STANFORD UNIVERSITY
DEPARTMENT OF COMPUTER SCIENCE
PROFILES OF RESEARCH PROJECTS

Founded in 1965, the Department of Computer Science is a center for research and education at the undergraduate and graduate levels. Strong research groups exist in areas of artificial intelligence, robotics, foundations of computer science, scientific computing, and systems. Basic work in computer science is the main research goal of these groups, but there is also a strong emphasis on interdisciplinary research and on applications that stimulate basic research.

Fields in which interdisciplinary work has been undertaken include chemistry, genetics, linguistics, physics, medicine and various areas of engineering, construction, and manufacturing. Close ties are maintained with researchers with computational interests in other university departments. In addition, both faculty and students commonly work with investigators at nearby research or industrial institutions. The main educational goal is to prepare students for research and teaching careers either in universities or in industry.

ARTIFICIAL INTELLIGENCE

Artificial intelligence consists of a number of related research projects with both basic and applied research objectives. Current projects include basic research in artificial intelligence and formal reasoning, expert systems, large knowledge bases, agent-based architectures, image understanding, robotics, machine learning, mathematical theory of computation, program synthesis and verification, natural language understanding, parallel architectures, design/manufacturing, and portable LISP systems.

The Artificial Intelligence faculty perform their research in collaboration with a number of Stanford laboratories and centers. Some of these are the Robotics Laboratory, the Center for the Study of Language and Information (CSLI), the Center for Integrated Systems (CIS), the Stanford Integrated Manufacturing Association (SIMA), Center for Integrated Facility Engineering (CIFE), and the Knowledge Systems Laboratory (KSL).

FACULTY: Thomas O. Binford, Professor (Research), Ph.D. University of Wisconsin, 1965.

AFFILIATIONS: Robotics Laboratory, CIS.

GROUP: Machine Perception

RESEARCH: Medical imaging; image processing; manufacturing: geometric design and tolerancing; precision machining (with Prof. Dan DeBra); AI vision systems; geometric modeling and display; mobile robots

STUDENTS: Leonid Frants, Wally Mann, Thilaka Sumanaweera, Sheng-Jyh Wang

FACULTY: Edward Feigenbaum, Professor of Computer Science, Ph.D. Carnegie-Mellon, 1960, Scientific Director of the Heuristic Programming, and
Richard Fikes, Professor (Research), Ph.D. Carnegie-Mellon, 1968., Co-Scientific Director of the Heuristic Programming Project.

RESEARCH SCIENTISTS: Robert Engelmore, Yumi Iwasaki, Thomas Gruber

GROUP: Knowledge Systems Laboratory/Heuristic Programming Project

RESEARCH: The How Things Work (HTW) Project is:

- building explicit computational models of physical devices, physical phenomena, and engineering techniques, and mechanisms for making use of these models in order to perform such tasks as diagnosis, design, and prediction;
- developing a general framework for modeling physical devices that supports reasoning about their designed structure, intended function, and actual behavior;
- building intelligent assistants for acquiring knowledge about engineered devices to facilitate cooperative design, training, and the construction of knowledge based systems;
- developing techniques for reasoning with both qualitative and quantitative information and for reasoning with multiple models (at different levels of abstraction and from different perspectives).

A closely related research effort on Knowledge Sharing Technology (KST) aims to enable the reuse of knowledge bases and knowledge systems, by:

- developing a standard knowledge interchange language and tools to automate the translation process,
- developing standard domain vocabularies in a portable form, and
- developing tools for defining and combining ontologies.

Basic research in knowledge representation, in order to facilitate knowledge sharing and modeling of physical devices, is a vital component of the HTW and KST efforts.

SENIOR STUDENTS: A. Levy, J. Mohammed, P. Nayak.

FACULTY: Michael R. Genesereth, Associate Professor, Ph.D. Harvard, 1978.

AFFILIATIONS: Director of the Logic Group, Robotics Laboratory, CIFE.

GROUP: The Logic Group

RESEARCH ASSOCIATES: Terril Hurst, Pierre Huyn, Reed Letsinger, Narinder Singh

RESEARCH: The general subject of research is the theory of symbolic systems, i.e. dynamic systems capable of storing and manipulating symbolic information about the world. This includes both object-oriented systems (in which real world objects are represented as structured objects with slots and values) and Fregean systems (in which information is represented in the form of algebraic expressions and logical statements).

Our current emphasis within this general area of research is the study of deliberate systems (sometimes called *agents*). Deliberate systems are symbolic systems with information about themselves and their environments -- systems that, in effect, contain their own documentation. Deliberate systems can tell you what they know and what they do; they can accept new information and new tasks in the midst of operation and regulate their behavior accordingly.

The central (in fact, defining) role of symbolic information in the group's research places strong emphasis on the following areas of basic research.

- Knowledge representation. First order logic and logics that are expressively superior to first order logic (as needed, for example, in the representation of knowledge about knowledge). The use of highly perspicuous languages to express information in these logics, e.g. graphs, charts, tables, etc. Encoding of knowledge about the physical work within these languages, with special attention to electromechanical systems.
- Automated reasoning. This includes work on both deductive reasoning (e.g. resolution) and

nondeductive reasoning (e.e. induction and analogical reasoning). The Group has a long list of publications in the area of reasoning efficiency. There is a special concentration on controlling search in theorem proving and in synthetic reasoning tasks (e.g. runtime planning, hardware design, and software design) and on the reformulation of knowledge bases to obtain efficiency for specific reasoning methods.

- Theory of rational action. Topics in individual rationality include the tradeoffs between computation and action (preprogramming, runtime deliberation, and interleaved planning and execution) and action in the face of limited information and resources. Topics in collective rationality concern principles of cooperation and competition, with and without communication.

In concert with this basic research, substantial effort is devoted to the implementation of results in specific application areas. Due to the theoretical nature of the group's research, this work is not essential to establish the correctness of our results, but the effort pays off in validation of the importance of the results and as a way of discovering new problems on which to work.

In our application, we have a strong emphasis on integrative efforts, i.e. building complete systems rather than just their parts. This helps us to assess the relative importance of various aspects in the design of such systems. At present, we are concentrating on the following applications:

- Mobile Robots. The goal of this project is the design of mobile robots capable of navigating Stanford buildings and the Stanford campus. Map building and planning are key technologies.

- Designworld. Designworld is an automated engineering system for small-scale electromechanical devices. The design for a device is entered into the system via a multi-media design workstation (Helios). To assist in the design effort, the system provides capabilities for simulation, analysis, verification, and debugging. As the design progresses, the system develops plans for manufacture and maintenance of the device and provides feedback to the designer (on cost, reliability, and so forth). When the design is finished, the system configures a robotic assembly cell (Robotworld) to manufacture working units. As they are produced, the units are tested, reworked as necessary, and released. If subsequent problems are encountered, the units can be returned to the robotic cell for diagnosis are repair. The main application of the system is the engineering of products in low volume, e.g. on-off devices, design prototypes, and individualized variations on standard designs. Such products can be produced rapidly (in terms of total time from concept to working unit) and at low cost (since the manufacture and maintenance are automated). (Joint with Jean-Claude Latombe.)

- Autonomous Construction Robots. A project similar to Designworld in scope but with application to the construction of large-scale facilities. The long range goal of this project is the construction of a small building on the Stanford campus. (Joint with Jean-Claude Latombe and Paul Teicholz.)

- Agent-Based Software Engineering. The goal her is the creation of a technology to allow automatic interoperation among heterogeneous, distributed systems. Our approach to solving this problem is based on the notion of deliberate systems. The user starts up programs or plugs machines together; the interlinked systems then exchanged their documentation and use this information to coordinate their efforts. The ever increasing proliferation of networked programs makes this a very important application. We already have an internet with many thousands of interlinked computers; and with the advent of ubiquitous computing (interlinked computers in cars and highways, toasters and microwave ovens), this number should grow astronomically.

- Stanford Theorem Prover. Proof development and checking and automatic proof generation. Theorem proving strategies and systems architecture.

- National Knowledge Service. The goal here is a computer system to provide easy access to technical knowledge. The user of the system should be able to retrieve specific data (e.g.

the performance curves of a semiconductor), definitions (e.g. the axioms defining an algebraic ring), and physical laws (e.g. Maxwell's equations). the user of the system should be able to search for specific information, browse by subject matter, and ask hypothetical questions (e.g. what happens if I put 220 volts across my GE toaster?). Our work on this project is not so much concerned with the problems of scale in building such a system; instead, the emphasis is on representation and reasoning and, to a lesser extent, user interface.

SENIOR STUDENTS: Todd Feldman, Don Geddis, Ofer Matan, Illah Nourbakhsh, Scott Roy, Vishal Sikka, John Woodfill

FACULTY: Barbara Hayes-Roth, Senior Research Scientist, Ph.D. University of Michigan, 1974.

AFFILIATIONS: KSL, CIFE, CIS.

GROUP: Knowledge Systems Laboratory (Director, Prof. Edward Feigenbaum)

RESEARCH PROJECT: Adaptive Intelligent Systems

We are investigating adaptive intelligent systems that perform multiple tasks while interacting with other dynamic entities (e.g., people, processes, other computer agents) in real time. Such systems must perceive, reason, and act continuously over extended periods of time. They must integrate diverse knowledge and reasoning skills. They must allocate their own computational resources effectively to achieve the most important objectives in a timely fashion. Current research topics include: architectural foundations; selective attention; reactive behavior; model-based diagnosis, prediction, and explanation; real-time planning and replanning; learning from experience; and global control of loosely-coupled tasks. We are studying these issues in the context of several intelligent monitoring applications that provide an experimental workbench for exploring, demonstrating, and evaluating theoretical ideas. These are: the Guardian system for monitoring intensive care patients (with Drs. Adam Seiver, David Gaba, Thomas Feeley and Julie Barr of the Medical School); a system for monitoring semiconductor manufacturing equipment (with Profs. Krishna Saraswat, Gene Franklin, and Robert Dutton of Electrical Engineering); and a system for monitoring power plant equipment (with Profs. Ray Levitt and Paul Teicholz of Civil Engineering). Our goal is to develop a generic AI architecture, with associated knowledge representations and reasoning mechanisms, to support the larger class of adaptive intelligent systems. We are also interested in developing reusable knowledge modules for use in new application of the architecture.

PH.D. STUDENTS: Richard Washington, Michael Wolverton, David Ash, Janet Murdock, Vlad Dabija

M.S. STUDENTS: Alex Macalalad, Henny Sipma, Michael Hollander

PROJECT PROGRAMMER: Lee Brownston

OTHER PROJECT MEMBERS: Paul-Andre Tourtier (visiting Ph.D. student), Serdar Uckun and Philippe Morignot (post-doctoral fellows)

FACULTY: Oussama Khatib, Associate Professor of Computer Science and (by courtesy) Mechanical Engineering, Ph.D. Ecole Nationale Superieure de L'Aeronautique et de l'Espace (France) 1980.

AFFILIATION: Robotics Laboratory, SIMA, CIFE.

GROUP: Manipulation Project -- Robotics Laboratory

RESEARCH: Robot control architectures, object-level manipulation, multi-arm cooperation, sensor-based strategies and compliant motion primitives, real-time collision avoidance, robot programming and processing environments, integrated planning and control systems, and design and development of a new generation of force-controlled robot manipulator and mini-manipulator systems. One of the primary objectives of this research is the development of a general framework for task-oriented sensor-based robot control with emphasis on its connections with planning systems.

STUDENTS: Sanford Dickert, Bob Holmberg, Sean Quinlan, David Williams, Ramin Zabih.

FACULTY: Jean-Claude Latombe, Associate Professor and Director, AI and Robotics Division, Docteur-Ingenieur Grenoble, 1972; Docteur d'Etat Grenoble, 1977.

AFFILIATIONS: Director, Robotics Laboratory, CIFE, CIS, SIMA.

GROUP: Robot Reasoning

RESEARCH: Develop autonomous robot systems, which allow the user to specify what he wants done rather than how to do it. Robots include, but are not limited to mobile vehicles, manipulator arms, and combinations of them. Current activities focus on automatic Spatial Reasoning. Our goal is to provide robots with general-purpose capabilities for reasoning about their tasks and their workspace, and for planning and controlling their motions and sensing acts appropriately. Motions include motions of mobile vehicles and manipulator arms, coordinated motions of multiple effectors, and operations of various-end effectors (e.g. dextrous hands). Three applications of the concept of autonomous robot are considered: construction/assembly tasks both on earth and in space, office/clean room automation, integration of design and manufacturing.

SENIOR STUDENTS: David Zhu (CS), Randy Wilson (CS), Shashank Shekhar (ME)

FACULTY: John McCarthy, Charles M. Pigott Professor of Engineering, Professor of Computer Science and (by courtesy) Electrical Engineering, Ph.D. Princeton, 1951.

GROUP: Commonsense Reasoning

RESEARCH: Developing the theory of commonsense knowledge and reasoning, which is concerned with designing formal languages for expressing commonsense knowledge and with representing the process of commonsense reasoning by logical deduction. This theory plays a central part in AI. Its mathematical foundation is formed mainly by the concepts and ideas of logic which were originally proposed for the formalization and study of proofs in mathematics. There are remarkable similarities between commonsense reasoning and mathematical proofs, but there are also important differences. One distinctive feature of commonsense reasoning is the use of default assumptions, when an argument is justified by the absence of certain information among available facts. Default reasoning is non-monotonic, in the sense that extending the given knowledge base may invalidate some of the previously used default assumptions and thus force us to retract some of the conclusions made before. The study of non-monotonic reasoning is a new and intensively developing area of applied logic related to the theory of logic programming and deductive databases.

FACULTY: John McCarthy

RESEARCH ASSOCIATES: Carolyn Talcott and Ian Mason

ASSOCIATED RESEARCHERS: Jussi Ketonen, Richard Weyhrauch, Lou Galbiati

GROUP: Formal Reasoning and Program Proving

RESEARCH: General research interests include semantics and design of programming languages, mathematical theory of computation (proving properties of programs), and mechanization of formal reasoning. The basic theoretical work in programming language theory is the study of mathematical properties of computation mechanisms and the computation structures needed to support them. Of particular interest are languages with function and control abstractions, objects with state, and concurrency. Recent work includes developing inference systems for proving program equivalence in such languages, and specification of classes of programming language algebras that include operational, denotational, and static semantics. The work on formal reasoning includes development of ‘direct logic’, extracting bounds information from mechanized proofs, and analysis of structures and tools needed to support and implement mechanized formal reasoning.

FACULTY: Mark Musen, Assistant Professor of Medicine and (by courtesy) Computer Science, M.D. Brown, 1980; Ph.D. Stanford, 1988.

AFFILIATION: Medical Computer Science Group within the KSL.

PROJECT: PROTEGE

RESEARCH: Automated generation of task-specific knowledge-acquisition tools for expert systems. Languages for modeling problem-solving methods. Architectures for knowledge sharing and reuse. Visual languages for knowledge acquisition. Adaptive user interfaces. Medical expert systems.

SENIOR STUDENTS: Keith Campbell, Amar Das, John Egar, Gretchen Purcell, Yuval Shahar.

FACULTY: H. Penny Nii, Senior Research Scientist, M.S. Stanford, 1973.

AFFILIATION: Heuristic Programming Project, Knowledge Systems Laboratory.

RESEARCHER: Sanjay Bhansali, Research Scientist

STAFF PROGRAMMER: Nelleke Aiello

GROUP: Knowledge Systems Laboratory

RESEARCH: Knowledge-Assisted Software Engineering (KASE) Project

The objective of the research is to build a mixed-initiative design environment for software designers and analysts using knowledge-based techniques. The research focus is on building an intelligent agent that helps software designers and analysts design and redesign relatively large systems, rather than on code generation. Current research tasks include:

- to design and develop a scheme for representing generic architectures for different classes of tasks;
- to discover methods and techniques for customizing a generic architecture into application-specific architectures;
- to develop methods of acquiring domain models using generic architectures;
- to develop a language to express architectural and domain constraints;
- to develop an efficient constraint management system that can detect constraint violations and suggest remedial actions;
- to develop a graphic user interface that uses multi-perspective views of software architectures; and
- to develop the above capabilities within an environment that support incremental, opportunistic design process.

SENIOR STUDENTS: Liam Peyton, Raymonde Guindon

FACULTY: Nils Nilsson, Kumagai Professor of Engineering, Professor of Computer Science, Ph.D. Stanford, 1958.

AFFILIATIONS: CIS, CSLI, CIFE, Robotics Laboratory, Symbolic Systems Program.

GROUP: ’Bots

RESEARCH: We are interested in programming autonomous agents (such as mobile robots) so that they can function in a real, dynamic world---possibly cooperating, communicating, and negotiating with other agents. We seek to integrate learning and planning mechanisms into agent architectures so that these mechanisms can modify hand-coded programs that control the agent’s behavior. We are exploring a behavior formalism based on ‘‘teleo-reactive trees’’ and delayed-reinforcement learning mechanisms for these trees. We are also interested in ‘‘active perception’’ as a means for focusing sensory processing.

SENIOR STUDENTS: Anton Schwartz, Andrew Kosoresow.

FACULTY: Nils Nilsson

SENIOR RESEARCH SCIENTIST: Matt Ginsberg

GROUP: PRINCIPIA

RESEARCH: We work on the application of formal methods to representation and reasoning problems in artificial intelligence. This work typically involves multivalued or probabilistic logics, modal operators, and applications to real-time reasoning and to planning. We also work on the development of new search techniques. Recent results here include the development of iterative broadening and dynamic backtracking, and experimental evaluation of these techniques on a variety of toy and practical problems.

SENIOR STUDENTS: Current Ph.D. students: Adnan Darwiche, Don Geddis, Will Harvey, Ari Jonsson and Scott Roy. Current Masters student: Hernan Epelman-Wang.

FACULTY: Stanley J. Rosenschein, Consulting Professor, Ph.D. Pennsylvania, 1975, Director of Research, Teleos Research; CSLI.

GROUP: Situated Automata

RESEARCH: The situated-automata project is engaged in a long-term program of research aimed at developing theoretical foundations and design methods for sophisticated embedded computer systems. Work to date has concentrated on mathematical theories of information and new programming techniques and languages, as well as experimental robotic applications intended to test the theories and design methods.

A key element in our understanding of embedded systems is the way that information from the environment is picked up by a system through its sensors and converted to a form ultimately usable for controlling action, and for this reason the situated-automata project also includes work on machine perception. This work focuses on computational theories of early vision, using ‘‘sign-correlation’’ techniques for binocular stereo and the measurement of optical flow fields. A prototype hand/eye tracking system has been developed and is being used for experiments in intelligent reactive control. Experimental work on reinforcement learning in embedded systems is also underway. A formal model has been developed, together with criteria for comparing learning algorithms on the basis of correctness, convergence, and computational resources.

FACULTY: Yoav Shoham, Assistant Professor, Ph.D. Yale, 1986.

AFFILIATIONS: Robotics Laboratory, SIMA, CSLI, CIFE, CIS.

GROUP: The Artificial Agents Group, aka Nobotics, located at the Robotics Laboratory

RESEARCH: Our group is interested in a theory of artificial agents and their environment. Among the themes investigated:

- What is a useful and precise way of ascribing mental qualities to machines?

- What are the social rules of artificial agent societies?

The research style emphasizes both mathematical rigor and experimentation. The theoretical tools include temporal, epistemic and nonmonotonic logics, as well as continuous mathematics. Applications include robotics, civil engineering and communication systems.

STUDENTS: Ronen Brafman, Nita Goyal, Ronny Kohavi, Eyal Mozes, Becky Thomas, Alvaro del Val

FACULTY: Jay M. Tenenbaum, Consulting Professor, Ph.D. Stanford, 1971.

AFFILIATIONS: President, Enterprise Integration Technologies Corp. (EIT); Co-Director, MCC Enterprise Integration Initiative; CIS; SIMA.

IN AFFILIATION WITH: Jeff Pan, Jay Glicksman, and CIS Faculty and Staff.

GROUP: Manufacturing Knowledge Systems

RESEARCH: We are exploring the application of knowledge systems in complex, real-world environments such as engineering and manufacturing. Current research is focusing on:

- Virtual Factories: The use of computer simulation to eliminate the time and cost of physically prototyping both products and processes. Through computational prototyping, one may (a) rapidly analyze and control dynamic factory performance, (b) predict the results of product and process changes prior to implementation, enabling continuous improvement and faster development, and (c) provide accurate and efficient models for process diagnosis and control.

- Autonomous Agent Architectures: Development of distributed, modular, scalable architectures for integrating people, systems and activities. Such architectures consist of agents representing people, software/tools, equipment, databases, etc. The agents communicate using a common business language. They register their capabilities and interests with an enterprise-wide information service, which is then responsible for routing messages to the appropriate parties. Systems structured in this way are inherently reliable and flexible. They also suggest new approaches to system integration. For example, rather than viewing software as a ‘‘package’’ to buy, install, and maintain, software can be viewed as a service, available as needed over a network. (interest in distributed systems)

- Intelligent Agents: Autonomous agents that are programmed (by end-users) to watch over some element of production, responding to events with actions such as shutting down a machine, starting up a program, or sending a message to another agent. Agents are self scheduling, yet cooperate to achieve the common goal. Ultimately, thousands can be dispatched to watch over every aspect of production, ensuring that information flows rapidly and that nothing slips through cracks. Such radically distributed models of factory control have interesting analogies in hopes of developing factories that are highly reliable and able to adapt and learn continuously. (interest in distributed AI, distributed systems, user interfaces, learning).

SENIOR STUDENTS: John Mohammed and Janet Murdock

FACULTY: Jay M. Tenenbaum, Mark Cutkowky, Assistant Professor, ME

GROUP: SHARE

RESEARCH: SHARE is a distributed environment for concurrent engineering, where human and computational agents cooperate by exchanging information and services. SHARE is being used to explore a variety of interesting computer science issues, such as:

- knowledge sharing and reuse (including knowledge representation standards) : building large, sharable engineering knowledge bases that include both formal representations of knowledge and unstructured representations such as hypermedia; overcoming the single-student barrier in AI systems, where a student typically spends several years developing a knowledge base that is useful only for one application and is so baroque that only that student can maintain it.

- electronic medium: using multimedia ‘‘electronic engineering notebooks’’ to access shared knowledge and facilitate cooperative work.
- distributed problem solving: extending classical planning methodologies to domains where the planner must cooperate with other planners and specialized domain experts--human and computerized.
- incremental re-design: maintaining dependencies linking design requirements to artifacts through decisions. Each decision is justified in terms of requirements, other decisions, engineering knowledge, assumptions, simulation, test results, and so forth. The problem is rationalizing design decisions step by step as a design evolves and understanding how to fix the design when any of the factors it depends on change -- and doing all this in a distributed, man-machine environment where the design representation is maintained by multiple agents, each concerned with a particular subsytem, engineering discipline, or lifecycle stage.
- using computer vision techniques to create CAD and process models from physical artifacts and to recognize features in CAD models.
- monitoring the actual execution of process plans on machine tools using a variety of sensors (e.g., vibration, acoustic emission, and infrared imagery) and a combination of neural networks and expert systems to interpret the resulting data.

SENIOR CS STUDENTS: Randy Wilson, Andrew Philpot.

FACULTY: Jay M. Tenenbaum

GROUP: Electronic Commerce

RESEARCH: (New Project) Create the foundation (technology and policies) for a national information infrastructure, supporting manufacturing, design, engineering, and analysis. Participate in EINet, a prototype currently under development by MCC and five other major manufacturing consortia (NCMS, Sematech, ITI, CAMI and Iacocca Institute). Ultimately, EINet will encompass thousands of design, engineering, prototyping, and acquisition services -- all accessible through the Internet, and serving to seamlessly link manufacturers, customers, suppliers, subcontractors, development partners and distributors. To support this vision of electronic commerce, it will be necessary to enhance the Internet along several dimensions. These include security, ease of use (e.g., directory and navigation aids), functionality (e.g., multimedia, standards, translation), and remittance (pay for use of data and services). Numerous policy issues relating to commercial use (e.g., advertising, scaling of traffic, standard trade agreements) must also be addressed. (ideal for CS students who also have an interest in business).

FACULTY: Terry Winograd, Professor, Ph.D. MIT, 1970.

PROJECT: People, Computers and Design project at the Center for the Study of Language and Information (CSLI)

RESEARCH: Design of computer systems for cooperative work. Professor Winograd’s focus is on developing the theoretical and practical background needed to incorporate human contextual elements into the design and analysis of computer systems. His primary research emphasis is on the development of a ‘‘language/action perspective’’ in which current and potential software and hardware devices are analyzed and designed in the context of their embedding in work and communication. The language/action perspective grew out of his earlier work in artificial intelligence, but it shifts the focus of attention away from the mental and the individual, to the social activity by which people generate the space of cooperative actions in which they work---and to the technology that is the medium for those actions. The People, Computers, and Design project is funded by an NSF grant to develop a series of innovative courses in human-computer interaction, in which students from a variety of disciplines learn through theoretical discussion and group projects to design computing systems from the perspective of how they will be interpreted and used.

Scientific Computing

The research in scientific computing involves two closely related aspects: development of mathematically based theory to solve particular problems, and implementation of appropriate computer algorithms. Particular emphasis is placed on numerical accuracy of a computation; additional considerations are made of algorithm design, computational efficiency, data structures and parallel procedures. Seminars are held almost every week of the year with members of the Stanford community and industrial researchers. Close cooperation and collaboration with government and industrial laboratories is maintained, and there is an active visitors program.

Research is done in collaboration with the Scientific Computing and Computational Mathematics Program.

FACULTY: Gene H. Golub, Fletcher Jones Professor of Computer Science and (by courtesy) Electrical Engineering, and Director, Scientific Computing Division, Ph.D. University of Illinois, 1959.

Affiliation: Chairman, Scientific Computing and Computational Mathematics (SCCM); CIS.

RESEARCH: My current research is mainly concerned with the design and analysis of algorithms arising in linear algebra; my book (with Charles Van Loan) reflects my interests in that area. Recently, I have been concerned with matrix methods in signal processing, updating/downdating of orthogonal polynomials and the solution of large sparse systems of equations, mainly arising in the discretization of pde's. Each of these problems is motivated by applications. I have a peripheral interest in parallel algorithms and my recent work describes the implementation of a variety of algorithms.

FACULTY: **Joseph Oliger**, Professor, Ph.D. Uppsala, 1973.

Affiliation: Director, Center for Large Scale Scientific Computation (CLaSSiC Project); SCCM.

CO PRINCIPAL INVESTIGATORS AND ASSOCIATED FACULTY: J. Ferziger (ME); R. Street (CE); J. Keller (Math); G. Homsy (ChemE) T. Hughes (ME)

GROUP NAME: CLaSSiCS, Center for Large Scale Scientific Computing at Stanford.

RESEARCH: This is an interdisciplinary project focused upon problems associated with the computer simulation or analysis of physical processes. The system aspect of interaction of problem formulation (mathematical), algorithm formulation, language and system implementation, and hardware design is central to this project. Current projects are under way in each of these areas and their interaction. The identification and implementation of parallel processes is a current focus. Other projects involve the development of programming environments for scientific computations.

SENIOR STUDENTS: Steve Suhr, Amala Krishna.

FACULTY: Andrew M. Stuart, Assistant Professor of Computer Science and Mechanical Engineering, D. Phil Oxford, 1986.

AFFILIATIONS: Numerical Analysis. Joint appointment with Mechanical Engineering, Division of Applied Mechanics.

RESEARCH: Numerical Analysis of Dynamical Systems; Numerical Analysis of Dissipative Partial Differential Equations; Applied and Numerical Analysis of Mathematical Models of Dissipative Phenomena; Numerical Methods for various model problems of interest to NASA.

STUDENTS: Tony Humphries, Gabriel Lord, Jeremy Smith, Tianbing Wang.

Systems

The systems area encompasses both experimental and theoretical work growing out of topics in operating systems and compilers, computer communication and networks, architecture, programming languages and environments, distributed systems, VLSI design and fabrication, graphics, reliability and fault tolerance, system specification and verification, and user interfaces. A large concentration of systems research takes place within the Computer Systems Laboratory, a joint laboratory of the Departments of Computer Science and Electrical Engineering.

FACULTY: Charles Bigelow, Associate Professor of Art and Computer Science (Teaching), B.A. Reed, 1967, C.A.S. Harvard University, 1991.

RESEARCH: Digital typographic designs. My interests include the design of fonts, symbols, and icons for user interfaces and for electronic publishing, and the design of electronic and printed documents.

FACULTY: David Cheriton, Associate Professor, Ph.D. University of Waterloo, 1978.

AFFILIATIONS: CIS, CSL.

PROJECT: Distributed Systems

RESEARCH:

- Parallel architecture, systems and programming, including hardware development of a multiprocessor machine, parallel operating system, parallel language and compiler, plus various parallel and distributed applications.
- High-performance Communication: very fast wide-area networks, network interfaces, multiprocessor gateways, multicast, distributed systems architectures and protocols.
- High-performance cluster operating systems, supporting distributed execution, process migration, scheduling, atomic transactions and distributed shared memory.

SENIOR STUDENTS: Kieran Harty, Bob Kutter

FACULTY: Giovanni De Micheli, Associate Professor, Electrical Engineering and (by courtesy) Computer Science, Ph.D. University of California, Berkeley, 1983.

Affiliations: CIS, CSL.

PROJECT: SYNTHESIS OF DIGITAL CIRCUITS AND SYSTEMS

RESEARCH: Professor Giovanni De Micheli is interested in Computer-Aided Design methods and tools for VLSI circuits. At present he is interested in automated synthesis systems, with particular emphasis on optimization techniques, and in circuit verification methods. His research is centered on the development of algorithms, the study of their properties as well as their implementation. Recently he engaged research in system-level design automation, including hardware-software codesign techniques.

FACULTY: David Dill, Assistant Professor, Computer Science and (by courtesy) Electrical Engineering, Ph.D. Carnegie Mellon, 1987.

AFFILIATIONS: CIS, CSL.

PROJECT: Finite-state Concurrent Systems

RESEARCH: Specification and verification of finite-state concurrent systems. Specifically:

- verification and synthesis of real-time systems.
- protocol specification and verification
- hardware verification
- automata theory and temporal logic

SENIOR STUDENTS: Steve Nowick, Howard Wong-Toi, Elizabeth Wolf, Han Yang, Alan Hu, Seungjoon Park, Norris Ip

FACULTY: David Dill

PROJECT: Asynchronous Circuits (joint project with **Prof. Teresa Meng**)

RESEARCH: Explorations in asynchronous circuit design, including:

- automatic synthesis from high-level specifications
- safe design methods
- alternative design styles
- fault modeling and testability
- automatic verification
- performance analysis

SENIOR STUDENTS: Steve Nowick, Peter Beerel, Ken Yun, Chris Myers

FACULTY: Michael J. Flynn , Professor of Electrical Engineering, Ph.D. Purdue University, 1961.

PROJECT: CASIS--Center for Aeronautics and Space Information Sciences

RESEARCH: CASIS is a large interdisciplinary multi-faculty program to support evolving information technology for application in NASA's space program. Currently CASIS is organized into four major thrust areas: computer architecture and compilers to support high performance architectures, neural networks and control, high-speed data networks, and data presentation technology, including graphics and database management.

OTHER FACULTY INVOLVED: Profs. John Cioffi, Hector Garcia-Molina, Bert Hesselink, Tom Kailath, Monica Lam, Marc Levoy, Allen Peterson, Fouad Tobagi, Gio Wiederhold.

SENIOR STUDENTS AND ASSOCIATES: Qian Lin, Ramin Samadani, Jim Burr

FACULTY: Michael J. Flynn

PROJECT: Studies in Concurrent Architectures

RESEARCH: To develop tools to predict the performance of various processor combinations using actual programs as a workload. These tools have lead earlier to the development of the Architect's Workbench (the AWB) which has been widely distributed and used for advanced microprocessor floorplanning.

A shared memory version of the AWB, called the PAWB, is being developed for study and use in the design of shared memory multiprocessors.

Another aspect of this program is studies in massively parallel or fine-grained parallel processing. In this area we have been particularly interested in particle or cellular solutions to problems in high-turbulence gas dynamics as an alternative to the numerical solution of the Navier-Stokes equations.

Senior students: Fung Lee, Andrew Zimmerman, Bill Lynch.

FACULTY: Michael J. Flynn

OTHER FACULTY: Profs. Giovanni De Micheli, Bob Dutton, Fabian Pease, Bruce Wooley

PROJECT: SNAP--Subnanosecond Arithmetic Processor

RESEARCH: The SNAP program was developed to create a basic technology for high-speed arithmetic which would provide significant performance advantages in the area of floating point arithmetic and arithmetic co-processors. The design goals of SNAP pushed the state of the art in both chip and system technologies. It consists of multiple chips on an active substrate (wafer). Each chip represents a functional unit of a high-speed arithmetic co-processor. The functional units originally identified include (together with their target execution delays): floating-point adder, 2-3 ns, floating-point multiplier, 5 ns, floating-point divide, 10-15 ns, register file, 1 ns, control unit, and integer unit, 1-2 ns.

Finally, we have an objective of a one-ns repetition rate for communication between chips and for pipeline rates across chips. During the past year we have completed a unique test chip based upon the concept of wave pipelining, which uses delays intrinsic to the circuits themselves as storage elements in a high-speed pipelined functional implementation. Wave pipelining is a target technology that we believe will allow us to achieve extremely high pipelining rates with functional units. Simulations and measurements show that two or more pipelined waves are possible, thus doubling the normal clock frequency without additional registers.

Senior Students: John Beale, Gary Bewick, C. Chao, Ben Chui, Nhon Quach, Eric Schwarz, Paul Song

FACULTY: Hector Garcia-Molina, Professor, Ph.D. Stanford, 1979.

Affiliations: CIS, CSL.

SENIOR RESEARCH SCIENTIST: Arthur Keller, Ph.D. Stanford, 1985.

Affiliations: KBMS, CSL, CIFE, CIS.

GROUP: Transaction Processing

RESEARCH: Current algorithms for distributed computing and distributed data management tend to assume a relatively small number of tightly coordinated computers working in synchrony. These algorithms (e.g., two phase commit, multicast, load balancing, remote data access, remote procedure call) work well in some environments. However, they can lead to significant delays, low throughput, or simply lack of service when computers and networks fail, when computers become disconnected, or when the number of system components grows. For example, a computer holding a copy on an object X may be unable to access it because some other, unreachable node might be updating X.

Our goal is to develop new distributed object management and transaction processing algorithms that are "elastic:" (1) Algorithms that permit data values to diverge in a controlled way when necessary. For example, instead of telling an application that the value of X is unavailable, they may report that it is unknown with certainty, but that it is guaranteed to be in a particular range of values. (2) Defensive algorithms that do not put blind trust in the decisions of other participants. For example, if a remote node asks the local node to lock certain objects, these locks will not be given unconditionally and for an indefinite period, as in conventional distributed two phase locking. (3) Opportunistic algorithms that exploit whatever communications bandwidth is available. If there is bandwidth, tight consistency is provided; if not, operation is more decoupled.

FACULTY: Anoop Gupta, Assistant Professor, Ph.D. Carnegie-Mellon, 1985.

AFFILIATIONS: CIS, CSL.

GROUP: DASH (Directory Architecture for Shared Memory)

OTHER FACULTY: **Prof. John Hennessy**

RESEARCH: Design of general-purpose scalable parallel computer architectures -- mechanisms for cache consistency; ways to reduce or tolerate memory latency; hardware/software support for efficient synchronization of processes; design of interconnection networks for multiprocessors; understanding the role of locality. To test our ideas, we are currently building a scalable directory-based shared-memory multiprocessor using very high performance individual nodes. In the software area, working on multiprocessor scheduling and resource allocation issues and on the design of a concurrent object-oriented programming language. In the applications area, studying parallelism in VLSI CAD and scientific domains. The applications work is also used for evaluating our hardware and software efforts.

SENIOR STUDENTS: Wolf-Dietrich Weber, Ed Rothberg, Andy Tucker.

FACULTY: John Hennessy , Willard R. and Inez Kerr Bell Professor of Engineering, Professor of Computer Science and Electrical Engineering, and Director, Systems Division, Ph.D. State University of New York at Stony Brook, 1977.

AFFILIATIONS: Director of CSL, CIS

PROJECT: Multiprocessor Architecture

RESEARCH: The objective of this research is to explore the design and implementation of large-scale, general-purpose multiprocessors. The end-goal will be to build a large-scale multiprocessor as a follow-on to the current DASH project.

FACULTY: John Hennessy

PROJECT: Parallel processing compiler and language technology

RESEARCH: The goals of this project are to explore programming language design and implementation focusing on programming and efficiency issues for large-scale multiprocessors.

FACULTY: Mark Horowitz, Associate Professor of Electrical Engineering and (by courtesy) Computer Science, Ph.D. Stanford, 1984.

AFFILIATIONS: CIS, CSL.

PROJECT: Multiprocessor Architecture

RESEARCH: The objective of this research is to explore the design of large-scale, general-purpose, shared-memory multiprocessors. This is a large project involving a number of faculty members. My research interest in this project is looking at how implementation issues affect the architecture and design of the multiprocessor. We have been looking at how to organize the design of the directory controller that maintains cache consistency, and how to use high-bandwidth I/O techniques to build a high-speed interconnection network

FACULTY: Mark Horowitz

PROJECT: Low Power Integrated Circuit / Processor Design

RESEARCH: The objective of this research is to explore methods of reducing the power required by high-performance processors and signal processors. Initial work will focus on getting a better understanding of where power is dissipated in current ICs and to explore the speed / power tradeoffs possible in the microarchitecture of a machine. In this research we assume that the power supply voltage used for the chip is not fixed, and can be adjusted to provide the right speed/power tradeoff. Assuming different chips can use different supply voltages leads to a number of circuit challenges (supply generation, and communication between the chips for example) that will be addressed in this work. This work is also investigating the feasibility of placing a tracking clock generator on the chip, to allow the chip to be run at the silicon limit, rather than the worst-case spec limit. Using the tracking clock generator both provides a higher peak performance and a better speed/power product.

FACULTY: Mark Horowitz

PROJECT: Superconducting Memory Design

RESEARCH: The objective of this research is to explore the design of memory devices for superconducting electronics. There has been some recent work in superconducting electronics that has renewed interest in this technology. One of the principle problems with this technology is the difficulty in making dense/fast memories. This work will investigate superconducting memory design to understand what is difficult with conventional architectures, and see if some novel structures (like the kinetic inductance cell) improve the situation. In addition we plan on looking at hybrid semi/super technologies to see if the resulting hybrid system has advantages of a purely semiconductor or superconductor approach

FACULTY: Monica Lam, Assistant Professor of Computer Science, Ph.D. Carnegie Mellon University, 1987.

AFFILIATIONS*:* CSL, CIS.

GROUP: Compiling for Parallel Machines

OTHER FACULTY: Prof. John Hennessy

RESEARCH: To build a compiler that automatically parallelizes sequential code to run on parallel machines. The target architectures include superscalar machines, shared memory multiprocessors and distributed memory machines. We do experimental research; the research consists of the entire cycle of designing, implementing and evaluating new concepts. We also use compiler technology to evaluate parallel computer architectures.

STUDENTS: Saman Amarasinghe, Jennifer Anderson, Rob French, Amy Lim, Dror Maydan, Todd Mowry, Brian Murphy, Karen Pieper, Mike Smith, Steve Tjiang, Bob Wilson.

FACULTY: Monica Lam

GROUP: JADE -- Language for Parallel Machines

RESEARCH: To develop a programming language for writing parallel applications. We have been exploring the concept of expressing parallelism in a data-oriented fashion. Prototypes have been built for both shared and distributed memory machines.

STUDENTS: Martin Rinard, Daniel Scales

FACULTY: Marc Levoy, Assistant Professor of Computer Science and Electrical Engineering, Ph.D. University of North Carolina at Chapel Hill, 1989.

AFFILIATIONS: CSL, CIS.

GROUP: Computer Graphics

RESEARCH: Computer graphics, scientific visualization, and interactive techniques. Professor Levoy's interests within computer graphics include shading models, visibility algorithms, texture generation and display, strategies for anti-aliasing and photorealism, and other aspects of image synthesis. His recent research has focused on algorithms and architectures for visualizing multidimensional sampled data. He is also interested in computer animation, high-performance graphics architectures, parallel algorithms as they apply to graphics, and exotic user interface technologies such as eye tracking, head tracking and head-mounted displays.

SENIOR STUDENTS: Brian Curless, Alice Yu

FACULTY: David Luckham, Professor (Research) of Electrical Engineering, Ph.D. MIT, 1963.

AFFILIATIONS: CIS, CSL.

GROUP: Program Analysis and Verification Group

RESEARCH ASSOCIATES: Doug Bryan, Sriram Sankar, Walter Mann.

GROUP MEMBERS: Prof. Sigurd Meldal (from Univ. Bergen), Dr. Frank Belz of TRW, and other visitors from afar and from industry.

RESEARCH: Design of new computer languages for specification and prototyping of large distributed time-critical systems that include both hardware and software components. Also the group is designing and implementing various programming environment tools to support applications of specifications in developing systems, particularly:

- runtime checking of consistency between behavior and specifications,
- formal analysis of specifications,
- automated checking of hardware simulations, and
- construction of self-checking systems on multi-processors.

The group is also collaborating with other Stanford CSD faculty members, and research groups at other universities and in industry. Research Topics:

- design of a new prototyping language for large distributed realtime systems,
- specification languages for concurrency and timing,
- formal proof rules for verifying behavioral and timing properties of concurrent systems,
- applications of multi-processor systems
- theory and formalisms for software systems architecture.

SENIOR STUDENTS: John Kenney, Neel Madhav, Larry Augustin.

FACULTY: Edward J. McCluskey Professor of Computer Science and Electrical Engineering, Sc.D. MIT, 1956.

AFFILIATIONS: Director, Center for Reliable Computing, CIS, CSL.

GROUP: Center for Reliable Computing

RESEARCH:

Broad objective is design and production of high quality digital systems. Specific current projects:

- Design and implementation of a high-level synthesis system that will accept a functional-level system specification written in enhanced VHDL. This specification includes functional, performance, reliability, testability, fault tolerance and maintainability attributes. An appropriate implementation will be produced automatically.
- Studies of the characteristics (failure modes) of actual hardware failures (reliability and

transient failures in CMOS and BiCMOS) and procedures for detecting these failures. We are especially interested in failures that are not detected as single-stuck faults.

- Development of new, novel design-for-testability techniques for detecting failures that are not detected by standard single-stuck fault test methods.

SENIOR STUDENTS: LaNae Avra, Piero Franco, Hong Hao

FACULTY: Teresa Meng, Assistant Professor of Electrical Engineering and (by courtesy), Computer Science, Ph.D. U.C. Berkeley, 1988

AFFILIATIONS: CSL, CIS, ISL.

PROJECT: Asynchronous Design for Parallel Processing Architectures

RESEARCH: The objective of this research is to provide an interconnect design synthesis methodology which facilitates a modular design approach without compromising the global performance. The goal is to support fast-prototyping of parallel processing systems with a minimal amount of design effort by providing an asynchronous interconnect strategy which allows the overall system performance to be improved by individual optimization of computer modules. This allows the equivalent of the ‘‘object oriented’’ software methodology to be applied to hardware design. The main tasks of this effort will be the development of the theory for asynchronous circuit synthesis from a high-level specification, with emphasis on timing and performance among parallel computing hardware objects, and the application of this design methodology to physical implementations of multi-processing systems.

FACULTY: Teresa Meng

PROJECT: Portable Video

RESEARCH: The proposed research aims to make portable digital video possible through development in both the algorithms and a hardware prototype. We propose to develop compression algorithms capable of maintaining consistent visual quality for image and video signals transmitted in a wireless communication environment. These algorithms will therefore display reliability and recovery in situations of severe packet loss and can sustain effective video service in a network in which the available bandwidth is constantly changing. Further, because of portability constraints, they will be optimized for efficient hardware implementations and for low power requirements.

The chosen hardware architecture for the optimized algorithms will be implemented and the effectiveness measured. A chip set will make use of advanced packaging techniques to implement the critical components in the optimized compression algorithm will be designed and tested. The compression module will be incorporated into a real-time video system for the demonstration of visual quality under the effects of packet loss and variable channel capacity.

FACULTY: Teresa Meng

PROJECT: Asynchronous Circuits (joint with **Prof. David Dill**)

FACULTY: Kunle Olukotun, Assistant Professor of Electrical engineering and (by courtesy) Computer Science, PH.D. The University of Michigan, 1991

AFFILIATIONS: CSL, CIS

PROJECT: Automating Architectural Design

RESEARCH: The goal of this research is to significantly reduce the time and effort that it takes to design and evaluate new computer implementations. To achieve this goal we will develop the capability to explore more of the computer design space in much less time than is possible today. The approach taken will be to combine sophisticated computer performance evaluation with logic synthesis and layout synthesis to create an environment for automatically estimating the cost and performance of candidate computer designs. The key challenge of this research will be to formalize the high level architectural design process so that it can be integrated with existing tools for low level design. Much of the design exploration at the architectural level will be done by manipulating a hardware description language (HDL) representation of the architecture. To make sophisticated performance evaluation practical, techniques that will speedup the simulation of HDLs by two orders of magnitude must also be developed.

FACULTY: Mendel Rosenblum, Assistant Professor of Computer Science, Ph.D. U.C. Berkeley, 1992

AFFILIATIONS: CSL

PROJECT: High Performance Storage Systems

RESEARCH: This research project is focused on the hardware and system software used to manage permanent storage in computer systems. This includes the techniques used to allocate and transfer data from storage devices, the design of the I/O controllers and channels used to transfer data, caching techniques for improving performance, and file system interfaces used by application programs. The project will use a combination of analytic analysis, simulations and prototype implementations to investigate ideas.

FACULTY: Fouad Tobagi, Professor of Electrical Engineering and (by courtesy) Computer Science, Ph.D. UCLA, 1974.

AFFILIATIONS: CSL, CIS, Stanford Center for Telecommunications.

PROJECT: Design and Analysis of High-Speed Integrated Services Communications Networks

Spurred by the high transmission capacity of fiber optics technology, applications requiring higher bandwidth than possible in present networks have emerged. Because of the high data rates involved as well as the diversity of the applications characteristics, the creation of future high-speed integrated service networks is a challenging task. We have been engaged in research in several interrelated areas pertaining to the design of such networks and to the characterization of the applications: (i) architecture design and hardware implementation of large-size fast packet switches, with emphasis on added functionalities (e.g., support of bursty traffic); (ii) performance evaluation of high-speed networks in the presence of bursty traffic, using measures of quality relevant to the applications (e.g., signal-to-noise ratio for still image transmission or for video). (iii) design of dynamically reconfigurable optical networks; and (iv) study of distributed computation and communications on a large scale.

STUDENTS: F. Chiussi, W. Wang, B. Paul-Dubois-Taine, C. Noronha, I. Dalgic.

Theory

Faculty in the theory division seek greater understanding of fundamental computational techniques and their inherent limitations. Research includes the development of new sequential and parallel algorithms, computational problems in databases, computational geometry, design and analysis of programs and programming languages, and supporting studies in combinatorial, logical, and algebraic mathematics.

FACULTY: Robert W. Floyd, Professor, Computer Science, B.A., 1955; B.S. University of Chicago, 1958.

RESEARCH: Eclectic. Recent subjects include analysis of algorithms (Quicksort, coalesced hashing), digital halftones, greedy table search, syntax error recovery, fixpoint theory, random sampling, Fibonacci arithmetic, Chaitin-Kolmogorov complexity. I have directed several distinguished dissertations, including those of Manna, Tarjan, Rivest, Earley, and Deutsch; subjects include correctness and termination of algorithms, graph planarity testing, table searching with partial information, parsing of context-free languages, and automated deduction.

FACULTY: Andrew Goldberg, Assistant Professor and (by courtesy) Operations Research, Ph.D. MIT., 1987.

GROUP: Design and Analysis of Algorithms

RESEARCH: Design and analysis of efficient algorithms for combinatorial optimization problems. Parallel and distributed algorithms. Lower bounds and complexity theory.

SENIOR STUDENTS: Michael Kharitonov, Tomasz Radzik, Robert Kennedy, Steve Murray.

FACULTY: Leonidas J. Guibas, Professor, Ph.D. Stanford, 1976.

AFFILIATION: Theory, Robotics, CSL.

GROUP: Computational Geometry and Computer Graphics

RESEARCH:

- THE DESIGN AND ANALYSIS OF GEOMETRIC ALGORITHMS: Computational geometry is, in its broadest sense, the study of geometric problems from a computational point of view. At the core of the field is a set of techniques for the representation of geometric objects and for the design and analysis of algorithms manipulating these representations.
- Specific current areas of interest include:
  - Parametric searching: The goal here is to generically transform algorithms that solve a decision problem (‘‘Is the largest distance between any two points of a given set S of n points in 3-d less than or equal to d?’’) to algorithms that solve the corresponding optimization problem (‘‘What is the largest distance...?’’).
  - Output-sensitivity in degenerate situations: Over the last five or six years several techniques have been developed to deal with degeneracies (e.g. collinear points, etc.) in geometric computations. These often work by effecting or simulating a perturbation of the input that destroys the degeneracies. Such perturbations, however, are likely to increase the output size of the result, and thus the complexity of computing it. We are investigating ways to handle degeneracies while remaining faithful to the complexity of the original unperturbed output.
  - Model-based recognition in computer vision: In computer vision many problems arise where the goal is to recognize a shape among a class of possible given shapes. What makes this a challenging problem is that the image of an object may be scaled, rotated, or otherwise transformed from the original prototype, it may be partially occluded, the input data may be noisy, etc. We are especially interested in the additional complexity introduced by objects having a small number of degrees of freedom (such as, scissors), and in the use of recognition techniques for localization and navigation. (Additional contact: Prof. Rajeev Motwani)
  - Line hulls of objects in 3-D: Roughly speaking, the line hull of an object is the smallest object containing the original that cannot be distinguished from it with probes involving infinite straight lines moving through free space. The line hull produces all the same shadows as the original object and is a useful concept for vision algorithms trying to extract shape from rotation. We are currently investigating ways to represent the line hull and algorithms for computing it.
  - Tail estimates for randomized geometric algorithms: A very powerful technique for developing simple and efficient algorithms in geometry has been the use of randomization. This is especially so in the area of randomized incremental

constructions, where a geometric structure (such as a convex hull, or a Voronoi diagram) is built by inserting the objects defining it in a random permutation. Such methods invariably have very good expected behavior. In this project we are developing mathematical techniques to understand the full running time distribution of these methods.

- Upper bound theorems for general algebraic varieties: One of the most useful combinatorial tools in geometry is the Upper Bound Theorem. This roughly states that the convex polytope determined by intersecting n halfspaces in d dimensions has complexity as most $n^{d/2}$. We are interested in developing techniques for proving similar results for more general algebraic objects. If successful, this would give us a tool for settling many basic but still open questions in combinatorial geometry.

- EFFICIENT MODELING AND REASONING ABOUT PHYSICAL SYSTEMS: The goal of this research is to develop new techniques for simulating physical phenomena, or for reasoning about physical objects. Our main application areas are computer graphics and robotics, so we are most concerned with modeling well the geometric and optical properties of various environments. Our goal is to attack problems involving ‘‘many’’ objects that would be prohibitively expensive if approached in a straightforward manner. We plan to use clever combinatorial and geometric techniques to find approximate solutions within reasonable time bounds.
    - A new formulation of the global illumination problem: We have a new formulation for the problem of computing the steady-state light energy flow in a scene, as well as a new algorithm for computing a good approximation to this flow. Classical radiosity in computer graphics deals only with surface reflections and even for these it cannot properly handle directional reflections off specular elements. Our method is based on a hierarchical volume subdivision of the scene and on ‘‘three-way’’ light interactions (these ideas are intimately coupled). As a result, both participating media (such as fog), and directional reflections (light from A bouncing off B to reach C) can be accurately modeled in our framework. A combinatorial argument shows that the number of significant triplet interactions we have to account for is much less than the total possible.

SENIOR STUDENT: G.D. Ramkumar

FACULTY: Joseph Halpern, Consulting Professor, Ph.D. Harvard, 1981., Research Staff Member, IBM Almaden Research Laboratory.

GROUP: Reasoning about Knowledge and Probability

RESEARCH: Systems of agents (where the ‘‘agents’’ can be robots, processors in a distributed system, negotiators in a bargaining system, etc.) interact through communication and then must make decisions. A good way of analyzing this interaction turns out to be in terms of knowledge, considering how the knowledge changes as a result of communication and what levels of knowledge are required to perform various actions. The goal is to provide a general framework for modeling knowledge, and to use that framework to analyze and better understand interacting systems. Another line of research is to investigate how to use this information to form principled degrees of belief.

SENIOR STUDENTS: Adam Grove, Daphne Koller.

FACULTY: Zohar Manna, Professor, Ph.D. Carnegie-Mellon, 1968.

GROUP MEMBERS: Richard Waldinger (Consulting Professor of Computer Science, and Staff Scientist, SRI International); Amir Pnueli (Visiting Professor of Computer Science, and Professor, Weizmann Institute)

RESEARCH TOPICS: Program Verification: The Logic of Programs; Program Synthesis: Deriving

Programs from Specifications; Theorem Proving: Automated Deduction; Reactive Systems: Concurrency, Real-time, Hybrid; Temporal Logic: The Logic of Time; Automated Planning.

RESEARCH:

- Automated Deduction: Techniques and Applications
  The development of an interactive theorem prover and its application to software engineering and artificial intelligence. Applications include program synthesis and verification, and robotic and commonsense planning. The deductive system combines unification-based methods such as resolution with proof by mathematical induction. Proofs are restricted to be sufficiently constructive so that programs and other information may be extracted.

  A situational logic has been employed for the synthesis of imperative programs and for planning applications. Research is underway leading to the synthesis of concurrent programs and plans. Domain-dependent inference rules are being introduced to give the system particular facility with application theories.

- Reactive, Real-Time, and Hybrid Systems
  A reactive system is a system whose role is to maintain an ongoing interaction with its environment, such as concurrent and real-time programs, embedded systems, communication networks, air-traffic control systems, avionic and process control programs, operating systems, and many others. One of the most promising and widely accepted approaches to the specification and verification of reactive systems is that of temporal logic. Extensions to the temporal methodology for reactive systems that are currently investigated include real-time and hybrid systems. The real-time extension considers the timing aspects of reactive systems and contains a specification language and a proof methodology for specifying and proving the timing properties of systems.

  Hybrid systems are systems consisting of a non-trivial mixture of discrete and continuous components, such as a digital controller that controls a continuous environment, control of process and manufacturing plants, guidance of transport systems, robot planning, and many similar applications. The extended methodology will provide a (temporal) language for specifying the desired behavior of such systems, methods for their rigorous development, and proof techniques, that will enable verifying that a developed program, interacting with an appropriate continuous environment, meets its specification. We also examine approaches to the automatic verification (model checking) of hybrid systems.

SENIOR STUDENT: Eddie Chang

FACULTY: John C. Mitchell, Associate Professor, Ph.D. MIT, 1984.

AFFILIATION: CSL

GROUP: Programming Language Theory and Design

RESEARCH: Type systems for programming languages, object-oriented programming, logical methods for reasoning about programs, development of programming language design tools. Two language design projects are a type and module system for object-oriented programming and a revision of the programming language ML. Theoretical projects include linear logic, semantic domains based on recursion theory, revised treatments of nontermination, and integration of algebraic and type-theoretic methods.

FACULTY: Rajeev Motwani, Assistant Professor, Ph.D. University of California, Berkeley, 1988.

RESEARCH:

Algorithms and Complexity Theory.

Algorithms and data structures, with emphasis on randomization and approximation algorithms. Complexity theory and lower bounds. Computational geometry and applications to vision/robotics.

FACULTY: Serge A. Plotkin, Assistant Professor, Ph.D. MIT., 1988.

RESEARCH: Design and analysis of parallel algorithms and data structures with emphasis on efficiency in terms of both time and the number of required processors. Theoretical issues that arise in the context of computation over a distributed network of processors. In particular, issues that arise because of the costs involved in moving information from one processor to another and because the processors are asynchronous. Design of efficient parallel and sequential algorithms for combinatorial optimization problems.

FACULTY: Vaughan Pratt, Professor and Director, Computer Science Theory Division, Ph.D. Stanford, 1972.

RESEARCH SCIENTIST: Rob van Glabbeek

INDUSTRY VISITOR: Jeremy Gunawardena

INDUSTRY AFFILIATES: Vadim Kotov, Lucy Cherkasova

GROUP: Concurrency Modeling

RESEARCH: Understanding existing models of concurrency, developing new models, and investigating their properties. Application of the results to the design and definition of concurrent programming languages and their associated tools.

FACULTY: Jeffrey D. Ullman, Professor and Chairman of Computer Science and (by courtesy) Electrical Engineering, Ph.D. Princeton, 1966.

VISITING SCIENTIST: Shinichi Morisita

GROUP: NAIL (Not Another Implementation of Logic)

RESEARCH: We are developing the technology for a new generation of database systems. These systems, like relational systems, use languages that are based in logic. However, unlike SQL and similar languages, we use a general-purpose logic language called GLUE. Work of the group includes design and implementation of an experimental system. We also develop the optimization algorithms needed for the new system, including ‘‘magic-sets’’ and other techniques for transforming logical rules into forms that allow fast evaluation. We investigate the use of constraints to communicate between distributed databases and are examining efficient ways to enforce constraints. Finally, we are interested in applications of the NAIL system to problems in engineering and other disciplines, and we have developed several prototypes.

FUNDING: NSF, ARO

STRUCTURE: Informal structure. Weekly meetings for discussing progress, important papers, crazy ideas, etc.

SLOTS AVAILABLE: One

SENIOR CS STUDENTS: Hakan Jakobsson, Ashish Gupta, Alberto Torres.