Graduate students are an essential part of the research productivity of the KSL. Currently 41 students are working with our projects centered in Computer Science and another 12 students are working with the MCS/MIS programs in Medicine. Of the 41 working in Computer Science, 25 are working toward Ph.D. degrees, and 16 are working toward M.S. degrees. A number of students are pursuing interdisciplinary programs and come from the Departments of Engineering, Mathematics, Education, and Medicine.

Because of the highly interdisciplinary and experimental nature of KSL research, two special degree programs have been established:

The **Medical Information Sciences (MIS)** program is an interdepartmental program approved by Stanford University in 1982. It offers instruction and research opportunities leading to the M.S. or Ph.D. degree in medical information sciences, with an emphasis on either medical computer science or medical decision science. The program, directed by Ted Shortliffe and co-directed by Larry Fagan, is formally administered by the School of Medicine, but the curriculum and degree requirements are coordinated with the Dean of Graduate Studies and the Graduate Studies Committee of the University. The program reflects our local interest in the interconnections between computer science, artificial intelligence, and medical problems. Emphasis is placed on providing trainees with a broad conceptual overview of the field and with an ability to create new theoretical and practical innovations of clinical relevance.

The **Master of Science in Computer Science: Artificial Intelligence (MS:AI)** program is a terminal professional degree offered for students who wish to develop a competence in the design of substantial knowledge-based AI applications but who do not intend to obtain a Ph.D. degree. The MS:AI program is administered by the Committee for Applied Artificial Intelligence, composed of faculty and research staff of the Computer Science Department. Normally, students spend two years in the program with their time divided equally between course work and research. In the first year, the emphasis is on acquiring fundamental concepts and tools through course work and and project involvement. During the second year, students implement and document a substantial AI application project.

## Academic and Research Achievements

The primary products of our research are scientific publications on the basic research issues that motivate our work, computer software in the form of the expert systems and AI architectures we develop, and the students we graduate who continue AI research in other academic and industrial laboratories.

The KSL has averaged publishing more than 45 research papers per year in the AI literature, including journal articles, theses, proceedings articles, and working papers. In addition, many talks and invited lectures are given annually. In the past few years, 11 major books have been published by KSL faculty, staff, and former students, and several more are in progress. Those recently published include:

- *Heuristic Reasoning about Uncertainty: An AI Approach*, Cohen, Pitman, 1985.
- *Readings in Medical Artificial Intelligence: The First Decade*, Clancey and Shortliffe, Addison-Wesley, 1984.
- *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, Buchanan and Shortliffe, Addison-Wesley, 1984.
- *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*, Feigenbaum and McCorduck, Addison-Wesley, 1983.
- *Building Expert Systems*, F. Hayes-Roth, Waterman, and Lenat, eds., Addison-Wesley, 1983.

- *System Aids in Constructing Consultation Programs: EMYCIN*, van Melle, UMI Research Press, 1982.
- *Knowledge-Based Systems in Artificial Intelligence: AM and TEIRESIAS*, Davis and Lenat, McGraw-Hill, 1982.
- *The Handbook of Artificial Intelligence*, Volume I, Barr and Feigenbaum, eds., 1981; Volume II, Barr and Feigenbaum, eds., 1982; Volume III, Cohen and Feigenbaum, eds., 1982; Kaufmann.
- *Applications of Artificial Intelligence for Organic Chemistry: The DENDRAL Project*, Lindsay, Buchanan, Feigenbaum, and Lederberg, McGraw-Hill, 1980.

Our laboratory has pioneered in the development and application of AI methods to produce high-performance knowledge-based programs. Programs have been developed in such diverse fields as analytical chemistry (DENDRAL), infectious disease diagnosis (MYCIN), cancer chemotherapy management (ONCOCIN), pulmonary function evaluation (PUFF), machine fault diagnosis (DART), VLSI design (KBVLSI/PALLADIO), and molecular biology (MOLGEN). Some of these programs rival human experts in solving problems in restricted domains. A number of projects have developed generalized software tools for representing and using knowledge; of these, EMYCIN, AGE, MRS, and BB1 are available to outside research groups. Some of our systems and tools (e.g., DENDRAL, PUFF, UNITS, and EMYCIN) are now also being adapted for commercial development and use in the burgeoning AI industry.

Following our lead in work on biomedical applications of AI and the development of the SUMEX-AIM computing resource, a nationally recognized community of academic projects on AI in medicine has grown up.

Central to all KSL research are our faculty, staff, and students. These people have been recognized internationally for the quality of their work and for their continuing contributions to the field. KSL members participate extensively in professional organizations, government advisory committees, and journal editorial boards. They have held major managerial posts and conference chairmanships in both the American Association for Artificial Intelligence (AAAI) and the International Joint Conference on Artificial Intelligence (IJCAI).

Several KSL faculty and former students have received significant honors. In 1976, Ted Shortliffe received the Association of Computing Machinery Grace Murray Hopper award. In 1977, Doug Lenat received the IJCAI Computers and Thought award, and in 1978, Ed Feigenbaum received the National Computer Conference Most Outstanding Technical Contribution award. In 1981, Ted Shortliffe's book *Computer-Based Medical Consultation: MYCIN* was identified as the most frequently cited work in the IJCAI-81 proceedings. In 1982, Doug Lenat won the Tioga prize for the best AAAI conference paper while Mike Genesereth received honorable mention. In 1983, Ted Shortliffe was named a Kaiser Foundation faculty scholar, and Tom Mitchell received the IJCAI Computers and Thought award. In 1984, Randy Davis and Doug Lenat were named among the 100 most promising U.S. scientists under 40 by a prestigious scientific panel assembled by Science Digest. Also in 1984, Ed Feigenbaum was elected a fellow of the American Association for the Advancement of Science (AAAS), and he and Ted Shortliffe were elected fellows of the American College of Medical Information.

## KSL Research Environment

**Funding** -- The KSL is supported solely by sponsored research and gift funds. We have had funding from many sources, including DARPA, NIH/NLM, ONR, NSF, NASA, and foundations and industry. Of these, DARPA and NIH have been the most substantial and long-standing sources of support. All, however, have made complementary contributions to establishing an effective overall research environment that fosters interchanges at the intellectual and software levels and that provides the necessary physical computing resources for our work.

**Computing Resources** -- Under the Symbolic Systems Resources Group, the KSL develops and operates its own computing resources tailored to the needs of its individual research projects. Current computing resources are a networked mixture of mainframe host computers, Lisp workstations, and network utility servers, reflecting the evolving hardware technology available for AI research. Our host machines include a DEC 2060 and 2020 running TOPS-20 (these are the core of the national SUMEX biomedical computing resource) and a VAX 11/780 running UNIX. Our growing complement of Lisp machines includes more than 25 Xerox 1100's, a Xerox Dorado, a Symbolics LM-2, eight Symbolics 3600's, and five Hewlett-Packard 9836's. Network printing, file, gateway, and terminal interface services are provided by dedicated machines ranging from VAX 11/750's to microprocessor systems. These facilities are integrated with other computer science resources at Stanford through an extensive Ethernet and to external resources through the ARPANET and Tymnet. Funding for these resources comes principally from DARPA and NIH.

# Appendix B

# Resource Operations and Usage Data

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

1. Overall resource loading data (page 294).
2. Relative system loading by community (page 295).
3. Individual project and community usage (page 298).
4. Network usage data (page 305).
5. System reliability data (page 305).

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 severe changes in most of the graphs. This is due to many reasons which I will mention briefly here ;

1. Even though the TENEX operating system used on the KI-10 was a forerunner of the current Tops20 operating system, the Tops20 system is still different from TENEX is many ways. Tops20 uses a radically different job scheduling mechanism, different methods for computing monitor statistics, different I/O routines, etc. In general, it can not be assumed that statistics measured on the TENEX system correlate one to one with similar statistics under Tops20.
2. The KL-10 processor on the 2060 is a faster processor than the KI-10 processor used previously. Hence, a job running on the KL-10 will use less CPU time than the same job running on the KI-10. This aspect is further complicated by the fact that the SUMEX KI-10 system was a dual processor system.
3. The SUMEX-AIM Community was changing during the time of the transfer to the 2060. The usage of the GENET community on SUMEX had just been phased out. This part of the community accounted for much of the CPU time used by the AIM community. Since the purchase of the 2060 was partially funded by the Heuristic Programming Project (HPP), an additional number of HPP Core Research Projects started using the 2060, increasing the Stanford communities usage of the machine. And finally, the move to the 2060 occurred during a pivotal time in the community when more and more projects were either moving to their own local timesharing machines, or onto specialized Lisp workstations. It also was the time for the closure of many long time SUMEX-AIM projects, like DENDRAL and PUFF/VM.

Any conclusions reached by comparing the data before and after February, 1983 should be done with caution. The data is included in this years annual report mostly for casual comparison.

Also, it should be noted that monthly statistics are not available for this past year because of problems with the accounting program at this writing. The appropriate average data quantity for the year is shown instead for each month so the graphs appear to be "flat" in the area corresponding to the current period.

*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.

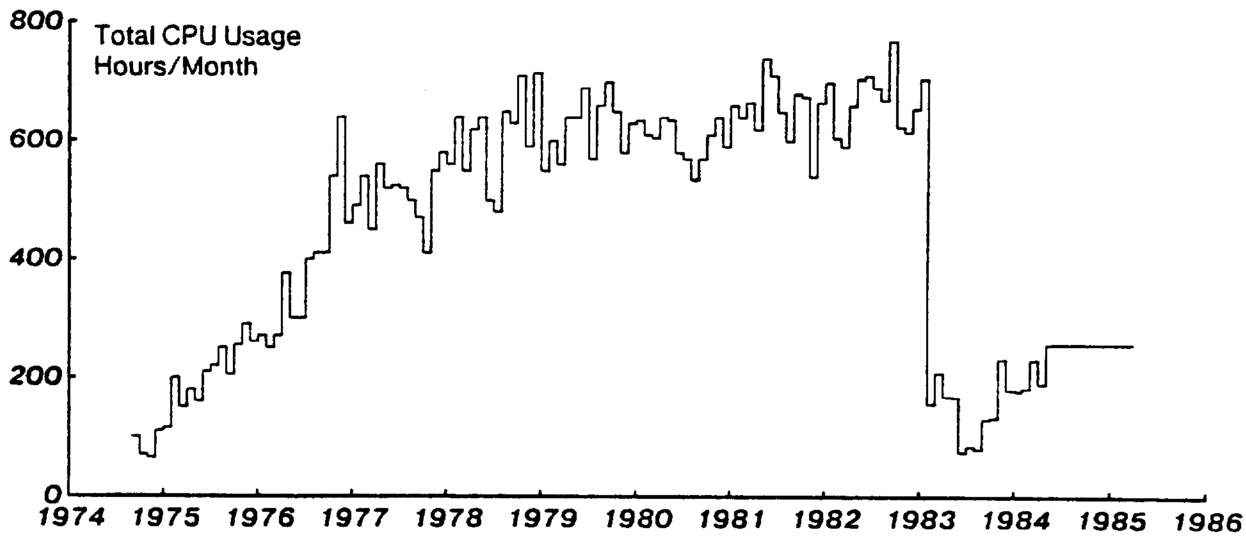


**Figure 14:** Total CPU Time Consumed by Month

*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 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 resources and terminal connect time for each of these three communities relative to their respective aliquots is shown in the plots in Figure 15 and Figure 16. As mentioned on page 293, these plots include both KI-10 and 2060 usage data.

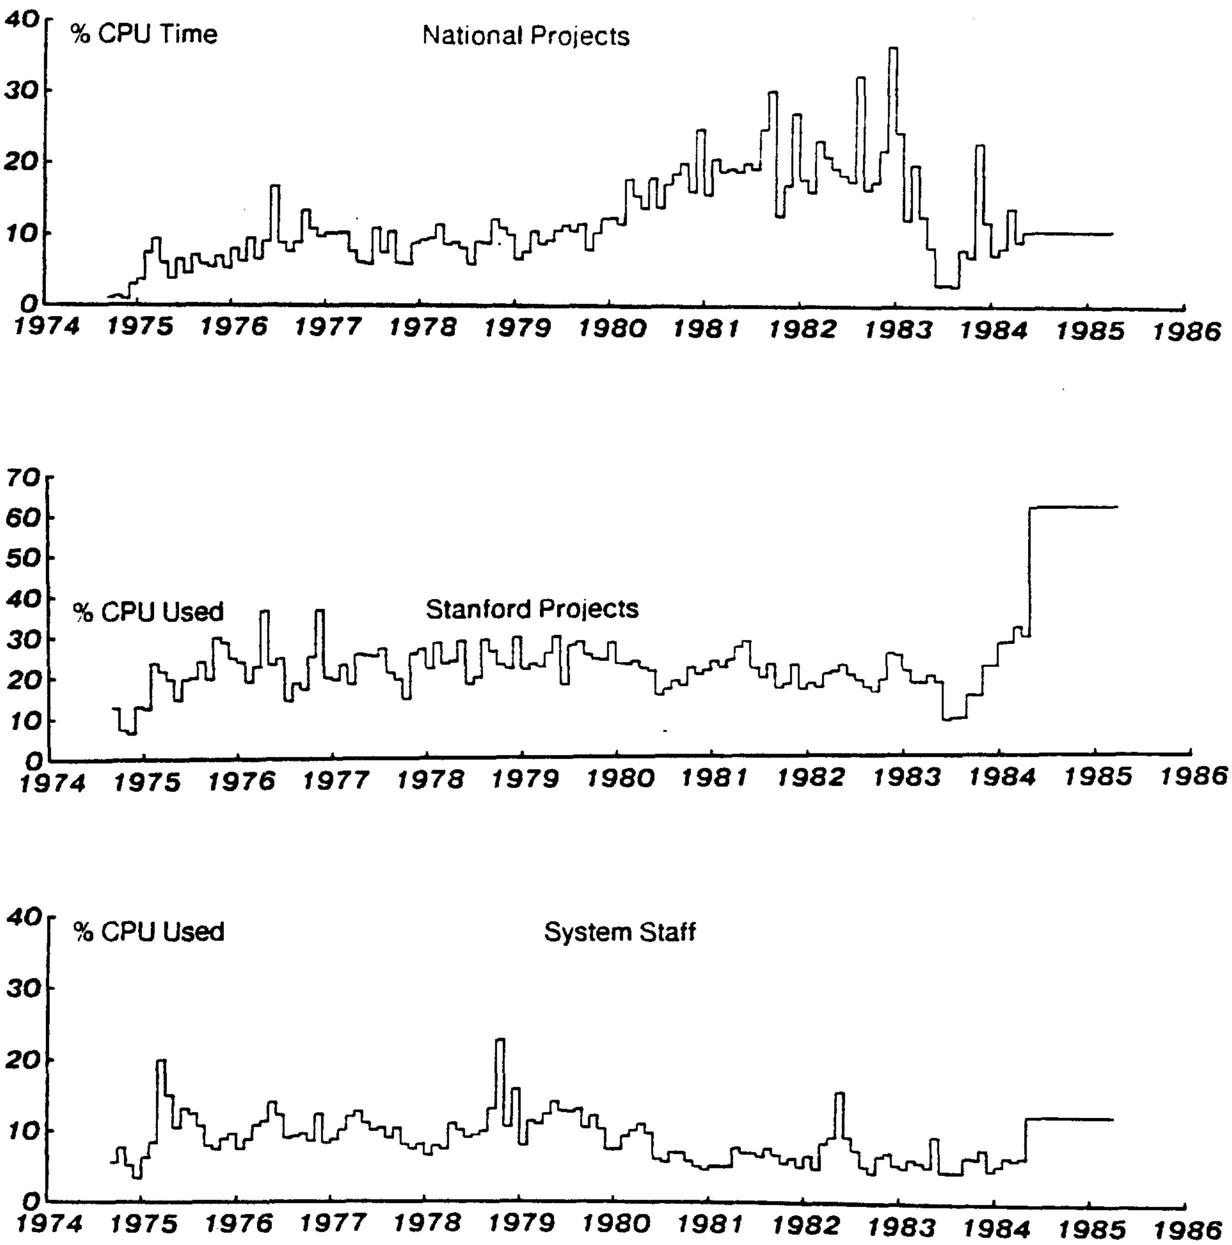


**Figure 15:** Monthly CPU Usage by Community

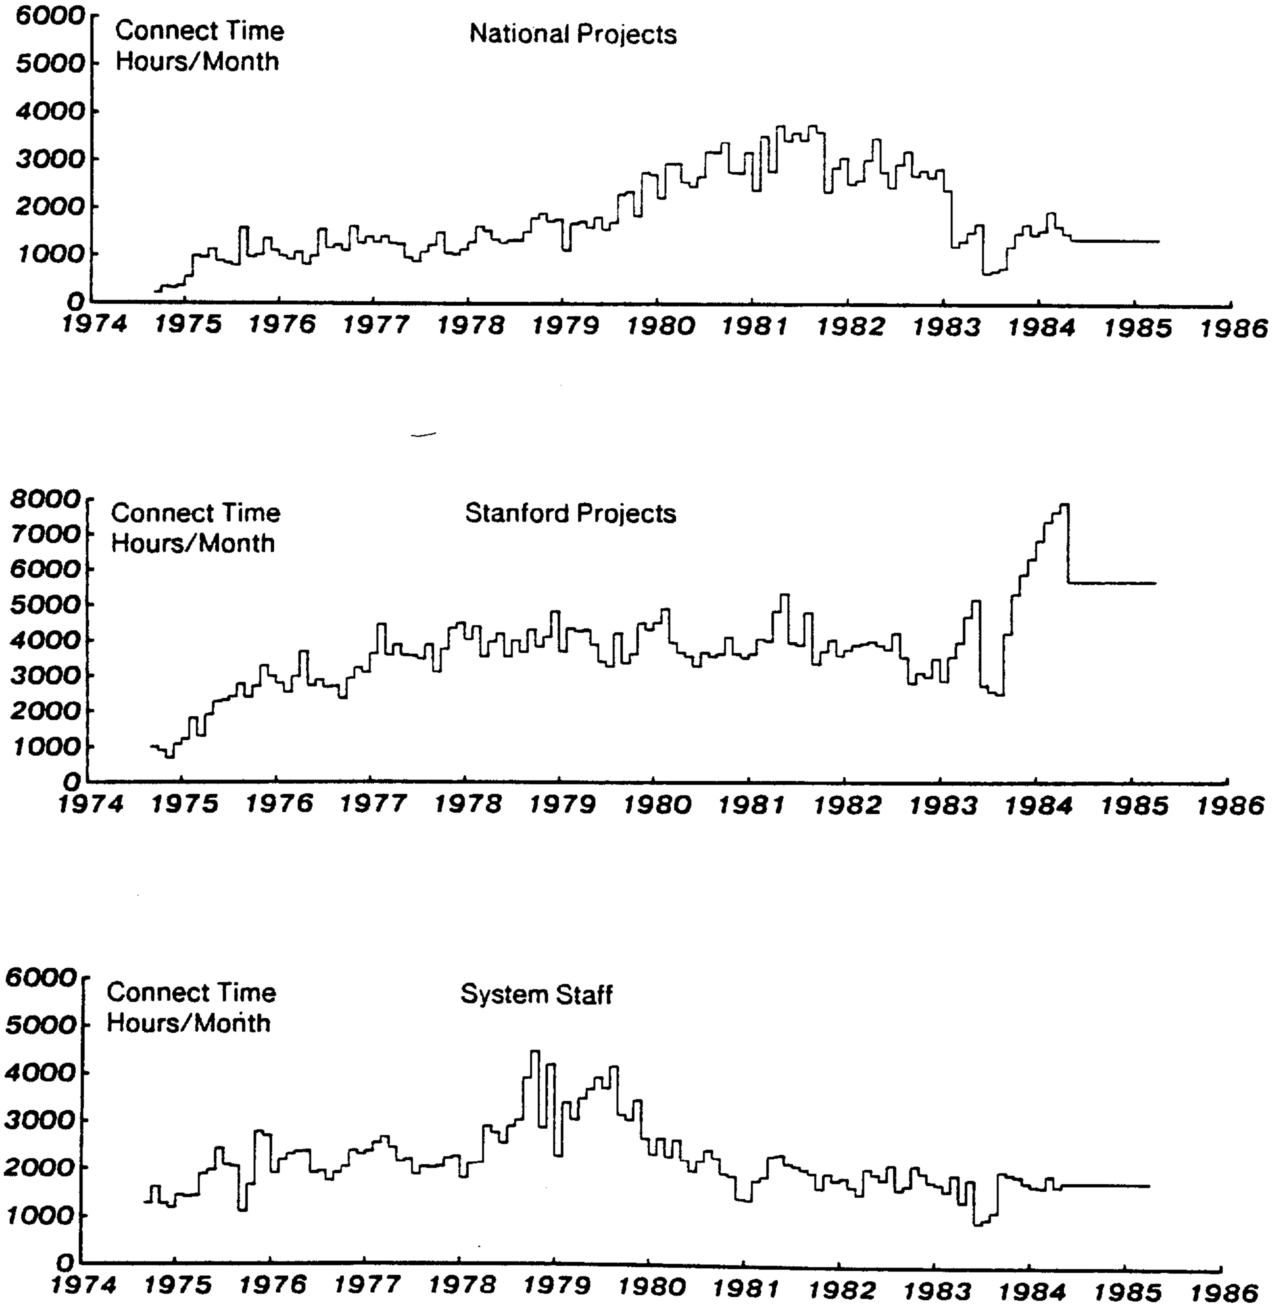


**Figure 16:** Monthly Terminal Connect Time by Community

*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, 1984 and April 30, 1985, on the SUMEX-AIM DECsystem2060 system.

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, 1984 and May, 1985.

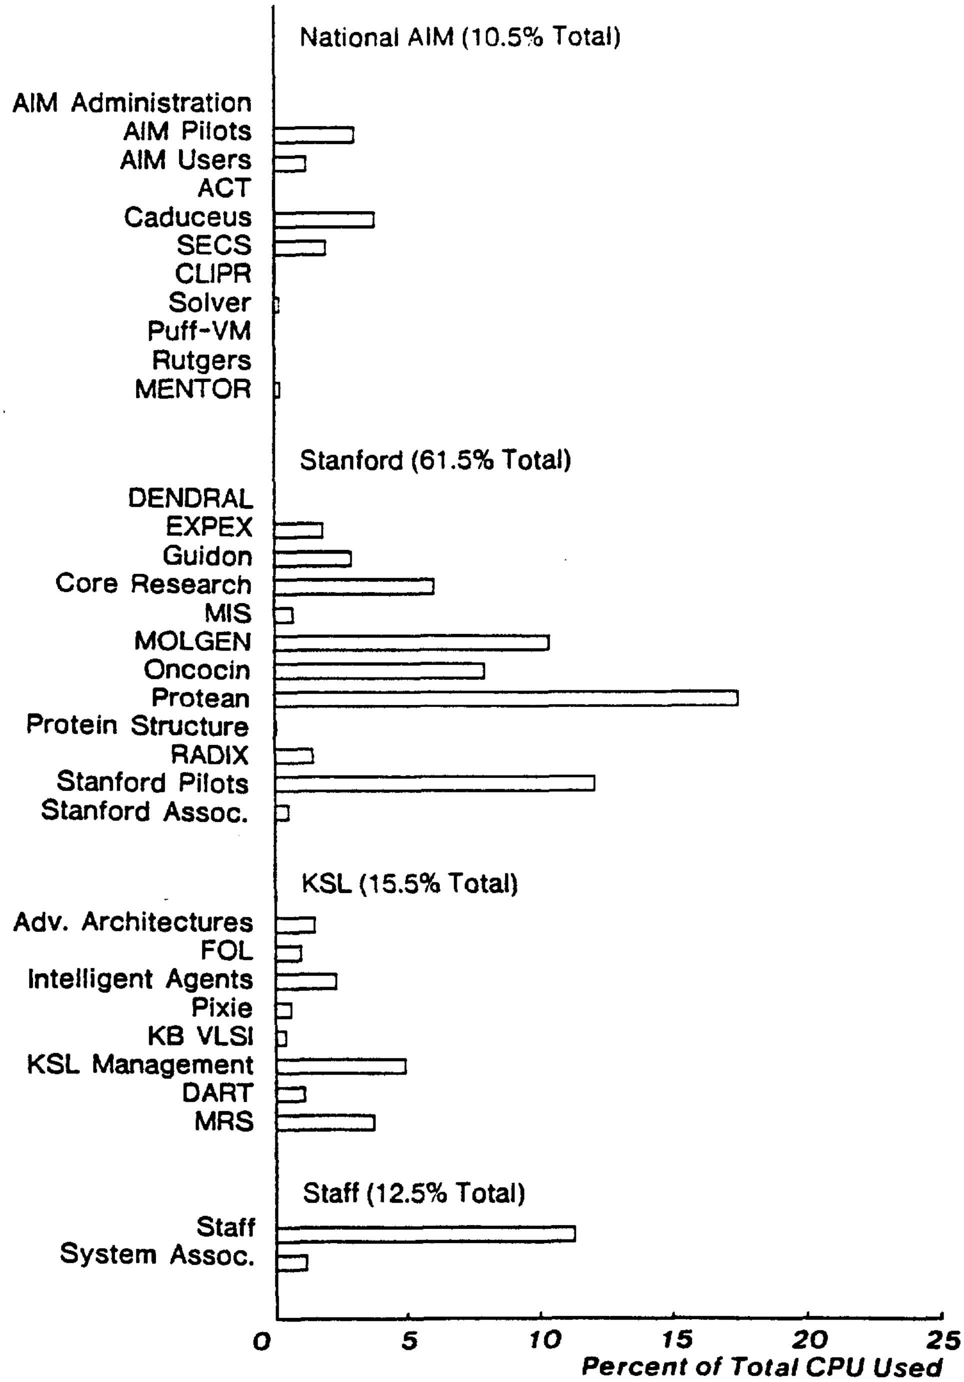


**Figure 17:** Cumulative CPU Usage Histogram by Project and Community

**Resource Use by Individual Project - 5/84 through 4/85**

| *National AIM Community* | CPU (Hours) | Connect (Hours) | File Space (Pages) |
|---|---|---|---|
| 1) CADUCEUS<br>"Clinical Decision Systems<br>Research Resource"<br>Jack D. Myers, M.D.<br>Harry E. Pople, Jr., Ph.D.<br>University of Pittsburgh<br>NIH 5 R24 RR-01101-08<br>7/80-6/85 $1,607,717<br>7/84-6/85 $354,211<br>NIH 5 R01 LM03710-05<br>7/80-6/85 $817,884<br>7/84-6/85 $210,091<br>NIH New Invest 5 R23 LM03889-03<br>Gordon E. Banks, M.D.<br>4/82-3/85 $107,675<br>4/84-3/85 $35,975 | 86.72 | 1809.97 | 8028 |
| 2) CLIPR Project<br>"Hierarchical Models<br>of Human Cognition"<br>Walter Kintsch, Ph.D.<br>Peter G. Polson, Ph.D.<br>University of Colorado<br>NIMH 5 R01 MH-15872-14-16 (Kintsch)<br>7/84-6/87 $145,500<br>7/84-6/85 $40,500<br>NSF (Kintsch)<br>8/83-7/86 $200,000(*)<br>IBM (Polson)<br>David Kieras<br>University of Arizona<br>1/85-12/85 $250,000(*) | 1.14 | 119.94 | 129 |
| 3) SECS Project<br>"Simulation & Evaluation<br>of Chemical Synthesis"<br>W. Todd Wipke, Ph.D.<br>U. California, Santa Cruz<br>NIHEHS ES02845-02<br>4/82-3/85 $257,801<br>4/84-9/85 $89,140<br>Evans & Sutherland Corp.<br>Equipment gift<br>Value $95,000<br>Stauffer Chemical Co.<br>$6,000 | 45.14 | 5542.39 | 12230 |