Computing and Information Science at Cornell

2000–2001 Annual Report
Education and Research

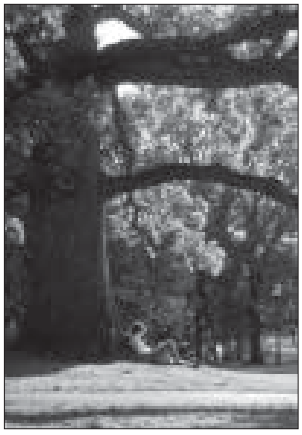

Table of Contents

Message from the Dean

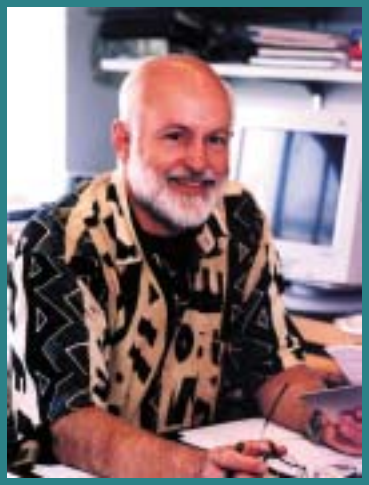

Robert L. Constable

This has been a busy and productive year in the short life of the Cornell initiative in Computing and Information Science (CIS). We hired three more faculty with appointments in the Faculty of Computing and Information (FCI), Paul Ginsparg, joint with Physics; Phoebe Sengers, joint with Science & Technology Studies; and Hod Lipson, joint with the Sibley School of Mechanical & Aerospace Engineering. These are outstanding appointments which I talk more about in this message and who appear in this report in the section on new faculty (see page 11). In addition we have assembled more components of the FCI, and we have continued those activities of the CIS initiative that started before the FCI came into being, namely our teaching with the Johnson Graduate School of Management in e-Business and our collaboration with the School of Electrical and Computer Engineering in the Computer Systems Laboratory.

The current structure of Computing and Information Science (CIS) and its relation to the FCI is provided in the image below.

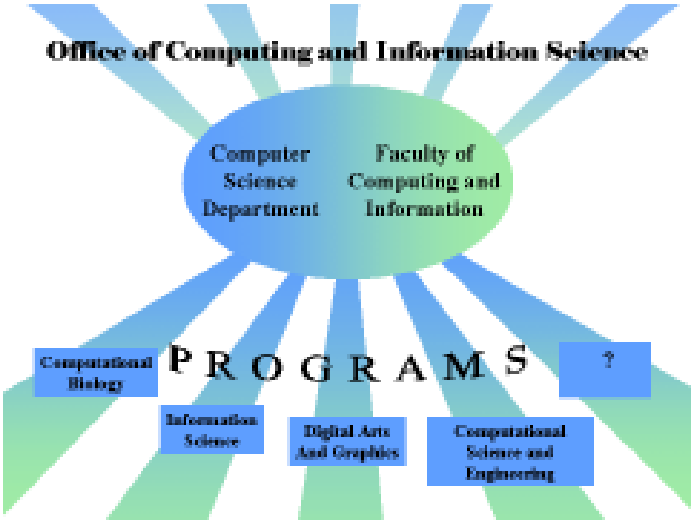

As background to the new academic programs in the Faculty of Computing and Information, I want to quote from an article I wrote called "Universities in the Information Age," which tries to set the FCI in context:

> The Information Revolution is changing every human activity, especially the way we learn, teach, discover and communicate knowledge. Underlying the change is a powerful new way of thinking called the "algorithmic method." Those who know the algorithmic method bring new categories of thought to intellectual tasks; they know what part of the task can be delegated to a computer or how much of it could be accomplished in partnership with a computer. The algorithmic method and computing technology have already dramatically transformed science, and will soon transform the arts and humanities as well. Teaching this far-reaching new mode of thought will become part of a modern liberal education. Before we see widespread effects, however, the basic ideas must become infused more deeply in university curricula and academic structures. The Computing and Information Science (CIS) initiative at Cornell is our attempt to infuse these ideas broadly.

The CIS initiative is developing along the lines laid out in the report *Cornell in the Information Age* (available at www.cis.cornell.edu/information.htm). In particular, Cornell has formed a new academic unit called the Faculty of Computing and Information (FCI) as called for in the report. Its mission is to develop new academic programs to complement the computer science major and to stimulate research in computing and information broadly. The Office of the Dean for CIS has been given substantial new funding to develop the FCI and the Department of Computer Science (CS). Some of these funds have already been used to create new faculty positions for Computer Science while others have been used to build academic programs in computing and information.

Building the FCI shares some characteristics of building a large software system. The design document, *Cornell in the Information Age*, is informed by experience and some "theory" of organizations, but to know whether it works, many details have to be filled in, just as in completing system modules or filling in steps of a large proof. In the end, we have to "run the structure" for several years to see how it works. But it is much more clear now that all the parts fit together and function. Material in this report illustrates important details of the structure.

This year we assembled more parts of the FCI, and we hired more FCI faculty. This is being done with a committee called the *FCI Founders*, consisting of computer science professors and distinguished professors in other units with interest and experience in computing and information science. The Founders have identified the first four areas in which we will build academic programs and hire new faculty. They are Information Science, Computational Biology, Computational Science and Engineering, and Digital Arts and Graphics.

The Information Science program and the Computational Biology program are each well under way. Basic facts about these programs are available on the CIS home page, www.cis.cornell.edu. They are also summarized in this report just below. Computational Science and Engineering has developed a set of short courses, and the FCI has hired Dr. Andrew Pershing of the Earth & Atmospheric Sciences department to teach them. Information about these courses will also be on the CIS web site and the Department of Computer Science web site.

The impact of these programs is being felt almost as soon as they are in place. They have guided CS and other departments in planning course offerings, and they have attracted the interest of many students who will help us implement this emerging system. In terms of an analogy between social systems and computer systems, it is designed so that "new code" can be added while it is running -- but some of the documentation, like this report, is done in the summer when the load on the system is light.

New funds have allowed Cornell to hire three excellent faculty into FCI programs this year, all of whom are highlighted in the section on new faculty. We have attracted Paul Ginsparg as a full professor to FCI with a joint appointment in Physics. Paul is well known for his revolutionary work in scientific publishing embodied in his e-print archive for physics, astronomy, mathematics and computer science. Paul's system is called arXiv.

Jointly with the School of Mechanical & Aerospace Engineering, the FCI hired assistant professor Hod Lipson who works in computer aided design of mechanical systems, graphics, and evolutionary design of systems.

Jointly with the Department of Science and Technology Studies, the FCI hired assistant professor Phoebe Sengers. Her computing expertise is in AI, especially the study of agents and human computer interfaces. Phoebe's Ph.D. is also in the field of Cultural Studies, and she looks at how computer systems embody cultural values. She will be part of the Information Science program.

The FCI also inagurated a Distinguished Lecture Series this year. The first two speakers to present were William A. Wulf, President of the National Academy of Engineering, who spoke on "Some Challenges for Computer Science and Engineering in the 21st Century," and Neil Gershenfeld, Director of the Physics and Media Group, MIT Media Lab and Co-director of the Things that Think research consortium, who spoke on "Things that Think."

The CIS initiative has been helping the Johnson Graduate School of Management build a program in e-Business. This year Fred B. Schneider worked with Richard Conway, from the Johnson School, to design and implement the CS contribution to an immersion course in e-Business. Ken Birman and Johannes Gehrke lectured in this course along with Fred. The CS contribution was well appreciated by the Johnson School students.

Another piece of the CIS initiative is the collaboration between the Department of Computer Science and the School of Electrical and Computer Engineering (ECE) in the Computer Systems Laboratory. CIS provides some lab space that supports investigation in common topics of interest in computer systems. We hope to eventually build new laboratory facilities closer to the Duffield Systems Lab on the third floor of Upson Hall. Students in the lab come from both CS and ECE. The CSL web page (www.csl.cornell.edu) lists fifteen courses offered in this area by over eight faculty from the two units.

The Department of Computer Science continues to play a broad role in the university, now in partnership with the FCI. The Department of Computer Science started its life being managed as a central university resource reporting to the Office of the Provost. Once it created undergradu-

ate programs, they were offered in both Arts & Sciences and in Engineering. Both of these majors are thriving, and the Department of Computer Science plays an active role in the educational programs of both colleges. Its Engineering BS program and its Master of Engineering program are the largest in the college. CS co-teaches courses with Electrical and Computer Engineering and with Civil and Environmental Engineering. Its major in Arts & Sciences is very high quality -- it produced two of the current group of Cornell CS professors. CS is also one of the most active departments in the Cognitive Studies Program, with Joe Halpern as one of the co-directors and two other CS faculty on the executive committee. CS co-teaches courses with Mathematics, Physics, Linguistics and Psychology, and will add a course with Economics and one with Science and Technology Studies soon.

The Department of Computer Science is starting to work more closely with the College of Agriculture and Life Sciences through the FCI programs in computational biology and information science. There are joint research projects and joint seminars. CS faculty will play an important role as this college starts a new department in biological statistics and computational biology.

In addition to its responsibilities in the large undergraduate colleges, CS plays a leading role along with the FCI founders in the CIS initiative. Most of its faculty are active in FCI working groups or other CIS initiatives. Nine CS faculty helped eight university units in recruiting. Being a university department is both demanding and exciting. The CIS faculty are exposed to a variety of stimulating problems and ideas, from systems engineering to digital arts. The FCI and the CS have discovered that teaching people to use the computer as a medium of expression is just as challenging as using it as a problem solving tool.

The Computing and Information Science initiative is being energized by a great deal of university support. The CIS Task Force and the Department of Computer Science helped sketch a broad faculty vision. University president Hunter R. Rawlings has not only encouraged and inspired the effort, he has also helped shape it, helped recruit faculty with us and continues to raise new endowment. Provost Biddy Martin has helped us find new opportunities at Cornell, and she has spent time connecting us with specific programs, especially in support of Digital Arts and Graphics. The University librarian, Sarah Thomas, has been a strong partner as have several college deans, notably Susan Henry, Dean of the College of Agriculture and Life Sciences; Phil Lewis, Dean of the College of Arts and Sciences; Bob Swieringa, Dean of the Johnson Graduate School of Management, and Porus Olpadwala, Dean of the College of Architecture, Art, and Planning. I look forward to another exciting year as we continue to develop Cornell's initiative in Computing and Information Science.

Robert L. Constable

You can access the CIS Dean's web page at
http://www.cs.cornell.edu/cis-dean/.

Message from the Chair

Charles Van Loan

It is a watershed time for the Department of Computer Science. We have emerged from a difficult period in our history characterized by open faculty lines and exploding enrollments. We are stronger than ever before because the CIS structure enables us to address boldly every teaching and research challenge that comes our way. We are now set to move in several directions, each of which represents a different academic adventure, a different path to the sea. I am reminded of the five great river watersheds that grace our part of the Northeast. The Hudson, the Susquehanna, and the Delaware each get their start in beautiful upstate New York. The raindrops that fall here in Ithaca are rewarded with a trip down the St. Lawrence after stopovers in Cayuga Lake and Lake Ontario, while those that fall just to our west find their way into the Gulf of Mexico via the Allegheny, the Ohio, and the Mississippi. Upstate New York is a watershed paradise, and I like to think that Ezra Cornell appreciated the metaphor when he chose the location for this great university. Our job in the department is not to cogitate at the river's edge, but to be part of its source with innovative research and teaching. Quite simply our mission is to affect positively the quality of life downstream in the information age.

To that end there is no better way to ensure the continued flow of remarkable scholarship than to hire creative young people. Gun Sirer (University of Washington) and Golan Yona (Hebrew University) joined the faculty in January, adding new dimensions to our efforts in systems and computational biology respectively.

Rich Caruana (Carnegie Mellon University), Daisy Fan (Cornell), Thorsten Joachims (Dortmund), and Jayavel Shanmugasundaram (Wisconsin) will be new faculty colleagues this fall. Rich works in machine learning and data mining and has interests in medical decision making and bioinformatics. Daisy is a civil engineer who will continue her research in computational science while she helps run our large CS 100 operation. Thorsten works in machine learning and intelligent agents with a focus on Support Vector Machines and learning with text. Jai's research interests include internet data management, database systems, and transaction processing in emerging system architectures.

In January, two additional faculty members come on board. Jeanna Neefe Matthews (UC Berkeley) has research interests in file systems and storage systems. Radu Rugina (UC Santa Barbara) does research on pointer analysis, parallelizing compilers, and parallel computing.

With these additions the department will become the home to twelve assistant professors—an all-time high. Working to ensure that these young scholars realize their potential in our environment is easily one of the best parts about being the chair. But my job would be virtually impossible without the leadership of my senior colleagues who provide support in so many important ways. Our wildly successful faculty recruiting effort this year is a tribute to the environment that they have created. Research entities like the Information Assurance Institute and the Intelligent Information Systems Institute are increasingly important when it comes to attracting the top CS scholars to Cornell.

Before I leave the topic of new faculty, I would like to mention the return to the department of John Hopcroft. John has been gone for nearly ten years serving the university as an associate dean and dean in the College of Engineering. It will be great to have a Turing Award winner back in Upson Hall.

There are many new developments on the teaching front. During the year we started eight new courses. For freshman our one-credit minicourse *Great Ideas from Computer Science* (CS 150) gives snapshots of the field through a sequence of eight expository lectures. As a follow-up to CS 130,

our introductory web course, Graeme Bailey taught a prototype of *Intermediate Web Design* (CS 230). CS 130 and 230 focus on the client side and the server side respectively. Professor David Mermin from Physics taught *Quantum Information Processing* (CS 483). The theory of quantum computation offers striking new perspectives on computation and information, as well as on the quantum theory itself. Bill Arms taught the first edition of *Information Discovery* (CS 430). This course looks at the methods used to search for and discover information in digital libraries and web information systems. Al Demers introduced a much-needed upper level undergraduate course on *Programming Languages and Logics* (CS 411). At the graduate level we were able to expand our offerings in the theory area with Jon Kleinberg's *Advanced Design and Analysis of Algorithms* (CS 683). The course emphasizes algorithmic problems in a range of areas including networks, large datasets, lattices, and the design of heuristics. Mats Rooth from the Department of Linguistics started *Computational Linguistics* (CS 324) and *Introduction to Natural Language Processing* (CS 474).

This coming year Lillian Lee will teach a new course entitled *Computation, Information, and Intelligence* (CS 172). This will be a non-programming, freshman-level introduction to computer science using methods and examples from artificial intelligence. We think that the novel approach of this course will help attract more women to the CS major. Eva Tardos will teach a new graduate-level course on *Approximation and Network Algorithms* (CS 684). Another course making its debut this coming year will be our new *Java Practicum* (CS 212), which will serve as a "large project" companion to the second course in our Java sequence (CS 211). Dave Schwartz will be starting *Advanced Unix Programming and Tools* (CS 214) in the spring semester. Rich Caruana will introduce a new course on *Empirical Methods in Machine Learning and Data Mining* (CS 578) in the fall.

A number of broad curriculum reviews/initiatives will take place during the coming year. The breadth and depth of the machine learning group will make it possible for us to put together a remarkable curriculum in this important area. Work on this will begin in the fall. Likewise, the systems group plans to revamp its set of courses working closely with colleagues in the School of Electrical and Computer Engineering. We expect to see curriculum gaps filled and increased coordination between the two units via the cross-listing of critical courses. We look to expand our commitment to the Academic Excellence Workshop (AEW) idea by offering several of the CS 100 AEWs on North Campus and by starting to offer AEWs in CS 211. Finally, long-time colleague David Gries will be visiting us during the fall semester and will be exploring how we might develop a self-paced introduction to Java tailored to incoming freshman (and others) who are trained in C++.

Expansion of faculty and curriculum creates space problems. However, thanks to the department's space committee and Pat Musa, we will be able to manage for another twelve months. Nevertheless, a major challenge during the coming year will be to find new offices and laboratories to accommodate our growth.

The department's relationship with other academic units will be much enhanced because of new faculty appointments that were made possible by the FCI. In particular, the hiring of Paul Ginsparg in Physics, Phoebe Sengers in Science and Technology Studies, and Hod Lipson in Mechanical and Aerospace Engineering are extremely exciting developments for us. Following the example of Mats Rooth, who joined Linguistics this past year, and Geri Gay in the Department of Communications, these FCI-supported faculty will be the basis for new research alliances and new cross-listed courses. The close proximity of these faculty to the department's research engine is a hallmark of Cornell's CIS initiative.

In connection with on-campus relationships, we note that several of our faculty served on search committees in other units. Joe Halpern and Johannes Gehrke worked with Operations Research and Industrial Engineering; Dexter Kozen with Mathematics; Bob Constable with Science and Technology Studies; Greg Morrisett and myself with Electrical and Computer Engineering; Ron Elber for the Tri-Institional initiative; Fred Schneider, Joe Halpern, and Johannes Gehrke with the business school.

Other activitity that is serving to cement our relationship with other units includes Ramin Zabih's joint appointment with the Department of Radiology at Weill Medical College in New York City and the intersession teaching of Ken Birman, Johannes Gehrke, and Fred Schneider in the Johnson Graduate School of Management.

Our faculty continue to render the very highest levels of service to the university. Bill Arms chaired the Provost's Advisory Committee on Distance Learning and is a Director of eCornell. Tom Coleman is Director of the Cornell Theory Center (CTC) and the Financial Industry Solution Center down in Manhattan. Bob Constable is Dean for Computing and Information Science. Ron Elber is the Head of the National Institutes of Health Resource in the CTC. Joe Halpern co-chairs the cognitive studies program. Ken Birman is a member of the advisory council for the Cornell Research Foundation, and Head of the Responsible Conduct of Research Committee for the university.

In alumni and corporate relations we are moving rapidly to organize our efforts. The CS major is only 20 years old, and so we are just now beginning to see the majority of our graduates move into the prime of their careers. Successful Cornell CS gatherings in both Boston and Palo Alto were held during the year. Our Alumni Weekend breakfast here

on campus had more attendees than ever before. Dan Jenkins will be orchestrating additional alumni events of this sort in the future. Relatedly, Marcy Rosenkrantz will now be managing corporate relations, and this will enable us to track more effectively the many connections that we have in the industrial sector.

Several faculty were honored at the national and international level making us all extremely proud to have such wonderful colleagues in our midst. Eva Tardos was elected to the American Academy of Arts and Sciences for her pioneering work in the algorithms area. Joe Halpern received both a Fulbright Scholarship and a Guggenheim Fellowship. Juris Hartmanis won the Lielo Medalo from the Latvian Academy of Sciences. Jon Kleinberg received the National Academy of Sciences Award for Initiatives in Research. Greg Morrisett was honored by Carnegie Mellon University with the Allen Newell Medal for Research Excellence. Greg also received a Presidential Early Career Award for Scientists and Engineers and was honored at a White House ceremony last fall. Johannes Gehrke received an IBM Faculty Development Award for the second year in a row.

The CS faculty continue to be campus leaders in the classroom. Dexter Kozen won the Stephen and Margery Russell Teaching Award in the College of Arts and Sciences. Johannes Gehrke, Juris Hartmanis, and Greg Morrisett won College of Engineering teaching awards.

There are many more exciting things that I could write about. But it is far better for you to read the firsthand accounts that follow. My colleagues are the ones who made everything happen!

Charles Van Loan

You can access the Chair's web page at
http://www.cs.cornell.edu/chair/ .

Computer Systems lab with ECE

The Cornell Computer Systems Laboratory (CSL) brings together faculty with common interests from the School of Electrical and Computer Engineering and the Department of Computer Science at Cornell.

The Computer Systems area encompasses both experimental and theoretical work, growing out of topics in computer architecture, parallel computer architecture, operating systems and compilers, computer protocols and networks, programming languages and environments, distributed systems, VLSI design and fabrication, and system specification and verification.

Graduate students are admitted to either ECE or Computer Science. Usually students with primary interest in computer architecture, multiprocessor design, VLSI, CAD, and circuit design enroll in ECE, while students with interest in compilers, operating systems, and programming environments enroll in CS. There are no rigid student classifications; ECE students can have a thesis advisor in Computer Science and vice-versa. Indeed, the interdisciplinary composition of the research teams is a strength of the Computer Systems Lab.

You can visit the CSL website at:
http://www.csl.cornell.edu.

Information Assurance Institute

The Cornell Information Assurance Institute (IAI) is intended to support a broad spectrum of research and education efforts aimed at developing a science and technology base to enhance information assurance and networked information systems trustworthiness—system and network security, reliability, and assurance. IAI is also intended to foster closer collaborations involving Cornell and Air Force Research Lab (AFRL) researchers.

Fred B. Schneider is the director of the institute, and Marcy Rosenkrantz is the associate director.

AFRL researchers are able to participate in Cornell research projects, facilitating technology transfer as well as exposing Cornell researchers to problems facing the Air Force. Cornell researchers are able to become involved in AFRL projects and have access to unique AFRL facilities. The Institute thus makes both Cornell and AFRL more attractive places to work, facilitating recruitment of higher caliber personnel at each site.

A series of technology-transition activities dominated Cornell-AFRL interactions this year. Also under the auspices of the IAI, Cornell researchers were involved in the development of the Air Force Joint Battlespace Infosphere. Various other technical collaborations are being explored—in the use of "gossip protocols," in language-based security policy enforcement technology, and in data mining from networks of sensors.

We are also exploring the expansion of IAI to include a few key industrial partners. These partnerships will not only support and leverage IAI activities but will add another important perspective to the problems IAI researchers attack and the solutions they investigate.

For more information, see:
http://www.cis.cornell.edu/iai/.

Information Assurance Institute Director, Fred B. Schneider

Charles Holland, Director of Information Systems, Office of the Undersecretary of Defense (Science and Technology), speaking at the IAI inaugural

The Intelligent Information Systems Institute

The Intelligent Information Systems Institute (IISI) began operation in December 2000. The mission of the IISI is threefold: To perform and stimulate research in compute- and data-intensive methods for intelligent decision making systems; to foster collaborations between Cornell researchers, our sponsors, and the scientific community; and to play a leadership role in the research and dissemination of the core areas of the institute. The institute is funded by AFRL/AFOSR. Carla Gomes is the director of the institute and Marcy Rosenkrantz is the associate director. The Scientific Advisory Board of the institute consists of Robert Constable (Cornell), Nort Fowler and Charles Messenger (AFRL/IF), and Neal Glassman (AFRL/AFOSR).

Intelligent Information Systems Director, Carla P. Gomes

The IISI supports basic research within the Department of Computer Science and promotes a cross fertilization of approaches from different disciplines, including Operations Research, Mathematics, Statistics and Physics. Areas of investigation within the IISI are: Search and Complexity, Planning and Scheduling, Knowledge Representation, Data Mining and Information Retrieval, Reasoning under Uncertainty, Natural Language Processing, Machine Learning, and Agent Technology. Current Cornell members of the institute are Raffaello D'Andrea (dynamics and control), Claire Cardie (natural language understanding and machine learning), Johannes Gehrke (database systems and data mining), Carla Gomes (artificial intelligence and operations research techniques for combinatorial problems and multi-agent technology), Joseph Halpern (knowledge representation and uncertainty), Mark Heinrich (active memory, simulation methodology), Lillian Lee (statistical methods for natural language processing), David Shmoys (algorithms for large-scale discrete optimization), Bart Selman (knowledge representation, complexity, and multi-agent systems), Chris Shoemaker (large scale optimization and modeling), Evan Speight (distributed computing, computer architectures), and Stephen Wicker (intelligent wireless information networks).

Through the IISI, a number of research projects involving direct collaborations between Cornell and AFRL/IF researchers were initiated during the last six months. These projects cover a variety of topics, such as probabilistic decision making, architectures for active memory systems, multi-agent sensor networks, and visualization of reasoning and search methods. The IISI also hosted a hands-on workshop on foundations and complexity of multi-agent systems. As one of the outcomes of the workshop, a team of researchers from Cornell University, Stanford University, and the University of Washington is developing a tuneable benchmark suite for the design and evaluation of new algorithms for combinatorial auctions. The IISI is also sponsoring the AAAI Symposium on Uncertainty within Computation, the 2001 Conference on Empirical Methods in Natural Language Processing (EMNLP 2001), Language Technologies 2001, and Association for Computational Linguistics (NAACL 2001).

To further its research mission, the IISI hosts a large number of visiting scientists. During the first year, this group has included researchers from AFRL/IF, AT&T Labs, Hebrew University, Microsoft Research, Stanford University, Technion, University of Leida, University of Lisbon, University of Minnesota, University of Washington , Washington University, and York University.

For additional information, see:
www.cis.cornell.edu/iisi.

Marcy Rosenkrantz, Associate Director of the IAI and IISI

CIS Electronic-commerce immersion with the JGSM

The Johnson Graduate School of Management (JGSM), in collaboration with Cornell's Computer Science Department, initiated an immersion course in electronic commerce (E-business Immersion NBA 613).

A $1 million grant from Corning, Inc. to JGSM enabled the school to expand the program.

Professor Richard Conway led the course. Professor Fred B. Schneider coordinated the Computing and Information (CIS) involvement. The goal was to give students an in-depth understanding of all aspects of doing business in the networked economy. A JGSM/CIS collaboration means that students receive a solid education in the enabling technologies that form the backbone of e-commerce.

JGSM immersion courses represent a new model of management education (introduced by Conway years ago);it replaces static case-based training with integrated, experiential, just-in-time learning. In an immersion course, students work on real-world problems under real-world time pressures. The result is savvy students who combine a sound theoretical background with a hands—on sophistication—a "been there, done that" kind of confidence and the "big picture" perspective required for success in the business world today.

With Intel's support, we have created an E-commerce laboratory for the E-business immersion as well as to support a broad collection of e-commerce courses in JGSM. The lab comprises a group of powerful workstations for individual students and several powerful servers to support a data management infrastructure for a fictional e-commerce company (including a web server, an industry-strength database system, web-database connectivity, and database application software).

We also taught a second-year elective course, NBA 601 "The Software Infrastructure of Electronic Commerce;" The course introduced students to computer security, distributed systems, fault-tolerance, software construction, reasoning under uncertainty, and database and data-mining technology.

For further information, see:
http://www.cs.cornell.edu/courses/NBA602/2001sp/

Cornell Theory Center

The Cornell Theory Center (CTC) is Cornell's high-performance computing and interdisciplinary computational research center, serving over 150 faculty research groups across campus and at Weill Medical College. CTC has pioneered the use of industry-standard computational clusters running Windows as a productive large-scale computing environment, keeping Cornell at the forefront of computational science and engineering. The center currently operates a complex that includes more than 600 processors. One part of the cluster complex, Velocity+, is dedicated to strategic applications that require large-scale parallel computing to achieve results. Among these applications are computational materials, genomics, and structural biology. Several Computer Science faculty are involved in these activities, including Keshav Pingali in computational materials and Ron Elber in both structural biology and genomics.

Another CTC research focus is computational finance, an activity headed by CTC director and computer science professor Thomas Coleman. Professor Robert Jarrow of the Johnson Graduate School of Management also plays a leadership role in this effort. Projects include investigating new optimization algorithms for large-scale portfolio analysis and value-at-risk calculations. Much of CTC's computational finance work takes place at CTC-Manhattan, which is located across from the New York Stock Exchange.

CTC brings high-performance computing into the undergraduate curriculum through a NASA-funded project, which provides students in CEE 479 & MAE 491 with access to EduCluster and the CAVE. EduCluster is a 16-processor cluster dedicated to student applications; the CAVE is an immersive 3-D virtual reality environment. CTC's three-wall CAVE allows scientists to "immerse" themselves in their application. The project, titled Advanced Interactive Design Environment," will have seniors using these resources to design a piece of NASA's future Reusable Launch Vehicle.

The Theory Center has helped a number of organizations on campus acquire and operate smaller clusters dedicated to their research fields. For example, CTC is home to the cluster used by the Cornell Institute for Social and Economic Research (CISER) for extensive census data mining and analysis. Running parallel SAS, a statistical software package, the CISER cluster is also used in social and economic classes across campus.

CTC has done pioneering work in science communication outreach and informal education through its Virtual Worlds SciCentr, which consists of a series of multi-user virtual environments. This project has engaged several interdisciplinary teams of undergraduate programmers, designers, and content developers in the creation of interactive exhibits. A number of team members come from Computer Science. SciCentr brings CTC into interaction with a number of faculty affiliated with FCI, in the fields of Communication, Fine Arts, Theatre Arts, Music, and Architecture.

Rich Caruana

Assistant Professor
Department of Computer Science

http://www.cs.cmu.edu/~caruana/
Ph.D. Carnegie Mellon University, 1997

I work in machine learning and data mining, medical decision making and bioinformatics, feature selection, missing values, inductive transfer, artificial neural networks, memory-based learning. I will be joining the Department of Computer Science at the start of the fall 2001 semester.

HONORS/AWARDS

Invited speaker, AAAI Workshop on Learning from Imbalanced Data (July, 2000). "Methods for Learning from Imbalanced Data."

Invited speaker, Research Jamboree, University of Edinburgh, Scotland (May, 2000). Three talks: "Unabridged and Multitask Learning," "Machine Learning, Medical Decision Making, and Explanation," and "Semi-supervised Clustering."

PUBLICATIONS

"Predicting Cesarean Delivery: Decision Tree Models." To appear in the *American Journal of Obstetrics and Gynecology* (2001). With Cynthia J. Sims, Leslie Meyn, Rich Rao, R. Bharat, Tom Mitchell, and Marijane Krohn.

"An Evaluation of Machine Learning Methods for Predicting Pneumonia Mortality." *Artificial Intelligence in Medicine* 9:107-138 (1997). With G. F. Cooper, C. F. Aliferis, R. Ambrosino, J. Aronis, B. G. Buchanan, M. J. Fine, C. Glymour, G. Gordon, B. H. Hanusa, J. E. Janosky, C. Meek, T. Mitchell, T. Richardson, and P. Spirtes.

"Multitask Learning." *Machine Learning* 28:41-75, Kluwer Academic Publishers (1997).

"Experience with a Learning Personal Assistant." *Communications of the A CM* (1994). With Tom Mitchell, Dayne Freitag, John McDermott, and David Zabowski.

"Fifteen Useful Tricks with Extra Outputs." In Neural Networks: Tricks of the Trade, G. B. Orr, and K–R. Muller, editors, Springer-Verlag (1988).

PATENTS

"Using Active Monitor Illumination for 3-D Active Imaging." Patent disclosure filed April, 2000. With Rahul Sukthankar, Keiko Hasegawa, and Matt Mullin.

"Iterated K-nearest neighbor Method and Article of Manufacture for Filling in Missing Values." United States Patent 6,047,287, Assignee: Justsystem Pittsburgh research Center, Pittsburgh, PA. Filed May 5, 1998, granted April 4, 2000.

K-Y. Daisy Fan

Assistant Professor
Computer Science Department

kdf4@cornell.edu
Ph.D. Cornell University, 2001

My research interests include the application of systems analysis techniques for water resources and environmental problems. I will be joining the Department of Computer Science at the start of the fall 2001 semester . With Dave Schwartz, I will be running CS 100 and developing the academic excellence workshops that are associated with that very important course.

HONORS/AWARDS

Graduate Teaching Assistant Award, Department of Computer Science, Cornell University (2000).

New York State Section American Water Works Association Russell L. Sutphen Scholarship.

John E. Perry Teaching Assistant Prize, School of Environmental Engineering, Cornell University (1999).

PUBLICATIONS

"Regression Dynamic Programming for High-dimensional Continuous-state Problems." In preparation for submittal to *Operations Research*.

"Regression Dynamic Programming for Multiple Reservoir Control." Proceedings of the 2000 ASCE Joint Conference on Water Resources Engineering and Water Resources Planning and Management, Minneapolis, MN (2001).

Paul Ginsparg

Professor
FCI, joint with Physics

http://xxx.lanl.gov/blurb/p96unesco.html
Ph.D. Cornell, 1981

I received my A.B. degree summa cum laude in Physics from Harvard University in 1977 and my Ph.D. in Physics from Cornell University in 1981 (Quantum Field Theory, thesis advisor: Kenneth G. Wilson), where I was supported as an NSF graduate fellow and A.D.White fellow. I was in the Harvard Society of Fellows from 1981-84, and an assistant professor in the Harvard University Physics department from 1984-90, where I was supported as an A.P. Sloan Fellow and as a DOE Outstanding Junior Investigator. I was a Technical Staff Member in the Los Alamos National Laboratory Theoretical Division from 1990-2001. I have also held visiting positions at C.E.N. Saclay, France, Princeton University, Stanford Linear Accelerator Center, the Institute for Advanced Studies, Princeton, the Institute for Theoretical Physics at UC Santa Barbara, the Mathematical Science Research Institute at UC Berkeley, and Hebrew University, Jerusalem. In 1991 I inititated the "e-print arXiv" as a new form of communications research infrastructure for physics. I have served on many committees and advisory boards, including most recently an NRC/CODATA committee on "Transborder flow of Scientific Data," an NAS/NRC committee on "Future of Universities," an AAAS study committee on "Transition from Paper," and an NSF committee on "Knowledge Networks and Distributed Intelligence Initiative," and I currently serve on the APS global "Task Force on Electronic Information Systems," on the NIH's "PubMed Central National Advisory Committee," on the Open Archives Initiative Steering Committee (and founder), on the French "Centre pour la Communication Scientifique Directe" technical steering committee (and founder), and on the APS "Publications Oversight Committee." I have also given numerous invited keynote talks and colloquia, including recently at meetings "Future of Mathematical Communication" at the University of Minnesota; "Electronic Publishing in Science" at UNESCO HQ in Paris; "50th Anniversary of the Development of the ENIAC computer" at the University of Pennsylvania; "The impact of electronic publishing on the academic community" in Stockholm, Sweden; "Alternative Models for Scholarly Publishing in Higher Education" at UC Berkeley; Chautauqua meeting of the National Computational Science Alliance at the University of KY, Lexington; joint meeting of the Medical Library Association and the Canadian Health Library Association, Vancouver, BC; "The Impact of Barrier-free Access and New Technologies on Biomedical Publishing" at the New York Academy of Medicine; and "Open Archives European Open day," at Max Planck Library in Berlin.

AWARDS/HONORS

P.A.M. (Physics Astronomy Math) award from the Special Libraries Association. "Honors work which demonstrably improves the exchange of information in physics, math or astronomy" (1998).

Fellow of the American Physical Society. Cited "For his work relating to chiral symmetry on the lattice, for fundamental contributions to string theory, and for establishment and development of the revolutionary 'Los Alamos E-Print Archive'" (November, 2000).

PUBLICATIONS

"A Remnant of Chiral Symmetry on the Lattice." *Phys. Rev.* D25, 2649 (1982). And K. G. Wilson.

"Curiosities at 'c=1'."*Nucl. Phys.* B295 [FS21], 153 (1988).

"2D Gravity + 1D Matter." *Phys. Lett*. B240, 333 (1990). And J. Zinn-Justin.

"First Steps towards Electronic Research Communication." *Computers and Physics* 8 (4):390 (July/August, 1994).

"Winners and Losers in the Global Research Village." In proceedings of 'Electronic Publishing in Science,' Sir R. Elliot and D. Shaw, editors. Held at UNESCO HQ, Paris, ICSU Press (1996).

Thorsten Joachims

Assistant Professor
Department of Computer Science

http://www-ai.cs.uni-dortmund.de/PERSONAL/joachims.eng.html
Ph.D. Dortmund, 1997

My core scientific interests lie in machine learning and statistical learning theory, with its main applications in the fields of text mining and intelligent information agents. This application domain will be of increasing importance, since human attention and the speed with which we can process information are natural bottlenecks that limit our ability to make informed decisions. While data-base and data-mining techniques can already assist users when analyzing structured data, the problem of text mining, the automatic analysis of text-based information, is still largely unsolved. Therefore, from an application perspective, my goal is to develop systems and agents that focus, enhance, and accelerate our ability to access large quantities of natural language information. Central to such systems are, for example, tasks like text and speech classification, information extraction, and abstract generation.

To investigate the fundamental challenges in these tasks, I am particularly interested in machine learning approaches. This is because most tasks dealing with natural-language based information are inherently difficult to formalize and solve manually. In text classification, for example, there is no formal language that lets us describe common categories without reference to human-like background knowledge. Machine learning approaches have shown their ability to overcome the lack of background knowledge by exploiting statistical regularities of word usage patterns that are related to humanly defined concepts. The success of such learning methods is already reflected in their high commercial demand, and I am convinced that machine learning approaches can further push the limit in understanding text-based information. Therefore, from a methodological perspective, my goal is to develop and understand machine learning approaches that fit the properties of text-mining problems.

PUBLICATIONS

"Knowledge Discovery and Knowledge Validation in Intensive Care." *Artificial Intelligence in Medicine* 19(3):225-249 (2000). With K. Morik, M. Imhoff, P. Brockhausen and U. Gather.

"Aktuelles Schlagwort: Support Vector Machines." *Künstliche Intelligenz* 33(4):54-55 (1999).

"Browsing-assistenten, Tour Guides und Adaptive WWW-server." *Künstliche Intelligenz* 28(3):23-29 (1998). With D. Mladenic.

"Making Large-scale SVM Learning Practical." In Advances in Kernal Methods – *Support Vector Learning* , Schölkopf et al., editors, chapter 11, 169-184, MIT Press (1999).

"Text Categorization with Support Vector Machines: Learning with Many Relevant Features." In Proceedings of the European Conference on machine Learning (ECML), Chemnitz, Germany (1998). Ranked 12 in NEC Research Index most accessed publications (August, 2000).

Hod Lipson

Assistant Professor
FCI, joint with Mechanical and Aerospace Engineering

http://www.mit.edu/~hlipson
Ph.D. Technion, 1998

Engineers design by combining knowledge and resources to make products that achieve some functionality. Despite the fact that design is the basis of engineering, this process of problem solving by *synthesis* is not understood well, and is still taught, to a large extent, as an art. While we have elaborate computational models for analysis, we still have no computational model of synthesis. I believe understanding this process holds the key to future competitiveness, and presents a largely unadressed challenge across both engineering and computer science.

My research interests are in the area of computational design, information systems and fabrication—at the intersection of engineering and computer science. I am interested in understanding the synthesis process of design and emulating it computationally, and I focus on the ideas of self-organization and self-replication as new paradigms of design, fabrication and learning. I search for ways to harness all these areas to make the future CAD/CAM systems. I look both at the human design process and at natural design and fabrication as two sources of inspiration, and I build working systems to test my theories.

HONORS/AWARDS

Technology R eviews one of ten most promising technologies of the future, 2001.
TIME Magazine 's Annual 2001.
Shaping the Future Award, EXPO '2000.
Fischbach Postdoctoral Scholarship, 1998.
CIRP International F. W. Taylor Medal, 1997.
Charles Clore Scholarship Award for Academic Excellence, 1996.
Miriam and Aaron Gutwirth Memorial Fellowship, 1996.

PUBLICATIONS

"Automatic Design and Manufacture of Robotic Lifeforms." *Nature* 406:974-978 (2000). With J. B. Pollack.
"Clustering Irregular Shapes using High Order Neurons." *Neural Computation* 12(10):2331-2353 (2000). With H. T. Siegelmann.
"Conceptual Design and Analysis by Sketching." *Journal of AI Design and Manuf acturing* 14:391-401 (2000). With M. Shpitalni.
"3D Conceptual Design of Sheet Metal Products by Sketching." *Journal of Materials Pr ocessing Technology* 103(1):128-134 (2000). With M. Shpitalni.
"Identification of Faces in a 2D Line Drawing Projection of a Wireframe Project." *IEEE Transactions on P attern Analysis and Machine Intelligence* (*PAMI*) 18(10):1000-1012 (1996). With M. Shpitalni.

Jeanna Neefe Matthews

Assistant Professor
Department of Computer Science

http://www.clarkson.edu/~jnm/
Ph.D. Berkeley, 2000

My research interests include file systems, storage systems and more generally operating systems and distributed systems.

AWARDS/HONORS

Intel Foundation Graduate Fellowship Award, 1999-2000.
Cal VIEW Fellow Award, for excellence in teaching, 1998-1999.
National Science Foundation Graduate Research Fellowship, 1994-1998.
Award Paper, 1995 ACM Symposium on Operating System Principles.

PUBLICATIONS

"Improving the Performance of Log-structured File Systems with Adaptive Methods." Proceedings of the Sixteenth ACM Symposium on Operating Systems Principles, 238-251 (October, 1997). And D. Roselli, A. Costello, R. Wang, and T. Anderson.
"Serverless Network File Systems." *ACM Transactions on Computer Systems* (February, 1996). With T. Anderson, M. Dahlin, D. Patterson, D. Roselli, and R. Wang.
"Serverless Network File Systems." Award paper. In Proceedings of the Fifteenth ACM Symposium on Operating Systems Principles, 109-126 (December, 1995). With T. Anderson, M. Dahlin, D. Patterson, D. Roselli, and R. Wang.
"A Case for Network of Workstations: NOW." *IEEE Micr o* (February, 1995). With T. Anderson, D. Culler, D. Patterson, and the NOW Team.
"An Exploration of Network RAM." UC Berkeley Technical Report, UCB/CSD-98-1000 (December, 1994). With E. Anderson.

Radu Rugina

Assistant Professor
Department of Computer Science

http://www.cag.lcs.mit.edu/~rugina/
Ph.D. University of California, Santa Barbara, 2001

Program analysis automatically extracts information that is critical for understanding, maintaining and debugging the program; for checking properties about the program; or for applying various transformations to the program. Although it has traditionally been considered part of the programming languages and compilers community, program analysis has applications in virtually all areas of computer science, and these applications can deliver substantial benefits to scientists in these fields. Research in program analysis is interesting, relevant and important.

In the last few years I developed new techniques for pointer analysis and symbolic analysis of accessed memory regions. These techniques can analyze general programs, including programs that use recursion, multithreading, and manipulate pointers. I also interacted with researchers in other fields to apply these techniques to problems in computer architecture, parallel computing, and software engineering. Concrete results include the automatic parallelization of sophisticated divide and conquer problems, the static detection of array bounds violations and data races in multithreaded C programs that heavily use pointers and pointer arithmetic, and the use, by computer architects, of pointer analysis results to map programs onto the MIT RAW machine and hardware circuits.

In the future, I intend to extend this research to include software reliability and computer security. In these fields, there is a need to automate the process of checking important properties required to guarantee the functionality or safety of the program. For large, complex pieces of software, checking these properties manually, by humans, is a difficult and error-prone task. However, using formal verification techniques based on program annotations requires programmers to substantially change the way they write software. In this context, program analysis is an appealing alternative. It can be used to develop automatic tools to solve these kinds of problems, but without requiring changes of programming style. I intend to develop deep program analysis techniques and focus on application of these techniques for software reliability and security. I believe that to successfully solve these problems, the analysis techniques have to be general, not restricted only to a certain class of programs, like scientific applications. These program analyses also have to be based on general, formally correct frameworks, not on ad-hoc techniques.

AWARDS/HONORS

Tuition Fellowship from UC Santa Barbara between 1996 and 2001.

Merit Fellowship from Polytechnica University of Bucharest between 1991 and 1996.

Scholarships from the European Union TEMPUS Program for several projects at the Technical University of Eindhoven in Holland: April-June 1992, March-May 1993, and May-July 1996.

PUBLICATIONS

"Recursion Unrolling for Divide and Conquer Problems." In Proceedings of the13th International Workshop on Languages and Compilers for Parallel Computing, IBM T.J. Watson Research Center, Yorktown Heights, NY (August, 2000). With Martin Rinard.

"Symbolic Bounds Analysis of Pointers, Array Indices, and Accessed Memory Regions." In Proceedings of the ACM SIGPLAN 1999 Conference on Programming Languages Design and Implementation, Atlanta, GA (May, 1999). With Martin Rinard.

"Automatic Parallelization of Divide and Conquer Algorithms." In Proceedings of the ACM SIGPLAN 1999 Conference on Programming Languages Design and Implementation, Atlanta, GA (May, 1999). With Martin Rinard.

"Predicting the Running Times of Parallel Programs by Simulation." In Proceedings of the 12th International Parallel Processing Symposium and 9th Symposium on Parallel and Distributed Processing, Orlando, FL (April, 1998). With Klaus E. Schauser.

Phoebe Sengers

Assistant Professor
FCI, joint with Science
and Technology Studies

http://www.cs.cmu.edu/afs/cs.cmu/user//phoebe/mosaic/my-home-page.html

Ph.D. Carnegie Mellon University, 1998

I am a computer scientist and a cultural theorist. I build intelligent, interactive, expressive information systems, like an artificial agent that makes emotionally expressive child-like drawings. The goal of my work is to analyze carefully what our systems currently unconsciously express and to develop technology that allows us to express new aspects of human experience. I use the tools of cultural theory as a way to understand our systems better and to generate technical ideas for new forms of technology. I am part of a growing community of critical technical practitioners.

I have done research on agents, avatars, virtual environments, and computer graphics at the GMD in Bonn, Germany. I am active in the Narrative Intelligence research community. Last year, I was a Fulbright Guest Researcher at the Center for Art and Media Technology (ZKM) in Karlsruhe. In August 1998, I graduated from Carnegie Mellon University, with a self-defined interdisciplinary Ph.D. in Artificial Intelligence and Cultural Theory (administered jointly by the Department of Computer Science and the Program in Literary and Cultural Theory).

HONORS/AWARDS

Lingua Franca, Tech Top 20 (July, 1999), named one of the "top 20 researchers changing the way we think about technology."

Fulbright Fellowship (September 1998-July, 1999).

AAAI Doctoral Consortium (August,1996).

Office of Naval Research Allen Newell Graduate Fellowship (October 1994-September 1997).

Member, National Research Council Study for Computers and Creativity.

PUBLICATIONS

Narrative Intelligence . Michael Mateas and Phoebe Sengers, editors. Advances in Consciousness Series. Amsterdam: John Benjamins Publishing Company, forthcoming.

"Narrative Intelligence." In *Human Cognition and Social Agent Technology* . Kerstin Dautenhahn, editor. Advances in Consciousness Series. John Benjamins Publishing Company, 2000.

"Fabrikation der Subjekte: Verdinglichung, Schizophrenie, und Kuenstliche Intelligenz." In *Netzkritik: Materialien zur Internet-Debatte* . Geert Lovink and Pit Schultz, editors. Berlin: Edition ID-Archiv (1997).

"Technological Prostheses: An Anecdote." *ZKP-4 Net Criticism Reader*. Geert Lovink and Pit Schultz, editors (1997).

"Fabricated Subjects: Reification, Schizophrenia, Artificial Intelligence." *ZKP-2 Net Criticism Reader*. Geert Lovink and Pit Schultz, editors (1996).

"Practices for Machine Culture: A Case Study of Integrating Artificial Intelligence and Cultural Theory." *Surfaces* VIII (1999).

"Madness and Automation: On Institutionalization." *Postmodern Cultur e* (May, 1995).

"Wallowing in the Quagmire of Language: Artificial Intelligence, Psychiatry, and the Search for the Subject." *Cultronix* (Summer, 1994).

Jayavel Shanmugasundaram

Assistant Professor
Computer Science Department

http://www.cs.wisc.edu/~jai/index.html
Ph.D. Technion, 1998
Ph.D. University of Wisconsin, 2001

My broad research agenda is to build software systems that can serve as the infrastructure for creating and deploying Internet-based business applications (also referred to as e-business applications). The need for building such infrastructure systems arises because currently, in order to build e-business applications, application developers have to program against relatively low-level interfaces and write a lot of special purpose code. This plight of application developers is analogous to that of programmers in the early days of computing, who had to write assembly language programs without the aid of software systems such as compilers and operating systems. My research goal of building software infrastructure systems for e-business applications is thus motivated by the need to provide developers with higher levels of abstraction.

PUBLICATIONS

"Accessing Extra-database Information: Concurrency Control and Correctness." *Information Systems , An International Journal* 23(7):439-462 (1998). With Narain Gehani, Krithi Ramamritham, and Oded Shmueli.

"Efficiently Publishing Relational Data as XML Documents." In Proceedings of the Conference on Very Large Databases (VLDB), Cairo, Egypt (September, 2000). With Eugene Shekita, Rimon Barr, Michael Carey, Bruce Lindsay, Hamid Pirahesh, and Berthold Reinwald.

"Relational Databases for Querying XML documents: Limitations and Opportunities." In Proceedings of the Conference on Very Large Databases, Edinburgh, Scotland (September, 1999). With Kristin Tufte, Gang He, Chun Zhang, David DeWitt, and Jeffrey Naughton.

"Compressed Data Cubes for OLAP Aggregate Query Approximation on Continuous Dimensions." In Proceedings of the ACM SIGKIDD Conference on Knowledge Discovery and Data Mining, San Diego, CA (August, 1999). With Usama Fayyad, and Paul Bradley.

"Efficient Concurrency Control for Broadcast Environments." In Proceeding s of the ACM SIGMOD Cnference on the Management of Data, Pittsburgh, PA (May, 1999). With Arvind Nithrakashyap, Rajendran Sivasankaran and Krithi Ramamritham.

"Use of Recurrent Neural Networks for Strategic Data Mining of Sales Information." International Management Resources Association (IMRA) Conference, Hershey, PA (May, 1999). With Maram V. Nagendra Prasad, Sanjeev Vadhavkar, and Amar Gupta.

PATENTS

"Multi-dimensional Database and Data Cube Compression for Aggregate Query Support on Numeric Dimensions." United States patent application filed April 22, 1999.

"Using an XML Query Language to Publish Relational Data as XML." United States patent application filed March 21, 2000.

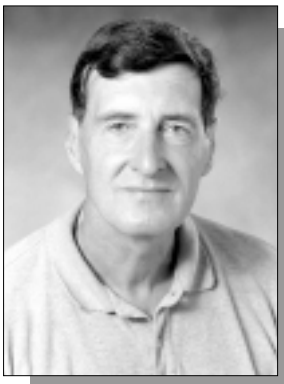

William Y. Arms

Professor

wya@cs.cornell.edu
D.Phil. University of Sussex
U.K., 1973

My research interests concentrate on web information systems, digital libraries and electronic publishing. These fields integrate methods from many disciplines, so that the work ranges from technical topics, such as distributed computing and information representation, to the economic and social aspects of change. My book *Digital Libraries* was published by the MIT Press in winter 2000.

This year we received a major grant to integrate many separate projects into the NSF's new digital library for science, mathematics, engineering and technology education. This is likely to be the largest and most heterogeneous digital library yet attempted. Cornell's multidisciplinary team combines computer science, librarian and user interfaces design expertise.

One of my principal interests is the change in scientific publication as online materials replace printed journals as the primary means of creating, storing, and distributing research information. I have recently completed a period as chair of the ACM Publications Board, am a member of the MIT Press Management Board, and am a member of a strategic planning committee of the American Physical Society.

As part of the NSF-funded Prism project, I am working with the Library of Congress to develop methods for long-term preservation of materials on the Web.

UNIVERSITY ACTIVITIES

Chair, Provost's Advisory Committee on Distance Education.
Director, eCornell.
Member of the Faculty Senate, the Faculty Advisory Board on Information Technology, and the University Library Board.

PROFESSIONAL ACTIVITIES

Publications Board, Association for Computing Machinery.
National Research Council study *Issues for Science and Engineering Researchers in the Digital Age*.
MIT Press, member of the Management Board and editor of the series on Digital Libraries and Electronic Publishing.

LECTURES

The Impact of the Internet on Research Universities. National Science Foundation (April 13, 2001).
The National Science Digital Library Program. Coalition for Networked Information (April 9, 2001).
Quality Control in Scholarly Publishing. What are the Alternatives to Peer Review? Keynote, workshop on the Open Archives initiative and peer review journals in Europe, Geneva (March 22 – 24, 2001).
Minerva: The Web Preservation Project. Library of Congress (February 2, 2001).
Strategies for Collecting and Preserving Open Access Materials on the Web. Federal Library and Information Center Committee, Washington D.C. (December 7, 2000).
The Web as an Open Access Digital Library. Closing address, 2000 Kyoto International Conference on Digital Libraries: Research and Practice, Kyoto, Japan (November 15, 2000).
Open Access to Digital Libraries. Must Research Libraries be Expensive? Keynote address, European Conference on Digital Libraries, ECDL2000, Lisbon, Portugal (September 18, 2000).

PUBLICATIONS

"Uniform Resource Names: Handles, PURLs, and Digital Object Identifiers." *Communication of the ACM*, 44(5):68 (May, 2001).
"Collecting and Preserving the Web: The Minerva Prototype." *RLG DigiNews* 5(2) (April 2001). With Roger Adkins, Cassy Ammen, and Allene Hayes. (http://www.rlg.org/preserv/diginews/diginews5-2.html#feature1)
"An Architecture for Reference Linking." Technical Report TR 2000-1820, Computer Science Department, Cornell University (October, 2000). With Donna Bergmark, and Carl Lagoze.
"Automated Digital Libraries. How Effectively can Computers be used for the Skilled Tasks of Professional Librarianship?" *D-Lib Magazine* 6(7/8) (July/August, 2000). (http://www.dlib.org/dlib/july00/arms/07arms.html)
"Economic Models for Open-access Publishing." *IMP* (March, 2000). (http://www.cisp.org/imp/march_2000/03_00arms.htm). *Digital Libraries*. MIT Press, ISBN 0-262-01180-8, 2000.

Graeme Bailey
Professor

bailey@cs.cornell.edu
Ph.D. University of Birmingham U.K., 1977

Originally working in low-dimensional topology and combinatorial group theory, through an odd mixture of circumstances I have become actively involved in research in mathematics and medicine. One of two ongoing research projects in this area is the modeling of lung inflation, together with a research group at the Class One Trauma Center at Upstate Medical Univ., Syracuse, NY. This is in the early stages of a program to extend to various pathologies affecting elasticity and aimed towards effective clinical treatments. The group, now having made some significant advances in answering questions that had remained unsolved for over 30 years, is now in the process of trying to obtain reliable mathematical models. This involves building computer simulations of dynamic packing results under constrained perturbations and deformations. The other project is in understanding deformations of transmembrane proteins used in cell-signaling processes. This is a carefully constrained version of the protein-folding problems that have been exciting the mathematical biology community in recent years; the application of a topological viewpoint in collaborating with molecular phmacologists and structural biologists has already yielded some intriguing insights.

HONORS/AWARDS

Kenneth A. Goldman '71 Excellence in Teaching Award, 2000.

UNIVERSITY ACTIVITIES

Adjunct Professor; Mathematics.
Member, Fellowship Selection Committees: Rhodes, Marshall, Churchill, and Fulbright.
Member, WCHI-Development and Transition Committee.
Member, Donlon Fellows Development.
Member, Master of Engineering Committee.
Member, Cornell EMS.
Faculty Advisor, Judo Club.

Kenneth P. Birman
Professor

ken@cs.cornell.edu
http://www.cs.cornell.edu/ken/
Ph.D. University of California, Berkeley, 1981

My research is concerned with reliability and security in modern networked environments. This work has three broad themes.

Our main focus is on a new system called "Spinglass" (http://www.cs.cornell.edu/Info/Projects/Spinglass). The idea is to explore a class of reliable multicast protocols that are extremely scalable and provide unusually stable throughput under stress. We believe that stable throughput is a common requirement in demanding critical settings, but few reliable protocols have this property. By scalability, we mean that a system which works with ten computers should also be usable with ten thousand of them.

Spinglass involves two subprojects. One, called Astrolabe, is concerned with a new way to represent data in a network. Astrolabe is like a network-wide database in which each computer or component contributes a live tuple. As data change, Astrolabe propagates the updates. The system uses a form of dynamically materialized view to continuously compute summaries of the picture of the network as a whole. This results in a powerful new tool for distributed monitoring, management, control, and live collaboration. Robbert van Renesse is the leader on this work, and we're collaborating with Al Demers and Johannes Gehrke on aspects related to databases and data mining.

The second big part of Spinglass is concerned with reliable multicast. We've developed a scalable multicast protocol that gives probabilistic consistency guarantees, and we are finding ways to apply this in practical settings. Graduate students looking at these questions include Indranil Gupta, who is developing algorithms that make direct use of probabilistic guarantees; Rimon Barr, who is looking at adapting these tools for mobile networks; and Ranveer Chandra and Venugopalen Ramasubramanian, who are

investigating ad-hoc mobile networking. Kate Jenkins and Ken Hopkinson are looking at applications that arise when using these protocols in real-world settings arising from the restructuring of the electric power grid.

Our third big activity is joint work with Bob Constable's Nuprl project, and involves the use of formal methods to prove properties of reliable communication protocols, such as those used in Isis, Horus, and Ensemble.

Our project is funded primarily by DARPA, with some additional funding from the Electric Power Research Institute and Microsoft Research. The project is directed by myself, R. van Renesse, and W. Vogels. R. Bhoedjang is visiting as a post-doc for a few years, and developing Intrusion Detection software to make use of Astrolabe.

HONORS

Stephen '57 and Marilyn Miles Excellence in Teaching Award, 2000.

UNIVERSITY ACTIVITIES

Committees: Founding Committee for Faculty of Computing and Information Science; University Conflicts of Interest; Chairman of the Responsible Conduct of Research Committee, Engineering College Policy Committee; IP Advisory Council for the Cornell Research Foundation.

LECTURES

Next Generation Internet: *Unsafe at any Speed?*

—. Keynote Speaker: ISDCS '01 (April, 2001).

—. University of Rochester (November, 2000).

—. *IBM* T.J. Watson Research Center (March, 2000).

—. Keynote: Middleware 2000.

PUBLICATIONS

"Technology Requirements for Virtual Overlay Networks." *IEEE Systems, Man and Cybernetics*: Special issue on Information Assurance (March, 2001).

"Next Generation Internet: *Unsafe at Any Speed?" IEEE Computer*, Special issue on Infrastructure Protection (Fall, 2000).

"Technology Challenges for Virtual Overlay Networks." *IEEE Systems, Man, and Cybernetics* Information Assurance and Security Workshop, West Point, New York (June 6-7, 2000).

"Optimized Group Rekey for Group Communications Systems." Networked and Distributed Systems Security 2000, San Diego, California. (V). Extended version available as Cornell University, Computer Science TR99-1764. With Ohad Rodeh, and Danny Dolev.

"A Dynamic Light-weight Group Service." *Journal of Parallel and Distributed Computing* 60:1449-1479 (2000). With Luis Rodrigues, Katherine Guo, and Paulo Verissimo.

"A Gossip Protocol for Subgroup Multicast." International Workshop on Applied Reliable Group Communication (WARGC 2001), Phoenix, AZ (April, 2001). With Kate Jenkins.

"Providing Efficient, Robust Error Recovery through Randomization." International Workshop on Applied Reliable Group Communication (WARGC 2001), Phoenix, AZ (April, 2001). With Zhen Xiao.

"Anonymous Gossip: Improving Multicast Reliability in Ad-hoc Networks." International Conference on Distributed Computing systems (ICDCS 2001), Phoenix, AZ (April, 2001). With Ranveer Chandra and Vanogupalen Ramasubramanian.

"A Randomized Error Recovery Algorithm for Reliable Multicast." IEEE Infocom 2001 AK (April, 2001). With Zhen Xiao.

"Using Epidemic Techniques for Building Ultra-scalable Reliable Communications Systems." Workshop on New Visions for Large-scale Networks: Research and Applications, Vienna, VA (March, 2001). With Werner Vogels, and Robbert van Renesse.

"Optimizing Buffer Management for Reliable Multicast." Submitted to the 2nd Annual Workshop on Networked Group Communication (NGC 2000), Palo Alto, CA (November 8-10, 2000). With Zhen Xiao, and Robbert van Renesse.

"Throughput Stability of Reliable Multicast Protocols." ADVIS' 2000, Dokuz Eylul University, Izmir, Turkey (October 25-27, 2000). With Oznur Ozkasap.

"A Probabilistically Correct Election Protocol for Large Groups." DISC 2000, Toledo, Spain (October 4-6, 2000). With Indranil Gupta, and Robbert van Renesse.

PUBLICATIONS - LANDMARK

"Bimodal Multicast." *ACM Transactions on Computer Systems* 17(2):41-88 (May, 1999). With M. Hayden, O. Ozkasap, Z. Xiao, M. Budiu, and Y. Minsky.

"A Probabilistically Correct Leadership Election Protocol for Large Groups." DISC-2000, Nov 2000, Toledo Spain. With I. Gupta and R. van Renesse.

Martin Burtscher

Assistant Professor
Member of the School of Electrical and Computer Engineering and the Graduate Field of Computer Science
burtscher@csl.cornell.edu
http://www.csl.cornell.edu/~burtscher/
Ph.D. University of Colorado at Boulder, 2000

My research interests are high-performance microprocessor architecture, instruction-level parallelism, and compiler optimizations. In particular, I am exploring hardware- and software-based value prediction, data compression, and latency reduction techniques.

The constantly widening speed gap between CPUs and memory is becoming more and more of a performance-limiting factor. In fact, current high-end microprocessors already spend a substantial amount of time waiting for memory accesses. To speed up program execution, the CPU needs to process useful instructions while waiting for the memory. One way of providing a processor with useful work is to predict what it will have to do next. Many commodity microprocessors already contain branch predictors to boost their performance, and it is likely that more predictors will be needed to meet the continuing demand for ever-faster CPUs. Designing, evaluating, and improving such predictors is an important focus of my research.

Ongoing projects include locating novel domains that can benefit from prediction, adding compiler support to aid and simplify the prediction hardware, devising means to reduce predictor sizes and power consumption without compromising performance, discovering as-of-yet unobserved patterns to build new predictors, and using value-prediction techniques to enhance branch-prediction accuracy and data-compression rates.

LECTURES

Designing a High-performance Load-value Predictor. Hewlett-Packard Company (December, 2000).

The Evolution of a High-performance Load-value Predictor. Lockheed Martin Corporation (November, 2000).

The Design of a High-performance Load Value Predictor. Compaq Computer Corporation (October, 2000).

Hybridizing and Coalescing Load-value Predictors. International Conference on Computer Design (September, 2000).

Predictability and Exploitability of Load Values. Microsoft Research (June, 2000).

PUBLICATIONS

"Hybridizing and Coalescing Load-value Predictors." International Conference on Computer Design, Austin, TX (September, 2000) :81-92. With B. Zorn.

"Exploring Last n Value Prediction." International Conference on Parallel Architectures and Compilation Techniques, Paris, France (October, 1999):66-76. With B. Zorn.

"Prediction Outcome History-Based Confidence Estimation for Load-value Prediction." *Journal of Instruction-level Parallelism*1 (May 1999). http://www.jilp.org/vol1/ With B. Zorn.

CS faculty on retreat

Claire Cardie

Associate
Professor

cardie@cs.cornell.edu
http://www.cs.cornell.edu/home/cardie/
Ph.D. University of Massachusetts
Amherst, 1994

My primary research areas are Natural Language Processing (NLP) and Machine Learning (ML) where we focus on developing corpus-based techniques for understanding and extracting information from natural language texts. In particular, my group investigates the use of machine learning techniques as tools for guiding natural language system development and for exploring the mechanisms that underlie language understanding. Our work encompasses three related areas: (1) machine learning of natural language, (2) the use of corpus-based NLP techniques to aid information retrieval (IR) and summarization systems, and (3) the design of user-trainable NLP systems that can efficiently and reliably extract the important information from a document.

In the past year or so we have made progress on both the natural language processing and machine learning aspects of our research. First, we have extended our approach to partial parsing of natural language texts to operate effectively in a weakly supervised learning framework. The original approach, developed with graduate students Scott Mardis and David Pierce, combines corpus-based grammar induction with a very simple pattern-matching algorithm and an optional constituent verification step. In evaluations on a number of large-scale partial parsing tasks involving on-line text, the approach produces partial parsers that are both fast and accurate.

Unfortunately, however, large amounts of expensive, human-annotated data are required for training. In new work with David Pierce, we investigate the use of weakly supervised learning algorithms for partial parsing that require only a small set of labeled training instances. In particular, we examine the learning behavior of co-training, a weakly supervised learning paradigm in which the redundancy of the learning task is captured by training two classifiers using separate views of the same data. This enables bootstrapping from a small set of labeled training data via a large set of unlabeled data. For noun phrase bracketing, we find that co-training reduces by 36 percent the difference in error between co-trained classifiers and fully supervised classifiers trained on a labeled version of all available data. However, degradation in the quality of the bootstrapped data arises as an obstacle to further improvement. To address this, we propose a moderately supervised variant of co-training in which a human corrects the mistakes made during automatic labeling. Our analysis suggests that corrected co-training and similar moderately supervised methods may help co-training scale to other natural language learning tasks that typically require large amounts of training data.

In a joint project with CoGenTex Inc. and the University of Montreal, we have begun to extend existing corpus-based learning algorithms for information extraction to acquire a broader set of extraction patterns. We have implemented a new learning algorithm, Autoslog-XML, that can extract linguistic entities beyond just noun phrases. Autoslog-XML is a semi-supervised algorithm for locating useful extraction patterns from unrestricted text. The learned extraction patterns have been employed to support domain-specific multi-document summarization in the natural disasters domain. We have also begun to investigate the application of the extraction pattern learning algorithm for Korean texts. Graduate students Vincent Ng and Kiri Wagstaff are part of this joint research effort.

The techniques described above can, in turn, be used to support a number of practical, end-to-end text-processing tasks in addition to multi-document summarization. For example, we are using the corpus-based partial-parsing techniques as the primary linguistic component in a new system for general-knowledge question answering. The system combines techniques for standard information retrieval, query-dependent text summarization, and shallow syntactic and semantic sentence analysis. In a series of experiments, we examined the role of each statistical and linguistic knowledge source in the question-answering system and find that even very weak linguistic knowledge can offer substantial improvements over purely IR-based techniques for question answering, especially when smoothly integrated with statistical preferences computed by the IR subsystems.

Finally, in machine learning research with Kiri Wagstaff, we are investigating the use of prior knowledge in the form of user-supplied constraints to improve the performance of clustering algorithms.

UNIVERSITY ACTIVITIES

College Scholar Advisor.
College Scholar Committee Chair.
Independent Major Advisor.
Cognitive Studies Concentration Advisor.

Member: Faculty of Computing and Information Founders; Faculty of Computing and Information Working Group for Information Science; Faculty of Computing and Information Working Group for Crosscutting Education Programs; Independent Major Advisory Board; Provost's Advisory Group of Women in Science and Engineering; Cognitive Studies Steering Committee; Computer Science Department Ph.D. admissions committee; Department of Computer Science space committee; Cognitive Studies Selection Committees for Summer Fellowships, for Continuing Fellowships, and for Incoming Fellowships.

PROFESSIONAL ACTIVITIES

Secretary: North American Association for Computational Linguistics (2000-2001).

Secretary: SIGNLL, Association for Computational Linguistics special interest group on Natural Language Learning (1999-2001).

Editorial Board: *Machine Learning* (1999-2001), Semantic Web Journal.

Action Editor: *Journal of Machine Learning Research* (2000-2002).

Program co-chair: Fourth Computational Natural Language Learning Workshop (CoNLL 2000).

Member: DARPA/NSF Question and Answering Roadmap Committee; DARPA/NSF Summarization Roadmap Committee; DARPA Translingual Information Detection, Extraction, and Summarization (TIDES) Evaluation Committee.

PROGRAM COMMITTEES:

Eighteenth International Conference on Machine Learning (ICML) (2001).

Seventeenth International Joint Conference on Artificial Intelligence (IJCAI) (2001).

First International Conference on Knowledge Capture (K-CAP) (2001).

39th Annual Meeting of the Association for Computational Linguistics (2001).

Second Meeting of the North American Chapter of the Association for Computational Linguistics (2001).

Fifth Computational Natural Language Learning Workshop (CoNLL) (2001).

Area Chair: The 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (2001).

Reviewer: Special issue of Computational Linguistics on anaphora resolution.

Executive Board: SIGDAT, Special Interest Group of ACL for Linguistic Data and Corpus-based approaches to NLP.

NSF Review Panel: Knowledge and Cognitive Systems (2000); Human-Computer Interaction (2001).

LECTURES

Noun Phrase Co-reference for Information Extraction. Logic and AI Seminar, University of Maryland (April, 2001).

Machine Learning for Information Extraction from Unrestricted Text. Alan Perlis Symposium, Yale University (April, 2001).

Rapidly Portable Translingual Information Extraction and Interactive Multi-document Summarization. DARPA TIDES Meeting (February, 2001).

Overview of Cornell University Projects in Natural Language Understanding. America On-line (December, 2000).

PUBLICATIONS

"Limitations of Co-training for Natural Language Learning from Large Datasets." Proceedings of the 2001 Conference on Empirical Methods in Natural Language Processing (EMNLP-2001), 1-9. Association for Computational Linguistics (2001). With David Pierce.

"Multi-document Summarization via Information Extraction." Proceedings of the First International Conference on Human Language Technology Research (2001). With Michael White, Tanya Korelsky, Vincent Ng, David Pierce, and Kiri Wagstaff.

"Issues, Tasks and Program Structures to Roadmap Research in Question & Answering" (Q&A). DARPA/NSF (2000). With J. Burger, V. Chaudhri, R. Gaizauskas, S. Harabagiu, D. Israel, C. Jacquemin, C. Lin, S. Maiorano, G. Miller, D. Moldovan, B. Ogden, J. Prager, E. Riloff, A. Singhal, R. Shrihari, T. Strzalkowski, E. Voorhees, and R. Weischedel.

"Using Clustering and SuperConcepts within SMART: TREC 6." *Information Processing and Management* 36(1):109-131 (2000). With C. Buckley, M. Mitra, and J. Walz.

"A Cognitive Bias Approach to Feature Selection and Weighting for Case-based Learners." *Machine Learning* 41:85-116 (2000).

"Examining the Role of Statistical and Linguistic Knowledge Sources in a General-knowledge Question-answering System." Proceedings of the Sixth Applied Natural Language Processing Conference (ANLP-2000), 180-187. Association for Computational Linguistics / Morgan Kaufmann (2000). With V. Ng, D. Pierce, and C. Buckley.

"Clustering with Instance-level Constraints." Proceedings of the Seventeenth International Conference on Machine Learning, 1103-1110. Morgan Kaufmann (2000). With Kiri Wagstaff.

"Integrating Case-based Learning and Cognitive Biases for Machine Learning of Natural Language." *Journal of Experimental and Theoretical Artificial Intelligence* 11:297-337 (1999).

"The Role of Lexicalization and Pruning for Base Noun Phrase Grammars." Proceedings of the Sixteenth National Conference on Artificial Intelligence, 423-430. AAAI Press / MIT Press (1999). With David Pierce.

"Noun Phrase Co-reference as Clustering." Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Very Large Corpora, 82-89. Association for Computational Linguistics (1999). With Kiri Wagstaff.

Error-driven Pruning of Treebank Grammars for Base Noun Phrase Identification." Proceedings of the Annual Conference of the Association for Computational Linguistics and COLING-98, 218-224. Association for Computational Linguistics (1998). With David Pierce.

"Empirical Methods in Information Extraction." *AI Magazine* 18(4):65-79 (1997).

[On sabbatic leave 2001-2002]

L. Paul Chew

Senior Research Associate

chew@cs.cornell.edu

http://www.cs.cornell.edu/Info/People/chew/chew.html

My primary interest is in geometric algorithms with an emphasis on practical applications. These practical applications have included placement, motion planning, shape comparison, vision, sensing, mesh generation, molecular matching, and protein shape-comparison.

In recent work on protein shape, Klara Kedem and I have developed a new kind of "consensus shape" for protein families. This is an analog of the consensus string that is sometimes used for multiple alignment of proteins. The idea is based, in part, on our previous work on the Unit-vector Root Mean Square (URMS) distance for protein shapes. The consensus shape is a pseudo-protein with useful properties. It is a pseudo-protein because it fails to have certain characteristics of real proteins. In particular, for the consensus shape, the spacing between successive alpha carbons is variable, with small distances in regions where the members of the protein family exhibit significant variation and large distances (up to the standard spacing of four angstroms) in regions where the family members agree. Despite this nonprotein-like characteristic, the consensus shape does preserve structural information. If all members of a protein family exhibit a geometric relationship between corresponding alpha carbons then that relationship is preserved in the consensus shape. In particular, distances and angles that are consistent across family members are preserved. Thus, the consensus shape provides a compact summary of the significant strucutural information for a family. We are exploring the use of the consensus shape (1) as a tool for improved protein threading for use in protein structure prediction and (2) as a tool for automating the division of proteins into families and subfamilies.

My work on mesh generation has been motivated by the finite element method for finding approximate solutions to partial differential equations. The first step of this method is to create a mesh, i.e. to divide the given problem region into simple shapes called elements (usually triangles or quadrilaterals in 2D, tetrahedra or hexahedra in 3D). A number of algorithms have been developed to automate this process, but most of them don't guarantee the quality of the resulting mesh (e.g., a triangle may cross a region boundary or there may be some flat triangles, leading to poor error bounds). I developed efficient techniques for producing meshes of guaranteed quality for problems in the plane and for curved surfaces. The triangles produced are close to equilateral in shape; all region boundaries are respected; and the user can control the element density, producing small elements in "interesting" regions and large elements elsewhere.

I extended this work to produce tetrahedral meshes for three-dimensional problems. The major difficulty here is to avoid producing "slivers": tetrahedra with nicely shaped faces but with near-zero volume. I showed that slivers can be avoided by choosing each new mesh point randomly within a small specified volume. The randomness helps; a good mesh point—one that does not form any slivers—can be found in constant expected time.

This work is being used in a large, multi-disciplinary project: developing adaptive software for field-driven simulations. In particular, we focus on computational fracture

mechanics and reactive, multiphase fluid flows. Our goal is to develop principles for building software systems that can adapt to changing conditions. These conditions include changes in the desired physics (e.g. we may need to change our physics model when we discover that vibration is significant), changes in the desired algorithms (e.g. we may change our solution technique depending on how quickly we are converging toward a solution), and changes in the computing environment (e.g. additional processors may become available or we may lose processors due to processor failure). Other Cornell researchers working on this project are Keshav Pingali, Steve Vavasis, Paul Stodghill, and Tony Ingraffea (Civil Engineering), along with participants at Mississippi State University, Ohio State University, Clark-Atlanta University, and the College of William & Mary.

PROFESSIONAL ACTIVITIES

Associate Editor: *Pattern Recognition: the Journal of the Pattern Recognition Society*.

Referee: *Journal of the Association for Computing Machinery*, *SIAM Journal on Scientific Computing*.

PUBLICATIONS

"Parallel FEM Simulation of Crack Propagation-challenges, Status, and Perspectives." Proceedings of Irregular 2000 (2000). With B. Carter, C.S. Chen, G. Heber, A.R. Ingraffea, R. Krause, C. Myers, P.A. Wawrzynek, K. Pingali, P. Stodghill, S. Vavasis, N. Chrisochoides, D. Nave, and G.R. Gao.

"Fast Detection of Common Geometric Substructure in Proteins." *Journal of Computational Biology* 6(3):313-325 (1999). With D. Huttenlocher, K. Kedem, and J. Kleinberg.

"Unit-vector RMS (URMS) as a Tool to Analyze Molecular Dynamics Trajectories." *Proteins: Structure, Function and Genetics* 37, 554-564 (1999). With K. Kedem and R. Elber.

Thomas F. Coleman

Professor
Director: Cornell Theory Center
Director: Financial Industry Solutions Center (FISC)

coleman@cs.cornell.edu
http://www.tc.cornell.edu/~coleman/
Ph.D. University of Waterloo, 1979

Our research is concerned with the design and understanding of practical and efficient numerical algorithms for continuous optimization problems. Our primary emphasis is the development of algorithms for large-scale optimization, especially as applied to the area of computational finance.

With colleagues Yuying Li, Peter Mansfield, Arun Verma, and Shirish Chinchalkar, we are developing a variety of tools and methods for computational finance in the areas of portfolio management and options pricing (and hedging). Several Ph.D. students in the Center for Applied Mathematics are also involved in this work: Cristina Patron, Yohan Kim, and Changhong He. Yohan Kim completed his Ph.D. dissertation in May, 2001: *Estimation of Smooth Volatility Functions in Option Pricing Models*.

UNIVERSITY ACTIVITIES

Director: Cornell Theory Center.

Director: Financial Industry Solutions Center (55 Broad Street, NYC).

Member: Engineering Dean Search Committee; Cornell Task Force on Genomics; Program Committee for the Center for Applied Mathematics.

PROFESSIONAL ACTIVITIES

Chair: SIAM Activity Group on Optimization.

Co-organizer: The 11th Annual Derivatives Securities Conference, New York (April 27-28).

Organizer: FISC Spring 2000 Workshop Series.

Program Committee: Automatic Differentiation 2000, INRIA, France.

Member: Advisory Board, Brookhaven Center for Data Intensive Computing; Scientific Program Committee.

Member: Organizing committee for the thematic year at the Fields Institute: Numerical and Computational Challenges in Science and Engineering (2001-2002).

Organizer: Proposal for an IMA Special Year on Optimization, Minneapolis, MN (2002-2003); Seventh SIAM Conference on Optimization, Toronto, Ontario (2002).

Editorial Board: *Applied Mathematics Letters; SIAM Journal of Scientific Computing; Computational Optimization and Applications*, Comm. on Applied Non-linear Analysis, Mathematical Modeling and Scientific Computing.

Referee/Reviewer: *Mathematical Programming; Computational Optimization and Applications; SIAM Journal of Optimization; SIAM Journal of Scientific Computing;* Department of Energy, NSF.

LECTURES

Dynamic Hedging and a Deterministic Local Implied Volatility Function. Eleventh Annual Derivatives Conference, New York (April 28, 2001).

RISK 2001 Europe, Paris France (April 10-11, 2001).

Introduction to Computational Finance in Matlab: A 1-day workshop at FISC-Japan, Tokyo (March 22, 2001).

A Newton Method for Option Valuation. U. Singapore, Singapore (December 12, 2000).

—. City University of Hong Kong, Hong Kong (December 22, 2000).

PUBLICATIONS

"ADMIT-1: Automatic Differentiation and MATLAB Interface Toolbox." *ACM Transactions on Mathematical Software* 22:150-175 (2000). With Arun Verma.

"Efficiency Improvements for Pricing American Options with a Stochastic Mesh: Parallel Implementation." *Financial Engineering News*, 1-2 (December, 2000). With Thanos Avranidis, Yuriy Zinchenko, and Arun Verma.

[On sabbatic leave 2001-2002]

Robert L. Constable

Professor
Dean for Computing and Information Science

rc@cs.cornell.edu
http://www.cs.cornell.edu/home/rc/
Ph.D. University of Wisconsin-Madison, 1968

I direct the PRL research group. For the past three years we have been building a system that we call a *logical programming environment* (LPE). It provides substantial automation in the design, coding, verification, and evolution of large software systems. Generally an LPE will integrate programming languages and logics. In our case we integrate the ML programming language and a programming logic based on type theory. Reasoning about ML programs is founded on type theoretic semantics for ML. The LPE also integrates a compiler and a theorem prover. We use the latest version of Nuprl as the prover.

Over the past year we have deployed our prototype LPE to support the Ensemble group communication system that Ken Birman, Robbert van Renesse, and their students have built. The LPE provides automatic code transformations that improve performance in a way that is guaranteed not to introduce errors. The LPE also supports the design and coding of new *adaptive protocols* that are part of the Spinglass project.

The Nuprl 5 system is a major re-implementation of Nuprl 4; its design is based on communicating processes that synchronize around a logical database that we call "the Library." The Library stores definitions, theorems, conjectures, proofs, algorithms, proof tactics, and commentary that is linked to the formal mathematics. We released the first version of Nuprl 5 last summer, with a debut at CADE in June involving the whole design and implementation team (Stuart Allen, Rich Eaton, Christoph Kreitz, Lori Lorigo, and me). Nuprl 5 supports multiple proof engines and multiple editors. It also integrates an automatic theorem prover, called the JProver, built by C. Kreitz, Jens Otten, and Stephan Schmitt.

The Library includes over ten thousand theorems. Many of these are used in system verification, but a large number are from general mathematics. These general theorems are a valuable resource. We are funded by ONR to further develop and explore the concept of a formal digital library of formal constructive mathematics built around these theorems. We will generalize our mechanisms to allow many users to contribute to the Library using other theorem provers such as ACL2, HOL and PVS.

We are also working on a more experimental LPE called MetaPRL, which started with Jason Hickey's thesis and now involves many people but especially Aleksey Nogin and Alexsey Kopylov. This system is built entirely in OCaml and supports OCaml as its programming language. It is coded for efficiency as well as modularity, and at some tasks it is over two orders of magnitude faster than Nuprl 5. It also supports multiple programming logics. Nuprl 5 and MetaPRL share mathematics libraries.

The two theorem provers are used in a variety of other projects as well, including the creation of formal courseware by S. Allen, the translation of formal proofs into natural language by Amanda Holland-Minkley, the automatic analysis of the computational complexity of higher-order programs by Ralph Benzinger, and efficient reflection being designed and implemented by Eli Barzilay. We follow the work of Greg Morrisett and his students on new ML compilers and on typed assembly language. We plan to use the LPE to broadly support research on language-based security in the department and at the new Information Assurance Institute, including the work of Dexter Kozen, Andrew Myers, and Fred Schneider.

UNIVERSITY ACTIVITIES

Dean: Computing and Information Science.

Committees: Applied Math Policy; Cognitive Studies Executive Committee; Theory Center Executive Committee.

PROFESSIONAL ACTIVITIES

Advisory Council: Princeton University Department of Computer Science.

Editor: *Journal of Logic and Computation*; *Formal Methods in System Design*; *Journal of Symbolic Computation*.

Director: NATO Summer School, Marktoberdorf, Germany.

General Committee Member: LICS.

LECTURES

How Computers Think. Dean's Forum, Faculty of Natural Sciences, Ben-Gurion University (June, 2001).

Formal Complexity Classes: An Approach to Automating Computational Complexity Analysis. Cornell University (May, 2001). With Ralph Benzinger.

How Nuprl Reasons. University of Delaware (May, 2001).

Taking a MEGABYTE: Cornell in the Information Age. University of Delaware (May, 2001).

How Nuprl Reasons. Yale University (February, 2001).

Taking a MEGABYTE: Cornell in the Information Age. Cornell University (February, 2001).

Representing the Faculty of Computing and Information. Cornell University (September, 2000).

How Computers Think. Cornell University (September, 2000).

Computer Science: Achievements and Challenges Circa 2000. Cornell University (March, 2000).

PUBLICATIONS

"An Experiment in Formal Design Using Meta-properties. In Proceedings of DARPA Information Survivability Conference and Exposition II (DISCEX 2001), IEEE Computer Society Press (June, 2001). With C. Kreitz, M. Bickford, and R. van Renesse.

"Protocol Switching: Exploiting Meta-properties." In Proceedings of the International Workshop on Applied Reliable Group Communication (WARGC 2001) (April, 2001). IEEE Computer Society Press. With C. Kreitz, X. Liu, R. van Renesse, and M. Bickford.

"Computational Complexity and Induction for Partial Computable Functions in Type Theory." In *Reflections: A Collection of Essays in Honor of Solomon Feferman*, Association for Symbolic Logic (2001). With K. Crary.

"Constructively Formalizing Automata." In *Proof, Language and Interaction: Essays in Honour of Robin Milner*, MIT Press, 213-238 (2000). With P.B. Jackson, P. Naumov, and J. Uribe.

"Nuprl's Class Theory and its Applications." In *Foundations of Secure Computation*, F.L. Bauer and R. Steinbruggen, editors. IOS Press, 91-115 (2000).

"Types in Logic, Mathematics and Programming." In *Handbook of Proof Theory*, S. R. Buss, editor. Elsevier Science B. V., 684-785 (1998).

"The Structure of Nuprl's Type Theory." In *Logic of Computation*, Springer-Verlag, (1997).

Implementing Mathematics with the Nuprl Development System. Prentice-Hall (1986). With S. Allen, et. al.

CIS Dean Robert L. Constable with Department of Computer Science Chairman, Charles Van Loan

Alan J. Demers

Professor

ademers@cs.cornell.edu
Ph.D. Princeton University, 1975

My current research concerns aspects of weakly-consistent data replication in databases and distributed systems.

With Ken Birman, Robert van Renesse, Johannes Gehrke, and others, I am studying randomized "gossip protocols." Such protocols are highly fault-tolerant and, when properly designed, extremely scalable as well. We are studying convergence properties of several flat and hierarchal versions of the basic protocols tailored to specific application requirements.

My particular focus is approximate evaluation of aggregate queries in such a system. We are studying age distributions of gossiped data in order to prove probabilistic bounds on the quality of aggregate query results. Alternatively, we can use this approach to bound the latency required to probabilistically guarantee a client-specified degree of consistency.

We are considering classes of "resource location" or "anomaly detection" problems, in which query results are site-specific and depend on distance and subsumption. Such problems can have efficient and highly scalable solutions using gossip partner choice distributions based on the distance between sites.

Finally, we are studying graph constructions for which flooding or deterministic gossip partner choices can be used, leading to reduced overhead while still retaining most of the desirable properties of randomized gossip.

The above is related to my previous work on the Clearinghouse and Bayou projects at Xerox PARC. I am also doing work supported by Oracle on asynchronous update-anywhere replication in a more traditional database setting. This involves algorithms for scheduling/reordering update propagation between sites to improve throughput while preserving eventual consistency and bounded inconsistency during propagation.

RECENT PAPERS

"Logarithmic Harary Graphs." ICDCS International Workshop on Applied Reliable Group Communication, Phoenix, Arizona (April, 2001). With K. Jenkins.

"Spatial Gossip and Resource Location Protocols." Proceedings of the 33d ACM Symposium on Theory of Computing, Crete (July, 2001). With D. Kempe, and J. Kleinberg.

Ron Elber

Professor

ron@cs.cornell.edu
http://www.cs.cornell.edu/ron/old_webpage-ron/index_ron.html
Ph.D. Hebrew University

My research is in the field of Computational Molecular Biology. We develop computer algorithms to study sequences, structures, dynamics, and function of proteins and apply these methods to a variety of biological problems. Our techniques are implemented in the systems MOIL and LOOPP, available on the web http://www.tc.cornell.edu/CBIO.

My current research directions include: mean field approaches for global optimization and structure prediction (Locally Enhanced Sampling). Structures are often determined by an optimization of an energy function. I introduced mean field approaches that modify the target function and make it more accessible to global optimization. We have applied these techniques to determine conformations of short peptides and to refine low-resolution structures of proteins. These approaches are implemented into MOIL.

We are also working on development of folding potentials using linear programming. An ideal folding potential assigns the lowest energy to the correct three-dimensional structure of a protein. All other structures must have higher energies. The design of folding potentials relies on considerable human intuition and many trials and errors. I developed an automated protocol that "learns" and improves the quality of the current potential energy. We used this protocol to prove that the widely used pairwise interaction

model cannot recognize exactly correct protein folds. Based on these studies, a novel threading algorithm was designed and implemented in the program LOOPP. In a threading algorithm a sequence is matched with a structure. In a recent article in *Science*, we published an intriguing application of this program. We suggested an evolutionary link between a gene that controls the size of the tomato fruit and a protein that participates in controlling cell growth and division. Malfunction of this protein causes cancer in humans (joint work with Steve Tanksley's group).

Another project concerns extending the time scale of simulations. One of the striking observations in dynamics of biological molecules is the extremely large time scale they covered. Initiation by light absorption of biochemical processes is very rapid (femtoseconds), while protein folding is slow (milliseconds to minutes). Current simulation approaches (Molecular Dynamics (MD)) are restricted to nanoseconds (10-9 seconds). I developed a stochastic path integral formulation that provides a numerically stable trajectory for almost any arbitrary time step. We apply the new algorithm to study activation of proteins (the R->T transitions in hemoglobin, microseconds) and to protein folding (folding of C peptide, tens of nanoseconds). The method provides systematic approximations to the dynamics and is more efficient than MD by orders of magnitude. It is available in MOIL.

PROFESSIONAL ACTIVITIES

Acting head: NIH resource for parallel computing at the Cornell Theory Center.

Committees: Statistical and Computational Genomics Committee; Computational Biology Committee for the collaborative efforts at Cornell, Rockefeller, and Sloan Kettering Institutes; Cornell Life Science Advisory Board; Planning Committee for Life Science and Technology Building; Theory Center committee.

National Committees: NIH study sections; NSF study section; Reviewer for the State of Texas; Chair: workshop on protein dynamics, Telluride (August, 2001).

LECTURES

Protein Recognition by Threading. DIMACS, Rutgers (March, 2001).

Parallel Computations of Trajectories. SIAM (March, 2001).

Long Time Dynamics of Proteins. University of Pennsylvania (February, 2001).

Long Time Dynamics and Protein Recognition by Threading. IBM Watson (February, 2001).

Long Time Dynamics of Biomolecules. Florida State, Computational Biophysics (January, 2001).

Protein Recognition by Threading. CUNY (December, 2000).

Protein Recognition by Threading. University of Maryland (October, 2000).

Long Time Dynamics of Biomolecules. NYU (October, 2000), M3.

PUBLICATIONS

"Protein Recognition by Sequence-to-structure Fitness: Bridging Efficiency and Capacity of Threading Models." Submitted to *Advances in Chemical Physics*, by invitation. With Jaroslaw Meller.

"The Enzymatic Circularization of a Malto-octaose Linear Chain Studied by Stochastic Reaction Path Calculations on Cyclodextrin Glycosyltransferase." *Proteins, Structure, Function and Genetics* 43:327-335 (2001). With Joost C.M. Uitdehaag, Bart A. van der Veen, Lubbert Dijkhuizen, and Bauke W. Dijkstra.

"Cloning, Transgenic Expression and Function of fw2.2: a Quantitative Trait Locus Key to the Evolution of Tomato Fruit." *Science* 289:85-88 (2000). With A. Frary, C. Nesbitt, A. Frary, S. Grandillo, E. van der Knaap, B. Cong, J. Liu, J. Meller, K.B. Alpert, and S.D. Tanksley.

"Distance Dependent Pair Potential for Protein Folding: Results from Linear Optimization." *Proteins, Structure, Function and Genetics* 41:40-46 (2000). With D. Tobi.

"Probing the Role of the Local Propensity in Peptide Turn Formation." *International Journal of Quantum Chemistry* 80:1125-1128 (2000). With D. Mohanty, and D. Thirumalai.

Professor Ken Birman, Associate Professor Bart Selman, and Professor Keshav Pingali

Geri Gay

Professor
FCI, joint with Communication

gkg1@cornell.edu
http://www.comm.cornell.edu/faculty/gay.html
Ph.D. Cornell University, 1985

I am the director of the Human Computer Interaction Group (HCI Group) and a professor of communication at Cornell University. The HCI Group is a research and development group whose members design and research the use of computer-mediated learning environments. My research interests focus on cognitive and social issues for the design and use of interactive communication technologies. Past research has explored navigation issues, knowledge management, mental models and metaphors, knowledge representations, collaborative work and learning, and system design.

I have received funding for my research and design projects from the National Science Foundation (NSF), the National Endowment for the Humanities (NEH), the Mellon Foundation, Intel, Microsoft, IBM, Getty, and several private donors. I teach courses in interactive multimedia design and research, computer-mediated communication, human-computer interaction, and the social design of communication systems.

HONORS/AWARDS

NYS Chancellor's Award for Excellence 2001.
Innovative Teaching Award 2000.

CURRENT PROJECTS

NASA and AT&T Advanced Technology for Learning projects.
Intel Museum Context Aware Computing Project
Intel and NSF studies on the use of wireless computing, covered in *Chronicle of Higher Education*, *USA Today*, *Newsweek*, *The New York Times*, *Globe and Mail*, and NPR.

LECTURES

American Educational Research Association, International Communication Association; ACM Multimedia; Japanese Private University Association.

PUBLICATIONS

Articles in *Journal of Computer-Mediated Communication*; *Journal of Research on Computing in Education*; *Journal of Educational Technology*; *Journal of Educational Computing Research*; *Journal of Educational Psychology*; *International Journal of HCI*; *ACM Digital Libraries*; *Journal of Information Technologies*.

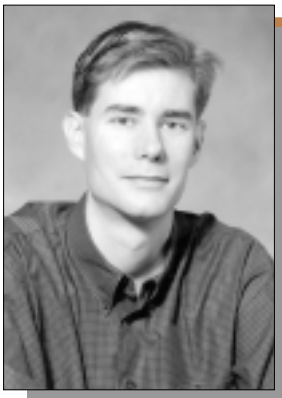

Johannes Gehrke

Assistant Professor

johannes@cs.cornell.edu
http://www.cs.cornell.edu/johannes/
Ph.D. University of Wisconsin-Madison, 1999

My primary research interest is in the development of new data mining and database technology. My group is currently involved in three projects: The Himalaya Data Mining Project, the Cougar Sensor Database System, and the Amazon Stream Processing Project.

In the Himalaya Data Mining Project we develop new data mining functionality, and we work on techniques to make the resulting data mining models more understandable to the user. As an example, consider classification trees, a data mining model that is supported in nearly all commercial data mining suites. In recent research we have shown that a large class of classification tree construction algorithms is biased (including most algorithms used in commercial tools), thus, users could draw incorrect conclusions from the resulting "incorrect" classification tree. Our methods can provably eliminate this bias from any existing split selection method. Other recent results include the fastest published algorithm for mining long market baskets, and new methods for mining long sequential patterns.

The widespread deployment of sensors and mobile devices is transforming our physical environment into a computing platform. There is now computing power on every device, and emerging networking techniques ensure that devices are interconnected and accessible from local or wide-area networks. This is a distributed database system of unprecedented scale. In the Cougar Sensor Database

System, we develop database technology for tasking, mining, and monitoring such a large number of distributed data sources. We have implemented the first generation of the Cougar Device Database System, where we leverage the processing power on the devices to push query processing directly to the data sources. Different query processing strategies allow us to balance resource usage, accuracy, and speed of query answers. Our current research focuses on distributed and fault tolerant query processing and meta-data management.

In many applications, for example in intrusion detection, sensor networks, and network management, data arrive in streams, and the large volume of such high-speed data streams makes storage and offline processing of the data infeasible. In the Amazon Stream Processing Project, we are developing query processing techniques for long running queries over infinite data streams. The main difficulty here is the new model of computation: Instead of being able to re-read data many times and to perform expensive offline computation on a static dataset, we need to compute query answers and maintain summary statistics in an online fashion. Our recent results include computation of correlated aggregates and quantiles over data streams.

HONORS/AWARDS

IBM Faculty Development Award (2000, 2001).
James and Mary Tien Excellence in Teaching Award (2001).

UNIVERSITY ACTIVITIES

Member: Space Committee, Department of Computer Science; Faculty Search Committee, School of Operations Research and Industrial Engineering; faculty Search Committee, The Computational and Statistical Genomics Trust.

PROFESSIONAL ACTIVITIES

Program Committees:
Twenty-sixth International Conference on Very Large Databases (VLDB), Cairo, Egypt (September, 2000). Demonstrations Committee.
Seventeenth IEEE International Conference on Data Engineering (ICDE 2001), Heidelberg, Germany (April 2001).
Fourth International Workshop on Parallel and Distributed Data Mining. Workshop held in conjunction with the 15th International Parallel and Distributed Processing Symposium, San Francisco, CA (April, 2001).
Twentieth ACM SIGMOD Conference (SIGMOD 2001), Santa Barbara, CA (May, 2001).
ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery (DMKD 2001) held in cooperation with SIGMOD 2001. Santa Barbara, CA (May, 2001). Workshop Co-chair.
Eighteenth International Conference on Machine Learning (ICML 2001), Williams College, MA (June, 2001).
Sixth ACM SIGKDD Conference (KDD 2001), San Jose, CA (August, 2001).
Twelfth International Conference on Software Engineering and Knowledge Engineering, Chicago, IL (Ju;ly, 2000).
Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2000). Boston, MA (August, 2000).
Editorial Boards:
Knowledge and Information Systems.
Journal of Database Management.

LECTURES

An Introduction to Data Mining. Air Force Research Laboratory, Rome, NY (September 12, 2000).
An Overview of Modern Data Mining Technology. Workshop at the Financial Industry Solutions Center (FISC) New York (November 8, 2000).
The Infrastructure of Electronic Commerce. Lectures on Database Technology and Data Mining. Johnson Graduate School of Management, Cornell University (January, 2001).
Honest Classification Trees. IBM Watson Research Center, Yorktown, NY (March, 2001).
Panel Manager for "Storage—A Crowded Place?" Panel at the 2001 Leadership in the Technology Marketplace Symposium, Cornell Johnson Graduate School of Management, Ithaca (April, 2001).
Querying the Physical World. DARPA Sensor Information Technology PI Meeting. St. Petersburg, FL (April, 2001).
Mining Very Large Databases. Invited talk at the 33rd Symposium on the Interface of Computing Science and Statistics, Costa Mesa, CA (June, 2001).

PUBLICATIONS

"RAINFOREST - A Framework for Fast Decision Tree Construction of Large Datasets. In *Data Mining and Knowledge Discovery* 4(2/3):27-162 (July, 2000). With Raghu Ramakrishnan, and Venkatesh Ganti.
"Querying the Physical World." In *IEEE Personal Communications*, special issue on Smart Spaces and Environments (October, 2000). With Philippe Bonnet, and Praveen Seshadri.

"Towards Sensor Database Systems." In Proceedings of the Second International Conference on Mobile Data Management, Hong Kong, China (January, 2001). With Philippe Bonnet, and Praveen Seshadri.

"DEMON: Mining and Monitoring Evolving Data." In *IEEE Transactions on Knowledge and Data Engineering* 13(1):50-63 (January/February 2001). With Venkatesh Ganti, and Raghu Ramakrishnan.

"MAFIA: A Maximal Frequent Itemset Algorithm for Transactional Databases." In Proceedings of the 17th International Conference on Data Engineering, Heidelberg, Germany (April, 2001). With Doug Burdick, and Manuel Calimlim.

"On Computing Correlated Aggregates Over Continual Data Streams." In Proceedings of the 2001 ACM SIGMOD International Conference on Management of Data, Santa Barbara, CA (May, 2001). With Flip Korn, and Divesh Srivastava.

"Query Optimization In Compressed Database Systems." In Proceedings of the 2001 ACM SIGMOD International Conference on Management of Data, Santa Barbara, CA (May, 2001). With Zhiyuan Chen, and Flip Korn.

"Bias Correction in Classification Tree Construction." In Proceedings of the Seventeenth International Conference on Machine Learning (ICML 2001), Williams College, MA (June, 2001). With Alin Dobra.

Carla Gomes

Research Associate
Director, Intelligent Information Systems Institute (IISI)

gomes@cs.cornell.edu
http://www.cs.cornell.edu/gomes
Ph.D. University of Edinburgh, 1993

My research interests are centered around the integration of methods from artificial intelligence and operations research for solving hard combinatorial problems. I consider applications in areas ranging from combinatorial design, planning and scheduling, reasoning, multi-agent systems, and machine learning. Recently, I have focused on randomized search techniques. In this work, I study so-called heavy-tailed distributions that characterize complete randomized search methods. A promising way of exploiting heavy-tailed behavior is by combining a suite of search methods into a portfolio, running on a distributed compute cluster. It can be shown that such portfolios dramatically reduce the expected overall computational cost, thereby allowing us to solve large, previously unsolved combinatorial problems. Another recent research direction (joint work with groups at the University of Washington and Microsoft Research) involves the use of machine learning techniques and Bayesian models to develop effective adaptive algorithmic strategies given bounded computational resources.

I also established and direct the newly formed Intelligent Information Systems Institute (IISI) at Cornell. The mission of the institute is to foster research in computation and data intensive methods for intelligent decision making systems. See www.cis.cornell.edu/iisi.

PROFESSIONAL ACTIVITIES

Director, Intelligent Information Systems Institute (IISI).

Guest Editor, *Journal of Knowledge Engineering Review*, Cambridge Press.

Guest Editor, *Artificial Intelligence Journal*.

Editorial Board, *Journal of Knowledge Engineering Review*.

Editorial Board, *International Journal AI Tools (IJAIT)*.

Co-chair, AAAI Symposium on Uncertainty in Computation, Boston, MA (2001).

Co-chair, AAAI Workshop on Leveraging Probability and Uncertainty in Computation, AAAI (2000).

Member, Advisory Committee International Scientists, Ministry of Science and Technology, Portuguese Government, Presidency of European Union (2000).

Member, Program Committee, SAT, Boston, MA (2001).

External Examiner, Ph.D. Thesis of Ramon Bejar, Univ. of Leida, Barcelona, Spain.

Reviewer for 7th Int. Joint Conference on Artificial Intelligence (IJCAI); *Journal of Automated Reasoning*; *Journal of Artificial Intelligence Research*; *Constraints: An International Journal*; *Discrete Applied Mathematics*.

LECTURES

Structure, Duality, and Randomization: Common Themes in AI and OR.

—. Broad Area Colloquium, Stanford University, Stanford, CA (November, 2000).

—. AI Seminar, SRI, Palo Alto, CA (November, 2000).

—. Research Seminar, NASA/Ames, Mountain View, CA (November, 2000).

—. Colloquium, Univ. of Lisbon, Lisbon, Portugal (March, 2001).

Survey of Information Retrieval and Knowledge Representation, 3 lectures at AFRL/IF, Rome, NY (November, 2000).

Vision and Directions for the Intelligent Information System Institute, AFRL/IF, Rome, NY (February, 2001).

Impact of Structure on Complexity. MURI /AFOSR meeting on Coop. Control of Distributed Autonomous Vehicles in Adversarial Environments, UCLA, Los Angeles, CA (May, 2001).

Structure and Complexity in the Virtual Transportation Company. Meeting on Taskable Agent Software Kit, Darpa, Sante Fe, NM (April, 2001).

Controlling Computational Cost. Meeting on Autonomous Negotiation Teams, DARPA, Lake Tahoe, CA (May, 2001).

PUBLICATIONS

"On the Intersection of AI and OR." *Journal of Knowledge Engineering Review* 16(1) (2001).

"Algorithm Portfolios." *Artificial Intelligence Journal* 126(2001). With B. Selman.

"A Bayesian Approach to Tackling Hard Computational Problems." Proc. 17th Conf. on Uncertainty and Artificial Intelligence (UAI-2001), Seattle, WA (2001). With E. Horvitz, Y. Ruan, H. Kautz, B. Selman, and M. Chickering.

"Balance and Filtering in Structured Satisfiable Problems. Proc. 17th Intl. Conf. on Artificial Intelligence (IJCAI-2001), Seattle, WA (2001). With H. Kautz, Y. Ruan, D. Achlioptas, B. Selman, and M. Stickel.

"Extending the Reach of SAT with Many-Valued Logics." *Electronic Notes in Discrete Mathematics* 9, Elsevier Science Publ. (2001). With R. Bejar, A. Cabiscol, C. Fernandez, and F. Manya.

"An Application of Randomization and Restarts in Proof Planning." Proc. of the 6th European Conference on Planning (ECP-01), Toledo, Spain (2001). With A. Meier , and E.Melis.

"Extending the Reach of Proof Planning by Randomization and Restart Techniques." *Future Directions in Automated Reasoning*, IJCAR Workshop, Siena, Italy (2001). With A. Meier and E. Melis.

"Generating Hard Feasible Schedules. Proc. of the 6th European Conference on Planning (ECP-01), Toledo, Spain (2001). With J. Argelich, R. Bejar, A. Cabiscol, C. Fernandez, and F. Manya.

"Heavy-tailed Behavior and Randomization in Proof Planning. Model-based Validation of Intelligence, AAAI 2001 Spring Symposium Series, Stanford, CA (2001). With A. Meier and E. Melis.

"Distribute Constraint Satisfaction in a Wireless Sensor Tracking System. Proc. Workshop on Distributed Constraint Reasoning (CONS-2), IJCAI-2001, Seattle (2001). With R. Bejar, B. Krishnamachari, and B. Selman.

"Heavy-tailed Phenomena in Satisfiability and Constraint Satisfaction Problems. *Journal of Automated Reasoning* 24(1/2):67-100 (2000). With B. Selman, N. Crato, and Henry Kautz). Results featured in *Science News* (2000).

"Artificial Intelligence and Operations Research: Challenges and Opportunities in Planning and Scheduling and Operations Research. *Journal of Knowledge Engineering Review* 15(1) (2000).

Donald P. Greenberg

Professor and Member of the FCI, the Johnson School of Management, the Department of Architecture, and the Graduate Field of Computer Science

dpg@graphics.cornell.edu

http://www.graphics.cornell.edu/people/director.html

Ph.D. Cornell University, 1968

I am one of the pioneers in the emerging field of computer graphics, having served as a leading researcher and teacher in the field since 1965. My research is primarily concerned with physically based image synthesis and with applying graphic techniques to a variety of disciplines. My specialties include color science, parallel processing, and realistic image generation. My application work now focuses on medical imaging, architectural design, perception, digital photography, and real-time photorealistic image generation.

Consistent with the interdisciplinary nature of the field of computer graphics, I am a member of Cornell's faculty in Johnson Graduate School of Management, the Department of Computer Science, and the Department of Architecture. In recent years I have taught courses in computer graphics, computer-aided architectural design, digital photography, and disruptive technologies.

I was the founding director of the NSF Science and Technology Center for Computer Graphics and Scientific Visualization, now in its eleventh year. I am the director of the Program of Computer Graphics and former director of the Computer-aided Design Instructional Facility at Cornell.

I have published more than 200 articles on computer graphics, and many of my students have been highly recognized in the field, including several who have received the SIGGRAPH Achievement Award and others who have received Hollywood Oscars.

In 1987, I received the ACM Steven Coons Award, the highest honor in the field, for my outstanding creative contributions in computer graphics. I also received the National Computer Graphics Association Academic Award in 1989. In 1997 I received the ASCA Creative Research Award in Architecture. An honorary doctoral degree from New Jersey Institute of Technology was presented to me in 1999.

PROFESSIONAL ACTIVITIES

Member: National Academy of Engineering.

Fellow: ACM and International Association of Medical and Biological Engineering.

LECTURES

Disruptive Histories of Our Future. Cornell JGSM Reunion 2001, Cornell University (June 9, 2001).

Art & Design for Living. Cornell University, Parents Visit (April 20, 2001).

Virtual Environments of the 21st Century. Cornell Association of Professors Emeriti, Cornell University (April 19, 2001).

Rendering History & Progress. STC Lecture, Program of Computer Graphics, Cornell University (April 17, 2001).

The Real Challenge for Architecture in a Virtual World. Architecture Department, University of Virginia, Charlottesville, VA (April 5, 2001).

Rendering. EAG 2001, Boston, MA (April 3, 2001).

Working Today on Tomorrow's Design Software. SOM Lecture Series, NY (February 15, 2001).

B-schools: A Case History of Our Future. MBA Leadership 2001 Conference, New Orleans, Louisiana (January 25, 2001).

Distance Learning. Cornell Club of Eastern Florida, Delray Beach, FL (December 6, 2000).

Great Ideas for Computer Science. Lecture, ComS 150, Computer Science, Cornell University (October 18, 2000).

Technology & Design Practices. Real-time Workshop on Technology and Design Practice, Program of Computer Graphics, Cornell University (October 8, 2000).

Tomorrow's Internet and Design Software. Proceedings: Emerging Information Technologies for Facilities, NAS, FCC, Washington D.C. (October 20, 2000).

How Do We Prepare Our Students for the Future? Architecture Summer School: Architecture Department, Cornell University (July 12, 2000).

PUBLICATIONS

"Spatiotemporal Sensitivity and Visual Attention For Efficient Rendering of Dynamic Environments." *ACM Transactions on Graphics* (2001). With Hector (Yang-Li) Yee, and Sumanta N. Pattanaik.

"Field Trip to Ithaca, N.Y.: Autodesk Development Team Refining Sketching Advantages of Architectural Studio." *Design Architecture.com* (June 4, 2001). With Wendy Talarico.

"Lighting the Way: A Conversation with Don Greenberg of Cornell's Program in Computer Graphics." *Cadence Web* (http://www.cadenceweb.com/features/interviews/greenberg.html), 1-2 (2001).

"Tomorrow's Internet and Design Software." Symposium: Emerging Information Technologies for Facilities, NAS, FCC, Washington D.C. (October 20, 2000).

"Once and Future Graphics Pioneer." *Architecture Week* (18):1-7 (September, 2000). With B. J. Novitski.

"Time-dependent Visual Adaptation for Realistic Image Display." Computer Graphics Proceedings, Annual Conference Series, *ACM SIGGRAPH*, 47-53(July, 2000). With Sumanta N. Pattanaik, Jack Tumblin, and Hector Yee.

"Toward a Psychophysically Based Light Reflection Model for Image Synthesis." Computer Graphics Proceedings, Annual Conference Series *ACM SIGGRAPH*, 55-64 (July, 2000). With Fabio Pellacini, and James A. Ferwerda.

"A Lab Ahead of its Time: Cornell Graphics Lab Sets High Standards." *Architectural Record*, 198-204 (June, 2000). With B. J. Novitski, and Moreno A. Piccolotto.

"Approximate Visibility for Illumination Computations using Point Clouds." *Program of Computer Graphics Technical Report*, Cornell University (June 1, 2000). With Philip M. Dutre, and Parag Tole.

"A System for 3D Conceptual Modeling for Architectural Design." *Program of Computer Graphics Technical Report*, Cornell University (January 3, 2000). With Moreno Piccolotto, Sebastian Fernandez, Kavita Bala, and Michael Malone.

"Interactive Direct Lighting in Dynamic Scenes." *Program of Computer Graphics Technical Report, Cornell University* (January 2, 2000). With Sebastian Fernandez, Kavita Bala, and Moreno Piccolotto.

Zygmunt J. Haas

Associate Professor
Member of the School of Electrical and Computer Engineering and the Graduate Field of Computer Science
haas@ee.cornell.edu
http://www.ee.cornell.edu/$\sim$haas/wnl.html
Ph.D. Stanford University, 1988

My research is in the area of mobile and wireless systems and networks. Selected examples of the projects that are conducted in my Wireless Network Laboratory (WNL) are: ad-hoc networks (routing, medium access control, security, etc.), quality of service, cross-layer protocol design, mobile web access, multicasting, and mobility management.

MEMBERSHIPS

Professional Societies: IEEE Communications Society; IEEE Vehicular Technology; Association for Computing Machinery (ACM); Special Interest Group on Mobile Communications (SIGMOBILE).

AWARDS/HONORS

Michael Tien '72 Award, Cornell College of Engineering, Excellence in Teaching Award (September, 2000).

UNIVERSITY ACTIVITIES

Member, Ad-hoc Tenure Promotion Review Committee, Computer Science Department, Cornell University; Computing Policy Committee (CPC), College of Engineering, Cornell University.

EE Policy Committee, School of Electrical and Computer Engineering, Cornell University (2000).

Member of the Committee on "Bits On Our Minds 2001 (BOOM '01).

PROFESSIONAL ACTIVITIES

Editorial Board: *IEEE Transactions on Networking*.

Editorial Board: *IEEE Communications* Magazine.

Organizer and chair of the session on "Outrageous Opinions" at MobiHoc '01, The ACM Symposium on Mobile Ad-hoc Networking & Computing, Long Beach, CA (October 4-5, 2001).

Member: IASTED (The International Association of Science and Technology for Development) Technical Committee on Telecommunication.

Member:ACM Mobicom Steering Committee.

Chairof the IEEE Technical Committee on Personal Communications (TCPC).

Vice-chair of the IEEE Technical Committee on Personal Communications (TCPC).

Editorial Board of the journal *Wireless Communications and Mobile Computing*, John Wiley & Sons.

Guest Editor: *Wireless Personal Communication Journal*, special issue on Multimedia Network Protocols and Enabling Radio Technologies.

Editorial Board: *ACM/Baltzer Wireless Networks.*

Editorial Board: *Journal of High Speed Networks*.

Program Committee: IEEE Symposium on Ad-hoc Wireless Networks (SAWN), in conjunction with IEEE GLOBECOM 2001, San Antonio, TX (November 25-29, 2001).

Program Committee member and session chair, Milcom'01, McLean, VA (October 28-31, 2001).

Committee member: Wireless Communications and Networking Conference 2002 (WCNC'02), Orlando, FL (March 18-21, 2002).

Program Committee: European Wireless 2002, Next Generation Wireless Networks: Technologies, Protocols, Services and Applications, Florence, Italy (February 26-28, 2002).

Program Committee: 11th IEEE Workshop on Local and Metropolitan Area Networks, Boulder, CO (May 18, 2001).

Program Committee: IEEE Wireless Networks and Mobile Computing Workshop, Phoenix, AZ (April 16-19, 2001).

Program Committee: ACM/IEEE MobiCom'2000, Boston, MA (August 6-11, 2000).

Program Committee: First IEEE Workshop on Mobile Ad HOC Networking and Computing Workshop (MobiHOC), Boston, MA (August 11-12, 2000).

NSF reviewer and panelist.

LECTURES

Research in the Wireless Networks Laboratory at Cornell. Department of Electrical Engineering, Columbia University (April 23, 2001).

Research in the Wireless Networks Laboratory at Cornell. Communication Systems Department, Swiss Federal Institute of Technology Lausanne (EPFL) (December 11, 2000).

PUBLICATIONS

"The Zone Routing Protocol: A Hybrid Framework for Routing in Ad-hoc Networks." *Ad-hoc Networks*, Charlie Perkins, editor. Addison Wesley (2001). With M. R. Pearlman.

"A Communication Infrastructure for Smart Environments — A Position Article." *IEEE Personal Communications* Magazine, special issue on "Networking the Physical World," 6-10 (October , 2000).

"Predictive Distance-based Mobility Management for PCS Networks." 2000 Cornell Summer Workshop on Information Theory (Bergerfest), Ithaca, NY (August 18-19, 2000). And B. Liang.

"On the Impact of Alternate Path Routing for Load Balancing in Mobile Ad-hoc Networks." First Annual IEEE/ACM Workshop on Mobile Ad-hoc Networking & Computing, MobiHOC'2000, Boston, MA (August 11, 200). With M. R. Pearlman, P. Scholander, and S. S. Tabrizi.

"A Decision-theoretic Approach to Resource Allocation in Wireless Multimedia Networks." Fourth International Workshop on Discrete Algorithms and Methods for Mobile Computing, DIALM 2000, Boston, MA (August 11, 2000). And J. Y. Halpern, L. Li, and S. B. Wicker.

"Securing Ad-hoc Networks." *IEEE Network Magazine* 13(6) (November/December 1999) With L. Zhou.

"Determining the Optimal Configuration for the Zone Routing Protocol." *IEEE Journal on Selected Areas in Communications (JSAC)*, issue on Ad-hoc Networks 17(8):1395-1414 (August, 1999). And M. R. Pearlman.

"The Dynamic Packet Reservation Multiple Access Scheme for Multimedia Traffic." *ACM/Baltzer Mobile Networks Applications* 4:87-99 (1999). With D. A. Dyson.

"Ad-hoc Location Management using Quorum Systems." *IEEE Transactions on Networking*, ACM/*IEEE Transactions on Networking* (April, 1999). And B. Liang.

"The Multiply-detected Macrodiversity Scheme for Wireless Cellular Systems." *IEEE Transactions on Vehicular Technology* 47(2) (May, 1998). And C-P. Li.

NEW PATENT APPLICATIONS

"ITAMAR — Independent Tree Ad-hoc Multicast Routing." Cornell Research Foundation, Application number: D-2823 — Haas. And M. S. Sajama.

"COCA: A Secure Distributed On-line Certification Authority." Cornell Research Foundation, Application number: D-2732A — Haas. With F. Schneider, L. Zhou, and R. van Renesse.

"Adaptive Power Control in Wireless Ad-hoc Networks." Cornell Research Foundation, Application number: D-2507 — Haas. And Miguel Sanchez.

"Routing and Mobility Management Protocols for Ad-hoc Networks." Cornell Research Foundation, Application number: D-2191 — Haas.

Joseph Y. Halpern

Professor, and Co-director: Cognitive Studies Program

halpern@cs.cornell.edu
http://www.cs.cornell.edu/home/halpern/
Ph.D. Harvard University, 1981

My research is concerned with representing and reasoning about knowledge and uncertainty in multi-agent systems. The work uses tools from logic (particularly modal logic and the idea of possible-worlds semantics), probability theory, distributed systems, game theory, and AI, and I like to think that it contributes to our understanding of each of these areas as well.

Some themes of my current research include: (1) applying ideas of decision theory to constructing algorithms in asynchronous distributed systems, database systems, and wireless systems, (2) providing foundations for useful qualitative notions of decision theory, (3) reasoning about security.

HONORS

Guggenheim Fellowship (for 2001-02).
Fulbright Fellowship (for 2001-02).

UNIVERSITY ACTIVITIES

Co-director: Cognitive Studies Program.
Chair, Admissions Committee, Department of Computer Science.

PROFESSIONAL ACTIVITIES

Fellow, American Association of Artificial Intelligence.
Editor-in-chief: *Journal of the ACM* (as of May, 1997).
Consulting Editor: *Chicago Journal of Computer Science*.
Editorial board: *Artificial Intelligence Journal*; *Information and Computation*; *Journal of Logic and Computation*.
Member: ACM Publications Board.
Chairman: ACM Preprint Repository.
Coordinator: CoRR (Computing Research Repository).
Member: LICS (IEEE Conference on Logic in Computer Science) Advisory Board.

President of Board of Directors: Corporation for Theoretical Aspects of Reasoning About Knowledge.
Program Chair, 16th Annual IEEE Symposium on Logic in Computer Science (2001).
Program Committee Member and Conference Chair, 8th Conference on Theoretical Aspects of Rationality and Knowledge (2001).

LECTURES

A Computer Scientist Looks at Game Theory.
—. Invited talk, Games 2000, Bilbao, Spain (July, 2000).
—.Invited talk, 4th Conference on Logic and Foundations of Game and Decision Theory, Torino, Italy (July, 2000).
Plausibility Measures and Default Reasoning. Amsterdam University (May, 2001).
Plausibility Measures: a General Approach for Representing Uncertainty Northwestern University (May, 2001).
Complexity, Logic, and Computation: A Symposium in Honor of Albert Meyer , Boston, MA (June, 2001).

PUBLICATIONS

"Axiomatizing Causal Reasoning." *Journal of AI Research* 12:317-337 (2000).
"A Note on Knowledge-based Programs and Specifications." *Distributed Computing* 13(3):145-153 (2000).
"First-order Conditional Logic Revisited." *ACM Transactions on Computational Logic* 1(2):175-207 (2000). With N. Friedman and D. Koller.
"Multi-agent Only Knowing." *Journal of Logic and Computation* 11(1):41-70 (2001). With G. Lakemeyer.
"A Logic for SDSI's Linked Local Named Spaces." *Journal of Computer Security* 9(1,2):47-74 (2001). With R. van der Meyden.
"A Decision-theoretic Approach to Reliable Message Delivery. *Distributed Computing* 14:1-16 (2001). With F. Chu.
"A Decision-theoretic Approach to Resource Allocation in Wireless Multimedia Networks." Proceedings of *Dial M for Mobility*, 86-95 (2000). With Z. Haas, L. Li, and S. B. Wicker.
"Conditional Plausibility Measures and Bayesian Networks." Proceedings of the Sixteenth Conference on Uncertainty in AI, 247-255 (2000).
"Minimum-energy Mobile Wireless Networks Revisited." Proceedings of the IEEE Conference on Communications (2001). With L. Li.
"A Logical Reconstruction of SPKI." Proceedings of the 14th IEEE Computer Security Foundations Workshop (2001). With R. van der Meyden.
"Review of 'Probability and Conditionals: Belief Revision and Rational Decisions.'" *Journal of Philosophical Logic* 100(2):277-281 (2000).
"The Unusual Effectiveness of Logic in Computer Science." *Bulletin of Symbolic Logic* 7(2):213-236 (2001). With R. Harper, N. Immerman, P. G. Kolaitis, M. Y. Vardi, and V. Vianu.
"Editorial: An Author's Bill of Rights and Responsibilities." *Journal of the ACM* 47(5):828-825 (2000).
"CoRR: A Computing Research Repository (with commentary)." *ACM Journal of Computer Documentation* 24(2):41-48 (2000).

LANDMARK PUBLICATIONS

Reasoning About Knowledge. MIT Press (1995). With R. Fagin, Y. Moses, and M. Y. Vardi.
"Knowledge and Common Knowledge in a Distributed Environment." *Journal of the ACM* 37(3):549-587 (1990). With Y. Moses. Awarded Godel Prize in 1997.
"Belief, Awareness, and Limited Reasoning." *Artificial Intelligence* 34:39-76 (1988). With R. Fagin. Conference version winner of MIT Press Publisher's Prize as best paper of the 9th International Joint Conference on Artificial Intelligence (1985).
"An Analysis of First-order Logics of Probability." *Artificial Intelligence* 46:311-350 (1990). Conference version winner of Publisher's Prize as best paper of the 11th International Joint Conference on Artificial Intelligence (1989).
"Plausibility Measures and Default Reasoning." Proceedings of the Thirteenth National Conference on Artificial Intelligence, 1297-1304 (1996). To appear in the *Journal of the ACM* . With N. Friedman. Commended for its excellence by the Committee on the "IGPL/FoLLI Prize for the Best Idea of the Year 1996."

Professor Juris Hartmanis (on right), with former Ph.D. student: now U. of Toronto Computer Science Professor A.B. Borodin

Juris Hartmanis

Walter R. Read Professor
of Engineering
Turing Award Winner

jh@cs.cornell.edu
Ph.D. California Institute of Technology, 1955

The strategic goal of my research is to contribute to the development of a comprehensive theory of computational complexity. Computational complexity, the study of the quantitative laws that govern computation, is an essential part of the science base needed to guide, harness, and exploit the explosively growing computer technology. My current research interests focus on understanding the computational complexity of chaotic systems and the classification of undecidable problems in complexity theory.

AWARDS

Recipient of the Grand Medal of the Latvian Academy of Sciences (Lielo Medalu) (2001).

PROFESSIONAL ACTIVITIES

Member: National Academy of Engineering.
Foreign Member: Latvian Academy of Sciences.
Fellow: American Academy of Arts and Sciences; New York State Academy of Sciences; AAAS.
Editor: *Springer-Verlag Lecture Notes in Computer Science*; *Journal of Computer and Systems Sciences*; *Fundamenta Informaticae*.
Advisory Board: *EATCS Monographs in Theoretical Computer Science*; Springer-Verlag; *International Journal of Foundations of Computing*.
Member: DIMACS External Advisory Committee.
Member: Santa Fe Institute Science Board.
Member: Santa Fe Institute Science Steering Committee.
Member: University of Cincinnati Computing Program Review Committee (November 1-3, 2000).

UNIVERSITY ACTIVITIES

Chair: Engineering College Nominating Committee.
Member: FCI Founders Committee.

LECTURES

Goedel, Undecidability and Automata Theory, Half Century of Automata Theory. University of Western Ontario (July 26, 2000).
Four lectures at Jyvaskyle University, Finland, (August 10-11, 2000).
—. Undecidability and Incompleteness Results in Theory of Computing
—. Succinctness and Minimality of Automata Description
—. Search for Limits of Feasible Computations
—. On the Complexity and Shape of Mathematical Proofs
Computational Complexity and Mathematical Proofs, University of Saarbruecken, Germany (August 30, 2000).
What can Computational Complexity Theory Say about Mathematics?
—. Iowa State University (October 23, 2000).
—. Cray Lecture Series, University of Minnesota (October 30, 2000).

PUBLICATIONS

"Computational Complexity and Mathematical Proofs." *Informatics— 10 Years Back, 10 Years Ahead."* Reinhard Wilhelm, editor. Springer-Verlag LNCS 2000, 251-256 (2001).

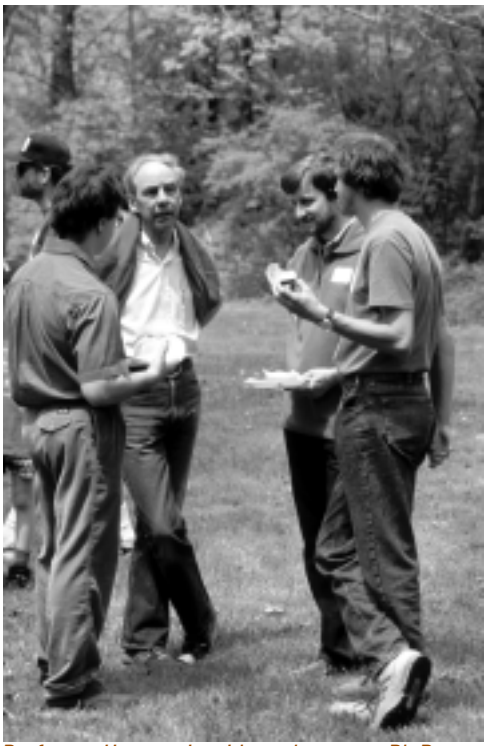

Professor Hartmanis with students at a Ph.D. picnic

Mark Heinrich

Assistant Professor
Member of the School of Electrical and Computer Engineering and the Graduate Field of Computer Science
heinrich@csl.cornell.edu
http://www.csl.cornell.edu/~heinrich/
Ph.D., Stanford University, 1998

My research is concerned with the design of active memory and I/O systems for next-generation servers and data-intensive computing. This work has focused on extending the cache coherence mechanism in modern servers to implement active memory operations—computation performed in the memory system on behalf of the microprocessor to speed up overall execution time. Coupled with this work is the exploration of the effect of new networking technologies (i.e. InfiniBand) on next-generation servers and the integration of active memory and I/O techniques with this networking technology. We have also shown that active memory machines and hardware cache-coherent distributed shared-memory (DSM) machines need much the same support, and, in fact, that by building our single-node active memory system we can also support a multiprocessor version of the machine that we call active memory clusters. Active memory clusters can achieve hardware DSM performance at the low cost of clusters.

In my work on Active I/O systems I am developing a smart InfiniBand switch (which can also be used in active memory clusters) that can support either normal or intelligent I/O devices, and offload computation from the microprocessor to minimize latency and reduce bandwidth requirements in the I/O system. This work also involves innovative operating system restructuring, including the filesystem and the network stack. The operating system in active I/O systems must be partitioned between the microprocessor and the active I/O devices. Our own operating system, SplitOS, is joint work between our research group and groups at Rutgers and Princeton.

In work on scalable cache coherence protocols, I am working on issues of fairness and robustness in scalable distributed shared-memory (DSM) machines. In addition, I am looking at the quantitative impact of many coherence protocol techniques by evaluating each technique on a flexible DSM prototype. Together with Martin Burtscher, I am also exploring predictive techniques in cache coherence protocols to minimize latency.

AWARDS/HONORS

Michael Tien '72 Excellence in Teaching Award (2000-2001).
IEEE Teacher of the Year Award (1999-2000).
NSF Faculty Early Career Development Award (2000-2004).

UNIVERSITY ACTIVITIES

Member: Intelligent Information Systems Institute; ECE Curriculum and Standards Committee; ECE Long-Range Recruiting Committee; ECE Experimental Systems Recruiting Committee; ECE Circuits & MEMS Recruiting Committee; CURIE Summer Program for Women in Engineering; Fields of Electrical Engineering, Computer Science.

PROFESSIONAL ACTIVITIES

Publicity Chair: International Symposium on High-performance Computer Architecture (PCA) (January, 2001).
Panelist: Mathematics, Information, and Computational Sciences Division within the Office of Advanced Scientific Computing Research in the Office of Science at the U.S. Department of Energy (DOE) (April, 2001).
Program Committee: Workshop on Caching, Coherence, and Consistency (WC3), held in conjunction with the ACM Conference on Supercomputing (June, 2001).
Program Committee: International Conference on Parallel Architectures and Compilation Techniques (PACT) (September, 2001).

LECTURES

Flash Forward: Better, Faster, Cooler. Cornell University Silicon Valley Event, hosted by Hunter Rawlings, San Mateo, CA (April, 2001).
Providing Hardware DSM Performance at Software DSM Cost. Seminar, University of Rochester (April, 2001).
Hardware DSM Performance at Software DSM Cost. Air Force Research Laboratory, Rome, NY (March, 2001).
A Case for Asynchronous Active Memories. ISCA Workshop on Solving the Memory Wall (June, 2000).

PUBLICATIONS

"FLASH vs. (Simulated) FLASH: Closing the Simulation Loop." In Proceedings of the 9th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 49-58 (November, 2000). With J. Gibson, R. Kunz, D. Ofelt, M. Horowitz, and J. Hennessy.
"Using Meta-level Compilation to Check FLASH Protocol Code." In Proceedings of the 9th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 59-70 (No-

vember, 2000). With A. Chou, B. Chelf, and D. Engler. "A Case for Asynchronous Active Memories." In Proceedings of the ISCA Workshop on Solving the Memory Wall (June, 2000). With R. Manohar.

Sheila S. Hemami

Associate Professor
Member of the School of Electrical and Computer Engineering and the Graduate Field of Computer Science
hemami@cs.cornell.edu
http://www.ece.cornell.edu/people/faculty/faculty_list.shtml#shei
Ph.D. Stanford University, 1994

The emerging information superhighway provides an example of the flexibility required of image and video compression and transmission techniques. Varying network capacities, differences in viewing devices, and a broad spectrum of user needs suggest the desirability of coding techniques that can efficiently span large quality and bandwidth ranges. Additionally, coded data must be robust to errors and loss of varying degrees across multiple network segments. For practicality, algorithms must be inexpensive to implement, in either hardware or software. My research interests broadly concern such communication of visual information. Particular topics of interest include multirate video coding and transmission, compression specific to packet networks and other lossy networks, and psychovisual considerations.

UNIVERSITY ACTIVITIES

College of Engineering Committee of Faculty Women.
Society of Women Engineers Faculty Advisor.
EE Curriculum and Standards Committee.
Peoplesoft Lab Oversight Committee.
EE General Recruiting Committee.
Summer presentation to students in "Inventing an Information Society" (July, 2000).

PROFESSIONAL ACTIVITIES

Reviewer, *IEEE Trans. Signal Processing; IEEE Trans. Image Processing, IEEE Trans. Circuits and Systems for Video Technology; IEEE Communications Letters; IEEE ICIP* 2001.
Associate Editor, *IEEE Trans. Signal Processing.*

AWARDS/HONORS

HKN C. Holmes MacDonald Outstanding Teaching Award.
Michael Tien '72 Cornell College of Engineering Teaching Award.
Fulbright Distinguished Lecturer, State of Morocco (2001).

PUBLICATIONS

"Perceptual Quantization for Wavelet-based Image Coding." Proc. IEEE Int. Conf. on Image Processing, Vancouver, B.C. (September, 2000). With M. G. Ramos.

"Subjective Quality Evaluation of Low Bit Rate Video." Proc. Human Vision and Electronic Imaging 2001, San Jose, CA (January, 2001). With M. A. Masry, W. Osberger, and A. M. Rohaly.

"Generalized Rate-distortion Optimizations for Motion-compensated Video Coding." IEEE Trans. Circuits & Systems for Video Tech. (September, 2000). With Yan Yang.

"Perceptually-based Robust Image Transmission over Wireless Channels." Proc. IEEE Int. Conf. on Image Processing, Vancouver, B.C. (September, 2000). With M. E. Buckley, M. G. Ramos, and S. B. Wicker.

"Robust Data Hiding using Psychovisual Thresholding." Proc. IEEE Int. Conf. on Image Processing, Vancouver, BC, Canada (September, 2000). With M. A. Masry, and M. G. Ramos.

Professors Claire Cardie and Gregory Morrisett

John E. Hopcroft

Joseph Silbert Dean of Engineering
Turing Award Winner
Professor

jeh17@cornell.edu
www.cs.cornell.edu/annual_report/99-00/hopcroft.htm
Ph.D. Stanford University, 1964

Since January 1994, I have been the Joseph Silbert Dean of the College of Engineering. Upon completion of my term as dean, which ends on June 30, 2001, I plan to return to research in the Department of Computer Science at Cornell. My research will center on the study of information capture and access. I have also been involved in the theoretical aspects of computing, especially analysis of algorithms, formal languages, automata theory, and graph algorithms. I have coauthored four books on formal languages and algorithms with Jeffrey D. Ullman and Alfred V. Aho.

I was formerly the associate dean for college affairs and the Joseph C. Ford Professor of Computer Science. After receiving an M.S. (1962) and Ph.D. (1964) in electrical engineering from Stanford University, I spent three years on the faculty of Princeton University. In 1967, I joined the Cornell faculty, was named professor in 1972 and served as chairman of the Department of Computer Science from 1987 to 1992. I am an undergraduate alumnus of Seattle University, and was honored with a Doctor of Humanities Degree, *Honoris Causa*, in 1990.

PROFESSIONAL ACTIVITIES

National Academy of Engineering.
Scientific Advisory Committee for The David and Lucile Packard Foundation;
John Von Neumann Medal Selection Committee for the Institute of Electrical and Electronics Engineers, Inc. (IEEE);
School of Engineering Advisory Committee, Hong Kong University of Science and Technology;
National Advisory Committee on Informatics Engineering, National College of Industrial Relations (Ireland).
Fellow: American Academy of Arts and Sciences; American Academy for the Advancement of Science; Institute of Electrical and Electronics Engineers (IEEE); (Charter) ACM.
Associate Editor (and Editor): *Journal of Computer and System Sciences.*
Editor (and Member, Executive Committee): *Algorithmica.*
Associate Editor: *Information Sciences.*
Editor: *International Journal of Computational Geometry and Applications.*

Klara Kedem

Professor

kedem@cs.cornell.edu
http://www.cs.cornell.edu/kedem
Ph.D. Tel Aviv University, 1989

My research area is Computational Geometry with applications to problems in computer vision and bio-information. The attempt to deal with practical problems (like shape comparison) by investigating their geometric nature, yields a better theoretical understanding of the problems and provides sound and efficient algorithms. Among the theoretical problems I work on are problems in geometric optimization, such as covering a set of points by a given number of shapes and facility location. I have investigated the minimum Hausdorff distance as a tool for measuring shape resemblance between images.

Many practical problems in the area of shape comparison seek a fully automated solution. The robustness of the minimum Hausdorff distance lends itself to such problems. In the last two years I have looked into shape comparison problems in three dimensions. Such problems arise in many life-science disciplines. In computational molecular biology I have come up with a new metric, the URMS, to measure substructure resemblance between proteins. This measure has been implemented and further applied to the analysis of molecular dynamics. It proves superior to previously used measures. With the department of life sciences at Ben Gurion University, I have been looking at dendrite shape comparison and classification. Here we applied a three dimensional Hausdorff distance for the problem and are in the midst of devising new measures.

PROFESSIONAL ACTIVITIES

Chair of the Computer Science Department, Ben-Gurion University of the Negev, Israel.

Editorial board: *Pattern Recognition Society Journal.*

Guest editor of *Computational Geometry: Theory and Applications*.

Workshop program committee: Workshop on Algorithmic Engineering (WAE 01) BRICS, University of Aarhus, Denmark (August 28-31, 2001).

Reviewer: Israeli-U.S. Bi-national Science Foundation; The Academia – Ministry of Science, Israel; *Pattern Recognition Journal*; *International Journal of Computational Geometry and Applications*.

GRANTS:

LSRT Consortium, "Optimal Layout Design of Large Scale Satellite Telephony Systems in Rural Regions" (2000-2001).

Israeli Defense Ministry, "BGU Bioinformatics Center for the Interpretation of the Human Genome" (2001-2002).

PUBLICATIONS

"Optimal Facility Location under Various Distance Functions." Accepted for publication in *International Journal of Computational Geometry and Applications*. With S. Bespametnik, M. Segal and A. Tamir.

"Improved Algorithms for Placing Undesirable Facilities." Accepted for publication in *Computers and Operations Research*. With M. J. Katz, and M. Segal.

Jon Kleinberg

Associate Professor

kleinber@cs.cornell.edu

http://www.cs.cornell.edu/home/kleinber/

Ph.D. MIT, 1996.

My research is concerned with algorithms that exploit the combinatorial structure of networks and information.

The information we deal with is taking on an increasingly networked structure; the World Wide Web serves as perhaps the most compelling example of this phenomenon. One direction I have pursued, motivated by this issue, is the development of algorithms that analyze the link structure of the Web to identify high-quality information resources — hubs and authorities relevant to broad topics. A second direction is the development of graph models that can provide insight into the structure of large networks. I have been studying an algorithmic framework in which to frame questions about certain 'small-world' properties of networks; and with Duncan Callaway, John Hopcroft, Mark Newman, and Steve Strogatz, I have recently looked at models of random graphs that evolve over time. A third line of research has been concerned with designing algorithms that facilitate the rapid spread of information in a network; with David Kempe and Al Demers, I have worked on the design of randomized 'gossip' protocols that provide efficient, decentralized mechanisms for such tasks.

Discrete optimization provides a collection of powerful techniques for approaching problems in this area; in particular, balancing constraints imposed by many individuals leads to a range of interesting new problems. With Christos Papadimitriou and Prabhakar Raghavan, I have been studying an optimization-based framework for clustering and data mining in which one has access to data from a large population; we consider the extent to which designing effective algorithms can be balanced with concerns about the privacy of each individual's data. With Amit Kumar, I have been looking at resource allocation problems with multiple users; since each user wants as much of the resource as possible, there are many competing objectives, and it is natural to seek the 'fairest' solution.

Finally, many of the core problems in computational biology require algorithms for analysis of large volumes of data; such data is being generated at an accelerating pace by experimental studies of genomes and proteins across a wide range of species. With Debra Goldberg, Susan McCouch, and David Liben-Nowell, I have been investigating mathematical and computational approaches to the analysis of evolutionary relationships among species at the genomic level and the connection between this type of analysis and the problem of comparative mapping. I have also developed algorithms for relating sequence information to protein structure, and with Ron Elber I am continuing to look at approaches to this issue based on threading methods.

AWARDS

National Academy of Sciences Award for Initiatives in Research (2001).

David and Lucile Packard Foundation Fellowship (1999-2004).

ONR Young Investigator Award (1999-2002).

NSF Faculty Early Career Development Award(1997-2001).
Alfred P. Sloan Research Fellowship (1997-1999).

UNIVERSITY ACTIVITIES

Member: Cornell Faculty of Computing and Information (FCI) Founders.

PROFESSIONAL ACTIVITIES

Member: National Academies Computer Science and Telecommunications Board study on Fundamentals of Computer Science.

Program Committees: IEEE Symposium on Foundations of Computer Science, 2001; International World Wide Web Conference (2001); ACM International Conference on Knowledge Discovery and Data Mining (2001); Workshop on Algorithms and Data Structures (2001); Workshop on Randomization and Approximation Techniques in Computer Science (2001).

Lecturer: Lipari Summer School on On-line Algorithms (2000).

NSF Review Panel member.

LECTURES

Information Networks: Models and Algorithms.
—. America Online (December, 2000).
—. Packard Fellowship Annual Meeting (September, 2000).
—. Santa Fe Institute Workshop on Complex Interactive Networks (August, 2000).

Small-World Phenomena and the Dynamics of Information.
—. Workshop in Honor of Allan Borodin's 60th Birthday (June, 2001).
—. Snowbird Learning Workshop (April, 2001).

Structure and Content in World-Wide Web Search. Invited plenary lecture at the Meeting of the North American Chapter of the Association for Computational Linguistics (June, 2001).

Detecting a Network Failure. IEEE Symposium on Foundations of Computer Science (November, 2000).

On-Line Algorithms for Routing and Search. Lipari Summer School (July, 2000).

PUBLICATIONS

"Spatial Gossip and Resource Location Protocols." Proc. 33rd ACM Symposium on Theory of Computing (2001). With D. Kempe and A. Demers.

"Provisioning a Virtual Private Network: A Network Design Problem for Multicommodity Flow." Proc. 33rd ACM Symposium on Theory of Computing (2001).With A. Gupta, A. Kumar, R. Rastogi, and B. Yener.

"On the Value of Private Information." Proc. 8th Conf. on Theoretical Aspects of Rationality and Knowledge (2001). With C. Papadimitriou, and P. Raghavan.

"Adversarial Queuing Theory." *Journal of the ACM* 48(1):13-38 (2001). With A. Borodin, P. Raghavan, M. Sudan, and D. Williamson.

"Universal-stability Results and Performance Bounds for Greedy Contention-resolution Protocols." *Journal of the ACM* 48(1):39-69 (2001). With D. M. Andrews, B. Awerbuch, A. Fernandez, F. T. Leighton, and Z. Liu.

"Navigation in a Small World." *Nature* 406:845 (2000).

"Detecting a Network Failure." Proc. 41st IEEE Symposium on Foundations of Computer Science, 231-239 (2000).

"Fairness Measures for Resource Allocation." Proc. 41st IEEE Symposium on Foundations of Computer Science, 75-85 (2000). With A. Kumar.

"Algorithms for Constructing Comparative Maps." *Comparative Genomics: Empirical and Analytical Approaches to Gene Order Dynamics, Map Alignment and the Evolution of Gene Families* (David Sankoff and Joseph H. Nadeau, editors), 243-261 (2000). With D. Goldberg and S. McCouch.

"The Syntenic Diameter of the Space of N-chromosome Genomes." *Comparative Genomics: Empirical and Analytical Approaches to Gene Order Dynamics, Map Alignment and the Evolution of Gene Families* (David Sankoff and Joseph H. Nadeau, editors), 185-198 (2000). With D. Liben-Nowell.

"Allocating Bandwidth for Bursty Connections." *SIAM J. Computing* 30(1):191-217 (2000). With Y. Rabani, and E. Tardos.

"Node-disjoint Paths on the Mesh, and a New Trade-off in VLSI Layout." *SIAM J. Computing* 29(4):1321-1333 (2000). With A. Aggarwal, and D. Williamson.

"Approximation Algorithms for Classification Problems with Pairwise Relationships: Metric Labeling and Markov Random Fields." Proc. 40th IEEE Symposium on Foundations of Computer Science, 14-23 (1999). With E. Tardos.

"Efficient Algorithms for Protein Sequence Design and the Analysis of Certain Evolutionary Fitness Landscapes." *Journal of Computational Biology* 6(3-4):387-404 (1999).

"Authoritative Sources in a Hyperlinked Environment." *Journal of the ACM* 46(5): 604-632 (September ,1999).

Dexter Kozen

Joseph Newton Pew, Jr.
Professor of Engineering

kozen@cs.cornell.edu
http://www.cs.cornell.edu/kozen/
Ph.D. Cornell University, 1977

My research interests include the theory of computational complexity, especially complexity of decision problems in logic and algebra, program logic and semantics, and computational algebra. Recent work includes new polynomial-time algorithms for type inference in type systems with subtypes and recursive types; algorithms solving systems of set constraints as used in program analysis; a unification algorithm for set constraints and a new constraint logic programming language based on set constraints; development of the theory of rational spaces and their relationship to set constraints; an algorithm for decomposition of algebraic functions; a new polynomial-time algorithm for resolution of singularities of plane curves; efficient algorithms for optimal transmission of encoded video data; optimality results for digital interleavers; and expressiveness, complexity and completeness results for Kleene algebras with tests. Recently I have begun to investigate algorithms for efficient code certification and the application of Kleene algebra with tests to the verification of compiler optimizations.

HONORS

Class of 1960 Scholar, Williams College.

Stephen and Margery Russell Distinguished Teaching Award, College of Arts and Sciences, Cornell.

UNIVERSITY ACTIVITIES

Member: Undergraduate Admissions Committee, College of Engineering; University Arbitration Panel.

Faculty Advisor: Cornell Men's Rugby Football Club; Johnson Graduate School of Management Rugby Football Club; Cornell Women's Rugby Football Club.

PROFESSIONAL ACTIVITIES

Program Committee Member: Foundations of Software Science and Computation Structure; Mathematical Foundations of Computer Science.

Editorial Board Member: *Journal of Relational Methods in Computer Science; Theory of Computing Systems.*

Supervisory Board: Centre for Basic Research in Computer Science (BRICS), Aarhus University; Goedel Prize Committee.

LECTURES

Language-based Security. Department of Computer Science, Dartmouth College, Hanover, NH (March, 2000).

On the Completeness of Propositional Hoare Logic. RelMiCS 5 Conference, Quebec City, Canada (January, 2000).

A Computer Scientist's View of Admissible Sets. Tarski Centenary Conference, Warsaw, Poland (May, 2001).

PUBLICATIONS

"Certification of Compiler Optimizations using Kleene Algebra with Tests." In *Proc. 1st Int. Conf. Computational Logic* (CL2000), v. 1861 of *Lecture Notes in Artificial Intelligence,* J. Lloyd, V. Dahl, U. Furbach, M. Kerber, K.-K. Lau, C. Palamidessi, L. M. Pereira, Y. Sagiv, and P. J. Stuckey, editors, London, Springer-Verlag , 568-582 (July 2000).

"On the Completeness of Propositional Hoare Logic." *Proc. of the 5th International Seminar Relational Methods in Computer Science (RelMiCS 2000)*, 195–202 (January, 2000). With J. Tiuryn.

"On Hoare Logic and Kleene Algebra with Tests." *Trans. Computational Logic* 1(1):60-76 (July, 2000).

"A Note on the Complexity of Propositional Hoare Logic." *Trans. Computational Logic* 1(1):171-174 (July, 2000). With E. Cohen.

Dynamic Logic. MIT Press, Cambridge, MA (2000). With D. Harel and J. Tiuryn.

"Myhill-Nerode Relations on Automatic Systems and the Completeness of Kleene Algebra." In *Proc. 18th Symp. Theoretical Aspects of Computer Science*, A. Ferreira and H. Reichel, editors. Dresden, Germany (February, 2001). Lecture Notes in Computer Science, v. 2010, Springer-Verlag, 27-38.

"Intuitionistic Linear Logic and Partial Correctness." *Proc. 16th Symp. Logic in Computer Science*. IEEE (June, 2001) With J. Tiuryn.

Dean Krafft

Senior Research Associate
Director of Computing Facilities

dean@cs.cornell.edu
http://www.cs.cornell.edu/home/dean/
Ph.D. Cornell University, 1981

I serve both as a researcher and an administrator in the Department of Computer Science at Cornell. In my guise as an administrator, I manage the Computing Facilities Support group and worry about a number of issues including computer security, networking, and building web services. Most recently, as part of the Nomad research project (http://www.nomad.cornell.edu), I have been leading a campus-wide pilot of megabit-speed wireless LAN technology. This will be rolled out by Cornell Information Technologies as a Cornell campus service in the fall 2001.

On the research side, I am part of the Cornell Digital Libraries Research Group (CDLRG - http://www.cs.cornell.edu/cdlrg). A major focus of our effort is on interoperability issues for digital libraries. As part of that broader thrust, I am a Co-principal Investigator on the NSF-funded National Science Digital Library Project at Cornell (http://www.SiteForScience.org). My own particular interests focus on ensuring the availability in the digital world of pre-digital published and manuscript materials, as well as related issues on copyright, the public domain, and public access to older and out-of-print materials.

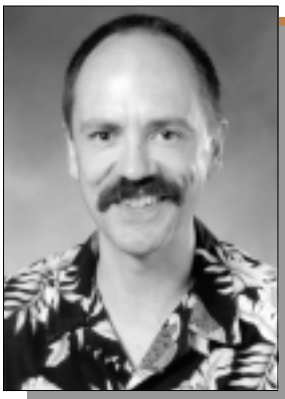

Christoph Kreitz

Senior Research Associate

kreitz@cs.cornell.edu
http://www.cs.cornell.edu/home/kreitz
Ph.D FernUniversitaet Hagen, 1984

My primary research interest is the application of automated deduction to the design, verification, and optimization of software systems. My current research aims at developing a Logical Programming Environment for the construction of reliable and efficient distributed systems. In collaboration with the Nuprl and Ensemble research groups, I have built semantics-based tools for the automatic optimization of protocol stacks in the Ensemble group communication system. More recently, Robbert van Renesse, Mark Bickford and I have developed a generic switching protocol for the construction of adaptive systems and proved it correct with the Nuprl proof development system. For this purpose, we introduced the concept of meta-properties and used them to characterize communication properties that can be preserved by switching. We also identified switching invariants that an implementation of the switching protocol must satisfy in order to work correctly. The verification efforts revealed a variety of implicit assumptions that are usually made when designing communication systems and uncovered minor design errors that would otherwise have made their way into the implementation.

I am also interested in the development of automatic proof procedures for classical and non-classical logics. Together with former students from the Technical University of Darmstadt, I work on proof search methods based on matrix-characterizations of logical validity, a very compact representation of the search space. We have developed a uniform proof search procedure for classical logic, intuitionistic logic, various modal logics, fragments of linear logic, and for inductive specification proofs. We have also developed a uniform algorithm for transforming the machine-found matrix proofs into sequent proofs. In the past year, we implemented JProver, a first-order intuitionistic theorem prover that creates sequent-style proof objects, and connected it as external proof engine to the interactive proof assistants Nuprl and MetaPRL. The combination of these systems gives a user the full expressive power of the proof assistant when dealing with complex proofs and verifications, while at the same time taking advantage of well-understood and efficient proof techniques for subproblems that only depend on first-order reasoning.

PROFESSIONAL ACTIVITIES

Program Commitee: IJCAR Workshop on Verification (2001).

Referee: *Handbook of Automated Reasoning; Journal of Symbolic Computation; Journal of Functional Programming.*

LECTURES

The NuPRL Open Logical Environment. Scottish Theorem Proving Meeting, St. Andrews, Scotland (July, 2000).

Matrix-based Inductive Theorem Proving, International Conference TABLEAUX-2000, St. Andrews, Scotland (July, 2000).

Building Reliable, High-performance Communication Systems from Components. Technical University of Darmstadt, Germany (July, 2000).

A Logical Programming Environment for Communication Systems. University of Potsdam, Germany (December, 2000).

Protocol Optimization in Ensemble/Spinglass. DARPA FTN PI Meeting, St. Petersburgh, (January, 2001).

Advances in Logical Programming Environments. DARPA PCES PI meeting, San Diego (February, 2001).

An Open Logical Programming Environment,.DARPA PCES PI meeting, St. Louis (May, 2001).

Formal Design of Reliable Software Systems. DARPA Information Survivability Conference and Exposition II, Anaheim (June, 2001).

Formal Design and Verification: Challenges and Prospects. Invited talk, workshop on verification, International Joint Conference on Automated Reasoning, Siena, Italy (June, 2001).

PUBLICATIONS

"J. Prover: Integrating Connection-based Theorem Proving into Interactive Proof Assistants." Proc. International Joint Conference on Automated Reasoning, LNAI, Springer (June 2001). With S. Schmitt, L. Lorigo, and A. Nogin.

"An Experiment in Formal Design using Meta-properties." Proc. DARPA Information Survivability Conference and Exposition II (DISCEX 2001), IEEE Computer Society Press (June 2001). With M. Bickford, R. van Renesse, and R. Constable.

"Protocol Switching: Exploiting Meta-properties." Proc. International Workshop on Applied Reliable Group Communication (WARGC 2001), IEEE Computer Society Press (April 2001). With X. Liu, R. van Renesse, M. Bickford, and R. Constable.

"A Uniform Procedure for Converting Matrix Proofs into Sequent-style Systems." *Journal of Information and Computation* 162 (1-2):226-254 (2000). With S. Schmitt.

"Matrix-based Inductive Theorem Proving." Proc. International Conference TABLEAUX-2000, LNAI 1847, Springer (July 2000). With B. Pientka.

Carl Lagoze

Digital Library Scientist

lagoze@cs.cornell.edu
http://www.cs.cornell.edu/lagoze/lagoze.html
MSE Wang Institute

Our group investigates the policies, organization, and architecture of distributed information spaces. The Web, and the massive amount of content that it makes available to us in our daily lives, provides the backdrop for our work. The goal of our research is to understand and prototype the services and organizational structures that we can build on top of this global information base in order to increase its functionality, integrity, and ease of use. We undertake this research with the recognition that any proposed solutions must balance the economy and speed of automated solutions against the often-irreplaceable expertise that comes from human intervention.

Within this context we examine a number of research areas:

Architectures for storage of and access to the multiple forms of digital content.

Policies and enforcement mechanisms that facilitate the preservation and secure management of distributed content.

The role of metadata, in its many forms, in the management of digital content.

Protocols for federating information across distributed repositories and services.

Architectures and services for interlinking amongst document references and citations.

We work in close collaboration with other researchers in the Cornell, national and international library, computer science, and Internet communities. Our research model is highly applied: building standards and systems and supporting the deployment of them to our collaborators in the global information community. This research model has produced Dienst, an architecture and protocol for creating distributed document repositories and services, the Open Archives Initiative Metadata Harvesting Protocol that provides access to metadata in a variety of forms, and the

Program Committee: Third Conference on Recent Advances in Natural Language Processing (RANLP) (2001).

Co-organizer: Text Learning: Beyond Supervision, a workshop at the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI).

Invited Participant: Institute for Mathematics and its Applications (IMA) Workshop on Mathematical Foundations of Natural Language Modeling (2000).

Member: NSF review panel.

Technical coordinator: joint study agreement between Cornell University and IBM Watson on text mining and document management, 1999-present.

Affiliated faculty: Intelligent Information Systems Institute.

Referee: *ACM Transactions on Information Systems.*

UNIVERSITY ACTIVITIES

Department of Computer Science liaison to the Department of Linguistics.

Member: Women in Science and Engineering Advisory Group.

Member: Women in Science and Engineering Term Chair Award committee.

Reader: Arts and Sciences Undergraduate Admissions.

Member: Field of cognitive studies.

Affiliated faculty: Cornell Information Science program.

LECTURES

Distributional Similarity: Models and Methods. Carnegie Mellon University (August, 2000).

Weakly-supervised Statistical Segmentation of Japanese. WhizBang! Labs (August, 2000).

Natural Language Technology (three-hour tutorial). Rome Air Force Lab, Knowledge Representation and Reasoning series (October, 2000).

Applications of EM Techniques. Institute for Mathematics and its Applications workshop on Mathematical Foundations of Natural Language Modeling (November, 2000).

On the Effectiveness of the Skew Divergence for Statistical Language Analysis. Eighth Meeting of Artificial Intelligence and Statistics (January, 2001).

Distributional Similarity and the Skew Divergence. Highland Technologies (April, 2001).

The Iterative Residual Rescaling Algorithm: An Analysis and Generalization of Latent Semantic Indexing.

—. University of Maryland (April, 2001).

—. Columbia University (May, 2001).

PUBLICATIONS

"Iterative Residual Rescaling: An Analysis and Generalization of LSI." Proceedings24th Annual International Conference on Research and Development in Information Retrieval (SIGIR) (2001). With Rie Kubota Ando.

"On the Effectiveness of the Skew Divergence for Statistical Language Analysis." Proceedings Artificial Intelligence and Statistics (2001).

"Mostly-unsupervised Statistical Segmentation of Japanese: Applications to Kanji." Proceedings First Conference of the North American Chapter of the Association for Computational Linguistics (NAACL) (2000). With Rie Kubota Ando.

"Measures of Distributional Similarity." Proceedings 37th Annual Meeting of the Association for Computational Linguistics (ACL) (1999).

"Similarity-based Models of Word Co-occurrence Probabilities." *Machine Learning* 34 (1999). With Ido Dagan, and Fernando Pereira.

"Fast Context-free Parsing Requires Fast Boolean Matrix Multiplication." Proceedings 35th Annual Meeting of the Association for Computational Linguistics and 8th Conference of the European Chapter of the Association for Computational Linguistics (ACL/EACL) (1997).

"Distributional Clustering of English Words." Proceedings 31st Annual Meeting of the Association for Computational Linguistics (ACL)(1993). With Fernando Pereira and Naftali Tishby.

Professor Graeme Bailey

Yuying Li

Senior Research Associate

yuying@cs.cornell.edu
http://www.cs.cornell.edu/home/yuying/yuying.html
Ph.D. Waterloo University, 1988

My general research interests include numerical optimization and scientific computation. In addition, I am interested in the application of optimization methods to medical, engineering, and financial problems.

My current interests focus on solving nonlinear constrained problems for which the gradient computation is expensive and inexact. These problems arise from many financial application problems.

LECTURES

Lung Nodule Segmentation Using Optimization. The Workshop on Mathematics in Image Processing 2000, The University of Hong Kong (December 14-16, 2000).

Reconstructing the Unknown Local Volatility Function. A special session of mathematics/optimization for finance at Optimization 2001, Aveiro, Portugal (July 23-25, 2001).

PUBLICATIONS

"A Trust Region and Affine Scaling Interior Point Method for Nonconvex Minimization with Linear Inequality Constraints." *Mathematical Programming* Series A, 88(1):1-32 (2000).

Rajit Manohar

Assistant Professor
Member of the School of Electrical and Computer Engineering and the Graduate Field of Computer Science
rajit@csl.cornell.edu
http://vlsi.cornell.edu/~rajit/
Ph.D., California Institute of Technology, 1998

My research is concerned with the design of efficient asynchronous computation structures in VLSI and the use of formal methods to guarantee the correctness of such structures.

In work on formal methods, I have developed a new way to analyze the correctness of a class of program transformations commonly used in asynchronous VLSI synthesis. The technique "decompiles" the circuit into a higher level programming language, and provides high-level information about the effect of applying the transformation.

The amount of power required by a processor is quickly becoming a design constraint. In work on energy-efficient computation, I have developed a new adaptive number representation that significantly reduces the amount of energy required during instruction execution in a general-purpose processor.

The presence of precise exceptions in a processor can complicate its design. In work with Mika Nystrom and Alain J. Martin, I developed a simple, distributed implementation of an algorithm that implements precise exceptions in asynchronous processors, along with a proof of its correctness using program transformations. An instance of this mechanism was used in the design of an asynchronous MIPS processor.

AWARDS/HONORS

Sonny Yau '72 Excellence in Teaching Award (2001).
IEEE Teacher of the Year Award (2001).
NSF Faculty Early Career Development Award (2000-2004).
Tau Beta Pi and Cornell Society of Engineers Excellence in Teaching Award (2000).

UNIVERSITY ACTIVITIES

Member: Graduate Admissions Committee.

Member: Fields of Electrical Engineering; Computer Science; Applied Math.

PROFESSIONAL ACTIVITIES

Program Committee : IEEE/ACM Symposium on Advanced Research in Asynchronous Circuits and Systems (March, 2001).

Workgroup Organizer: NSF Workshop on Neuromorphic Engineering (July, 2000).

LECTURES

Width-adaptive Data Word Architectures. 2001 Conference on Advanced Research in VLSI (March, 2001).

An Analysis of Reshuffled Handshaking Expansions. IEEE/ACM Symposium on Asynchronous Circuits and Systems (March, 2001).

Low Energy Adaptive Processors. Computer Science Colloquium, Cornell University, Ithaca NY (October, 2000).

A Case for Asynchronous Computer Architecture. ISCA Workshop on Complexity-effective Design (June, 2000).

PUBLICATIONS

"Width-adaptive Data Word Architectures." Proc. 19th Conference on Advanced Research in VLSI, 112-129 (2001).

"Precise Exceptions in Asynchronous Processors." Proc. 19th Conference on Advanced Research in VLSI, 16-28 (2001). With Mika Nystrom, and Alain J. Martin.

"An Analysis of Reshuffled Handshaking Expansions." Proc. 7th IEEE/ACM Symposium on Asynchronous Circuits and Systems, 96-105 (2001).

"A Case for Asynchronous Computer Architecture." Proc. ISCA Workshop on Complexity-effective Design (2000).

"A Case for Asynchronous Active Memories." Proc. ISCA Workshop on Solving the Memory Wall (2000). With Mark Heinrich.

Greg Morrisett

Associate Professor

jgm@cs.cornell.edu

http://www.cs.cornell.edu/home/jgm/

Ph.D. Carnegie Mellon University, 1995

My research interests are in the design, semantics, and implementation of programming languages. I am particularly interested in exploring how language technology can be used to build reliable, secure, and high-performance systems software. The unifying thread for all my research is the application of advanced semantic constructs in real-world applications.

Recently, I have concentrated on type systems and logics for enforcing security properties in low-level code. One byproduct of this research is the design of a type system called TAL (Typed Assembly Language) for the Intel x86 architecture and a suite of tools that can efficiently type check binary code. The TAL type system is sufficiently expressive that we can efficiently compile a variety of high-level safe languages, such as ML, Scheme, Safe-C, and Java, to type-correct assembly code. This technology provides a means to safely extend systems such as kernels, web browsers, or hand-helds without the overheads of a virtual machine or just-in-time compiler. Recent work on TAL includes extending the type system to provide enforcement of other security policies, such as resource constraints, that are outside the scope of traditional type systems. I am also studying how other language technologies, such as binary rewriting and inlined reference monitors, can be combined with types and logics to support a wider class of important security properties.

Another interest is in language, compiler, and runtime support for application-specific memory management. Though standard garbage collection techniques provide a safe and convenient programming model, many systems applications cannot tolerate the overheads introduced by a general-purpose collector. My research here focuses on a range of topics from fast conservative garbage collectors to advanced type systems for region-based memory management, to generalizations of linear type systems. The goal in all of this work is to provide the programmer with more control over memory management without sacrificing safety.

My other recent interests are in run-time code generation and modal type systems, efficient data representation, and rich forms of polymorphism. Many of these issues are being explored in the context of Cyclone, a next-generation systems language that is being developed jointly between researchers at Cornell and AT&T Laboratories.

AWARDS

Allen Newell Medal for Research Excellence (2001).

Ralph Watts Excellence in Teaching Award (2001).

Presidential Early Career Award for Scientists and Engineers (2000).

NSF Faculty Early Career Development Award (1999).

Sloan Fellow (1998).

UNIVERSITY ACTIVITIES

Ph.D. admissions committee (1998-2000).

ECE hiring committee (2000).

PROFESSIONAL ACTIVITIES

Editor: *Journal of Functional Programming*.

Associate Editor: *ACM Transactions on Programming Languages and Systems*.

Member: IFIP Working Group 2.8 on Functional Programming.

Participant: INFOSEC Research Council Study Group on Malicious Code; DARPA ISAT Study on Mobile Code.

Program Committee: Symposium on Principles of Programming Languages (PoPL 2001); International Conference on Functional Programming (ICFP 2000); Principles and Practice of Declarative Programming (PPDP 2000); International Symposium on Memory Management (ISMM 2000); Principles of Programming Languages (PoPL 2002); Workshop on Semantics, Applications, and Implementation of Program Generation (SAIG 2001); Workshop on Multi-language Infrastructure and Interoperability (BABEL 2001).

LECTURES

Language-based Security. Danish Technical Institute (ITU), Copenhagen, Denmark (June, 2001).

Towards Next Generation Low-level Languages. University of Minnesota, Minneapolis, MN, (April, 2001); Harvard University, Cambridge, MA (February, 2001).

Next-generation Low-level Languages. Workshop on Semantics, Program Analysis, and Computing Environments for Memory Management, London, England (January, 2001).

Mobile Code Security: An Overview. TARA Review, Air Force Research Laboratory (March, 2000).

The Role of Type Systems in Mobile Code Security. DARPA ISAT Study Group on Mobile Code (January, 2000).

Mobile Code Security. Air Force Scientific Advisory Board (December, 1999).

Advanced Type Systems for Low-level Languages. OpenSIG Conference (October, 1999).

Why Languages and Compilers Matter. INFOSEC Research Council Study Group on Malicious Code (October, 1999).

PUBLICATIONS

"Syntactic Type Abstraction." *Transactions on Programming Languages and Systems* 22(6)1037-1080 (November, 2000). With D. Grossman, and S. Zdancewic.

"Attacking Malicious Code: A Report to the INFOSEC Research Council." *IEEE Software* 17(5) (September/October, 2000) With G. McGraw.

"Alias Types for Recursive Data Structures." In *ACM Workshop on Types in Compilation*, Montreal, Canada (September, 2000). With D. Walker.

"Scalable Certification for Typed Assembly Language." In *ACM Workshop on Types in Compilation*, Montreal, Canada (September, 2000). With D. Grossman.

"Typed Memory Management via Static Capabilities." *Transactions on Programming Languages and Systems* 22(4):701-771 (July, 2000). With D. Walker, and K. Crary.

"Alias Types." *European Symposium on Programming*, Berlin, Germany (March, 2000). With F. Smith, and D. Walker.

"Type Structure for Low-level Programming Languages." *1999 International Colloquium on Automata, Languages, and Programming* . With K. Crary.

"From System F to Typed Assembly Language." *ACM Transactions on Programming Languages and Systems* 21(3):528-569 (May 1999). With D. Walker, K. Crary, and N. Glew.

"Principals in Programming Languages: A Syntactic Proof Technique." In the *1999 International Conference on Functional Programming*, Paris, France, 197-207 (September, 1999). With S. Zdancewic, and D. Grossman.

CS Chair, Charles Van Loan and IAI Director, Fred B. Schneider

Andrew Myers

Assistant Professor

andru@cs.cornell.edu
http://www.cs.cornell.edu/andru
Ph.D. Massachusetts Institute of Technology, 1999

Semantic information about programs and data, obtained from the programming language level, provides leverage for addressing difficult problems in computer systems. Programming language ideas can be applied effectively to problems in security, systems, and databases. I am particularly interested in using language-level information to improve security guarantees, performance, and transparency for distributed systems and mobile code.

One example of this approach is our current work on the problem of protecting confidential data. Current trends are making this problem both more important and more difficult. Computer systems are nearly completely connected via the Internet, allowing software to disseminate private information to almost any location. In addition, we increasingly use untrusted software; for example, downloaded software such as applets. Standard access-control mechanisms are inadequate because they do not control information propagation.

Static information control is a promising approach for confidentiality. Our programming language Jif judiciously extends Java with privacy annotations that facilitate static analysis of information flows within programs. We are applying the Jif programming model to distributed systems, in which not only programs but also hosts may be untrusted. Computations spanning a network require a protocol for the communicating hosts; we can show through static analysis of such a distributed program when its protocol releases unintended information to untrusted hosts participating in the computation. To develop a theory for the security of such programs in the presence of downgrading channels, we have introduced the idea of 'robust declassification,' which captures the idea that untrusted hosts are unable to exploit these channels.

PROFESSIONAL ACTIVITIES

Program Committee Member: 2001 IEEE Symposium on Security and Privacy; 18th ACM Symposium on Operating Systems Principles (SOSP).

LECTURES

Protecting Confidentiality Against Untrusted Programs and Hosts. Mini-Workshop on Mobile Objects/Code and Security, University of Tokyo, Japan (October, 2000).

Enforcing Confidentiality in Low-level Programs. SDI/LCS SeminarFest, Carnegie Mellon University, Pittsburgh, PA (June, 2001).

PUBLICATIONS

"Robust Declassification." Proceedings of the 14th IEEE Computer Security Foundations Workshop, Cape Breton, Nova Scotia, Canada (June, 2001). With Steve Zdancewic.

"Secure Information Flow and CPS. Proceedings of the 10th European Symposium on Programming, Genova, Italy (April, 2001). With Steve Zdancewic.

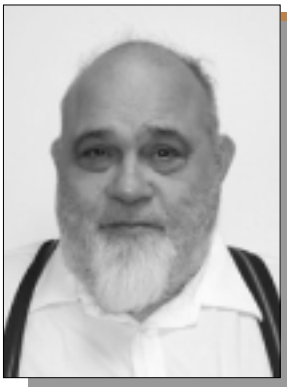

Anil Nerode

Professor
Member of the Department of Mathematics and the Graduate Field of Computer Science

anil@math.cornell.edu
www.math.cornell.edu/~anil/
Ph.D. University of Chicago, 1956

A decade in the works, the book *Automata Theory and its Applications*, by Bakhadyr Khoussainov and myself, was published in February, 2001 by Birkhauser (ISBN 0-8176-4027-2). This book makes available in one place material that is used all the time for decidability results in computer science. The basic theme is the correspondence between classes of automata and languages for finite automata, Buchi Automata, Rabin automata and the corresponding games and strategies for those games. It is suitable for a one semester graduate course or a two-semester undergraduate course. "Constructive Concurrent Dynamic Logic," with Duminda Wijesekera, was finished this year and will appear in the *Annals of Pure and Applied Logic* shortly. This was a project also of 10 years duration, primarily because of the number of cases involved in the intuitionistic treatment, which fortunately decreased from 120 to about 50, for the completeness theorem.

My principal project is a research monograph with engineer Wolf Kohn on the use of Finsler manifolds we associate with distributed optimal control problems throughout engineering to extract close-to-optimal controls in the form of finite automata. We dubbed this area Hybrid Systems in 1991, and many now work in it.

Kohn and I founded a research and development company, Hynomics, some years ago, with venture capital, in Seattle, and it is about to release to a client a prototype 25,000 agent distributed system for supply chain management, entirely based on new mathematical-computer science technology.

PUBLICATIONS

Foreword to: *Principles of Modeling and Asynchronous Distributed Simulation of Complex Systems* by S. Ghosh, IEEE Press (2000).

Keshav K. Pingali

Professor

pingali@cs.cornell.edu

http://www.cs.cornell.edu/Info/Projects/Bernoulli/

Ph.D. Massachusetts Institute of Technology, 1986

My research group works on programming languages and compiler technology for program understanding, restructuring, and optimization. Our goal is to develop the algorithms and tools that are required to raise the level of abstraction at which people program computers, freeing them from having to worry about low-level details of machine architectures, memory hierarchies, etc.

Our current focus is making programs adapt to hardware faults. This problem has been studied by the operating systems community, but proposed solutions such as message logging have a large run-time overhead. We are developing program analysis and transformation techniques to reduce this run-time overhead by exploiting information about program behavior that can be deduced at compile-time. We plan to deploy these techniques in the Adaptive Software Project, which is a multi-year, multi-institutional project funded by the NSF as part of its Information Technology Research (ITR) initiative.

We have continued our work on technology for optimizing the performance of programs running on machines with deep memory hierarchies. A recent break-through was the invention of fractal symbolic analysis, which is a powerful program analysis technique that permits compilers to restructure complicated programs far beyond the reach of conventional compilers that use dependence analysis. Traditional symbolic analysis is powerful but it is intractable for most programs. To circumvent this problem, fractal symbolic analysis analyzes a program and its transformed version by repeatedly simplifying these programs until symbolic analysis becomes tractable, ensuring that equality of simplified programs is sufficient to guarantee equality of the original programs. We have shown that this approach is adequate for restructuring codes like LU factorization with pivoting, which have not yielded to previous techniques in the literature. The fractal approach to analysis is likely to prove useful for proving other program properties.

We are continuing our work on next-generation generic programming techniques. In our approach, algorithm implementors use a different API than data structure designers, and the gap between these APIs is bridged by a compiler. One view of this approach is that it exploits restructuring compiler technology to perform a novel kind of template instantiation. We are demonstrating the usefulness of this new technology by deploying it in a system that generates efficient sparse codes from high-level algorithms and specifications of sparse matrix formats.

These ongoing projects build on our earlier work on restructuring compilation technology. Our group implemented one of the first compilers that generated code for distributed memory machines, starting from sequential shared memory programs. We introduced techniques called runtime resolution and owner-computes rule, which have now become standard in the area. Our work on linear loop transformations for enhancing parallelism and locality has been incorporated by Hewlett-Packard into its entire compiler product line. We also developed fast algorithms for program analysis problems such as computing the control dependence relation, the static single assignment form of a program, and dataflow analyses. Many of these algorithms have been incorporated into commercial and research compilers.

UNIVERSITY ACTIVITIES

Member: Cornell Theory Center Advisory Committee;

PROFESSIONAL ACTIVITIES

Program Committee: ACM Symposium on Programming Languages Design and Implementation (PLDI '01); International Conference on Supercomputing (ICS '01); EuroPar '01, Supercomputing 2001; HiPC 2001.

Consultant: Grammatech Inc.

Referee/Reviewer: *ACM TOPLAS; IEEE Trans. Computers; Journal of Parallel and Distributed Computing; Journal of Supercomputing; IEEE Computer; Software Practice and Experience*;

Editorial Board: *Int. Journal of Parallel Programming;* and *Discrete Mathematics and Theoretical Computer Science*.

LECTURES

Crash Recovery for Long-running Scientific Applications. IBM T. J. Watson Research Center, Yorktown Heights, NY (April, 2001).

Fractal Symbolic Analysis. Indian Institute of Technology, Kanpur (January, 2001).

Automatic Synthesis of Locality Enhancing Transformations for Imperfectly-nested Loop Nests. IBM T.J. Watson Research Center, Yorktown Heights, NY (November, 2000).

Synchronization on Blue Gene. IBM T.J. Watson Research Center, Yorktown Heights, NY (October, 2000).

PUBLICATIONS

"Tiling Imperfectly-nested Loop Nests." Supercomputing 2000. With Nawaaz Ahmed, and Nikolay Mateev.

"Automatic Generation of Block-recursive Codes." EUROPAR 2000, 125-134. With N. Ahmed.

"Left-looking to Right-looking and Vice Versa: An Application of Fractal Symbolic Analysis to Linear Algebra Code Restructuring." EUROPAR 2000, 155-164. With N. Mateev, and V. Menon.

"A Framework for Sparse Matrix Code Synthesis from High-level Specifications." Supercomputing 2000. With Nawaaz Ahmed, Nikolay Mateev, and Paul Stodghill.

"Landing CG on EARTH: A Case Study of Fine-grained Multithreading on an Evolutionary Path." Super computing 2000. With Kevin B. Theobald, Gagan Agrawal, Rishi Kumar, Gerd Heber, Guang R.Gao, and Paul Stodghill.

Robbert van Renesse

Senior Research Associate

rvr@cs.cornell.edu

http://www.cs.cornell.edu/home/rvr/

Ph.D. Vrije University, Amsterdam, 1989

I am currently interested in locating and aggregating data in a scalable manner. In the Astrolabe system, we construct a hierarchical distributed database. Only the leaves are writable. The internal nodes are generated at run-time using aggregation dictated by SQL queries that are installed on-the-fly. All internal communication is done using a peer-to-peer epidemic protocol. Astrolabe incorporates an hierarchical public key infrastructure for security.

My other interests include applying formal methods to systems problems and networking support for peer-to-peer applications.

PROFESSIONAL ACTIVITIES

Vice-president, Research: Reliable Network Solutions, Inc.

Technical Consultant to FAST.

PUBLICATIONS

"Scalable Fault-tolerant Aggregation in Large Process Groups." International Conference on Dependable Systems and Networks, Goteborg, Sweden (July, 2001). With Indranil Gupta, and Kenneth P. Birman.

"Spinglass, Scalable and Secure Communication Tools for Mission-critical Computing." DARPA Information Survivability Conference and Exposition II (DISCEX 2001), IEEE Computer Society Press (June, 2001). With Ken Birman and Werner Vogels.

"Protocol Switching: Exploiting Meta Properties." International Workshop on Applied Reliable Group Communication at the International Conference on Distributed Computing Systems (ICDCS), Phoenix, AZ (April, 2001). With Xioming Liu, Mark Bickford, Cristoph Kreitz, and Robert Constable.

"Using Epidemic Techniques for Building Ultra-scalable Reliable Communication Systems." Large Scale Networking Workshop: Research and Practice. Vienna, VA (March, 2001). With Werner Vogels, and Kenneth P. Birman.

"A Probabilistically Correct Leader Election Protocol for Large Groups." Proc. 14th International Symposium on Distributed Computing (DISC 2000) - *LNCS 1914*, 89-103, Toledo, Spain (October, 2000). With Indranil Gupta, and Kenneth P. Birman.

Mats Rooth

Professor
FCI, joint with Linguistics

mr249@cornell.edu
Ph.D. University of Massachusetts, 1995

Over the past several years, in collaboration with several colleagues, I have been developing an approach to research and applications in linguistics and computational linguistics that combines theoretical-linguistic formalisms, knowledge, and problem statements with numerical modeling and parameter estimation techniques. We have developed and implemented a grammatical formalism known as head-lexicalized probabilistic context free grammar and have employed it in large-scale experiments on German and English concerned with learning lexical information from text corpora. The methods have the status of a basic language technology, which can be applied in numerous ways in research and applications. My own interests relate mainly to scientific research in linguistics, particularly syntax, semantics and lexical semantics.

The work aims to devolop an experimental paradigm for work in linguistics in which lingustic events are observed in large text corpora using a robust parser and used to develop and test theories. Linguistic theory is placed in contact with observations on a massive scale using mixed numerical/symbolic comptutational models, and much relevant knowledge is learned from data using numerical optimization algorithms.

The linguistic representations assumed in this work are those of linguistic theory, for instance grammars, labeled trees and structured lexical entries. In the numerical modeling component, grammars and the lexicon are given a probabilistic interpretation, as defining a parameterized family of probability distributions. Many problems can then be solved using probabilistic optimization algorithms. For instance, inducing a lexicon from a text corpus amounts to maximizing a certain polynomial.

I also work on the semantics of human languages, using logical and denotational-semantic techniques. Currently, I am investigating connections between the syntax and semantics of elliptical constructions and the semantics and phonology of intonation.

UNIVERSITY ACTIVITIES

Administration of Computational Linguistics Lab.
Graduate Admissions Committee.
Phonetics Search Committee.
Founders' Committee, Faculty of Computing and Information.
Social Sciences Advisory Council.

PROFESSIONAL ACTIVITIES

2000 Workshop on Language Engineering, Johns Hopkins University. (Five-week workshop in which a reading comprehension prototype was implemented).
Editorial board: *Natural Language Semantics*.

LECTURES

Parse Forest Computation of Expected Governors. Colloquium presentation, University of Sussex (January, 2001).
Inducing a Semantically Annotated Lexicon. Symposium From Signals to Structured Communication, Cornell University (May 4 and 5, 2001).
Empty Domain Effects for Presuppositional and Non-presuppositional Determiners. Conference on Presupposition, Stuttgart (September, 2000).

PUBLICATIONS

"Parse Forest Computation of Expected Governors." Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics (2001). With Helmut Schmid.
"Inducing a Semantically Annotated Lexicon via EM-based Clustering." Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics (1999). With Stefan Riezler, Detlef Prescher, Glenn Carroll, and Franz Beil.

"Inside-outside Estimation of a Lexicalized PCFG for German." Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics (1999). With Franz Beil, Glenn Carroll, Detlef Prescher, and Stefan Riezler.

"Valence Induction with a Head-lexicalized PCFG." Proceedings of the 3rd Conference on Empirical Methods in Natural Language Processing (1998).

"On the Interface Principles for Intonational Focus." Proceedings of the 6th Conference on Semantics and Linguistic Theory (1996). With Glenn Carroll.

Fred B. Schneider

Professor
Director, Information Assurance Institute

fbs@cs.cornell.edu
http://www/cs.cornell.edu/People/fbs/
Ph.D. SUNY Stonybrook, 1978

Research this past year continued into techniques to support the construction of concurrent and distributed systems for high–integrity, mission–critical settings. Security and fault–tolerance are of paramount concern here.

Inlined reference monitors (IRMs) remain a promising approach for enforcing security policies, especially the fine–grained policies needed when the Principle of Least Privilege is employed to protect against hostile mobile code. This past year, we concentrated on putting the approach into practice. Cornell's Digital Library Project is looking to PSLang/PoET for imposing rights management and preservation policies on digital objects in their repository; we have prototyped a "digital university" to understand the enforcement issues there. And we started investigating the deployment of IRMs in the Windows operating system.

Our investigations into interactions between security and fault–tolerance (joint work with Lidong Zhou and Robbert van Renesse) also continued. We have deployed COCA, our replicated certification authority, to four sites on the Internet and made detailed performance measurements. The system supports on–line revalidation of name–key bindings and is designed to resist a broad collection of denial of service attacks.

UNIVERSITY ACTIVITIES

Sabbatical leave (2000-2001).
Founders Committee, Faculty of Computing and Information.
Engineering College Teaching Awards Committee.

PROFESSIONAL ACTIVITIES

Director: AFRL/Cornell Information Assurance Institute.
Editor: *Distributed Computing; Information Processing Letters; High Integrity Systems; Annals of Software Engineering; ACM Computing Surveys*.
Co–managing Editor: *Texts and Monographs in Computer Science*, Springer–Verlag.
Chairman: International Review of UK Computer Science Research.
Program committee: NORDSEC 2000 Fifth Nordic Workshop on Secure IT Systems—Encouraging Co–operation; 3rd Information Survivability Workshop (ISW–2000); Information/System Survivability Workshop 2001; 2001 Usenix Security Symposium; 2000 PODC Influential Paper Award.
Industrial Advisory Committees: JavaSoft Security Advisory Committee; deCode Genetics Security Advisory Board; Eweb University.Com Board of Advisors; JXTA Technical Advisory Council; CIGITAL Technical Advisory Board; Fast Search and Transfer Technical Advisory Board.
Other Advisory Committees: MITRE Corporation, Research and Technology committee of the Board of Directors; UK Dependability Interdisciplinary Research Collaboration (DIRC), Steering Committee.
IFIP Working Group 2.3 (Programming Methodology).

HONORS

Fellow, American Association for Advancement of Science (1992).
Fellow, Association for Computing Machinary (1994).
Professor–at–large, University of Tromsoe, Tromsoe, Norway (1996-2004).
Daniel M. Lazar Excellence in Teaching Award (2000).

LECTURES

The Case for Language Based Security. Invited Lecture. *Informatics—10 Years Back, 10 Years Ahead*. Saarland University, Saarbrucken, Germany (August, 2000).
In–lined Reference Monitors. Microsoft Research. Redmond, Washington (October, 2000).
Radio interview, "All Things Considered" (October 27, 2000).
The Case for Language Based Security. IFIP wg2.3. Santa

Cruz, CA (January, 2001).
The Design and Deployment of COCA. Distinguished lecture series. SUNY Stony Brook. Stony Brook, NY (February, 2001).
Fast P2P Possibilities. Tysil, Norway (February, 2001).
The Design and Deployment of COCA. Department of Computer Science. University of Tromso. Tromso, Norway (February, 2001).
The Case for Language Based Security. Keynote Address, ACM Southeast Conference 2001, Athens, GA (March, 2001).
The Design and Deployment of COCA. Department of Computer Science. University of Texas, Austin, TX (March, 2001).
The Case for Language Based Security. IBM Corporation Hawthorne, NY (April, 2001).
The Design and Deployment of COCA. AFOSR Principal Investigators Meeting. Ithaca, NY (May, 2001).
—. Information Assurance Institute Seminar Series. AFRL/IF Rome Research Site, Rome, New York (June, 2001).
The Case for Language Based Security. AFCEA Conference. Hamilton, NY (June, 2001).
Escaping the Ivy Tower: Transitioning Technology from a University. AFCEA Conference. Hamilton, NY (June, 2001).

PUBLICATIONS

"Enforceable Security Policies." *ACM Transactions on Information and System Security* 3(1):30-50 (February, 2000).
"Formalizations of Substitutions of Equals for Equals." *Millennial Perspectives in Computer Science*, Proceedings of the 1999 Oxford–Microsoft Symposium in honour of Professor Sir Antony Hoare, (Davies, Roscoe, and Woodcock, editors) Palgrave Publishers, Hampshire, England, 119-132 (November, 2000). With David Gries.
"Editorial: Time for Change." *Distributed Computing* 13(4):187 (November, 2000).
"A Language–based Approach to Security. " *Informatics—10 Years Back, 10 Years Ahead*. Lecture Notes in Computer Science 2000 (Reinhard Wilhelm, editor), Springer Verlag, Heidelberg, 86-101, (2000). And Greg Morrisett, Robert Harper.

David Schwartz
Assistant Professor

dis@cs.cornell.edu
http://www.cs.cornell.edu/dis/
Ph.D. State University of New York at Buffalo, 1999

My main area of research is in computational mechanics. My recent work focuses on applying mathematical techniques of deterministic uncertainty to enhance methods of structural engineering analysis. I am applying one such technique, called Interval Analysis, which adapts traditional numerical operations by replacing numbers with intervals that model uncertain values. This set-based form of structural analysis uses discrete mathematics to perform types of parametric studies, that traditional techniques cannot. However, the interval approach introduces numerical inaccuracies in solutions that map to infeasible structural behaviors. A major goal of this research involves improving the quality of the interval-based solutions to reduce these inaccuracies to produce a technique suitable for design engineers.

This year I restructured the introductory programming courses to update the technology and problem-solving tools being taught. For the first time, students were introduced to Maple in CS99, a preliminary course in programming that provides experience to students before they take other courses. MATLAB has also received greater focus in a new introductory programming course, CS100M. It has generated quite a bit of excitement from students and other universities because of the equal focus on both MATLAB and Java to help teach programming using problems of computational science. I have also been coordinating the Computer Science Department's "short courses," which are 1-credit courses taught by Ph.D. students. These courses focus on teaching computer languages and help the student instructors gain teaching experience. The instructors turn to me for advice on curriculum development, teaching methods, and navigating university policies. I have created a new short course, CS214: Advanced UNIX Tools, which will be taught for the first time in the next academic year.

Continuing my work from spring 2000, colleagues in the College of Engineering and I converted the pilot CS100 Academic Excellence Workshop (AEW) into a full-fledged program that runs in conjunction with CS100M. We continued to pilot AEWs for CS100J, which is the more traditional programming course. The AEW fosters cooperative learning, where students work in teams to solve challenging problems in a non-competitive environment. Students have responded enthusiastically, so we began designing a

new lab in spring 2001 to expand our capacity, especially to help the CS100J pilot grow. Since AEWs tend to entice students, especially women and minorities, we will be tracking engineering and computer science enrollments in the next academic year to measure the effectiveness of our AEW program in increasing diversity in the sciences. I look forward to working on the CS AEW with Daisy Fan, whom I recruited into the department.

UNIVERSITY ACTIVITIES

Faculty Advisor: Association of Computer Science Undergraduates.
Coordinator: Academic Excellence Workshop (AEW) for CS100 program.
Coordinator: departmental short-course advice and instruction.
Member: Student Experience Committee.
Member: Computing Policy Committee.

PROFESSIONAL ACTIVITIES

Textbook Reviewer: McGraw-Hill.

LECTURES

The Inexperienced Educator's Guide To Managing A Large Hierarchical Staff. Emerging Technologies for Industry and Education, Annual Meeting of the St. Lawrence Section, American Society Engineering Education (March, 2001).

Bart Selman

Associate Professor

selman@cs.cornell.edu
http://www.cs.cornell.edu/home/selman
Ph.D. University of Toronto, 1991

The focus of my research is on computation intensive methods in artificial intelligence, in particular fast general reasoning, search, and planning techniques. I also investigate the various sources of complexity in hard computational problems. In this work, I explore connections between computer science, artificial intelligence, and statistical physics. In addition, I study issues in problem representation, including the robustness of encodings, abstraction, compilation, and approximation methods. These issues are critical to the successful application of reasoning and search methods in realistic domains. In terms of applications, I consider challenge problems from planning, knowledge representation, multi-agent systems, and machine learning. Our planning system, Black Box, developed jointly with Henry Kautz of the University of Washington, is one of the fastest general purpose planning systems. Finally, in recent projects, I am exploring connections between machine learning methods and reasoning and planning techniques. In joint work with groups at the University of Washington and Microsoft Research, I study the use of Bayesian machine learning techniques for dynamic adaptive control of computational resources.

AWARDS

Fellow: American Association for Artificial Intelligence.
Alfred P. Sloan Research Fellowship (1999-2001).
NSF Faculty Early Career Development Award (1998-2002).

UNIVERSITY ACTIVITIES

Chair: Bits On Our Minds Science Fair.
Coordinator of the AI Seminar Series.
Cognitive Studies Undergraduate Committee.
Member: Fields of Cognitive Studies and Applied Mathematics.

PROFESSIONAL ACTIVITIES

Executive Council: American Association for Artificial Intelligence.
Editorial Board: *Constraints: An International Journal*; *Annals of Mathematics*; and *Artificial Intelligence*.
Advisory Board: *Journal of Artificial Intelligence Research* (*JAIR*).
Guest Editor: *Theoretical Computer Science*; *Discrete Applied Mathematics*; *Electronic Notes in Discrete Mathematics*; and *Annals of Mathematics and Artificial Intelligence*.
Organizer and program co-chair: SAT-2001: Workshop on Theory and Applications of Satisfiability Testing.
Program committee: AAI-2000 Workshop on Leveraging Probability and Uncertainty in Computation; 17th Intl. Conf. on Artificial Intelligence (IJCAI-2001); AAAI-2001 Symposium on Uncertainty in Computation.

Referee/Reviewer: *Artificial Intelligence Journal* (*AIJ*); *JACM Constraints: An International Journal*; *Journal of Artificial Intelligence Research* (*JAIR*); *NSF Journal of Automated Reasoning*; *Theoretical Computer Science*; *Science*.

LECTURES

Insights from Statistical Physics into Computational Complexity. Colloquium, California Institute of Technology (Caltech), Pasadena, CA (October, 2000).

Survey of Artificial Intelligence and Knowledge Representation. Six lectures at AFRL/IF, Rome, NY (October/November, 2000).

Understanding Complexity: Recent Developments and Directions. Colloquium, University of Minnesota, Computer Science, Minneapolis, MN (October, 2000).

Principled Analysis and Synthesis of Agent Systems. Meeting on Taskable Agent Software Kit, DARPA, Charlotte, NC (October, 2000).

Challenge Problems for Propositional Reasoning and Search. Distinguished Lecture Series, 7th Annual Knowledge Representation and Reasoning Lecture, York University, York, UK (November, 2000).

Understanding Complexity: Recent Developments and Directions. Colloquium, Leeds University, Leeds, UK (November, 2000).

Controlling Complexity: Structured Problems and Backbone Variables. Meeting on Autonomous Negotiation Teams, DARPA, Charlotte, NC (November, 2000).

Satisfiability Testing: Recent Developments and Challenge Problems Plenary Lecture, AAAI Spring Symposium on Answer Set Programming, Palo Alto, CA (March, 2001).

PUBLICATIONS

"A Bayesian Approach to Tackling Hard Computational Problems." Proc. 17th Conf. on Uncertainty and Artificial Intelligence (UAI-2001), Seattle, WA (2001). With E. Horvitz, Y. Ruan, C. Gomes, H. Kautz, and M. Chickering.

"Balance and Filtering in Structured Satisfiable Problems." Proc. 17th Intl. Conf. on Artificial Intelligence (IJCAI-2001), Seattle, WA (2001). With H. Kautz, Y. Ruan, D. Achlioptas, C. Gomes, and M. Stickel.

"Algorithm Portfolios." *Artificial Intelligence Journal* 126 (2001). With C. Gomes.

"Distribute Constraint Satisfaction in a Wireless Sensor Tracking System." Proc. Workshop on Distributed Constraint Reasoning (CONS-2), IJCAI-2001, Seattle (2001). With R. Bejar, C. Gomes, and B. Krishnamachari.

"Analysis of Random Noise and Random Walk Algorithms." Principles and Practice of Constraint Programming (CP 2000). *Lecture Notes in Computer Science* 894:278-290 (2000). With B. Krishnamachari, Xi Xie, and S. Wicker.

"Compute Intensive Methods in Artificial Intelligence." *Annals of Mathematics and Artificial Intelligence* 28(1) (2000).

"Heavy-tailed Phenomena in Satisfiability and Constraint Satisfaction Problems." *Journal of Automated Reasoning* 24(1/2):67-100 (2000). With C. Gomes, N. Crato, and Henry Kautz.

Associate Professor Bart Selman, Associate Professor Steve Vavasis, Chair Charles Van Loan, and CIS Dean Robert L. Constable

Professors Hartmanis and Birman

David B. Shmoys

Professor and Member of the School of Operations Research and the Graduate Field of Computer Science

shmoys@cs.cornell.edu

http://www.cs.cornell.edu/home/shmoys/shmoys.html

Ph.D. University of California, Berkeley, 1984

The primary focus of my research is on the design and analysis of efficient algorithms for discrete optimization problems, and in particular, on approximation algorithms for NP-hard problems.

Computational complexity theory provides a mathematical foundation for the intractability of many computational problems by proving that all NP-complete problems are equally hard. However, in practice, real-world inputs for some NP-hard optimization problems are straightforward to solve, whereas for others, even quite modestly sized inputs are beyond the limits of the most sophisticated methods. Analogously, from a theoretical perspective, for some NP-hard optimization problems it is possible to efficiently compute solutions that are guaranteed to be arbitrarily close to optimal, whereas for others, computing even a crude approximation to the optimum is also NP-hard. In fact, the extent to which such approximation algorithms exist for a problem provides a surprisingly accurate theoretical yardstick for its actual computational difficulty.

Our work has been motivated by the fact that certain linear programming relaxations have been shown to provide extremely good lower bounds on typical data. We provided a theoretical understanding of the strength of these bounds by designing algorithms that "round" the fractional solutions to these linear programs to nearby integer solutions without degrading the quality of the solution too much. We have been investigating a variety of clustering problems, and in joint work with Moses Charikar, Sudipto Guha, and Eva Tardos, we have obtained the first constant factor approximation algorithm for the k-median problem. We have also obtained improved approximation algorithms for the uncapacitated facility location problem (in joint work with Fabian Chudak) and for a variety of more general multi-level facility location problems (in joint work with Karen Aardal and F. Chudak, and with Nathan Edwards). For the former problem, not only is the theoretical performance of our algorithm surprisingly good, but it is more effective in practice than all previously known heuristic procedures for this problem. We have also been investigating the application of some of these rounding methods to problems arising in computational genomics, in joint work with Dan Brown and the group in plant sciences at Cornell, working under Steve Tanksley. The resulting software for computing genetic linkage maps has already seen widespread application within the genomics community.

UNIVERSITY ACTIVITIES

Member: Search Committee for Dean of the Graduate School.
Member: Genomics Task Force.
Member: FCI Working Group on Computational Biology.
Member: FCI Working Group on Computational Science.

PROFESSIONAL ACTIVITIES

Editor-in-chief: *SIAM Journal on Discrete Mathematics*.
Editor: *SIAM Journal on Computing*.
Co-editor: *SIAM/MPS Series on Optimization*.
Associate Editor: *Mathematics of Operations Research; Mathematical Programming; Journal of Scheduling*.
Guest Editor: *J. of Algorithms* (Special issue devoted to selected papers of the 11th Annual Symposium on Discrete Algorithms).

LECTURES

Approximation Algorithms for Facility Location Problems.
—. Invited plenary lecture. CO 2000, Greenwich (July, 2000).
—. Invited plenary lecture. CONF 2000, Saarbrucken, Germany, September 2000.
A Constant-factor Approximation Algorithm for the K-median Problem. International Symposium on Mathematical Programming, Atlanta, GA (August, 2000).
Selective Mapping: A Strategy for Optimizing the Construction of High-density Linkage Maps.
—. Tri-Institutional Workshop on Computational Biology, Cornell (September, 2000).
—. NIH Symposium - From Genes to Proteins and to Biological Function: Computational Approaches, Cornell (October, 2000).

PUBLICATIONS

"Selective Mapping: A Strategy for Optimizing the Construction of High-density Linkage Maps." *Genetics* 155:407-420 (2000). With T. J. Vision, D. G. Brown, R. T. Durrett, and S. D. Tanksley.
"Approximation Algorithms for Facility Location Problems." In: *Approximation Algorithms for Combinatorial Optimization*, K. Jansen and S. Khuller, editors, APPROX 2000, *Lecture Notes in Computer Science* 1913:27-33, Springer, Berlin (2000).

LANDMARK PUBLICATIONS

"Using Dual Approximation Algorithms for Scheduling Problems: Theoretical and Practical Results." *J. Assoc. Comput. Mach.* 34:144-162 (1987). With D.S. Hochbaum.

"An Approximation Algorithm for the Generalized Assignment Problem." *Math. Programming* 62:461-474 (1993). With E. Tardos.

"Approximation Algorithms for Fractional Packing and Covering Problems." *Math. Oper. Res.* 20, 257-301 (1995). With S.A. Plotkin and E. Tardos.

"Scheduling to Minimize the Average Completion Time: On-line and Off-line Approximation Algorithms." *Math. Oper. Res.* 22:513-544 (1997). With L. A. Hall, A. S. Schulz, and J. Wein.

"A Constant-factor Approximation Algorithm for the K-median Problem." Proceedings of the 31st Annual ACM Symposium on Theory of Computing, 1-10 (1999). With M. Charikar, S. Guha, and E. Tardos.

Emin Gün Sirer

Assistant Professor

egs@cs.cornell.edu

http://www.cs.cornell.edu/People/egs/

Ph.D. University of Washington, 2001

I am broadly interested in the design and implementation of distributed operating systems for modern networks.

My main focus is on operating system support for ad-hoc mobile networks. Ad-hoc networking is a newly emerging computing paradigm where disparate, mobile hosts form temporary alliances to accomplish certain tasks. They naturally arise in many real-world settings including disaster relief operations, smart vehicles, and distributed sensor networks. Such ad-hoc settings, characterized by constant change and limited resources, require a high degree of adaptability and resource awareness from applications. Current state of the art, however, requires all such functionality to be manually encoded by the application programmer.

We are currently building an operating system, called Magnetos, that provides automatic and transparent distribution of ad-hoc networking applications. Magnetos makes an entire ad-hoc network appear as a single Java virtual machine to applications. By transparently migrating application components from node to node, Magnetos can increase system longevity, reduce application delays, and decrease bandwidth consumption, even in the presence of conflicting applications. Consequently, Magnetos simplifies application design, increases application portability by obviating error-prone manual coding, and achieves an efficient placement of application components within the network that adapt to changes in resource availability and connectivity.

In conjunction with Magnetos, we are examining the design of secure, energy efficient operating systems for small, embedded computers. Newly emerging smart cards, which are capable of downloading and executing arbitrary code, are becoming part of our critical infrastructure as they are being deployed in financial and telecommunications networks. We are looking at different strategies for testing their security, as well as evaluating code verification techniques for such severely resource-constrained systems.

A third research direction looks at distributed systems at the opposite end of the spectrum; namely, the high-performance clusters that make up web sites. We are developing techniques for formally specifying and reasoning about the interfaces they expose to the rest of the network, as well as techniques for enforcing a security policy uniformly across a web site.

LECTURES

Distributed Virtual Machines: A New System Architecture for Networked Computers (March-May, 2000).

A New System Architecture for Ubiquitous Computing. EGSA Graduate Engineering Social Seminar Series (March, 2001).

PUBLICATIONS

"Comprehensive Synchronization Elimination for Java." *Science of Computer Programming*, special issue on Static Analysis (April, 2001). With Jonathan Aldrich, Craig Chambers, and Susan Eggers.

"Design and Implementation of a Distributed Virtual Machine for Networked Computers. Proceedings of the Seventeenth Symposium on Operating Systems Principles, 202-216, Kiawah Island, SC (December, 1999).

With Robert Grimm, Arthur J. Gregory, and Brian N. Bershad.

"Extensibility, Safety and Performance in the SPIN Operating System. Proceedings of the Fifteenth Symposium on Operating Systems Principles, 267-284, Copper Mountain, CO (December, 1995). With Brian N. Bershad, Stefan Savage, Przemyslaw Pardyak, Marc Fiuczynski, David Becker, Craig Chambers, and Susan Eggers.

Evan Speight

Assistant Professor
Member of the School of Electrical and Computer Engineering and the Graduate Field of Computer Science
espeight@csl.cornell.edu
http://www-ds.csl.cornell.edu/~espeight/
Ph.D., Rice University, 1998

My research centers around different aspects of distributed computing. In particular, my three areas of focus are software runtime systems for distributed computing platforms, improving the performance of commodity cluster-based shared memory systems, and providing location-independent data access through the concept of affinity directed mobility.

My work in software distributed system runtime development has resulted in the dissemination of the Brazos system to over 100 registered users worldwide. Brazos is a high performance parallel programming environment distinguished by its use of multithreading, selective multicast, a software-only implementation of scope consistency, and several adaptive runtime performance tuning mechanisms. Brazos supports both shared memory and message passing programming styles, and provides very efficient mechanisms for thread migration and checkpoint/recovery. Currently, my research on parallel programming runtime systems includes the development a version of the Message Passing Interface (MPI) that provides thread migration between nodes in a cluster for load balancing, fault tolerance, and higher performance.

The emergence of ubiquitous communication infrastructure and high performance, low power computing resources challenges us to explore a better alternative to the current fragmentation of data, applications, and devices that many users are faced with today. The Bifrost location independent computing project seeks to provide a flexible and comprehensive information access environment. The function of Bifrost is to provide location and device independent access to data. Data in Bifrost encompasses both information and the applications used to manipulate that information. Bifrost uses *affinity* between data, and between users and data, to make decisions about when and where to move data. We refer to this approach as *affinity directed mobility*.

The core research issues of the project are mobility management (how we move data and threads to support user and device mobility), data management (how we represent, access, update, and protect information), and application management (how we provide system-wide access to application data). In contrast to previous approaches to mobile data management that employ *hoarding* in anticipation of disconnection, Bifrost anticipates "connection elsewhere" instead of disconnection. Our rationale for this "almost always connected" approach is as follows: We project that in 3-5 years, the situation will exist in which users will be able to be in an "always connected" state, connected almost anytime and anywhere that they choose. Wireless access points (e.g., 802.11, infrared, and Bluetooth) will be common, even ubiquitous, in nearly every home and public place, as they are today in many situations including coffee shops, airports, and office buildings. Additionally, services such as Infared and Bluetooth will allow connection to the Internet in nearly any setting from almost anywhere. In this situation, the problem of *how* a user can have unilateral access to any piece of remote personal data, regardless of device or location, becomes increasingly important. Thus Bifrost seeks to provide a common data access model, and exploring the design space while providing a rich set of tools for the manipulation and sharing of personal data is a core focus of the proposed research.

UNIVERSITY OR PROFESSIONAL ACTIVITIES

Member: ECE Curriculum and Standard Committee.

Member: Faculty Advisory Board on Information Technologies.

Reviewer: Transactions of Parallel and Distributed Computing.

LECTURES

WSDLite: A Lightweight Alternative to Windows Sockets Direct Path. Fourth Windows Symposium, Seattle, WA (August, 2000).

Efficient Parallel Computing on Multiprocessor Clusters. EE Colloquium, Cornell University (March, 2000).

PUBLICATIONS

"WSDLite: A Lightweight Alternative to Windows Sockets Direct Path." Proceedings of the 4th Usenix Windows Symposium (August, 2000). With John Bennett, and Hazim Abdel-Shafi.

"Efficient User-level Thread Migration and Checkpointing on Windows NT Clusters." In *P*roceedings of the 3rd Usenix Windows NT Symposium (July, 1999). With Hazim Abdel-Shafi and John Bennett.

"Realizing the Performance Potential of the Virtual Interface Architecture." Proceedings of the International Conference on Super computing (ICS) (June, 1999). With John Bennett.

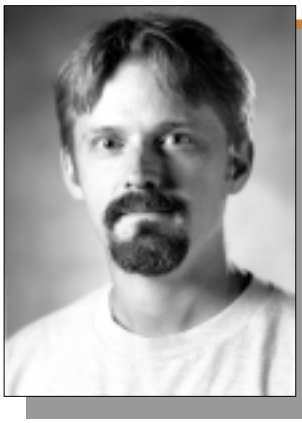

Paul Stodghill

Research Associate

stodghil@cs.cornell.edu
http://www.cs.cornell.edu/stodghil/
Ph.D. Cornell University, 1997

My research interest is in program transformations and program synthesis for computational science applications.

Increasingly, the frontiers of science are being explored using computer simulation and modeling. Creating these scientific applications poses a number of difficulties for computational scientists. First, these applications are often very complex and may contain hundreds of thousands, and even millions of lines of code. Furthermore, these applications are intended to run on the most advanced, high performance supercomputers. I am interested in developing software that helps to reduce the development cost of these applications. One way to do this is to raise the level of abstraction of programming languages and software libraries in order to bring them closer to the programmer's natural domain of discourse. In order to do this, specialized compilers and other program transformation systems are required in order to transform these abstractions into efficient executable programs.

An example of this is the work that I have done on sparse compilers with Keshav Pingali and his students. Matrix computations, such as those that perform linear algebra operations, are very naturally expressed using Fortran or Matlab style loops and arrays. However, many matrices that arise in scientific applications are sparse, which means that they consist almost entirely (>99%) of zero values. Modifying matrix algorithms to use special-purpose sparse matrix data structures is a complicated and error-prone process. We have developed software that allows the programmer to express their matrix computations and sparse matrix data structures separately and in their most natural representation. Then, we use program transformations to combine the algorithms and data structures into a single, efficient executable.

Another area in which I am working is that of fault-tolerance for scientific applications. Very few scientific applications would be considered "mission-critical," but many do run for very long periods of time (e.g., for weeks or months) or in environments in which failures are likely to occur (e.g., the desktop workstations within a department). We are currently investigating a number of different approaches for providing fault-tolerance for scientific applications, but one thing is already clear: different applications running on different computers call for different fault-tolerance solutions. A programmer would certainly like to have their application run in many different computing environments, but interfacing with a number of fault-tolerant systems is not usually realistic. In order to provide a single interface from the programmer's point of view, we are investigating a number of implementation techniques involving software library and program transformation technology. This work is being done as part of a project on Adaptive Software with researchers from other departments here at Cornell, Mississippi State University, the College of William and Mary, and other institutions.

Apart from my core research agenda, I have worked closely with computational scientists in order to develop novel, high-performance, scientific applications. Apart from being interesting in its own right, this work has given me insight into scientific programming that I have found extremely valuable in my research on program transformations. The most important project that I have been involved with was the Crack Propagation for Teraflop Computers (CPTC) project, whose goal was to develop fracture mechanics software that incorporated a number of recent advances in numerical analysis and computational geometry and which delivered very high performance on teraflop-scale computers. Other researches who worked on this project include, Keshav Pingali, Steve Vavasis, Paul Chew, the Cornell Fracture

Group headed by Tony Ingraffea (Civil Engineering), Gao Guong-Rong (ECE, University of Delaware) and Nikos Chrisochoides (CS, College of William and Mary).

LECTURES

Next-generation Generic Programming and its Application to Sparse Matrix Computations. ACM International Conference on Supercomputing, Sante Fe, NM (May, 2000).

Parallel FEM Simulation of Crack Propagation on the AC3 Velocity Cluster. ACM Second Workshop on Cluster Cluster-Based Computing, Sante Fe, NM (May, 2000).

Crack Propagation on Teraflop Computers. AC3 Meeting, Cornell University (June, 2000).

PUBLICATIONS

"A Framework for Sparse Matrix Code Synthesis from High-level Specifications." In *Supercomputing 2000*, Dallas, TX (November 4-11, 2000). With N. Ahmed, N. Mateev, and K. Pingali.

"Landing CG on EARTH: A Case Study of Fine-gained Multithreading on an Evolutionary Path." In Supercomputing 2000, Dallas, TX (November 4-11, 2000). With K. Theobald, G. Agrawal, R. Kumar, G. Heber, G. Gao, and K. Pingali.

"Next-generation Generic Programming and its Application to Sparse Matrix Computations." In International Conference on Supercomputing, 2000. With N. Mateev, K. Pingali, and V. Kotlyar.

"Parallel FEM Simulation of Crack Propagation on the AC3 Velocity Cluster." In The Second Workshop on Cluster Cluster-Based Computing, 2000. With G. Coulouris, G. Heber, D. Lifka, K. Pingali , D. Schneider, P. Wawrzynek, and J. Zollweg.

"Parallel FEM Simulation of Crack Propagation – Challenges, Status, and Perspectives." Irregular 2000. With B. Carter, et al.

Éva Tardos

Professor

eva@cs.cornell.edu
http://www.cs.cornell.edu/home/eva/eva.html
Ph.D. Eötvös University, Hungary, 1984

My research interest focuses on the design and analysis of efficient methods for combinatorial optimization problems and their applications to various fields. I am mostly working on problems that involve graphs or networks.

One general area of my research is designing fast algorithms that provide provably close-to-optimal results for NP-hard problems. Although research on polynomial time approximation algorithms started in the1970s soon after the discovery of NP-completeness, it has truly blossomed only in the past decade. Amazing progress has occurred both in our ability to design approximation algorithms, and in proving limits to approximability. Over the last years I have been working on different approximation algorithms on various basic combinatorial problems. I have worked on algorithms for various cut problems, and clustering type problems. These problems are motivated by applications arising in vision, networking and clustering.

I am also working on the interface of algorithms and game theory. Approximation algorithms provide a tool for understanding issues in game theory, such as evaluating and designing multi-agent games. Such games underlie many phenomena in networking. In joint work with Tim Roughgarden, I worked on understanding the quality of routing obtained if separate agents each make selfish routing decisions in a global network. Our goal is to understand the tradeoffs between introducing a global control mechanism, and the loss in quality obtained by letting each agent make selfish decisions. We use a model where each link in a network has a delay that is an arbitrary monotone and continuous function of the amount of flow on the link. We compare the quality of such a Nash equilibrium with a globally planned optimal routing whose goal is to minimize the sum of all delays. In the case of linear delay functions, we show that the global optimum can be at most a factor of four-thirds better than the Nash equilibrium. For more general delay functions, the value of the global optimum can be arbitrarily better than the Nash equilibrium even when the delay functions are rather simple. On the other hand, we show that for any monotone and continuous delay function the cost of a Nash equilibrium is at most as much as the optimal for the case in which each agent has to send twice as much flow. Roughly speaking, this shows that one can eliminate the need for introducing a global control mechanism at the cost of designing a network that can support twice as much flow.

AWARDS

Fellow: American Academy of Arts and Sciences (2001).
Fellow: ACM (1998).
Faculty of the Year: Association of Computer Science Undergraduates.

UNIVERSITY ACTIVITIES

Director of Graduate Studies: Computer Science Department, Cornell University
Member: Fields of Operations Research and Applied Mathematics
Member: FACTA, Faculty Advisory Committee on Tenure and Appointment.
Member: WISE, Advisory group on Women in Science and Engineering.

PROFESSIONAL ACTIVITIES

DIMACS External Advisory Board member, since 1996.
Co-organizer: DIMACS special year on Computational Intractability in1999-2001.
Area editor for Discrete Optimization: *Mathematics of Operations Research.*
Editor: *SIAM Journal of Computer Science*; *Chicago Journal of Theoretical Computer Science*; *Combinatorica*; *Journal of Interconnection Networks.*
Program committee: International Workshop on Randomization and Approximation Techniques in Computer Science (APPROX), 2000; ACM-SIAM Symposium on Discrete Algorithms (SODA) 2001; ACM Symposium on the Theory of Computing (STOC) 2001.

LECTURES

A Classification Problem Related to Multi-way Cuts. International Symposium on Mathematical Programming, Atlanta GA (August, 2000).
How Bad is Selfish Routing?
—. International Symposium on Mathematical Programming, Atlanta GA (August, 2000).
—. Cornell University, Department of Mathematics, VIGRE Interdisciplinary Colloquium (September, 2001).
—. Cornell University, Computer Science Department, Distinguished Lecture Series (October, 2000).
—. 41st Annual IEEE Symposium on the Foundations of Computer Science, Redondo Beach, CA (November, 2000).
—. University of Southern California, Computer Science Department, Distinguished Lecture Series (November, 2000).
Flow-based Algorithms for Some Metric Labeling Problems. International Symposium on Mathematical Programming, Atlanta GA (August, 2000).
A Constant-factor Approximation Algorithm for the K-median Problem. International Symposium on Mathematical Programming, Atlanta GA (August, 2000).
Classification with Pair-wise Relationships (invited talk). Horizons in Combinatorics, A Conference on Graph Theory, Vanderbilt, TN (May, 2001).
Sequence of four lectures on Approximation Algorithms. DONET summer school on Integer and Combinatorial Optimization, Utrecht, Netherlands (June, 2001).

PUBLICATIONS

"How Bad is Selfish Routing?" Proceedings of the 41st Annual IEEE Symposium on the Foundations of Computer Science (October, 2000). With Tim Roughgarden.
"Allocating Bandwidth for Bursty Connection." *SIAM Journal on Computing* 30(1):191-217 (February, 2001). With Jon Kleinberg, and Yuval Rabani.

LANDMARK PUBLICATIONS

"Approximation Algorithms for Classification Problems with Pair-wise Relationships: Metric Partitioning and Markov Random Fields." In the Proceedings of the 40th Annual IEEE Symposium on the Foundations of Computer Science, 14-23 (November, 1999). With Jon Kleinberg.
"A Constant-factor Approximation Algorithm for the K-median Problem. In the Proceedings of the 31st Annual ACM Symposium on the Theory of Computing, 1-10 (May, 1999). With Moses Charikar, Sudipto Guha, and David Shmoys.
"The Quickest Transshipment Problem." *Mathematics of Operations Research* 36-62 (February, 2000). With Bruce Hoppe.
Fast Approximation Algorithms for Multicommodity Flow Problems. *Journal of Computer and System Sciences* 50:228-243 (1995). With Tom Leighton, Fillia Makedon, Serge Plotkin, Cliff Stein, and Spyros Tragoudas.
"A Strongly Polynomial Algorithm for the Minimum Cost Circulation Problem." *Combinatorica* 5:247-255 (1985).

Tim Teitelbaum

Associate Professor

tt@cs.cornell.edu

http://www.cs.cornell.edu/Info/People/tt/Tim_Teitelbaum.html

Ph.D. Carnegie Mellon University, 1975

My research is concerned with the use of fine-grain dependence graphs for specification, development, and analysis of software and hardware systems. The objective is a new generation of tools that provide precise and complete information about the structure of complex systems. I am working to improve the performance and functionality of generic dependence-graph technology, and I am also exploring the use of the technology in various application domains.

Dependence-graph technology can be used in a program understanding system, where the graphs may include forward and backward links between each assignment statement and possible uses of the values stored by that assignment. Pointer analysis can be used so that indirect loads and stores through pointers are taken into account, as well as indirect function calls. Dataflow analysis can be used so that links between unrelated assignments and uses are excluded. Operations that highlight forward and backward slices show the impact of a given statement on the rest of the program (forward slicing) and the impact of the rest of a program on a given statement (backward slicing). Operations that highlight paths between nodes in the dependence graph (chops) show ways in which the program points are interdependent (or independent).

Uses of slicing and chopping include software development, maintenance and re-engineering of legacy code, test-data generation, security-assurance and safety-assurance inspection, and semantic interference checking in configuration management systems.

I am working with Ph.D. student Lyn Millett, who is studying program dependence-graphs and slicing of concurrent programs.

[On leave spring 2001.]

PROFESSIONAL ACTIVITIES

Co-founder and Chairman: GrammaTech, Inc.

Member: Science and Technology Study Group, Infosec Research Council (October 1999-May 2000).

LECTURES

Static-semantic analysis based on dependence graphs. SPAWAR, San Diego, CA (August, 1999).

—. Hewlett-Packard, Rancho Bernardo, CA (October, 1999).

PUBLICATIONS

"Issues in Slicing Promela and its Applications to Model Checking, Protocol Understanding, and Simulation." *International Journal on Software Tools for Technology Transfer* 2(4):343-349 (2000). With L. Millett.

"Program Slicing of Hardware Description Languages." In 10th IFIP WG10.5 Advanced Research Working Conference on Correct Hardware Design and Verification Methods (CHARM '99), Bad Herrenalb, Germany (September, 1999). With E.M. Clarke, M. Fujita, P.S. Rajan, T. Reps, and S. Shankar.

"A Case for Channel Analysis: Slicing Promela." Proceedings of the International Symposium on Software Engineering for Parallel and Distributed Systems, Los Angeles, CA, 52-61 (May, 1999). With L. Millett.

Associate Professor Claire Cardie with Theory Center Director Tom Coleman

Herbert Van de Sompel

Assistant Professor

herbertv@cs.cornell.edu
http://www.cs.cornell.edu/people/herbertv/
Ph.D. Ghent University, 2000

During the academic year 2000/2001, I focused on furthering my work in two areas of digital library research that I initiated in 1999: the Open Archives Initiative and open reference linking (OpenURL).

With the Open Archives Initiative, important progress was made through a generalization of the metadata harvesting specifications that were released as the Santa Fe Convention in early 2000. Those specifications were only applicable to preprint-related metadata, but nevertheless attracted the interest from communities outside of the preprint realm. It was decided to revise the specification in depth, in order to extend its applicability to metadata in general. The revision process took 5 months. In that timeframe, a meeting was organized to discuss issues involved in the generalization; several versions of a new protocol document were compiled; an alpha-testing group consisting of 14 international parties was assembled with the aim of testing and refining those versions. In January 2001, the Metadata Harvesting Protocol of the Open Archives Initiative was publicly released at a meeting in Washington DC. Carl Lagoze (Cornell University CS) and I were appointed Executives of the Open Archives Initiative, overseeing and coordinating the technical activities of the Initiative. The Digital Library Federation and the Coalition for Networked Information provide support for those activities. The Metadata Harvesting protocol attracts broad international attention, and several funding agencies explicitly or implicitly call for proposals that build on the protocol.

Important progress was also made in the realm of the open reference-linking track of my work. With Patrick Hochstenbach (Ghent University) and Oren Beit-Arie (Ex Libris USA), I had publicly released a draft OpenURL specification in January 2000. The OpenURL specification enables open reference linking by:

Allowing one to reference a scholarly object—as well as elements that describe the context in which the reference is provided—in a consistent way as *name=value* pairs on an HTTP GET or POST URL.

Allowing for the transfer of that reference—as well as of the contextual elements—to a linking server.

In December 2001, I filed for standardization of the OpenURL with NISO. Soon thereafter, the standardization request was granted, and a formal standardization committee was formed. Meanwhile, the draft OpenURL specification has already gained broad acceptance in the scholarly information industry, with leading companies such as ISI, EBSCO, Swets, SilverPlatter, etc. supporting it in their production systems. Also, an important OpenURL-based prototype was set up with DOI/CrossRef to demonstrate the feasibility of open reference linking based on DOIs.

PROFESSIONAL ACTIVITIES

Member: Steering Committee, Open Archives Initiative.
Executive: Open Archives Initiative.
Technical Committee: Open Archives Initiative.
Research Advisory Board: OCLC.
NISO Committee AX (OpenURL standardization)
Co-organizer: Open Archives Initiative Meeting, Washington DC (January, 2001).
Workshop on the Open Archives Initiative and peer-review journals in Europe, Geneva (March, 2001).

LECTURES

Interconnecting Distributed Scholarly Information Resources in a Context Sensitive Manner. Invited half-day workshop, Access 2000 conference, St. John's, Newfoundland, Canada (September, 2000).

The SFX Framework, the OpenURL and the Open Archives Initiative. Invited presentation, SLA Global 2000 conference, Brighton UK (October, 2000).

Invited presentation in the session "Digital Libraries and their Role in Knowledge Dissemination and Creation." ASIS 2000 Conference, Chicago IL (November, 2000).

Invited presentation in the session "Electronic Pre-print Initiatives: A Discussion on Comparative, Historical and Emerging Trends". At the ASIS 2000 Conference, Chicago IL (November, 2000).

The Roof is on Fire. Closing keynote at the fall 2000 meeting of the Coalition for Networked Information. San Antonio, TX (December, 2000).

The Open Archives Initiative. Keynote at the Workshop on the Open Archives Initiative and peer-review journals in Europe, Geneva (March, 2001).

The Open Archives Initiative and Scholarly Communication. Invited presentation, IATUL Conference, Delft (May, 2001).

PUBLICATIONS

The Open Archives Protocol for Metadata Harvesting, Herbert Van de Sompel and Carl Lagoze, editors (January, 2001). http://www.openarchives.org/OAI/openarchivesprotocol.htm.

"The Open Archives Initiative: Building a Low-barrier Interoperability Framework." JCDL2001 (2001). With C. Lagoze. http://www.cs.cornell.edu/lagoze/papers/oai-jcdl.pdf.

"Open Linking in the Scholarly Information Environment using the OpenURL Framework." *D-Lib Magazine* 7(3)(March, 2001). With Oren Beit-Arie. http://www.dlib.org/dlib/march01/vandesompel/03vandesompel.htm.

"Generalizing the OpenURL Framework beyond References to Scholarly Works." *D-Lib Magazine* 7(7/8)(July/August, 2001). With Oren Beit-Arie.

Charles Van Loan

Professor and Chair,
Department of Computer Science

cv@cs.cornell.edu
http://www.cs.cornell.edu/cv/
Ph.D. University of Michigan, Ann Arbor, 1973

I continue to work in the computational multilinear algebra area. This includes factorization approaches to various fast transforms, Kronecker product preconditioners, and Kronecker-constrained least squares problems.

This past year Adam Florence and I perfected our implementation of the fast Gauss transform (FGT) of Greengard and Strain. We also developed effective methods for total least squares and weighted least squares when the data matrix is a Kronecker product.

Applied Mathematics student Carla Martin and I completed work on a solver for linear systems of the form $(A-qI)x = b$ where A is the product of upper triangular matrices. Our method is an order of magnitude faster than the best of previous techniques.

UNIVERSITY OR PROFESSIONAL ACTIVITIES

Chair: Department of Computer Science.
Director of Undergraduate Studies: Department of Computer Science.
Member: FCI Founders.
Member: Core Curriculum Governing Board (Engineering).

LECTURES

The Ubiquitous Kronecker Product. Invited Lecture, SIAM Linear Algebra Meeting, Raleigh NC (October, 2000).

PUBLICATIONS

"The Ubiquitous Kronecker Product." *Journal of Computational and Applied Mathematics* 123:85-100 (2000).

"GEMM-based Level 3 BLAS: Algorithms for the Model Implementations." *ACM Transactions on Mathematical Software* 24:268-302 (1999). With P. Ling and B. Kagstrom.

Computational Frameworks for the Fast Fourier Transform, 273pp., SIAM Publications, Philadelphia, PA. (1992).

Matrix Computations, 3d ed., 694pp, Johns Hopkins University Press, Baltimore, MD (1996). With G.H. Golub.

Introduction to Scientific Computation: A Matrix-vector Approach Using Matlab, 2d Ed., 365pp, Prentice-Hall, Upper Saddle River, NJ (1999).

Graduate Studies Director Eva Tardos with Professor Johannes Gehrke, and Chairman Charles Van Loan at 2001 graduation ceremony

Stephen A. Vavasis
Associate Professor

vavasis@cs.cornell.edu
http://www.cs.cornell.edu/home/vavasis/
Ph.D. Stanford University, 1989

As computer hardware becomes more powerful, there is a corresponding growth in the demand for more efficient algorithms to solve large-scale scientific problems. My research is on the design and analysis of such algorithms. Two Ph.D. students, V. Howle and G. Jonsson (both of the Center for Applied Mathematics), completed their Ph.D.s working with me during the past year. Howle's thesis was on algorithms for modeling and simulation of AC electric power networks. Utility companies are interested in modeling the behavior of the network in the presence of a fault (closed circuit breaker). The governing equations, called the "swing equations," are nonlinear differential algebraic equations for the rotor angles of the generators in the system. We developed new, more accurate algorithms for the swing equations by solving a linear algebraic subproblem (complex-weighted least squares) more accurately. Work on geometry in scientific computing continues. Jonsson's thesis considers the problem of robust intersection of parametric patches with rays and planes. This problem arises in geometric modeling and mesh generation. Our results show that the problem can be solved accurately by transforming it to a generalized eigenvalue computation using the theory of resultants and other algebraic techniques.

Four new students have begun to work with me on other problems in scientific computing. We are looking at the following topics: the effect of domain shape on solution to boundary value problems, nonlinear diffusion equations, mesh generation for moving objects such as a heart, and the protein model-fitting problem of x-ray crystallography.

UNIVERSITY ACTIVITIES

Member: Graduate admissions committee, Applied Mathematics.
Graduate admissions committee, Computer Science.
Faculty Senate.
University Hearing Board.
BOOM 2001 Vice-chair.

PROFESSIONAL ACTIVITIES

Editorial Boards: *Journal of Global Optimization; SIAM Journal Matrix Analysis and Applications; SIAM Review; Math. Program.*
Vice-chair: SIAM Activity Group on Linear Algebra.
Referee: *SIAM J. Optimization, SIAM J. Matrix An. App., Linear Alg. App., Computational Geometry Theory and App.; Math. Progr.*; International Meshing Roundtable; *J. Complexity.*; NSF Panelist, Numeric, Symbolic and Geometric Computing Program.

LECTURES

Sparse Matrices and Automatic Differentiation in Semidefinite Programming, International Conference on Advances in Convex Analysis and Global Optimization, Samos, Greece (June 6, 2000).
—. 17th International Symposium on Mathematical Programming, Atlanta, GA (August 7, 2000).

PUBLICATIONS

"Quality Mesh Generation in Higher Dimensions." *SIAM J. Comput.* 29:1334-1370 (2000). With S. A. Mitchell.
"Accurate Solution of Weighted Least Squares by Iterative Methods." *SIAM J. Matrix An. App*, 22:1153-1174 (2001). With E. Y. Bobrovnikova.

Computer Science faculty on retreat

Werner Vogels

Research Associate

vogels@cs.cornell.edu
http://www.cs.cornell.edu/vogels
Ingenieur Hogere Informatica,
Haagse Hogeschool, 1989

My research explores the impact of scale on reliable distributed systems. The main focus is on the development of new network protocols and middleware, as well as on novel strategies to structure applications and support systems.

In the context of the Spinglass project, I am collaborating with Ken Birman and Robbert van Renesse on the development of a new generation of high-scalable reliable network protocols based on the principles of epidemic information dissemination. Although the research has already resulted in protocols for reliable multicast and group membership and failure detection, there are still many open questions such as the application of epidemic techniques to congestion control for many-to-many multicast protocols.

In the Galaxy project, the focus is on the distributed systems needs of enterprise cluster computing systems. In particular, I investigate the scalability problems that arise in these systems and try to provide solutions in a form that is directly applicable to current practical problems. Currently I am looking at the scalable cluster problems from three angles: foremost I try to address cluster management scaling problems by developing a framework for managing complete cluster farms with many different styles of cluster computing present in the farm. Secondly, I am investigating some of the issues that arise when providing support for the development of complex cluster aware applications. Most recently, I have started to address some of the more complex cluster system structuring issues such as meta-clusters and geographical distributed clusters.

I also remain active in the field of high-performance cluster communication. In the past, I collaborated with Thorsten von Eicken on building high-performance user-level network interfaces, which eventually resulted in the VIA industry standard for user-level network interfaces. More recently, I have started investigating the usage patterns of VIA based devices in large-scale clusters, especially with respect to hot-spot control and high-performance flow control.

PROFESSIONAL ACTIVITIES

Steering Committee: Usenix Windows Systems Symposium.

Program Committee: Usenix 4th Windows Systems Symposium, IEEE International Conference on Cluster Computing—Cluster 2000; 2001 IEEE Symposium on Applications and the Internet; International SRDS Workshop on Dependable System Middleware and Group Communication; Usenix 6th Conference on Object Oriented Tools and Systems (Program Committee); IEEE 2nd Workshop on Internet Applications—WIAPP'01; IFIP/ACM International Conference on Distributed Systems Platforms - Middleware 2001.

Program Chair: Distributed Systems Track of 2001IEEE Symposium on Applications and the Internet.

General Chair: IEEE 2nd Workshop on Internet Applications—WIAPP'01.

LECTURES

Cluster Computing Made Easy: New Tools for Scalable Servers and Services." Invited Lecture, Annual meeting of Advanced Cluster Computing Consortium (June 2, 2000).

PUBLICATIONS

"Spinglass, Scalable and Secure Communication Tools for Mission-critical Computing." Proceeedings of the 2001 DARPA Information Survivability Conference and Exhibition – II (June, 2001). With K. Birman, and R. van Renesse.

"Using Epidemic Techniques for Building Ultra-scalable Reliable Communication Systems." Proceedings of the Large Scale Networking Workshop: Research and Practice, Vienna, VA (March, 2001). With K. Birman and R. van Renesse.

"An Overview of the Galaxy Management Framework for Scalable Enterprise Cluster Computing." Proceedings of the IEEE International Conference on Cluster Computing: Cluster-2000, Chemnitz, Germany (December, 2000). With D. Dumitriu.

"Tree-saturation in the AC3 Velocity Cluster Interconnect." Proceedings of the 8th conference on Hot Interconnects, Stanford, CA (August, 2000). With D. Follett, J. Hsieh, D. Lifka, and D. Stern.

Golan Yona

Assistant Professor

golan@cs.cornell.edu
http://www.cs.cornell.edu/golan
Ph.D. Hebrew University, Jerusalem, 1999

My research focuses on computational molecular biology, with an emphasis on developing tools and methodologies for large-scale analysis of protein sequences and structures.

The goal of my research is to explore high-order organization within the space of all proteins and obtain a global view (a "road map") of the protein space. We hope that the global view will yield valuable insights about the nature and function of new genes and will lead to the discovery of high-level properties and principles in the protein space.

This interdisciplinary research is rooted in two different disciplines: computer science and molecular biology. Being on the borderline between the two disciplines, this study is related to fields of intensive research in both. It incorporates study and development of methods for metric embedding, unsupervised learning techniques, efficient graph algorithms and parallel applications, and efficient database management. On the computational biology side, it is involved with development of new algorithms and approaches for protein comparison, statistical models of protein families and study of the mapping from sequences to structures. A great emphasis is on developing novel machine-learning based techniques, both in the context of the study of the protein space and as general-purpose tools.

My study so far has resulted in two large databases that are being used by biologists to study new genes: *ProtoMap*: this database and its interactive web sites (http://protomap.cornell.edu and http://protomap.stanford.edu) are based on a graph-theoretic approach for the large-scale organization of protein sequences. The site offers an automatic hierarchical classification of all known protein sequences. It is a useful resource for the analysis of known as well as new protein sequences and for the study of relationships between protein families.

BioSpace: This database and its web sites (http://biospace.cornell.edu and http://biospace.stanford.edu) are based on a new unified framework for sequence and structure analysis. The result of this analysis is a preliminary global map of the protein space. The sites also provide all-atom three-dimensional models for over 160,000 proteins.

AWARDS

Burroughs-Welcome Fellowship from the Program in Mathematics and Molecular Biology (PMMB) (1998-2000).

Noted for Excellency in teaching, Faculty of Mathematics and Natural Sciences, Hebrew University, Jerusalem, Israel (1998).

Faculty of Mathematics and Natural Sciences Research award, Hebrew University, Jerusalem, Israel (1998).

Jewish National KKL Foundation Research award (1996).

Intel-Dean Prize (1996).

LECTURES

Towards a Complete Map of the Protein Space based on a Unified Sequence and Structure Analysis of all Known Proteins. NIH Symposium: From Genes to Proteins to Biological Function, Ithaca, NY (October, 2000).

Towards a Complete Map of the Protein Space based on a Unified Sequence and Structure Analysis of all Known Proteins. The Eighth International Conference on Intelligent Systems for Molecular Biology (ISMB), San Diego, CA (August, 2000).

A Unified Sequence-structure Classification of Protein Sequences: Combining Sequence and Structure in a Map of the Protein Space. Fourth Annual International Conference on Computational Molecular Biology (RECOMB), Tokyo, Japan (April, 2000).

A Unified Sequence-structure Classification of Protein Sequences: Combining Sequence and Structure in a Map of the Protein Space. Quantitative Challenges in the Post-genomic Sequence Era, San Diego, CA (January, 2000).

Methodologies for Target Selection in Structural Genomics. Innovative Computational Applications, San Francisco, CA (October, 1999).

Modeling Protein Families using Probabilistic Suffix Trees. Third Annual International Conference on Computational Molecular Biology (RECOMB), Lyon, France (April, 1999).

A Map of the Protein Space—An Automatic Hierarchical Classification of all Known Proteins. The Sixth International Conference on Intelligent Systems for Molecular Biology (ISMB), Montreal, Canada (June,1998).

PUBLICATIONS

"Within the Twilight Zone: A Sensitive Profile-profile Comparison Tool based on Information Theory." *Journal of Molecular Biology* (2001). With M. Levitt.

"Variations on Probabilistic Suffix Trees: Statistical Modeling and Prediction of Protein Families. *Bioinformatics* 17:23-43 (2001). With G. Bejerano.

"Towards a Complete Map of the Protein Space based on a Unified Sequence and Structure Analysis of All Known Proteins. In the *Proceedings of ISMB 2000*, 395-406, AAAI Press. With M. Levitt.

"A New Nonparametric Pairwise Clustering Algorithm based on Iterative Estimation of Distance Profiles. *Machine Learning* (2001). With S. Dubnov, R. El-Yaniv, Y. Gdalyahu, E. Schneidman, and N. Tishby.

"Methodologies for Target Selection in Structural Genomics. *Progress in Biophysical and Molecular Biology* 73:297-320 (2000). With M. Linial.

"A Unified Sequence-structure Classification of Protein Sequences: Combining Sequence and Structure in a Map of the Protein Space. *Proceedings of RECOMB 2000*, 308-317, ACM press. With M. Levitt.

"Comparison of Protein Sequences and Practical Database Searching. In *BioInformatics: Sequence, Structure, and Databanks*, edited by D. Higgins and W. Taylor. Oxford University Press (2000). With S. Brenner.

LANDMARK PUBLICATIONS

"ProtoMap: Automatic Classification of Protein Sequences, a Hierarchy of Protein Families, and Local Maps of the Protein Space. *Proteins: Structure, Function and Genetics* 37:360-378 (1999). With N. Linial, and M. Linial.

"Global Self-organization of All Known Protein Sequences Reveals Inherent Biological Signatures. *Journal of Molecular Biology* 268:539-556 (1997). With M. Linial, N, Linial, and N. Tishby.

Ramin Zabih

Associate Professor

rdz@cs.cornell.edu

http://www.cs.cornell.edu/rdz/

Ph.D. Stanford University, 1994

My research interests lie in computer vision and in medical imaging. I have worked on a variety of problems in early vision, including motion and stereo; many of these problems can be solved very accurately using algorithms based on graph cuts. In the last year I have been doing research on medical imaging in the Radiology Department at Cornell Medical School, where I have focused on improving the quality of MRIs. I have also investigated a number of applications of computer vision, including new methods for content-based access to databases of images, and have also developed some simple computer vision techniques to automate program debugging at Microsoft.

LECTURES

Fast Approximate Energy Minimization via Graph Cuts. University of Washington CS Colloquium (January, 2001).

PUBLICATIONS

"Visual Correspondence with Occlusions via Graph Cuts." Proc. International Conference on Computer Vision (ICCV) (July, 2001). With Vladimir Kolmogorov.

"An Efficient Real-time Navigator Algorithm: Motion Organized Simultaneous Acquisition with Interactive Control (MOSAIC)." Proc. IEEE Conference on Medical Imaging and Augmented Reality, Hong Kong (June, 2001). With Vladimir Kolmogorov, Yi Wang, Richard Watts, and Martin Prince.

"Exact Voxel Occupancy with Graph Cuts." Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2000). With Dan Snow, and Paul Viola.

LANDMARK PUBLICATIONS

"Color-spatial Indexing and Applications." *International Journal of Computer Vision* 35(3) (1999). With Jing Huang, S. Ravi Kumar, Mandar Mitra, and Wei-Jing Zhu.

"Fast Approximate Energy Minimization via Graph Cuts." Proc. International Conference on Computer Vision (ICCV) (1999). With Yuri Boykov, and Olga Veksler.

"Markov Random Fields with Efficient Approximations." Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (1998). With Yuri Boykov, and Olga Veksler.

"A Variable Window Approach to Early Vision." IEEE Transactions on Pattern Analysis and Machine Intelligence 20(12) (1998). With Yuri Boykov, and Olga Veksler.

Bits On Our Mind (BOOM) 2001

The following article was written by Melissa Globerman for the Cornell Chronicle. BOOM 2001 was co-sponsored by ECE.

On February 28, BOOM (Bits On Our Mind) 2001 came to the Engineering Quad for the fourth year, and the rumble could be felt strongly on three stories of Upson Hall. Forty-seven computer science–related exhibits highlighted student talent from across the university to show current research and new applications in digital technology. Projects ran the gamut from an instrument to transmit images of Mars in color for NASA's Athena mission, scheduled to launch in 2003, to an artificial intelligence prototype robot named Max, whose motor control is directed from the parallel port of an on-board laptop computer.

BOOM is a showcase of creativity from undergraduates and graduate students' work from engineering, fine arts, psychology, space sciences, and—of course—computer science. "Every year, we've become more interdisciplinary, but that's also because the interest in the event has grown tremendously all over campus," said Charles Van Loan, chair of the Department of Computer Science.

Two benefactors, alumnus Philip Young '62, B.M.E. '63, and the Microsoft Corporation, donated approximately $9,000 to the event to add swanky catering to the ambience and free t-shirts for the participants. Though Van Loan said it would be impossible to guess how many of the projects were already being developed for commercial production, scouts from corporations were invited to mingle with the students and browse their projects.

One project that garnered interest at the event will undoubtedly get more attention on East Hill later this spring. Three computer science undergraduates, working with the Human Computer Interaction research group in the Department of Communications, have created a Cornell campus tour guide that fits in the palm of a hand. Taking advantage of a campus pilot project in wireless networking for which eight buildings have been equipped with wireless transceivers, the students created an interactive messaging and mapping program. Using a personal digital assistant (PDA) with a global positioning system receiver attached, visitors can take a self-guided tour of the campus.

President Hunter R. Rawlings and students at BOOM 01

The PDA, which visitors will be able to rent at Day Hall, uses GPS to determine the user's location and then makes a beeping sound to alert them when information related to that location is available. "But what's cool about this program is that it creates social maps, so the user can then type in what they think about each of the various spots, rate them, and other users can then see those results," said Jenna Burrell, a senior in computer science and one of the exhibit's creators.

Other better known projects on campus also debuted their latest research. Members of Cornell's RoboCup team gave demonstrations of how their new omni-directional bots were able to capture their second robot soccer world championship title in Melbourne, Australia last summer.
The Cornell Hybrid Electric Vehicle team displayed photos of their most recent gem: the 2000 Chevrolet Suburban that they've converted to use a large battery pack as its energy storage buffer. The team, in its eighth year of competition, is almost ready for their second Future Truck competition.

Also, undergraduates from the Sibley School of Mechanical and Aerospace Engineering displayed the moonbuggy that they are readying for a NASA competition in Huntsville, Alabama this spring. A team of eight students constructed the vehicle—a four-wheel drive bicycle—for NASA's simulated lunar crater racing course. Junior Brett Lee explained that while the moonbuggy itself doesn't appear to be a digital technology project, the steering and suspension geometry necessary to create its most essential parts are assuredly "the products of high tech computer design and programs."

Whether the project is about designing an on-line coursework submission system, creating a program to instruct a computer to teach itself to play backgammon or checkers, or designing software to recognize human facial expressions, "this kind of forum gives the students a great opportunity to explain what it is they did and also get feedback on their work," said BOOM faculty advisor and associate professor of computer science Stephen Vavasis.

Learning opportunities notwithstanding, for participants and observers alike, BOOM was a blast of entertainment and a resounding success.

Corporate Interactions

Corporate Partnerships and Affiliations

The department's Corporate Partnerships and Affiliates programs support research collaborations and interactions with internationally respected scientists on a number of levels. Our programs are also designed to support education of our undergraduate and graduate students. The department offers opportunities for interaction and the support of education in such vital areas as:

- artificial intelligence
- bio-informatics
- computational methods for mechanical design and simulation
- computational molecular biology
- digital libraries
- distributed computing and fault tolerance
- formal specification and verification methodologies
- graphics
- information technology
- natural language, document classification, and retrieval
- networking databases
- parallel computing
- programming languages
- programming logics
- remote collaboration technologies
- scientific and numerical computing
- security
- supercomputing (through affiliation with CTC)
- theoretical computer science
- vision and image interpretation

Our Corporate Partnership program fosters strategic relationships with major sponsors of the department's activities. Corporate Partners are invited to participate directly in the technology development process, through on-campus representation, short term and extended visits, and consulting arrangements. Additional opportunities include access to technical reports, colloquia, seminars, the department's annual report, and resumes submitted by BA, BS, MEng, and Ph.D. candidates expecting to graduate. Our Corporate Affiliates program offers small companies many flexible arrangements to interact with our faculty and students. Such arrangements are designed on an individual basis to suit the goals of the companies and the department.

Department of Computer Science faculty and researchers continue to collaborate with corporate partners: Microsoft, Intel, GTE, Lockheed Martin, Lucent Technologies, IBM, McGraw Hill, and Sun Microsystems. GTE continued its support of the department's initiative in information technology with a five-year grant to support new faculty hires in this strategic area. Lockheed Martin provided support to the undergraduate and Ph.D. programs. Sun Microsystems made generous donations of equipment to the department. Lucent Technologies donated funds for a fellowship for Professor Jon Kleinberg's graduate student Amit Kumar.

Professor Johanne Gehrke continues to receive support from several companies for his research in new data-mining techniques. He received support from McGraw Hill to develop teaching materials and from Mastercard Corp. in support of his work in e-commerce. IBM gave professor Gehrke a Faculty Partnership award. Microsoft also awarded him $50,000 in support of his efforts in Online Data-mining Operators and to facilitate and foster collaboration between his research group and its Bay Area Research Center.

Professor E. Gun Sirer received funding from several sources to support his teaching and research efforts. Microsoft contributed equipment valued at $20,000 to be used for class projects by students in the spring 2001 CS414/415 course. In addition, Microsoft has committed to three years of support for his research in Assuring the Security of Components in the .NET Framework.

Intel provided major funding to the department for research and instruction. Gifts included $55,800 to fund a Graduate Fellowship for the academic year awarded to Stephanie Weirich, $2551 awarded to Professor Fred Schneider for equipment in support of his research in secure Information Assurance native code and $152,000 in equipment from Intel's Technology for Education 2000 grant.

Microsoft continued its generous support for research, instruction, and general support. Gifts included $75,000 to Werner Vogels's research in scalable enterprise cluster computing and to facilitate and foster collaboration between his group and its Bay Area Research Center. Microsoft also supported our migration to Win2000 and our student technology demonstration day, BOOM, with gifts totaling $42,000. In addition, it made gifts of Microsoft Office XP Professional software valued at more than $359,000.

Industrial Partners

AT&T
Hewlett Packard
IBM Corporation
Intel Corporation
Lockheed Martin
Microsoft
Sun Microsystems
Verizon

These industrial partnerships are recognized as a vital part of life in this department. We remain grateful for their ongoing support of our research and instructional activities.

Inquiries about corporate partnerships and affiliations may be addressed to:

Marcy E. Rosenkrantz, Coordinator of Corporate Relations
Department of Computer Science
4114 Upson Hall
Cornell University
Ithaca, NY 14853-7501
Telephone: (607) 255-9197 / Fax: (607) 255-4428
email: corprel@cs.cornell.edu

Corporate Gifts and Grants

The department is grateful for the support, including equipment and software, provided by our industrial partners.

Air Products & Chemicals	$2,500
Battery Management Corporation	1,000
Green Hills Software	1,800
IBM	40,000
Intel Corporation	58,351
Lockheed Martin	15,000
Lucent Technologies	45,000
McGraw Hill	11,000
Microsoft	551,400
Teradyne, Inc.	1,000
Verizon/GTE	50,000

The department is grateful for gifts from the following individuals:

- Mary Salton, a gift of $5000 for the Gerald Salton Reading Room in Upson Hall

- Phillip Young, a gift of $1000 to support BOOM, our student technology showcase.

The Office of Computing and Information Science wishes to acknowledge support from the following corporate partners:

Intel	$200,000
MasterCard International	77,114

University President Hunter Rawlings and CS students view the moonbuggy project at Boom 01.

Alumni and External Relations

This year marked a significant turning point in the history of alumni relations for computer science and computing related fields at Cornell. Alumni are networking in record numbers, and many have become actively involved with Cornell's efforts to support the development of Computer and Information Science across a wide range of disciplines.

Successful events were held in Boston and Palo Alto, which primarily included large gatherings of alumni and friends of the Department of Computer Science. The Boston event was a sponsored dinner hosted by the Dean of CIS, Robert L. Constable. Nearly three dozen alumni and friends attended. Corporate gifts were facilitated by **Neeraj Agrawal '95** of **Battery Management Corporation** and **Dan Proskauer '90** of **Teradyne.**

In Palo Alto over 50 alumni and friends attended an evening mixer that included Professor Dan Huttenlocher and **Charles Weiss '66** of **Oracle** (pictured below). Not only are these folks on the West Coast doing well, in spite of fluctuations in the market, they have an infectious optimism and justifiable excitement about things that are yet to come.

Charles Weiss '66 of Oracle and Professor Dan Huttenlocher.

At this year's alumni reunion breakfast in June, alumni and guests joined Computer Science Department Chairman Charlie Van Loan, and CIS Dean, Robert Constable. Among those attending was **George Joblove '76, MS '79**, special effects wizard and now senior vice president of technology at **Sony Pictures Imageworks**.

Cornell alum **Philip Young '62** made a substantial gift to support this year's student technology fair. The fair, which is called BOOM, was held in February, and featured 47 project presentations by students from Cornell. (See page 74 for more on BOOM 2001.)

Nikola Valerjev '96 was again instrumental in obtaining sponsorship from his company, **Green Hills Software**, for two Cornell programming teams to represent Cornell in the ACM Programming contests during the fall and spring semesters.

The **Degenfelder Family Scholarship** was awarded to **Marla Leahy '02**. It awards $5,000 to a student working at the boundary between computer science and biology. **Joseph R. Degenfelder '60** and his wife **Dr. Pauline Degenfelder '61** worked with Professor Ron Elber to establish an endowment for this special award.

Mary Salton, the wife of the late **Professor Gerard Salton**, made a generous gift to establish the Professor Gerard Salton Seminar and Reading Room. Gerry's ground- breaking research in information retrieval helped to found the current broad field of Information Science. The Department of Computer Science is pleased to honor his contribution. The Salton Reading Room, located on the 5th floor of Upson Hall, was dedicated in a ceremony on April 18th, 2001. In attendance to honor the event were Mrs. Mary Salton, wife of the late Professor Salton, and members of his family, as well as a number of friends and colleagues. Professor Edgar Rosenberg gave a moving tribute to the late professor, as did Professor Juris Hartmanis and Dean Robert Constable. The ceremonies ended with the cutting of the ribbon, by Chairman Charles Van Loan.

For more information about alumni or external relations in CIS or the Department of Computer Science, please contact Dan Jenkins at jenkins@cs.cornell.edu.

The late Gerard Salton.

Educational Statistics

Undergraduate Courses

COURSE NUMBER, TITLE, AND CREDIT HOURS	FALL 2000 INSTRUCTOR	STUD.	HOURS	SPRING 2001 INSTRUCTOR	STUD.	HOURS	TOTAL STUD.	TOTAL HOURS
099 Fund Prog Concepts - 2	Schwartz	33	66				33	66
100m Intro to Computer Prog. - 4	Schwartz	191	764	Yan	195	780	386	1544
100j Intro to Computer Prog. - 4	Teitelbaum	298	1192	Schwartz	196	784	494	1976
101 Intro to Cog. Science - 4	Spivey	10	40				10	40
113 Introduction to C -1	Vetsikas	32	32	Chen	31	31	63	63
114 Unix Tools - 1	Nogin	83	83	Holland-Minkley	57	57	140	140
130 Creating Web Doc. -3	Bailey	244	732				244	732
150 Gt. Ideas fr Comp Sci -1	Bailey/VanLoan	82	82				82	82
201 Cog Sci in Context - 4	B. Halpern	2	8				2	8
202 Transition to Java -1	Ezick	20	20	Marques	21	21	41	41
211 Comp & Prog - 3	Chew	285	855	Bailey	270	810	555	1665
213 C++ Programming - 2	Dinapoli	44	88	Dinapoli	36	72	80	160
221 Num. Mthds in CMB - 3	Elber	9	27				9	27
222 Intro Sci Computation - 3				Vavasis	198	594	198	594
230 Intermed Web Design - 3				Bailey	37	111	37	111
280 Discrete Structure - 4	Kleinberg	211	844	Artemov	167	668	378	1512
312 S & I of Comp Prog 4	Morrisett	107	428	Morrisett	116	464	223	892
314 Intro to DS& CO - 4	Yan	91	364	Manohar	149	596	240	960
381 Intro to Thry of Comp - 4				Artemov	153	612	153	612
409 Data Struct & Alg - 4				Halpern	15	60	15	60
411 Prog. Lang. & Logics - 4				Demers	28	112	28	112
412 Intr to Compilers & Trans - 3				Myers	32	96	32	96
413 Pract in Compilers & Trans-2				Myers	31	62	31	62
414 Sys Prog & Oper Sys - 3	Mosse	242	726	Sirer	64	192	306	918
415 Pract in Oper Sys - 2	Mosse	60	120	Sirer	65	130	125	250
417 Computer Graphics - 3				Greenberg	109	327	109	327
418 Pract Comp Graphics-2				Bala	82	164	82	164
421 Numerical Analysis -4	Vavasis	25	100				25	100
424 Comp Linguistics- 4				Rooth	5	20	5	20
430 Info Discovery -3				Arms	73	219	73	219
432 Intro Database Sys-3	Gehrke	141	423				141	423
433 Pract Database Sys -2	Gehrke	93	186				93	186
472 Foundations of AI -3	Cardie	130	390				130	390
473 Practicum in AI -2	Cardie	56	112				56	112
474 Intro Nat Lang Proc -4				Rooth	16	64	16	64
478 Machine Learning - 3				Yona	56	168	56	168
481 Intro. Thry Comp – Hnrs-4	Lee	37	148				37	148
482 Intro Analy of Alg -4				Tardos	220	880	220	880
483 Quantum Info Proc -2	Mermin	16	32				16	32
486 Applied Logic-4				Constable	16	64	16	64
490 Indp Study 1-4	Staff	30	99	Staff	38	126	68	225
Totals for Undergraduate Courses:		**2,725**	**8,573**		**2,324**	**7,674**	**5,048**	**16,245**

Graduate Courses

COURSE NUMBER, TITLE, AND CREDIT HOURS	FALL 2000			SPRING 2001			TOTAL	
	INSTRUCTOR	STUD.	HOURS	INSTRUCTOR	STUD.	HOURS	STUD.	HOURS
501 Softw Engr: Technol. & Techng -4	Arms	77	308				77	308
502 Comp.Methds for Dig Lib - 3				van de Smpl	30	90	30	90
504 App Sys Engineering I -3	Staff	5	15				5	15
514 Inter Comp Sys-4	Birman	54	216				54	216
515 Pract in Systems- 3	Bailey	29	87				29	87
517 Adv Computer Graphics - 3	Greenberg	21	63				21	63
522 Comp Tools & Methds for Fin -4				Li	33	132	33	132
574 Heuristic Methds for Opti -4				Shmkr/Slmn	42	168	42	168
611 Adv Programming Languages - 4	Myers	43	172				43	172
612 Comp Des for High-Perf Arch -4				Stodghill	12	48	12	48
614 Advanced Systems - 4				Birman	28	112	28	112
621 Matrix Computations -4	VanLoan	31	124				31	124
622 Num Optim. & Nonlinear Eq -4				Coleman	26	104	26	104
626 Comp Mol Bio -4				Elber	18	72	18	72
632 Adv Database Systems -4				Gehrke	21	84	21	84
671 Intro to Automated Reasoning -4	Constable	5	20				5	20
672 Adv Artificial Intelligence 4				Selman	43	172	43	172
676 Reasoning About Knowledge -4	Halpern	25	100				25	100
681 Analy of Algorithms-4	Kleinberg	43	172				43	172
682 Theory of Computing -3				Hartmanis	25	100	25	100
683 Adv. Des & Analy of Algor - 4				Kleinberg	36	144	36	144
686 Logics of Programs - 4				Kozen	15	60	15	60
709 Computer Sci Colloq-1	Staff	91	91	Staff	87	87	178	178
713 Topics in Cryptography -2	Bailey	23	46				23	46
715 Seminar on PRL -4	Constable	9	36	Constable	9	36	18	72
721 Topics in Numerical Analysis- 2	Vavasis	3	6				3	6
726 Prob & Persp in Comp Mol Biol-1	Elber	10	10	Elber	9	9	19	19
732 Seminar in Data Mining -4	Gehrke	8	32	Gehrke	11	44	19	76
734 Seminar in Replicated Data -4	Demers	8	32				8	32
754 Systems Research Seminar -1	Vogels	13	13	Vogels	12	12	25	25
772 Seminar in AI -4	Selman	12	48	Selman	8	32	20	80
773 Proseminar in Cog Studies I -4	Cardie	2	8				2	8
774 Proseminar in Cog Studies II - 4				Cardie	1	4	1	4
775 Seminar in Nat. Lang Underst-2	Lee	1	2	Lee	6	12	7	14
789 Sem in Thry of Algor & Com -4	Tardos	23	92	Tardos	20	80	43	172
790 Independent Rsch-var 15	Staff	98	590	Staff	78	453	176	1043
990 Doctoral Research -var 15	Staff	32	440	Staff	31	448	63	888
Totals for Graduate Courses:		**614**	**2,596**		**653**	**2,636**	**1,268**	**4,954**
Grand Total:		**3,339**	**11,169**		**2,976**	**10,308**	**6,316**	**21,199**

Degrees Granted

Doctor of Philosophy

AUGUST 2000

Eric G. Aaron
Chair: David Gries
Members: Robert Constable, Michael Spivey-Knowlton, Anil Nerode
Title: *Tactic-based Modeling of Cognitive Interference on Logically Structured Notation*

Marcos Kawazoe Aguilera
Chair: Sam Toueg
Members: Joseph Halpern, Michael Stillman
Title: *Broadening the Applicability of Failure Detectors*

Nawaaz Ahmed
Chair: Keshav Pingali
Members: Ramin Zabih, Adam Bojanczyk, Joyce Morgenroth
Title: *Locality Enhancement of Imperfectly-nested Loop Nests*

Daniel G. Brown
Chair: David Shmoys
Members: Michael Todd, L. Paul Chew, Ron Elber
Title: *Algorithmic Methods in Genetic Mapping*

Christopher Kirk Hawblitzel
Chair: Thorston von Eicken
Members: Gregory Morrisett, James Sethna
Title: *Adding Operating System Structure to Language-based Protection*

Scott Burten Hunter
Chair: Devika Subramanian
Members: David Gries, Carol Krumhansl
Title: *Safety Markov Decision Processes: Their Solution and Application to Agent Design*

Vera Maria Kettnaker
Chair: Ramin Zabih
Members: Daniel Huttelocher, Henry Shue
Title: *Stochastic Models for the Analysis of Traffic Video*

Scott Anthony Mardis
Chair: Claire Cardie
Members: Molly Diesing, Richard Zippel
Title: *Partial Parsing for Linguistic Relationship Identification*

Vijay S. Menon
Chair: Keshav Pingali
Members: Gregory Morrisett, Thomas Coleman
Title: *Symbolic Compilation Techniques for Array Computations*

JANUARY 2001

Jason Jonathan Hickey
Chair: Robert Constable
Members: Dexter Kozen, Roberto Bertoia
Title: *The MetaPRL Logical Programming Environment*

Xiaoming Liu
Chair: Robbert van Renesse
Members: Kenneth Birman, Zygmunt Haas
Title: *Building High-performance Adaptive Communication Systems from Components*

Lili Qiu
Chair: Robbert van Renesse
Members: Kenneth Birman, Zygmunt Haas
Title: *An Integrated Approach to Improving Web Performance*

David Patrick Walker
Chair: Greg Morrisett
Members: Dexter Kozen, Barbara Correll
Title: *Typed Memory Management*

Zhen Xiao
Chair: Ken Birman
Members: Zygmunt Haas, Joseph Halpern
Title: *Efficient Error Recovery for Reliable Multicast*

MAY 2001

Wei Tsang Ooi
Chair: Brian Smith
Members: Robbert van Renesse, Zygmunt Haas
Title: *Design and Implementation of Distributed Programmable Media Gateway Services*

Adam Florence
Chair: Charles Van Loan
Members: Steve Vavasis, Tim Healey
Title: *Computational Multilinear Algebra*

Ozan Hafizogularri
Chair: Robert Constable
Members: Greg Morrisett, Rebecca Harris-Warrick
Title: *A Typed Framework for Program Analysis*

Nicholas Howe
Chair: Daniel Huttelocher
Members: Claire Cardie, Michael Spivey Knowlton
Title: *Analysis and Representations for Automatic Comparison, Classification, and Retrieval of Digital Images*

Jia Wang
Chair: Robbert van Renesse
Members: Zygmunt Haas, Srinivasan Keshav, Balachander Krishnomurthy
Title: *Network Aware Client Clustering and Applications*

Yin Zhang
Chair: Robbert van Renesse
Members: Sam Toueg, Zygmunt Haas, Vern Paxon
Title: *Measurements and Analysis of End-to-End Internet Performance*

Lidong Zhou
Chair: Fred Schneider
Members: Zygmunt Haas, Robbert van Renesse
Title: *Towards Fault-tolerant and Secure On-line Services*

MS Special Degrees Conferred

AUGUST 2000

Rie Ando
Chair: Lillian Lee

Stephan Arthur Zdancewic
Chair: Andrew Myers

JANUARY 2001

Benjamin Atkin
Chair: Ken Birman

Eliyahu Barzilay
Chair: Robert Constable

Matthew D. Fleming
Chair: Ken Birman

Amit Kumar
Chair: Jon Kleinberg

Brian A. Meloon
Chair: J. Guckenheimer

MEng Degrees

AUGUST 2000

Roman Lobkovsky
Rui Wang

JANUARY 2001

Gaurav Agarwal
Douglas Burdick
Zhengyu Chen
Lizbeth Henson
Alex Holub
Chin-yu Hsu
Joe Huang
Jangwoo Kim
Yi Lin
Lori Lorigo
Brian Morgan
Benjamin Newton
Priya Rajan
Charles Shagong
Michael Shapiro
Anup Tamakuwala
Jasjeet Thind
Ngoclan Vu

MAY 2001

Andre Allavena
Andy Ang
Ali Anvari
Sean Byrnes
Kuan-Yuan Chang
Steve Chow
Brian Cody
Aditya Dada
Pantaleo de Candia
Valeriy Elbert
Michael Frei
Vadim Grinshpun
Daniel Hecht
Ernie Ho
Ranjita Jain
Ju Joh
Hyunjong Kang
Danquing Kong
Hsiang-Hao Kung
Allen Lamb
Lingling Li
Andrew Lin
Chien-Chih Liu
Nidhi Loyalka
Shyam Maniyedath
Omar Mehmood
Derek Messie
Eric Milkie
Ashish Motivala
Travis Ortogero
Li-wen Peng
Paradee Phempoonpanich
Joshua Pollak
Michael Priscott
Damien Raymond
Jason Rohrer
David Rollenhagen
Thibet Rungrotkitiyot
Daniel Ryazansky
Shrinivas Samant
Eric Strong
Sung-hsun Su
Yu-Ju Tu
David Wang
Anthony Watkins
Xiang Xie
Yukiko Yamashita
Donghui Yan
Hsien-Lung Yang
Puze Yang
Jinghao Zhang
Jingming Zhang
Qinghong Zhang
Ying Zheng
Wenjie Zhong
Xiaozheng Zhong
Maxim Zolotov

Bachelor of Arts

AUGUST 2000

Mark Stephen Gately
Wei Jing Guo
Eric Pohan Lee
Alexander Rakhlin

JANUARY 2001

Anit Agarwal
Jian-Ning Janet Cheng
Naveen Joshi
Mikhail Kobyakov
Nadine Latief
Eric Michael Strong

MAY 2001

Gregory Brian Artzt
Krista Bendig
Yuriy Berkovich
Jenna R. Burrell
Christopher Capobianco
Ruggiero Cavallo II
Shaun E. Chandran
William M. Chang
Christopher Hans Chattaway
Gregory Laurence Clinton
Fabiano Baroni Desouza
Vladimir Alexandrovitch Dizhoor
Alexander Druyan
Christina Nichole Dulitz
Steven John Engelbrecht
Robert D. Flint
Roger Houng Chiu Fong
Russell Douglas Greene
Rama C. Hoetzlein
Jason James Hofmann
Serena Kohli
Alexander Kordun
Simon Garrett Lang
Jennifer Anne Lazo
Mon Jed Liu
Ayan Kanti Mandal
Edward Munandar
Lyn Mien Ngooi
Jeffrey Michael Pine
Nikita O. Proskourine
Mahmood Reza
Bernard Lee Rissmiller
Ivan David Rosero
Vedran Rozic
Maria Kasiani Sammut
Alexandre Saverin
Samuel R. Scarano
David G. Seah
David Allen Shepperton

Wazila Z. Shikari
Joel James Skaliotis
Jared Tolman
Jeffrey M. Umetsu
Gilberto Rivera Vasquez
Amanda Marie Waack
Raymond James Wenderlich
Richard Ryan Williams
Ben Wong
John D. Woschinko
Mikhail Alexeivich Zatsman

Bachelor of Science

AUGUST 2000

William David Alonso
Herman Chen
John David Martinez Garza
Pratik Joshi
James Cremer Landis
Victor Y. Liu
Kun Shing Luk
Ankush Sahai
Mohan Sarovar

JANUARY 2001

Tareq Nisar Aryne
Michael John Cirello
Jeffrey Michael Derstadt
Valeriy Elbert
Ilia Gimelfarb
Sunny R. Gleason
Daniel Michael Hecht
Igor Kats
Chun Ho Leo Ku
Nidhi Loyalka
Jason A. Satran
Paras Hemant Shelawala
Traian Iavorov Stanev
Matthew Scott Taylor
Clyde Shi Ming Tsai
Allen Weipong Wang
David S. Weil
Brian Wesley Williams
Stephen Kai-Tung Yam
Tomi Yiu

MAY 2001

Michael Sitrin Adams
Jason William Adaska
Anish Aggarwal
Saif Reza Ahmed
Joseph Robert Aliperti
Chirag Rajnikant Amin
Thomas Jack Babinski
Megha Batra
Dmitriy Berman
Carson Grey Bloomberg
Brandon Roy Bray
Christopher Robert Bray
John H. Bull
Donald Chai
Waiki Chan
George Weikang Chang
Jesse Che-Chai Chang
John B. Chen
King Lun Choi
Jennifer Irene Chou
Andy Yu-An Chu
Andrew Chung
Regina Ruby Lee Clewlow
Brian Patrick Comerford
Robert Michael Cronin
James A.F. Curcio
Andrew Charles Cushman
Aparna Das
Brian G. Dekorte
David Michael Feldman
Michael Edwin Fettner
Daniel J. Flaccavento
Daniel J. Gagne
Aaron M. Gerlich
Rachel Eleanor Green
Cynthia Alexandra Greene
Wei Gu
Matthew Thomas Harren
David W. Hays
Abraham Samuel Heifets
Jeremy Chung-Kee Ho
Jeffrey Robert Hoy
Sandra Lynn Jablonka
Prasant Yedithi Jagannath
David Nai-Heng Kang
Kay Tiong Khoo
Theodore Won-Hyung Kim
Aleksey Kliger
Joseph Kong
Patrick Ian Kongsilp
Adam David Krauszer
Jeremy Martin Kubica
Joshua Ben Kulkin
Jay S. Kumar
Chee Yong Lee
Minsu Lee
Sze Ern Dawn Lee
James C. Lu
Christopher Michael Mancuso
Robert Anthony Maskaron
Eric Robert Mauskopf
Kyle George McKenna
Andrew Lawrence Mehler
Raj B. Merchant
Miroslaw Robert Mlot
Michael David Morse
Stewart Muñoz
Bernard Francis Murphy III
William Curtis Neff
Jill K. Newman
Gennady Nurik
Katherine Leigh Oliver
Dmitriy Patek
Joseph Robert Polastre
Daniel Scott Rabinovitz
Nathan Mathias Ramsey
Benjamin Andrew Ransford
Andrew Kamron Rostami
Patrick Brigham Ryan
Ahmad Saeed
James David Sampey
Robert C. Scanlon
David Gerard Seah
Kevin Yin-Sin Seng
Nikesh Keshav Shah
Piyanuch Silapachote
Nigel Adrian Singh
John Tam
Jeremy Russell Tavan
John D. Tennant
Charitha Tillekeratne
Fu Wah Ting
Hung Viet Tran
Schaun Lamont Valdovinos
Robert William Van Wicklen
Evangelos Vergetis
Grant Jenhorn Wang
Brian M. Weisberg
Devon Michael Welles
Mike Chia-Cheng Wey
Jonathan J. Wicks
Ian Thomas Withrow
William Robert Wolfrom
Sergio Wong
David W. Wu
Victoria Yampolsky
Judy Chao-Rong Yiu
Dennis C. Yueh
Tal M. Ziv

Awards

Faculty Awards

GRAEME BAILEY

Kenneth A. Goldman '71 Excellence in Teaching Award (2000).

ACSU Faculty of the Year (1999-2000).

KEN BIRMAN

Stephen '57 and Marilyn Miles Excellence in Teaching Award (2000).

JOHANNES GEHRKE

Engineering Teaching Award, Cornell (2001).

IBM Faculty Development Award, College of Engineering (2000, 2001).

James and Mary Tien Excellence in Teaching Award (2000).

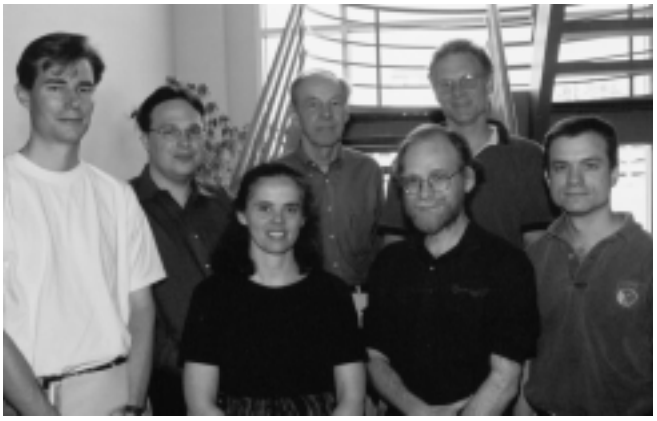

Award recipients Johannes Gehrke, Bart Selman, Juris Hartmanis, Dexter Kozen, Eva Tardos, Joe Halpern, and Greg Morrisett.

ZYGMUNT J. HAAS

Michael Tien'72 Award, Cornell College of Engineering. Excellence in Teaching Award (September, 2000).

JOSEPH HALPERN

Guggenheim Fellowship (2001-02) to be spent in Israel.

Fulbright Scholar (2001-02).

JURIS HARTMANIS

Lielo Medalu, Latvian Academy of Sciences. This is the highest Medal of Honor the Academy conveys.

Engineering Teaching Award, Cornell (2001).

Computing Research Association Distinguished Service Award (2000).

JON KLEINBERG

National Academy of Sciences Award for Initiatives in Research (2001). This prize recognizes innovative young scientists to encourage research likely to lead toward new capabilities for human benefit.

David and Lucile Packard Foundation Fellowship(1999-2004).

ONR Young Investigator Award (1999-2002).

NSF Faculty Early Career Development Award(1997-2001).

DEXTER KOZEN

Stephen and Margery Russell Distinguished Teaching Award, College of Arts and Sciences, Cornell (2001).

RAJIT MANOHAR

Sonny Yau '72 Excellence in Teaching Award (2000-2001).

IEEE Teacher of the Year Award (2000-2001).

NSF Faculty Early Career Development Award (2000-2004).

GREG MORRISETT

Engineering Teaching Award, Cornell (2001).

Allen Newell Medal for Research Excellence (2001).

Ralph Watts Excellence in Teaching Award (2001).

Presidential Early Career Award for Scientists and Engineers (2000).

FRED B. SCHNEIDER

Daniel M. Lazar Excellence in Teaching Award (2000).

Professor-at-Large, University of Tromsoe, Tromsoe, Norway (1996-2004).

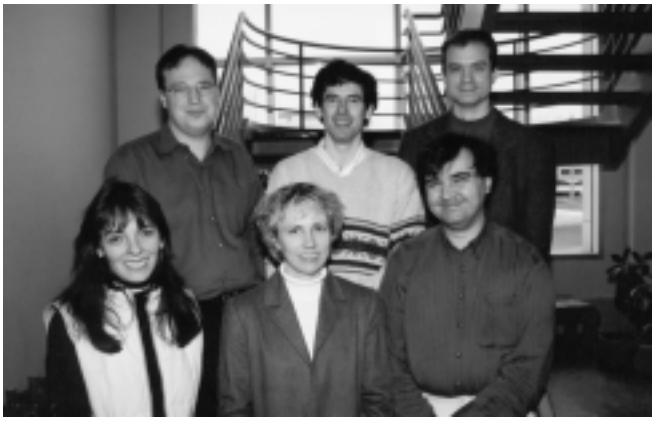

Award recipients Bart Selman, Jon Kleinberg, Greg Morrisett, Carla Gomes, Claire Cardie, and Ramin Zabih

BART SELMAN

Fellow: American Association for Artificial Intelligence.
NSF Faculty Early Career Development Award (1998-2002).

EVA TARDOS

Elected to the American Academy of Arts and Sciences (2001) for her contributions to the design of efficient algorithms.
Faculty of the Year , 2000-2001.

GOLAN YONA

1998-2000 Burroughs-Welcome Fellowship from the Program in Mathematics and Molecular Biology (PMMB).

Student Awards

Jonathan Marx Senior Prizes were awarded to **Joseph R. Aliperti** and **Joseph Robert Polastre**. These prizes recognize outgoing and energetic students who are community minded. The **Alan Marx Teaching Award** was given to **Brandon Roy Bray** for excellence in the support of undergraduate teaching. The Marx senior awards are provided in memory of **Jonathan Marx '85** and his father **Alan Marx, JD '61**.

The **Computer Science Prize for Academic Excellence** was given to **Richard Ryan Williams**. This award, given annually at graduation, recognizes a student who has achieved excellence in academics, promise in the area of research (typically, published work), and a commitment to the field of computer science in terms of teaching and/or other leadership activities.

The following undergraduates were recognized for their outstanding contributions as course consultants and undergraduate teaching assistants: **Brandon Bray, Young Cho, Bobby Chow, Rob Cronin, Ben Mathew, Doug Mitarotonda, Joe Polastre, Praveen Rajasethupathy**.

The **Lockheed Martin Awards** for Outstanding Achievement and Academic Service were awarded to: **Christian Hescock, Benedict Viglietta, Salman Arif, William Liu,** and **Steven Baker**. Among the selection criteria were academic achievement, research accomplishments, and/or service to the teaching mission of the Department of Computer Science.

Jeremy Kubica, Engineering Class of 2001, was named a **Merrill Presidential Scholar.** Recognizing the top 1 percent of each graduating class at Cornell, nominees were asked to name their most influential professor at Cornell. Jeremy selected Bart Selman, with whom he worked on an undergraduate research project.

Jeffrey M. Vinocur was awarded a **Goldwater Scholarship** in April 2001. The Barry M. Goldwater Scholarship and Excellence in Education Foundation was established by act of Congress in 1986 in recognition of the long government service of U.S. Sen. Barry Goldwater and to foster and encourage excellence in mathematics, science and engineering. Vinocur is applying to medical school and, in addition, hopes to earn his master's or doctorate and then pursue a career in research and clinical medicine, making use of his computer science education. At Cornell, he worked under Gregory Morrisett, professor of Computer Science, on type-safe programming languages, with the goal of providing a framework for designing crash-proof computers. He also has worked in the Neonatal-Physiology Research Laboratory at the MCP Hahnemann University School of Medicine in Philadelphia, investigating the biochemical mechanisms of injury to the developing brain. He will be presenting an abstract on his research at the Pediatric Academic Societies' 2001 annual meeting in Baltimore later this month.

Computer Science undergraduate **Marla Leahy '02** was awarded the **Degenfelder Family Scholarship** for academic year 2000-2001. This award, part of the Scholarship Challenge Campaign (Challenge 4 Campaign), will be awarded to a student in the Computational Biology concentration in academic year 2001-2002.

Cornell's Big Red team won its second consecutive championship in the fourth annual **RoboCup** tournament. Held in Melbourne Australia earlier this year, the Small Robots competition featured 23 international teams. According to the team's advisor, Professor Raffaello D'Andrea, new innovations in robot design and AI contributed to improved maneuverability and control. Student team members were **Bryan Audiffred '00. Michael Babish '99, Tobias Welge-Luessen '00, Josh Pollak '00, Saeed Saeed '00, Nicole Schlegel '01, Mark Schwager '00, and Will Stokes '03.**

Funded Research – Computing & Information Science / Computer Science

INVESTIGATOR	SPONSOR	TITLE	AWARD	PERIOD OF AWARD	
Arms	NSF	The NSDL Central System	799,085	9/15/00	8/31/01
Arms	Library of Congress	Collecting and Preserving Open-access Materials on the Web	72,214	1/1/00	6/30/00
Birman	NASA / JPL	Ensemble and Spinglass: Technologies for High Performance Cluster Management and Communication	240,000	5/1/00	9/30/01
Birman/ Constable	DARPA / AFRL	Spinglass Adaptive Probabilistic Tools for Advanced Networks	3,839,383	7/1/99	6/29/02
Cardie	NSF / POWRE	Integrating Natural Language Tools for Advanced Networks	68,695	7/15/00	12/31/01
Cardie	DARPA / ONR	Rapidly Portable Translingual Information Extraction and Interactive Multidocument Summarization	1,294,464	5/18/00	5/17/03
Coleman	DOE	Efficient Algorithms for Large-scale Constrained Optimization with Application to Inverse Problems	1,362,288	5/1/86	9/14/00
Constable	NSF	Creating and Evaluating Interactive Formal Course-ware for Mathematics and Computing	283,975	1/1/99	12/31/01
Constable	NSF	Enhancing Proof Assistant Systems	20,800	1/1/01	12/31/02
Constable	DARPA / AF	An Open Logical Programming Environment: A Practical Framework for Sharing Formal Models	2,271,068	6/2/98	9/30/02
Constable	DARPA / ONR	Building Interactive Digital Libraries of Formal Algorithmic Knowledge	1,938,148	5/1/01	4/30/04
Department	NSF	A Next Generation Computing and Communications Substrate	1,331,298	8/15/97	7/31/02
Elber	NSF	Kinetics of Ion Channels by Atomically Detailed Computer Simulations	465,743	4/1/00	3/31/03
Elber	NIH	Long Time Dynamics of Biomolecules	1,214,013	4/1/00	3/31/05
Elber	DARPA / ONR	Evolution of Protein Structures as a Tool to Predict Protein Shape and Protein Function	240,000	5/18/00	5/17/03
Gehrke	NSF	JAGUAR: JAVA in Next Generation Database Systems	370,000	9/1/98	8/31/01

INVESTIGATOR	SPONSOR	TITLE	AWARD	PERIOD OF AWARD	
Gehrke /Bonnet	AFRL	Flexible Decision Support in Device-saturated Environments- SenseIT	660,000	9/23/99	9/22/02
Gomes	AFRL	Compute-intensive Methods for Combinatorial Problems	351,094	3/1/99	7/31/01
Gomes	AFRL	Hybrid Approaches for Combinatorial Problems	304,835	3/1/99	2/28/02
Gomes	AFRL	Intelligent Information Systems Institute	3,100,000	12/1/00	11/30/03
Halpern	NSF A	Qualitative Framework or Reasoning Under Uncertainty	348,000	9/1/96	8/31/01
Halpern	NSF SG	ER: Decision-making in Complex Systems	90,000	9/1/00	2/28/02
Halpern	ONR	Semantic Consistency in Information Exchange	366,642	4/30/97	6/30/02
Halpern	ONR	Resource Bounded Knowledge and Security	200,000	3/1/00	2/28/02
Halpern	ONR	Reasoning About Noninterference	87,086	2/15/01	2/14/02
Halpern	ONR	Software Quality and Infrastructure Protection for Diffuse Computing	526,058	5/1/01	4/30/06
Kleinberg	NSF	(CAREER) Algorithmic Methods for Networks	200,000	4/1/97	3/31/02
Kleinberg	ONR	Algorithms for Networks and Link-structured Data	305,000	5/1/99	4/30/02
Kleinberg	Packard	Algorithmic Methods for Networks	625,000	10/13/99	10/12/04
Kleinberg /Lee/ Cardie /Selman	NSF ITR	The Construction and Analysis of Information Networks	450,000	9/1/00	8/31/03
Kozen	DARPA / NSF	Formal Methods for Software Certification	291,000	8/15/97	12/31/00
Lagoze	NSF	Security and Reliability in Component-based Digital Libraries	2,425,899	5/19/99	4/30/03
Lagoze	Digital Library Federation	Open Archives Initiative	84,800	1/1/01	12/31/01
Lagoze	NSF	Metadata for Resource Discovery of Multimedia Digital Objects Harmony	240,000	10/1/99	9/30/02
Lagoze	NSF	Integrating and Navigating E-print Archives through Citation-linking	291,650	10/1/99	9/30/02
Lagoze / Arms	CNRI	Digital Library Testbed Program	915,000	9/1/98	8/31/01
Morrisett	Sloan	Sloan Research Fellowship	35,000	9/16/98	9/15/02

INVESTIGATOR	SPONSOR	TITLE	AWARD	PERIOD OF AWARD	
Morrisett	NSF	Design, Applications, and Foundations of Safe Low Level Program Languages	205,000	3/1/99	2/28/03
Morrisett	AFOSR / PECASE	Next Generation Systems Languages	1,000,000	4/15/01	4/14/06
Pingali	NSF MATLAB	Extensions and Compiler Techniques for High Performance Computing	190,993	1/1/99	12/31/01
Pingali	NSF	Synthesis of Block-recursive Codes for Deep Memory Hierarchies	548,314	1/1/01	12/31/03
Schneider	DARPA/AF	AFRL/Cornell Information Assurance Institute	2,050,000	3/15/00	2/28/02
Schneider	AFOSR	CIPIAF for Information Assurance Institute	592,657	6/1/01	5/31/03
Schneider / Myers	DARPA/AF	Containment and Integrity for Mobile Code	2,197,784	6/16/99	6/15/02
Schneider/ Morrisett	AFOSR	Language-based Security for Extensible Systems	844,408	3/1/00	11/30/02
Schneider / Morrisett/ Kozen /Myers	ONR	Language-based Security for Malicious Mobile Code	4,247,977	7/1/01	6/30/06
Selman	Sloan	Sloan Research Fellowship	5,000	9/16/99	9/15/01
Selman	NSF	(CAREER) Compute Intensive Methods for AI	300,000	8/1/98	7/31/02
Selman /Gomes	AFRL	Principled Analysis & Synthesis of Agent Systems Using Tools from Statistical Physics	750,000	7/1/00	11/9/02
Selman /Gomes	DARPA/AF	Controlling Computational Cost: Structure, Phase Transitions and Randomization	1,621,041	4/27/00	8/26/03
Selman /Lee	DOD/ Expertology, Inc.	Expertise Location using Automatically Generated Network Models	30,000	10/15/00	10/14/01
Tardos	NSF	Algorithmic Issues in Communication Networks	249,559	7/1/97	6/30/01
Tardos	DARPA/ONR	Efficient Algorithms for Transportation in Dynamic Networks	256,212	9/8/00	9/30/02
Tardos	DARPA/ONR	Efficient Resource Management in High-speed Networks	801,548	4/1/98	9/30/01
Van Loan	NSF	New Applications and Algorithms that Involved the Kronecker Product	247,874	9/1/99	8/31/02

INVESTIGATOR	SPONSOR	TITLE	AWARD	PERIOD OF AWARD	
Yona	Bio & Life Sciences	Global Self-organization of the Protein Space: Towards a Map of the Protein Space	48,658	7/1/01	6/30/02
Zabih	NSF	Dynamic Contextual Recognition of Moving Objects	150,000	9/15/99	8/31/02

Submitted Grant Proposals

INVESTIGATOR	SPONSOR	TITLE	AWARD	PERIOD OF AWARD	
Arms	NSF	Interoperability Services Building on the Open Archives Initiative	493,635	9/1/01	8/31/04
Arms /Krafft/ Lagoze	NSF	Collaborative Project: Core Integration of the National SMETE Digital Library	1,562,500	9/1/01	11/30/02
Constable	NSF	ITR/SY: New Capabilities for Creating High-confidence Embedded Software	2,956,089	10/1/01	9/30/06
Gehrke	NSF	Interactive and Online Data Mining	235,000	9/1/01	8/30/04
Gehrke	NSF	Scalable Decision Tree Construction	463,663	9/1/01	8/31/04
Gehrke /Demers/Birman	NSF	ITR/SI: Intelligent Mining and Monitoring: Responding to the Data Tsunami	4,239,254	10/1/01	9/30/04
Halpern	AFOSR	Formulating and Reasoning about Security Policies	322,788	1/1/02	12/31/04
Kozen	NSF	Kleene Algebra	351,233	5/1/01	4/30/04
Lagoze	Univ of Virginia /Mellon	The Open Source FEDORA Repository Development Project	300,000	10/1/02	9/30/05
Pingali	NSF	Framework for Developing Complex Applications on High-end Petaflop-Class Machines	646,054	4/1/01	3/31/04
Pingali	NSF	ITR/SY: A New Framework for Program Optimization	791,443	9/1/01	8/3004
Selman	DOD/ Expertology, Inc.	Learning Social Network Models for Information Linkage	30,000	10/15/01	10/14/02
Sirer	Schlumberger Research Corp.	Assuring the Security of Java	36,795	5/1/01	12/31/01
Tardos	ONR	Algorithmic Issues in Network Design and in Information Access	481,714	10/1/01	9/30/03
Tardos/ Zabih	NSF	ITR/SY: Combinatorial Optimization Algorithms for Information Access	465,786	9/1/01	8/31/04

Collaborative Research at Cornell

INVESTIGATOR	SPONSOR	TITLE	AWARD	PERIOD OF AWARD	
Birman / Schneider/ECE	EPRI	Minimizing Failures While Maintaining Efficiency of Complex Interactive Networked Systems	625,000	1/1/99	12/31/03
Coleman/ Elber / CTC	NIH	Parallel Processing Resource for Biomedical Scientists	9,583,781	4/15/92	11/30/01
Coleman /C TC	SGI	SGI/Cornell Computational Finance Institute	1,433,810	3/23999	1/31/01
Coleman / CTC	TG Information Network	Financial Engineering and Tools	840,000	9/1/00	8/31/03
Coleman / CTC	NYS	Supercomputer Center Operational Support	1,200,000	5/1/00	4/30/01
Cardie / Zabih/ Cog Studies	NSF	Computational Aspects of Cognitive Science (Training Grant)	562,500	9/1/95	8/31/00
Elber/ Kleinberg/Chew Kedem MGB	NSF	Multiscale Hierarchical Analysis of Protein Structure and Dynamics	899,000	9/1/00	8/31/03
Halpern/ ECE	NSF	Predictive Sensor Assisted Wireless Multimedia Systems	862,465	9/15/97	8/31/00
Pingali / CTC	NSF	CISE Research Infrastructure: A Two-tier Computation and Visualization Facility for Multiscale Problems	1,500,000	10/1/99	9/30/04
Pingali / Vavasis/Chew/CTC	NSF ITR	Adaptive Software for Field-Driven Simulations	5,000,000	9/1/00	6/30/04
Pingali / Vavasis/CEE	NSF	Challenges in CISE: Crack Propagation on Tera Flop Computers	1,852,592	1/15/98	12/31/01
Pingali /C TC	NSF	Multiscale Modeling of Defects in Solids	1,500,000	10/1/98	9/30/01
Selman/ Gomes/ECE	DARPA	Self-configuring Wireless Transmission and Decentralized Data Processing for Generic Sensor Networks	165,281	4/18/00	4/17/03
Selman/ Gomes/MAE	AFOSR MURI	Cooperative Control in Uncertain Adversarial Environments	266,600	5/1/01	4/30/04

TOTAL EXPENDITURES, FISCAL YEAR 2000 - 2001: 10,255,536

Research Interests of the Faculty and Researchers

William Arms: Digital libraries and electronic publishing.

Graeme Bailey: Mathematical modeling, applications to medicine and biology, geometry, and parametrization spaces and connectivity.

Kenneth P. Birman: Reliability and security in modern networked environments.

Martin Burtscher: High performance microprocessor architecture, instruction-level parallelism, and compiler optimizations.

Claire Cardie: Natural language processing, machine learning, artificial intelligence.

Paul Chew: Computational geometry, computational biology.

Thomas F. Coleman: Numerical analysis, computational finance, scientific computing.

Robert L. Constable: Type theory, automated reasoning.

Alan Demers: Database systems, database replication, algorithms.

Ron Elber: Computational molecular biology, genomics.

Geri Gay: Cognitive and social issues for the design and use of interactive communicaion technologies.

Johannes Gehrke: Database systems, data mining, and mining and monitoring evolving data.

Carla Gomes: Artificial intelligence, operations research, and planning and scheduling.

Donald P. Greenberg: Realistic image synthesis, modeling, scientific visualization, computer-aided design, and image processing.

Zygmunt Haas: Wireless and mobile systems, including macrodiversity, interference-bounded dynamic channel allocation, mobile TCP/IP networks, location-independent access, personal communications services (PCS), and software for mobile systems.

Joseph Y. Halpern: Reasoning about knowledge and uncertainty in multi-agent systems, decision theory, logic, artificial intelligence, and security.

Juris Hartmanis: Computational complexity, complexity of chaotic systems.

Mark Heinrich: Design of active memory and I/O systems for next-generation servers and data-intensive computing.

Sheila S. Hemami: Communication of visual information, including multirate video coding and transmission, compression specific to packet networks and other lossy networks, and psychovisual considerations.

John E. Hopcroft: Robust geometric algorithms, modeling and simulation, and information capture and access.

Klara Kedem: Computational geometry.

Jon Kleinberg: Theory of computing, algorithms, computational biology.

Dexter Kozen: Theory of computational, proof-carrying code, computational complexity, analysis of algorithms, and program logics and semantics.

Dean Krafft: Digital libraries, information access.

Christoph Kreitz: Automated reasoning, program transformation, verification and synthesis.

Carl Lagoze: Digital libraries.

Lillian Lee: Natural language processing.

Yuying Li: Scientific computation, and numerical optimization.

Rajit Manohar: Design of efficient asynchronous computation structures in VSLI, and the use of formal methods to guarantee the correctness of such structures.

J. Gregory Morrisett: Programming languages, security, type systems, compilers.

Andrew Myers: Programming languages, compilers, distributed systems, runtime systems.

Anil Nerode: Automata theory, constructive concurrent dynamic logic.

Keshav Pingali: Programming languages, and parallel computing.

Robbert van Renesse: Distributed computing, fault-tolerance, real-time systems.

Mats Rooth: Scientific research in linguistics, particularly syntax, semantics and lexical semantics.

Fred B. Schneider: Concurrent and distributed systems, and computer and network security.

David Schwartz: Computational mechanics, applied mathematics, educational technology.

Bart Selman: Artificial intelligence, experimental computer science.

David B. Shmoys: Design and analysis of efficient algorithms, scheduling.

E. Gun Sirer: Design and implementation of distributed operating systems for modern networks with a main focus on ad hoc mobile networks.

Evan Speight: Software runtime systems for distributed computing platforms, commodity cluster-based shared memory systems, and affinity directed mobility.

Paul Stodghill: Program transformations, and program synthesis for computational science applications.

Eva Tardos: Design and analysis of algorithms, optimization, communication networks and combinatorics

Tim Teitelbaum: Programming languages, and systems and environments.

Charles Van Loan: Numerical linear and multilinear algebra.

Stephen A. Vavasis: Numerical analysis, and optimization.

Werner Vogels: Reliable distributed systems, with a focus on development of new network protocols and middleware, and novel strategies to structure applications and support systems.

Golan Yona: Computational molecular biology, with an emphasis on developing tools and methodologies for large-scale analysis of protein sequences and structures.

Ramin Zabih: Computer vision, multimedia, information technology

Editorial Activities of the Faculty

ACM Computing Surveys (Schneider, Editor)
ACM Transactions on Programming Languages and Systems (Morrisett, Associate Editor)
Algorithmica (Hopcroft, Editor and Member, Executive Committee)
Annals of Software Engineering (Schneider, Editor)
Annals of Mathemathics (Selman, Editorial Board)
Applied Mathematics Letters (Coleman, Editorial Board)
Artificial Intelligence Journal (Halpern; Selman, Editorial Board)
Chicago Journal of Theoretical Computer Science (Tardos, Editor; Halpern, Consulting Editor)
Combinatorica (Tardos, Editor)
Communication on Applied Non-linear Analysis (Coleman, Editorial Board)
*Computational Geometry: Theory and Application*s (Kedem, Guest Editor)
Computational Linguistics (Lee, Editorial Board)
Computational Optimization and Applications (Coleman, Editorial Board)
Constraints: An International Journal (Selman, Editorial Board)
D-Lib Magazine (Arms, Editor-in-chief)
Discrete Mathematics and Theoretical Computer Science (Pingali, Editorial Board)
Distributed Computing (Schneider, Editor)
Formal Methods in System Design (Constable, Editor)
Fundamenta Informaticae (Hartmanis, Editor)
High Integrity Systems (Schneider, Editor)
IEEE Transactions on Signal Processing (Hemami, Associate Editor)
Information and Computation (Halpern, Editorial Board)
Information Processing Letters (Schneider, Editor)
Information Sciences (Hopcroft, Associate Editor)
International Journal of Computational Geometry and Applications (Hopcroft, Editor)
International Journal Parallel Programming (Pingali, Editorial Board)
Journal of Logic and Computation (Constable, Editor)
Journal of Computer and Systems Sciences (Hartmanis, Hopcroft, Editors; Hopcroft, Associate Editor)
Journal of Functional Programming (Morrisett, Editor)
Journal of Global Optimization (Vavasis, Editorial Board)
Journal of Interconnection Networks (Tardos, Editor)
Journal of Logic and Computation (Constable, Halpern, Editorial Board)
Journal of Scheduling (Shmoys, Associate Editor)
Journal of Symbolic Computation (Constable, Editor)
Journal of the ACM (Halpern, Editor-in-chief)
Journal of Theoretical Computer Science (Tardos, Editor)
Machine Learning (Cardie, Editorial Board; Lee, Editorial Board)
Mathematical Modeling and Scientific Computing (Coleman, Editorial Board)
Mathematical Programming (Vavasis, Editorial Board, Shmoys, Tardos, Associate Editors)
Mathematics of Operations Research (Tardos, Area Editor, Shmoys, Associate Editor)
MIT Press Series on Digital Libraries and Electronic Publishing (Arms, Series Editor)
Natural Language Semantics (Rooth, Editorial Board)
Pattern Recognition Society Journal (Chew, Kedem, Editorial Board)
SIAM Journal on Computing (Shmoys, Tardos, Editors)
SIAM Journal on Discrete Mathematics (Shmoys, Editor-in-chief)
SIAM Journal of Matrix Analysis Applications (Vavasis, Editorial Board)
SIAM/MPS Series on Optimization (Shmoys, Co-editor)
SIAM Review (Vavasis, Editorial Board)
Springer-Verlag *Lecture Notes in Computer Science* (Hartmanis, Editor)
Springer-Verlag *Texts and Monographs in Computer Science* (Schneider, Co-managing Editor)

Faculty Personnel Changes

NEW FACULTY IN SPRING 2001

Gun Sirer, who is receiving his Ph.D. from the University of Washington, joined the faculty in January as an Assistant Professor. His research interests include extensible, distributed, and networked systems.

Golan Yona (Hebrew University, 1999) joined the faculty as an Assistant Professor in January. His research area is computational molecular biology.

NEW FACULTY IN 2001

Rich Caruana (CMU, 1997) works in machine learning and data mining, medical decision making and bioinformatics, feature selection, missing values, inductive transfer, artificial neural networks, memory-based learning. He joins the department in July.

Daisy Fan's (Cornell, Civil & Environmental Engineering, 2001) research interests include the application of systems analysis techniques for water resources and environmental problems.

Thorsten Joachims (Dortmund, 1997) works in machine learning and intelligent agents, with a focus on Support Vector Machines and machine learning with text. He will be joining the department in October.

Jeanna Neefe Matthews's (Berkeley, 2000) research interests include file systems, storage systems, and more generally, operating systems and distributed systems. She will join the department in January.

Radu Rugina (UCSB, 2001) is interested in pointer analysis, parallelizing compilers, and parallel computing. He will be joining the department in January.

Jayavel Shanmugasundaram's (Wisconsin, 2001) research interests include internet data management, database systems, and transaction processing in emerging system architectures. He will be joining the department in August.

DEPARTURES

Sam Toueg has resigned and has joined the faculty at the University of Toronto.

Praveen Seshadri has resigned and will remain at Microsoft, where he has been on leave for the last two years.

Thorsten von Eicken has resigned and will remain at Expertcity.com, where he has been on leave for the last two years.

CHANGES

Juris Hartmanis becomes emeritus, effective July 1, 2001.

Claire Cardie, Jon Kleinberg and Greg Morrisett were promoted to associate professor.

SABBATICALS AND LEAVES

Keshav Pingali, Fred Schneider, and Ramin Zabih are returning from academic year sabbatical leaves. Ramin will continue his affiliation with the Department of Radiology at the Cornell Medical School in New York City. He will be teaching on the Ithaca Campus during the fall 01 semester.

Claire Cardie, Tom Coleman, and Joe Halpern will be on sabbatical leave during the next academic year. Dan Huttenlocher will continue to be on leave from the department during the coming academic year.

Juris Hartmanis, Charles Van Loan, and Tim Teitelbaum

Faculty, Researchers, and Academic Visitors

PROFESSORS
William Arms
Graeme Bailey
Kenneth P. Birman
Thomas F. Coleman
Robert L. Constable
Alan Demers
Ron Elber
Donald Greenberg
Joseph Halpern
Juris Hartmanis
John E. Hopcroft
Daniel P. Huttenlocher
Klara Kedem
Dexter Kozen
Keshav K. Pingali
David B. Schmoys
Fred B. Schneider
Eva Tardos
Sam Toueg
Charles Van Loan

ASSOCIATE PROFESSORS
Claire Cardie
Jon Kleinberg
Greg Morrisett
Bart Selman
Tim Teitelbaum
Stephen Vavasis
Ramin Zabih

ASSISTANT PROFESSORS
Johannes Gehrke
Lillian Lee
Andrew Myers
David Schwartz
Emin Gun Sirer
Thorsten von Eicken
Golan Yona

LECTURERS
Gregery Buzzard
Ron DiNapoli
Matthew Morgenstern
Thomas Yan
Lidong Zhou

SENIOR RESEARCH ASSOCIATES
L. Paul Chew
Dean Krafft
Christoph Kreitz
Yuying Li
Robbert van Renesse

RESEARCH ASSOCIATES
Stuart Allen
Philip Bonnet
Saleh El Mohamed
Carla Gomes
Carl Lagoze
Koneshan Sivapathasundram
Paul Stodghill
Werner Vogels

RESEARCH STAFF
Donna Bergmark
Naomi Dushay
Rich Eaton
Lori Lorigo
Jaroslaw Meller
Sandy Payette
Carol Terizzi

POSTDOCTORAL ASSOCIATES
Kavita Bala
Ramon Bejar
Raoul Bhoedjang
Alfredo Cardenas
Tamara Galor
Avijit Ghosh
Anupam Gupta
Pavel Naumov
Koneshan Sivapathasundram
Veaceslav Zaloj

VISITING FACULTY
Sergei Artemov, Russian Academy of Sciences, Moscow, Russia
Bruno Codenotti, University of Pisa, Pisa, Italy
Daniel Mosse, University of Pittsburgh
Khalid Mughal, University of Bergen, Bergen, Norway
Herbert Van De Sompel, Open Archives Initiative
Walker White, University of Dallas

ACADEMIC VISITORS
Mathieu Baudet, Ecole Polytechnique, France
Tim Clark, Reliable Network Solutions
William Debany, Rome Air Force Research Lab
Dan Dumitriu, Reliable Network Solutions
Mark Dyson, Rome Air Force Research Lab
Walter Gadz, Rome Air Force Research Lab
Kevin Kwait, Rome Air Force Research Lab
Eva-Marie Luther, Technische Fachhochschule Berlin, Germany
Amy Magnus, Rome Air Force Research Lab
Andreas Meier, Universitaet des Saarlandes
Louis Pochet, Rome Air Force Research Lab
Leonard Popyack, Rome Air Force Research Lab
Jurek Tiuryn, University of Warsaw, Poland
Matthew Thomas, Rome Air Force Research Lab
Sharon Walter, Rome Air Force Research Lab
Richard Zippel, School of Computer & Media Science Interdisciplinary Center, Herzliya, Israel

Administrative and Technical Staff

Director of Administration: Pat Musa
Human Resources Manager: Susan Schwarz
Director of Corporate Relations: Marcy Rosenkrantz
Alumni Relations: Dan Jenkins
Assistant to the Chair: George Manning/Nora Balfour
Front Office Manager: Bonnie Maine

FINANCE OFFICE

Finance Manager: Claudia Wojcinski
Accounts Representative: Carol Ayer
Award Coordinator/HR Assistant: Laura Kratochvil
HR Assistant: Karla Consroe
Post Award Coordinator: Amy DeVaul
Pre-award Coordinator: Bonnie McCarthy

GRADUATE PROGRAMS

Assistant Director of Graduate Programs: Becky Stewart
Master of Engineering Program Coordinator: Stephanie Meik

UNDERGRADUATE EDUCATION PROGRAM

Assistant Director of Undergraduate Programs: Daniel Jenkins
Course Administrator: Laurie Buck
Reception and Records: Nicole Roy/Anna Salter

FACULTY ADMINISTRATIVE ASSISTANTS

Rosemary Adessa
Kathy Carpenter
Linda Competillo
Helene Croft
Juanita Heyerman
Tammy Howe
Cindy Robinson

Department Colloquium: Karla Consroe/Linda Competillo
Theory Seminar: Karla Consroe/Tammy Howe
Director of Computing Facilities: Dean Krafft
Systems Administrative Assistant: Cay Wilson
Systems Administrative Assistant: Mona Seamon

TECHNICAL STAFF

Network Technician: Bruce Boda
Network Technician: John Finley
Lead Consultant: Ellen Cramer/Jennifer Holleran
Consultant/Advisor: Eric Brinkman
Consultant/Advisor: Rob Collins
Consultant/Advisor: Robert O'Keeffe
Consultant/Advisor: Joseph McGuire
Senior Programmer/Analyst: Ellen Cramer
Senior Programmer/Analyst: Doug Flanagan
Senior Programmer/Analyst: Larry Parmelee
Programmer/Analyst: Dora Abdullah
Programmer/Analyst: Orlando Johnson
Programmer/Analyst: Dean Eckstrom
Info Tech Area Manager: William Holmes
Multimedia Editor: Una Moneypenny

Computer Science Department Staff

Departmental Computing Facilities

The department makes use of a mix of computing platforms, with about three-quarters of our research and instructional computing taking place on Microsoft's Windows NT/2000 operating system and Intel Architecture processors and the remaining quarter on Unix desktop and back-end servers. We have benefited greatly over the last year from several major equipment donations: Intel donated 25 1.5GHz Pentium 4 workstations and a server to upgrade our undergraduate teaching laboratory; Intel also donated 20 1.4GH Pentium 4 systems for research; and we have received 30 PocketPC-based PDAs, 802.11b wireless access cards, software, training, and books from Microsoft Corporation.

In the area of improved and upgraded infrastructure, during the past year, we:

- Begun the transition of our pilot 802.11b wireless networking infrastructure (Nomad) to a full service (RedRover) run by Cornell Information Technologies. The new infrastructure will provide close to 100 access points across campus when it is officially released at the beginning of the fall 2001 semester.
- Completed a major upgrade of departmental servers and desktop systems from Windows NT to Windows 2000, including an upgrade of our Exchange mail system from Exchange 5.5 to Exchange 2000.
- Added a new Linux Beowulf cluster with 12 1.2GHz AMD Athlon Thunderbird processors, and expanded two existing Linux clusters: one with an additional 40 processors and the other with 350GB of disk.
- Expanded departmental Linux support with three new Linux backend servers (CVS support, Apache/MySQL/PHP, and general Linux) and full support for RedHat Linux on research and teaching systems.

The department has over 1000 computers ranging from desktops to high-end parallel processing servers, over three terabytes of on-line disk storage, and a backbone network based on switched Gigabit Ethernet.

The department has a full-time computing facilities staff of fifteen. Dean Krafft serves as director, with programming support provided by Dora Abdullah, Jennifer Holleran, Dean Eckstrom, Doug Flanagan, Bill Holmes, Orlando Johnson, and Larry Parmelee; web development provided by Una Moneypenny; hardware support by John Finley and Bruce Boda; user consulting support by Rob Collins, Joseph McGuire, and Eric Brinkman; and systems administration by Cay Wilson and Mona Seamon. The staff provides full support for all the operating systems and standard software on our major computing platforms.

In addition to the resources directly owned and operated by the department, computer science students and researchers have access to a number of university facilities. The university provides extensive campus-wide networking, based on the TCP/IP protocols and implemented through a switched Gigabit Ethernet backbone connecting organizational Ethernets. National and international access is provided by three OC3 connections to NYSERNet and the global Internet. High-speed community access is available through Time-Warner's RoadRunner cable modem system and several DSL providers.

The department operates an undergraduate teaching laboratory of 15 Intel-donated 300MHz Pentium II systems, 30 Intel-donated 450MHz Pentium III systems, 30 Intel-donated 866MHz Pentium III systems, and 25 Intel-donated 1.5GHz Pentium 4 systems, all running Microsoft's Windows 2000. This lab provides support for a wide range of upper-level undergraduate courses and individual research projects. There is also a separate graphics teaching laboratory with 25 HP Visualize fx6+ workstations.

Finally, through the Cornell Theory Center and the Program of Computer Graphics, computer science researchers have access to a wide range of advanced parallel processing and supercomputer systems as well as advanced graphics and visualization systems.

The following list includes major computing equipment in the Department of Computer Science, owned either by Cornell or by the federal government.

DESKTOP MACHINES

10	Sun SunRay
214	Intel Pentium II Desktop PC
19	SUN UltraSparc 10
247	Intel Pentium III Desktop PC
4	SUN UltraSparc 5
46	Intel Pentium 4 Desktop PC
4	Apple G3/G4 PowerBook
5	Intel Celeron Desktop PC
3	Apple G3/G4 workstation
39	Intel Pentium II Laptop PC
25	HP VisualizeX 600MHz PIII
38	Intel Pentium III Laptop PC
1	SGI O2

128 Intel Celeron-based Laptop PC

BACK-END RESOURCES

9 Sun Ultra Enterprise 420/450 quad-processor
1 45-node Linux PIII dual-proc. cluster

SERVERS

5 Dell quad-processor Pentium III servers
2 SUN Ultra Enterprise 250 dual-processor servers
3 Intel quad-processor Pentium II servers
4 SUN Sparc-20/514 four-processor compute servers
4 Dell 8-way Pentium III servers
1 3.8TB SUN tape library
15 Intel Pentium III servers
1 12-node AMD 1.2GHz Athlon Beowulf cluster
3 Linux Pentium III servers

OTHER HARDWARE

9 Color Laser Printers
2 Cisco Catalyst 5000/5500 Fast Ethernet switches
57 B&W Laser Printers (HP/ Lexmark)
34 Cisco Catalyst 29xx Fast Ethernet switches
2 HP DesignJet 2500 poster printers
1 Cisco Catalyst 6509 Gigabit Ethernet switch
130 WindowsCE handheld/palm devices
7 Cisco Catalyst 1900 10/100 Ethernet switches

University President Hunter R. Rawlings and Students at BOOM 01

Professors Ken Birman and Keshav Pingali

Associate Professors Bart Selman and Steve Vavasis, and CS Chair Charles Van Loan with CIS Dean Robert L. Constable

Colloquium and Seminar Speakers

CS Distinguished Lecture Series

AUGUST 2000

David Mermin, Dept. of Physics, Cornell University. *How Quantum Mechanics Alters the Nature of Computation*.

SEPTEMBER

Rajit Manohar, Dept. of ECE, Cornell University. *Low Energy Adaptive Process.*

Charles Holland, AFRL/Cornell Information Assurance Institute (IAI) Inaugural Lecture. *Programs, Policy and Politics - Science and Technology in the National Interest.*

Rich Caruana, Computer Science Department, Carnegie Mellon University. *Multitask Learning Schedule*.

Siddhartha Chatterjee, CS Dept., Univ. of N. Carolina. *Fast Tree-Structured Computations and Memory Hierarchies Schedule*.

OCTOBER

Paul Edwards, Dept.of History, Univ. of Michigan, Joint CS and Science and Technology Studies Colloquium. *Systems, Networks, and Webs: Towards a History of Digital Convergence.*

Divesh Srivastava, AT&T Labs—Research. *Querying LDAP Directories Schedule*.

Eva Tardos, CS Dept., Cornell University. *How Bad is Selfish Routing?*

Ambuj Goyal, IBM. *Transactional Internet.*

NOVEMBER

Herbert Van de Sompel, Ghent University, Belgium. *SFX/OpenURL and the Open Archives Inititive: Achieving Interoperability in Digital Libraries via Low-barrier Standards.*

Tony Hey, University of Southampton, UK. *Feynman, Einstein, and Computers.*

Jeanna Neefe Matthews, Clarkson University. *Self-managing File Systems.*

Daniel Mosse, University of Pittsburgh. *Towards a View of Efficient Softer Real-time.*

Frans Kaashoek, MIT/LCS. *How to Design Flexible Software Systems or Applying the End-to-end Argument.*

JANUARY 2001

Balachander Krishnamurthy, AT&T Labs—Research. *What's New in Web Research?*

FEBRUARY

Mark Heinrich, School of Electrical and Computer Engineering, Cornell University. *Simulation vs. Reality: The Importance of Building Hardware*.

Stu Feldman, IBM, T. J. Watson Research Center. *Trends in E-Commerce and Challenges for Research*.

Greg Morrisett, Computer Science, Cornell University. *Towards Next-Generation Low-Level Languages*.

Herbert Van de Sompel, Ghent University, Belgium. *The OpenURL Framework for the Context-sensitive Provision of Service Links*.

MARCH

Jayavel Shanmugasundaram, University of Wisconsin. *XPERANTO: Bridging Relational Technology and XML*.

Venkatesan Guruswami, MIT. *List Decoding Of Error-Correcting Codes*.

Junghoo Cho, Stanford University. *Crawling the Web: Discovery and Maintenance of Large-Scale Web*.

Alexander Hartemink, MIT. *Principled Computational Methods for the Validation and Discovery of Genetic Regulatory Networks from Expression Data*.

Dieter van Melkebeek, Institute for Advanced Study. *Time-Space Tradeoffs for Satisfiability*.

APRIL

Paul Ginsparg, Los Alamos National Laboratory. *Creating a Global Knowledge Network*.

Andris Ambainis, UC, Berkeley. *Lower Bounds on Quantum Computing*.

Radu Rugina, MIT. *Program Analysis Techniques for Pointers and Accessed Memory Regions*.

Peter Manolios, University of Texas. *Combining Theorem Proving and Model Checking for the Verification of Reactive Systems*.

Thorsten Joachims, GMD. *The Maximum-Margin Approach to Learning Text Classifiers Methods, Theory, and Algorithms*.

Jovan Popovic, CMU. *Motion Design in Computer Animation*.

Yacov Yacobi, Microsoft Corp. Information Assurance Institute Invited Speaker. *A Dual Watermarking and Fingerprinting System*.

Igor Guskov, Caltech. *Meshes and Geometry Processing*.

MAY

Richard Han, IBM Watson. *Interacting Devices, Applications, and Users In a Pervasive Computing World*.

Publications

Technical Reports

Aaron, Eric. "Tactic-based Modeling of Cognitive Inference on Logically Structured Notation." TR00-1812. September 6, 2000.

Ahmed, Nawaaz. "Locality Enhancement of Imperfectly-nested Loop Nests." TR00-1811. August 24, 2000.

Ahou, Lidong. "Towards Fault-tolerant and Secure On-line Services." TR01-1840. May 11, 2001.

Barzilay, Eli. "Quotation and Reflection in Nuprl and Scheme." TR01-1832. January 8, 2001.

Bergmark, Donna, W. Arms, and C. Lagoze. "An Architecture for Reference Linking." TR00-1820. October 26, 2000.

Bergmark, Donna. "Automatic Extraction of Reference Linking Information from Online Documents." TR00-1821. November 30, 2000.

Bergmark, Donna, and C. Lagoze. "Reference Linking the Web's Scholarly Papers." TR01-1835. February 7, 2001.

Bickford, Mark, C. Kreitz, and R. van Renesse. "Formally Verifying Hybrid Protocols with the Nuprl Logical Programming Environment." TR01-1839. May 11, 2001.

Boykov, Yuri, and D. Huttenlocher. "A Graph Based Algorithm for Bayesian Recognition." TR00-1803. "August 28, 2000.

Chandra, Ranveer, V. Ramasubramanian, and K.P. Birman. "Anonymous Gossip: Improving Multicast Reliability in Mobile Ad-hoc Networks." TR01-1836. February 7, 2001.

Cheney, James, C. Lagoze, and P. Botticelli. "Toward a Theory of Information Preservation." TR01-1841. May 24, 2001.

Dushay, Naomi, and C. Lagoze. "Modeling Decisions for Digital Content." TR00-1807. July 19, 2000.

Gupta, Indranil. "Minimal CDMA Reducing Strategies in Power-controlled Ad-hoc Wireless Networks." TR01-1834. January 25, 2001.

Kolmogorov, Vladimir, and R. Zabih. "Computing Visual Correspondence with Occlusions via Graph Cuts." TR01-1838. March 5, 2001.

Kopylov, Alexi. "Dependent Intersection: A New Way of Defining Records in Type Theory." TR00-1809. August 14, 2000.

Kozen, Dexter. "Myhill-Nerode Relations on Automatic Systems and the Completeness of Kleene Algebra." TR00-1826. November 30, 2000.

Kozen, Dexter, and J. Tiuryn. "Intuitionistic Linear Logic and Partial Correctedness." TR00-1830. December 31, 2000.

Kozen, Dexter. "Automata on Guarded Strings and Applications." TR01-1833. January 25, 2001.

Lagoze, Carl. "Accomodating Simplicity and Complexity in Metadata: Lessons from the Dublin Core Experience." TR00-1800. June 30, 2000.

Lagoze, Carl, J. Hunter, and D. Brickley. "An Event-aware Model for Metedata Interoperability." TR00-1801. "June 30, 2000.

Maatev, Nikolay, K. Pingali, and P. Stodghill. "The Bernoulli Generic Matrix Library." TR00-1808. August 1, 2000.

Minky, Yaron, A. Tractenberg, and R. Zippel. "Set Reconciliation with Nearly Optimal Communication Complexity." TR00-1813. September 27, 2000.

Qiu, Lili, V. Padmanabhn, N. Venkata, and G.M. Voeljer. "On the Placement of Web Server Replicas." TR00-1806. July, 13, 2000.

Ramjee, R., L. Li, and T. Porta. "IP Paging Service for Mobile Hosts." TR00-1816. October 13, 2000.

Rodeh, Ohad, K. P. Birman, and D. Dolov. "The Architecture and Performance of Security Protocols in the Ensemble Group Communication System." TR00-1822. October 17, 2000.

Rodeh, Ohad, K. P. Birman, and D. Dolev. "Using AVL Trees for Fault Tolerant Group Key Management." TR00-1823. November 2, 2000.

Schneider, Fred B., G. Morrisett, and R. Harper. "A Language-based Approach to Security." TR00-1825. November 21, 2000.

Smith, Frederick, D. Grossman, G. Morrisett, L. Hornof, and J. Trevor. "Compiling for Runtime Code Generation (Extended Version)." TR00-1824. October 27, 2000.

Stoller, Scott D., and F.B. Schneider. "Automated Analysis of Fault-tolerance in Distributed Systems." TR00-1829. December 21, 2000.

Xiao, Zhen, and K. P. Birman. "A Randomized Error Recovery Algorithm for Reliable Multicast." TR00-1814. October 2, 2000.

Xiao, Zhen, K. P. Birman, and R. van Renesse. Optimizing Buffer Management for Reliable Multicast. TR00-1815. October 2, 2000.

Xiao, Zhen. "Efficient Error Recovery for Reliable Multicast." TR01-1831. January 3, 2001.

Zdancewic, Steve, and A. Myers. "Confidentiality and Integrity with Untrusted Hosts: Technical Report." TR00-1810. August 22, 2000.

Zhang, Yin, L. Qiu, and S. Kesav. "Speeding up Short Data transfers: Theory, Architectural Support, and Simulation Results. TR00-1799. July 19, 2000.

Zhang, Yin, and Lili Qiu. "Understanding the End-to-End Performance Impact of RED in a Heterogeneous Environment." TR00-1802. July 19, 2000.

Zhou, Lidong, F.B. Schneider, and R. van Renesse. "COCA: A Secure Distributed On-line Certification Authority." TR00-1828. December 11, 2000.

Student Lectures

Benzinger, Ralph. "Automated Higher-order Complexity Analysis and Certification." Invited talk, Implicit Computational Complexity Workshop 2000, affiliated with LICS 2000. Santa Barbara, CA (2000).

Cheney, James. "Compressing XML using Hierarchical Multiplexed PPM Models." 2001 IEEE Data Compression Conference, Snowbird, UT (March 27, 2001).

Florence, Adam. "An Improved Implementation of the Fast Gauss Transform." SIAM Applied Linear Algebra 2000 Conference, Raleigh, NC (October, 2000).

Gupta, Indranil, Robbert van Renesse, and Kenneth P. Birman. "A Probabilistically Correct Leader Election Protocol for Large Groups." Proceedings 14th International Symposium on Distributed Computing (DISC 2000) – LNCS 1914, 89-13, Toledo, Spain (October, 2000).

Gupta, Indranil. "Minimal CDMA Recoding Strategies in Power-controlled Ad-hoc Wireless Networks." Proceedings 1st International Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile Computing, San Francisco, CA (April, 2001).

Howe, Nicholas. "Digital Images from the Computer's Perspective." Guest lecture in course on digital art, NY University (November 3, 2000).

Howe, Nicholas. "Digital Images and Brain Images: Shall Ever the Twain Meet?" Guest lecture in cognitive studies proseminar, Cornell University (October 31, 2000).

Kumar, Amit. "Fairness in Resource Allocation." 41st IEEE Symposium on Foundations of Computer Science. Rodendo Beach, CA.

Nogin, Alexey. "Theorem Proving in Higher Order Logic." (TPHOL's) Conference (August, 2000).

Pellacini, Fabio. "Rederman: A Virtual Machine for Computer Graphics." Cornell University (October, 2000).

Pellacini, Fabio. "Computer Animation." Cornell University (September, 2000).

Pellacini, Fabio. "Towards a Psychophysically-based Light Reflection Model". PIXAR. July, 2000.

Qiu, Lili. "The Content and Access Dynamics of a Busy Web Server: Findings and Implications." ACM SIGCOMM 2000, Stockholm, Sweden (August, 2000).

Qiu, Lili. "On the Placement of Web Server Replicas." IEEE INFOCOM 2001, Anchorage, AK (April, 2001).

Roughgarden, Timothy. "Approximate K-MSTs and K-Steiner Trees via Primal-dual Method and Lagrangean Relaxation." 8th Conference on Integer Programming and Combinatorial Optimization (IPCO). (June 13-15, 2001).

Roughgarden, Timothy. "Stackelberg Scheduling Strategies." Theory Seminar, Department of Computer Science, Cornell University (April 23, 2001).

Roughgarden, Timothy. "How Bad is Selfish Routing?" Operations Research Seminar, GSIA, Carnegie Mellon University, Pittsburgh, PA (December 4, 2000).

Roughgarden, Timothy. "How Bad is Selfish Routing?" 41st Annual Symposium on Foundations of Computer Science (FOCS), Redondo Beach, CA (November 12, 2000).

Roughgarden, Timothy. "How Bad is Selfish Routing?" Theory Seminar, Department of Computer Science, Cornell University (October 23, 2000).

Roughgarden, Timothy. "How Bad is Selfish Routing?" 17th International Symposium on Mathematical Programming, Atlanta, GA (August 10, 2000).

Vetsikas, Ioannis. "Online Auctions: Software Agents Will Do Your Bidding." AI Seminar, Cornell University (November, 2000).

Wagstaff, Kiri. "Clustering with Knowledge-based Constraints." Doctoral Consortium, 17th National Conference on Artificial Intelligence. Austin, TX (July 31, 2000).

Wagstaff, Kiri. "Clustering with Knowledge-based Constraints." Invited Talk, Stuart Russell's research group, Berkeley, CA (September 13, 2000).

Weirich, Stephanie. "Type-safe Cast: Functional Pearl." International Conference on Functional Programming (September, 2000).

Weirich, Stephanie. "Resource Bound Certification." Harvard University (February, 2001).

Weirich, Stephanie. "Encoding Intensional Type Analysis." European Symposium on Programming. (April, 2001).

Weirich, Stephanie. "Polytypic Programming and Intensional Type (Constructor) Analysis." Working Group 2.8 (Functional Programming) (April, 2001).

Student Papers

Ando, Rie. "Latent Semantic Space: Iterative Scaling Improves Inter-document Similarity Measurement." In Proceedings of SIGIR '2000, 216-223.

Batu, Tugkan, Lance Fortnow, Ronitt Rubinfeld, Warren Smith, and Patrick White. "Testing that Distributions are Close." In Proceedings of 41st Annual Symposium on Foundations of Computer Science, 2000.

Benzinger, Ralph. "Automated Complexity Analysis of Nuprl Extracted Programs." *Journal of Functional Programming*. Cambridge University Press, 2001.

Benzinger, Ralph. "Automated Higher-order Complexity Analysis." *Theoretical Computer Science*, Elsevier, 2001.

Bronevetsky, Greg. "Circle Menus." 2001 Engineering Graduate Research Symposium (March, 2001).

Chen, Zhiyuan, Johannes Gehrke, and Flip Korn. "Query Optimization in Compressed Database Systems." ACM SIGMOD International Conference on Management of Data, Santa Barbara, CA, (May 21-24, 2001).

Chen, Zhiyuan, H.V. Jagadish, Flip Korn, Nick Koudas, S. Muthukrishnan, Raymond Ng, and Divesh Srivastava. "Counting Twig Matches in a Tree." In Proceedings of the 17th International Conference on Data Engineering, Heidelberg, Germany, 595-604 (April 2-6, 2001).

Cheney, James. "Compressing XML Using Hierarchical Multiplexed PPM Models." 2001 IEEE Data Compression Conference, Snowbird, UT (March, 2001).

Chudak, F.A., T. Roughgarden, and D. P. Williamson. "Approximate k-MSTs and k-Steiner trees via the primal-dual method of Lagrangean relaxation." 8th Conference on Integer Programming and Combinatorial Optimization (IPCO) (June, 2001).

Coury, D.V., J. S. Thorp, K. M. Hopkinson, and K.P. Birman. "Agent Technology Applied to Adaptive Relay Setting for Multi-Terminal Lines." 2000 IEEE PES Summer Meeting, 16-20, Seattle, WA (July, 2000).

Coury, D.V., J. S. Thorp, K. M. Hopkinson, and K.P. Birman. "Improving the Protection of EHV Teed Feeders Using Local Agents." IEEE 7th International Conference on Developments in Power System Protection, Amsterdam, The Netherlands (April 9-12, 2001).

Decoste, Dennis, and Kiri Wagstaff. "Alpha Seeding for Support Vector Machines". Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 35-349 (August 20-23, 2000).

Dobra, Alin, and Johannes Gehrke. "Bias Correction in Classification Tree Construction." International Conference on Machine Learning (June 28, 2001).

Dutre, Philip, Parag Tole, and Donald Greenberg. "Approximate Visibility for Illumination Computations using Print Clouds." Technical Report PCG-00-1, Program of Computer Graphics, Cornell University. Available at: http://www.graphics.cornell.edu/pubs/2000/DTG00.html.

Ferwerda, J., F. Pellacini and D. Greenberg."A Psychophysically-based Model of Surface Gloss Perception." Proc. of Human Vision and Electronic Imaging VI. San Jose, CA (January, 2001).

Gale, Amy, and Rod Downey. "On Genericity and Ershov's Hierarchy." *MLQ – Math. Log. Quart.* 47(2):161-182 (2001).

Grossman, Daniel, and Greg Morrisett. "Scalable Certification for Typed Assembly Language." In 2000 ACM SIGPLAN Workshop on Types in Compilation, Montreal, Canada (September, 2000).

Gupta, Indranil, Robbert van Renesse, and Kenneth P. Birman. "A Probabilistically Correct Leader Election Protocol for Large Groups." In Proceedings 14th International Symposium on Distributed Computing (DISC 2000) – LNCS 1914, 89-13, Toledo, Spain (October, 2000).

Gupta, Indranil. "Minimal CDMA Recoding Strategies in Power-Controlled Ad-hoc Wireless Networks." In Proceedings 1st International Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile Computing, San Francisco, CA (April, 2001).

Jenkins, K, K. Hopkinson, and K. Birman. "Reliable Group Communication with Subgroups." 2001 IEEE International Workshop on Applied Reliable Group Communication within the International Conference on Distributed Computing Systems, 16-19, Phoenix, AZ (April, 2001).

Kempe, David, Jon Kleinberg, and Alan Demers. "Spacial Gossip and Resource Location Protocols." In Proceedings of STOC 2001.

Kumar, Amit, and Jon Kleinberg. "Fairness in Resource Allocation." 41st IEEE Symposium on Foundations of Computer Science.

Li, L., J. Y. Halpern, P. Bahl, Y. M. Wang, and R. Wattenhofer. "Analysis of Cone-based Distributed Topology Control Algorithms for Wireless Multi-hop Networks." ACM Symposium on Principles of Distributed Computing (PODC) (August, 2001).

Li, L. and J. Halpern. "Minimum Energy Mobile Wireless Networks Revisited." Proc. IEEE International Conference on Communications (ICC) (June, 2001).

Liu, Xiaoming, and Robbert van Renesse. "Fast Protocol Transition in a Distributed Environment." ACM Symposium on Principles of Distributed Computing (PODC '00) (July, 2000).

Liu, Xiaoming, Robbert van Renesse, Mark Bickford, Christoph Kreitz, and Robert Constable. "Protocol Switching: Exploiting Meta Properties." International Workshop on Applied Reliable Group Communication, (IWARGC '01).

Ng, Vincent, M. White, T. Korelsky, C. Cardie, D. Pierce, and K. Wagstaff. "Multi-document Summarization via Information Extraction." Proceedings 1st International Conference on Human Language Technology Research (April, 2001).

Nogin, Alexey. "Theorem Proving in Higher Order Logic.". Proceedings (TPHOL's) Conference (August, 2000).

Ooi, Wei Tsang, and Robbert van Renesse. "An Adaptive Protocol for Locating Media Gateways." Eighth ACM International Multimedia Conference. Los Angeles, CA.

Ooi, Wei Tsang, Robbert van Renesse, and Brian Smith. "Design and Implementation of Programmable Media Gateways." Tenth International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV 2000). Chapel Hill, NC.

Padmanabhan, Venkata N., and Lili Qiu. "The Content and Access Dynamics of a Busy Web Site: Findings and Implications." Proceedings of ACM SIGCOMM 2000, Stockholm, Sweden (August, 2000).

Pucella, R. "An Approach to the Implementation of Overlapping Rules in Standard ML." Proceedings of the 1st International Workshop on Rule-based Programming (RULE2000), Montreal, Canada, 2000.

Pucella, R. "The Design of a COM-oriented Module System." Proceedings of the Joint Modular Languages Conference (Zurich, Switzerland). Lecture Notes in Computer Science 1897:104-118, Springer-Verlag (2000).

Qiu, Lili, Venkata N. Padmanabhan, and Geoffrey M. Voelker. "On the Placement of Web Server Replicas." In Proceedings of IEEE INFOCOM 2001, Anchorage, AK (April, 2001).

Qiu, Lili, George Varghese, and Subhash Suri. "Fast Firewall Implementations of Software and Hardware-based Routers." Extended Abstract in Proceedings of ACM SIGMETRICS 2001, Cambridge, MA (June, 2001).

Roughgarden, T. "Stackelberg Scheduling Strategies." 33rd Annual ACM Symposium on Theory of Computing (STOC) (July, 2001).

Roughgarden, T. and E. Tardos. "How Bad is Selfish Routing?" 41st Annual Symposium on Foundations of Computer Science (FOCS), 93-102 (November, 2000).

Vetsikas, Ioannis, Ramon Bejar, Carla Gomes, Henry Kautz, and Bart Selman. "Structure and Phase Transition Phenomena in the VTC Problem." DARPA Workshop (Task Project).

Xiao, Zhen, and Ken Birman. "A Randomized Error Recovery Algorithm for Reliable Multicast." Proceedings of IEEE Infocom (April, 2001).

Xiao, Zhen and Ken Birman. "Providing Efficient, Robust Error Recovery Through Randomization." In International Workshop on Applied Reliable Group Communication (April, 2001).

Wagstaff, Kiri, Claire Cardie, Seth Rogers, and Stefan Schroedl. "Constrained K-means Clustering with Background Knowledge." Proceedings of the 18th International Conference on Machine Learning. (June 28-July 1, 2001).

Wattenhofer, R., L. Li, P. Bahl, and Y. M.Wang. "Distributed Topology Control for Power Efficient Operation in Multihop Wireless Ad-hoc Networks." In Proc. IEEE INFOCOM (April, 2001).

Weirich, Stephanie. "Encoding Intensional Type Analysis."(October, 2000). European Symposium on Programming (ESOP '01), Genova, Italy (April, 2001).

Weirich, Stephanie. "Type Safe Cast: Functional Pearl." In Proceedings of the 5th ACM SIGPLAN International Conference on Functional Programming (ICFP '00), Montreal, Canada (September, 2000).

White, Mike, Tanya Korelsky, Claire Cardie, Vincent Ng, David Pierce, and Kiri Wagstaff. "Multi-document Summarization via Information Extraction." Proceedings of the 2001 Human Language Technology Conference (March 18-21, 2001).

FCI Programs

Computational Biology

Genomic databases, protein databanks, MRI images of the human brain, and remote sensing data on landscapes contain unprecedented detailed information about biological systems that are transforming the way that we do almost all of biology. Problems investigated by computational biologists span a wide spectrum including topics as diverse as the genetics of disease susceptibility, comparing whole DNA genomes to uncover the secrets of evolution, predicting protein structures and understanding their motions and interactions, designing new therapeutic drugs, mathematically modeling the complex signaling mechanisms within cells, predicting how ecosystems will respond to climate change, and designing recovery plans for endangered species. The computational biologist must have skills in mathematics, statistics, and the physical sciences as well as in biology. A key goal in training is to develop the ability to relate biological processes to mathematical models that can be solved computationally.

Cornell faculty work primarily in four subareas of computational biology: biomolecular structure, bioinformatics and data mining, ecology and evolutionary biology, and statistical and computational methods for modeling biological systems. These include the computational study of topics such as DNA databases, protein structure and function, computational neuroscience, biomechanics, population genetics, and management of natural and agricultural systems. Beyond the basic core skills in mathematics, physical sciences and biology, the computational biology program of study requires additional coursework in mathematics, computer programming, a "bridging" course aimed at connecting biology to computation, and an advanced course where the theoretical/computational component of one aspect of biology will be studied.

The FCI-created undergraduate program of study in computational biology encourages students to gain fundamental skills and understanding that will allow them to focus on specific subareas and problems later in their careers. Computational biology is a new emerging area that has applications as broad as biology itself. The problems of interest, as well as the tools available to study them, will undoubtedly change during the four years of an undergraduate program. The program is an excellent preparation for students who wish to specialize in one of these computational areas in graduate school. There is great, and increasing, demand for research scientists and technical personnel who can bring mathematical and computational skills to the study of biological problems.

Computational Molecular Biology (CMB) is an interdisciplinary field that brings together numerous diverse research areas. A separate and isolated program in CMB will have difficulties in maintaining excellence in all fields, in teaching the diverse tools, and in providing the breadth of research topics that form the core of CMB. We therefore propose a different model of a multi-field program in Computational Molecular Biology. Computational Molecular Biology is a program that crosses several fields. For example, to meet the program conditions, a Ph.D. candidate in computer science can have supplementing studies in molecular biology. Alternatively, a Ph.D. student in the biophysics field can have supplementing studies in computer science and meet the CMB requirements. Hence, the students of this program may come from diverse fields such as molecular biology and genetics or computer science, creating the diverse community of researchers that we seek in CMB.

Computational Science and Engineering

Many of the faculty in engineering and the sciences engage in research that is computationally driven. Computational science and engineering (CS&E) at Cornell continues to be as strong as ever. Critical to the overall environment is the Cornell Theory Center, whose Velocity Cluster supports lines of inquiry that require intensive, large-scale computation.

This year the CS&E subgroup of the FCI accomplished several things. First, it created four "FCI short courses" to be offered during the coming academic year:

CS 401 Applied Scientific Computing with MATLAB
CS 402 Scientific Visualization with MATLAB
CS 403 Development of Scientific Computing Programs in a Unix Environment
CS 404 Survey and Use of Libraries for Scientific Computing

These practical, 4-week courses are directed at beginning graduate students across the campus. With FCI support, the Department of Earth and Atmospheric Sciences hired

Dr. Andrew Pershing, who will serve as the instructor. We expect to be able to offer these courses and other courses as "CIS" courses in future years.

The CS&E subgroup also put together a website that publicizes the CS&E activities at Cornell. Faculty and courses are presented in a way that is useful to current and incoming students.

Finally, the CS&E subgroup helped orchestrate the hiring of Professor Hod Lipson in the School of Mechanical and Aerospace Engineering. This appointment will do a great deal to enrich the connection between the FCI and the College of Engineering.

Digital Arts and Graphics

Graphics is approaching a crossroads. University departments around campus depend on computer technology; architecture can't be produced without the computer; art graduates don't *have* to use computers, but many are discovering new creative outlets using them. Computational graphics, which has been dominated by the entertainment industry, will change dramatically in the next decade or two. We have the process, technology and knowledge to effect this change correctly—simulate the physical world with images that are faithful— and we believe that studies that integrate computer technology in the arts and graphics arenas will become an avenue of importance. The focus is on many areas, including, but not limited to music, art, and arcitecture. As technology becomes more integrated into their artistic enterprises, Cornell students from many disciplines will gain an understanding of the computer fundamentals that will help them grow as artists.

Information Science

Information Science at Cornell is an interdisciplinary program in the FCI that allows graduate and undergraduate students to study new theories, models, concepts and design principles that incorporate an understanding of both social and technical information systems. The field of information science combines aspects of computer science and human-computer interaction with an examination of the social, economic, political, and legal contexts in which information systems function. Students in the information science program will obtain an understanding of the core topics of study emerging in this new and quickly growing field: the design and analysis of computing applications, information infrastructures, and human-centered systems; the legal, economic, and ethical issues that surround the construction of information systems; and the ways in which information technology is transforming society. Specific topics emphasized in the information science program include electronic communication; knowledge networking; collaboration within and across groups, communities, organizations, and society; the web and web information systems; natural language processing; computational techniques in the collection, archiving and analysis of social science data; information privacy; methods of collecting, preserving, and distributing information; information system design; cognition and learning; and human interface design and evaluation.