Computing and the Arts at Cornell

CIS

CREDITS

Co-editors: Una Moneypenny and Nora Balfour
Faculty Advisor: Charles Van Loan

CONTACT INFORMATION

Office of Computing and Information Science
Cornell University
4132 Upson Hall
Ithaca, NY 14853-7501
607 255-9188
www.cis.cornell.edu

Department of Computer Science
Cornell University
4130 Upson Hall
Ithaca, NY 14853-7501
607 255-7316
www.cs.cornell.edu

Information Science

Annual Report

2003

C O N T E N T S

CORNELL

Education and Research

Note: Explanations of acronyms used in this report appear on the foldout inside the back cover

Message from the Dean for Computing and Information Science

Robert L. Constable

To many who reflect on the role of universities in modern society, it is clear that in the Information Age, society will come to rely even more on universities for leadership. These institutions must educate a citizenry that faces increasingly global issues—health, the food supply, the environment, globalization of business, jurisdiction of world government, the international sharing of intellectual property, and, increasingly, dependence of global institutions on a fragile software infrastructure.

Society will value those institutions that provide global leadership in coping with these issues. To lead in the Information Age, university administration, faculty, and staff will need to understand the capabilities and technologies of computing and information science, and bring that understanding to bear on the most pressing global problems—whether sequencing the SARS virus, building a protective digital skin for the planet, framing a coherent set of ideas and laws to manage digital intellectual property, or ensuring a reliable worldwide information resource.

Cornell has made it clear at the highest levels that it intends to remain a leader in computing and information science as the importance of this discipline increases. The university's academic units are presented on the university Web site under the heading "Academics". Among these units is the Faculty of Computing and Information Science (CIS); it is part of Cornell's response to leadership in the Information Age. Its mission is to create more capability in computing and information science by recruiting faculty, building academic programs, expanding research, and informing policy.

Academically, CIS shares attributes with the Graduate School, in that its budget and administration support academic programs that reside in several colleges, and thus it operates in close coordination with the schools and colleges. Administratively, it shares properties of the schools and colleges: it is led by a dean, is independently budgeted, and is engaged in faculty recruitment. In research, it has institutes and coordinates with research centers.

CIS–supported academic majors and the schools and colleges at Cornell with which they are affiliated are listed below:

COMPUTER SCIENCE
(Arts and Sciences, Engineering)
established 1972

COMPUTATIONAL BIOLOGY
(Agriculture and Life Sciences; Arts and Sciences)
established 2001

INFORMATION SCIENCE*
(Agriculture and Life Sciences; Arts and Sciences)
established 2003

Two other CIS programs are in the works:

COMPUTATIONAL SCIENCE AND ENGINEERING
(Graduate School)

DIGITAL ARTS AND GRAPHICS
(Architecture, Art, and Planning; Arts and Sciences)

CIS currently supports forty-eight faculty members, who are listed in this report. All are affiliated with at least one academic program, several with more than one. Most CIS–funded faculty members are in CS, all of whom are also appointed in Engineering; therefore CS is included in both CIS and the College of Engineering.

It is remarkable that these resources sustain broad programs outside the CS major and the CS graduate field. This is possible because of the strong coherence among the programs—a coherence that will be apparent in this report—and because of the participation of several other units—seventeen currently—that derive value from partnering actively with CIS, and thus contribute courses and activities.

One of the images associated with CIS is that of a woven tapestry. The vertical threads—the warp—are Cornell's ten colleges and schools in Ithaca, plus the Weill Medical College in New York City. The horizontal threads—the weft—are CIS and the Graduate School.

Our mission is to add strength by connecting units together by making clear and distinct the strength and "color" of the CIS programs. This report will focus your attention on our segments of the warp. You will see that this tapestry is a dynamic, living entity.

**with a minor available in Engineering, Human Ecology, Industrial and Labor Relations*

Message from the CS Chair

Charles Van Loan

At Cornell and other research universities, departments and research areas were once much more closely aligned. Physics was handled by faculty in the Department of Physics, anthropology was the purview of faculty in the Department of Anthropology, and so on. In those "Wild West" days, the department chair's job was like herding researchers on the open range! With widely spaced academic homesteads, the primary responsibility was to guarantee "good grazing" up to and including the horizon. Chairs in fledgling subjects such as computer science had the additional problem of defining the horizon. My predecessors did an excellent job in that regard.

Fast-forward to the age of multidisciplinary research, where it is impossible to say where one scholarly area ends and the next begins. As everyone knows, broad research agendas tend not to align with the university's grid of academic departments. Well-known administrative solutions that address this problem include the joint faculty appointment, the cross-listed course, and the on-campus multidisciplinary research center. At Cornell, we are fortunate to have a fourth device that can also be used to track critical research trends—the graduate-field system. In this system we define the set of allowable thesis advisors for a given student by field rather than by department, a distinction that makes our approach to multidisciplinary research friendly and effective. The system's inventors made Cornell a stronger research university because they challenged a department-centric view of graduate education. Thanks to their vision, our Ph.D. students see only the open prairie, even though it has long since been partitioned into a patchwork of administrative territories.

Multidisciplinary research is also forcing us to rethink how we deliver undergraduate education, because on this campus we insist upon the tight coupling of research and the undergraduate mission. Narrow definitions of "college" discourage the creation of new undergraduate programs and the flourishing of others when the subject matter fails to align with the university's subdivision into colleges. The Faculty of CIS addresses this issue in part by overseeing CS in a way that does not diminish the colleges of Engineering or Arts and Sciences, where our undergraduate majors reside. It is an administrative innovation that generalizes the concept of college so that membership issues are driven by intellectual considerations, just as they are in the graduate-field system. If all goes according to plan, Cornell undergraduates in computing and information science will likewise see just the open prairie.

New structures like the Faculty of CIS make departments all the more important. Departments are *the* critical social unit within academia. The loftiest university-level strategic plan depends upon how well the participating departments hire, mentor, and promote their faculty and how outward-looking they are in terms of curriculum. Life on the prairie is defined by life on the homestead and *the department* is the homestead.

In looking over our particular academic domicile, I am happy to report that we are stronger and more secure in our campus mission than ever before. Compared to last year, our research expenditures are up about 30 percent and the number of outside units that have representation in the field of computer science has doubled. (Psychology, Mechanical and Aerospace Engineering, and Science and Technology Studies join Electrical and Computer Engineering, Operations Research and Industrial Engineering, and Mathematics.) The number of departments that cross-list courses with CS has increased from a handful to about a dozen. This track record reflects our commitment to the university's strategic plan for computing and information science and confirms that the CIS structure has been a success.

Thinking about faculty, we have two new professors. Paul Francis (networks) and Uri Keich (bioinformatics) bring new strength to our systems and computational biology groups. We have a Sloan Fellowship award winner (Johannes Gehrke). We have a record number of assistant professors (fourteen). It's youth and creativity up and down the hallways of Upson!

Joe Halpern became a fellow of the ACM, Bart Selman and Don Greenberg became fellows of the AAAS, and Fred Schneider received an honorary doctorate from the University of Newcastle-upon-Tyne in England. Congratulations to these senior faculty members!

David Gries has returned to the faculty and will be serving as the associate dean for undergraduate education in the College of Engineering. We have joint appointments with the JGSM (Dan Huttenlocher) and the Weill Medical College (Ramin Zabih). There is outreach to other universities through the tri-institutional program for computational biology (Ron Elber) and an ITR grant concerned with high-performance code generation for scientific and engineering applications (Keshav Pingali, Steve Vavasis, Paul Chew). We have CS leadership in the CTC (Tom Coleman), the FISC (Tom Coleman), the IISI (Carla Gomes), the IAI (Fred Schneider), the PCG (Don Greenberg), and the NSDL (Bill Arms). These multidisciplinary adventures are supported by the department's commitment to collegiality and core CS research.

Saddle up. It's Big Sky Country!

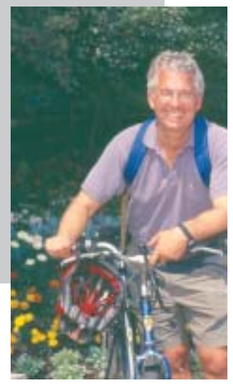

Computing and Information Science Highlights

Computer Systems Lab with ECE

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

The field of computer systems is both experimental and theoretical, having grown out of 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 CS. Usually students with primary interest in computer architecture, multiprocessor design, VLSI, computer-aided design (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 CS and vice-versa. Indeed, the interdisciplinary composition of the research teams is a strength of the Cornell Computer Systems Laboratory.

For further information, see *http://www.csl.cornell.edu.*

Digital Libraries and the National Science Digital Library (NSDL)

For ten years, Cornell's digital libraries research group has carried out research into architectures, protocols, services, and policies that facilitate the creation, management, accessibility, and longevity of distributed information. In particular, the group has had a focus on interoperability—the challenge of building coherent services from many heterogeneous, independently managed digital libraries. Recent achievements include the Open Archive Initiative Protocol for Metadata Harvesting (OAI–PMH), which enables technically inexperienced groups to share information, and the FEDORA mechanisms for the storage, manipulation, access management, and dissemination of digital library content, when the parties are more sophisticated technically.

These problems go far beyond conventional computer science research and, through the new Information Science program, the group works closely with colleagues who have expertise in human–computer interaction, electronic publishing, information preservation, evaluation, and software engineering.

The NSDL is a long-term program of the National Science Foundation (NSF) to build a digital library of all digital resources that could benefit education in the sciences. The NSF has funded almost one hundred independent projects, with one central project to integrate them into a single library. Following a successful demonstration at Cornell, the central grant has been awarded to a collaboration between the University Center for Atmospheric Research, Columbia University, and Cornell, with Cornell taking the technical lead.

The NSDL is simultaneously a production library, a testbed for digital-libraries research, and a source of new research challenges. For example, Donna Bergmark received the Vannevar Bush award for a paper describing her research into methods for automatic selection of materials for the NSDL, combining selective Web crawling with methods from classical information retrieval.

For further information, see *http://www.nsdl.org/*.

The Information Assurance Institute

The U.S. Air Force Research Laboratory (AFRL)/Cornell Information Assurance Institute (IAI) supports a broad spectrum of research and education efforts aimed at developing a science-and-technology base that can enhance information assurance and networked information–systems trustworthiness—system and network security, reliability, and assurance. IAI is also intended to foster closer collaborations among Cornell and AFRL researchers. Fred B. Schneider is the director.

AFRL researchers participate in Cornell research projects, facilitating technology transfer and exposing Cornell researchers to problems facing the Air Force; Cornell researchers 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.

Under the auspices of IAI, Cornell researchers are now involved in the development of the Air Force's Joint Battlespace Infosphere (JBI). Various other technical collaborations are also being explored—in the use of "gossip protocols", in language-based security policy–enforcement technology, and in data mining from networks of sensors.

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

The Intelligent Information Systems Institute

The mission of the IISI, founded in December of 2000, is threefold: To perform and stimulate research in compute- and data-intensive methods for intelligent decision-making systems; to foster collaborations within 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/U.S. Air Force Office of Scientific Research (AFOSR). Carla Gomes is the director of the institute. The Scientific Advisory Board of the institute consists of Robert Constable (Cornell), Nort Fowler and Charles Messenger (Information Directorate of the AFRL [AFRL/IF]), and Neal Glassman and Juan Vasquez (AFRL/AFOSR).

The IISI supports basic research within CIS, promoting a cross-fertilization of approaches from different disciplines, including computer science, engineering, operations research, economics, mathematics, statistics, and physics. Areas of research within the IISI are: search and complexity, planning and scheduling, large-scale distributed networks, data mining and information retrieval, reasoning under uncertainty, natural-language processing, machine learning, multi-agent systems, and combinatorial auctions.

Current IISI members at Cornell are Raffaello D'Andrea (dynamics and control), Claire Cardie (natural-language understanding and machine learning), Rich Caruana (machine learning, data mining, and bioinformatics), Carmel Domshlak (modeling and reasoning about preferences and uncertainty, combinatorial search and optimization, AI applications), Johannes Gehrke (database systems and data mining), Carla Gomes (artificial intelligence and operations research), Joseph Halpern (knowledge representation and uncertainty), Juris Hartmanis (theory of computational complexity), Mark Heinrich (active memory and simulation methodology), John Hopcroft (information capture and access), Thorsten Joachims (machine learning and information retrieval), Jon Kleinberg (algorithm design—networks and information), Lillian Lee (statistical methods for natural-language processing), Bart Selman (knowledge representation, complexity, and multi-agent systems), Phoebe Sengers (intelligent systems in human and social content; human computer interaction), David Shmoys (algorithms for large-scale discrete optimization), Chris Shoemaker (large-scale optimization and modeling), Evan Speight (distributed computing and computer architectures), Steve Strogatz (complex networks in natural and social science), and Stephen Wicker (intelligent wireless-information networks).

Several research projects that involve direct collaborations between Cornell and AFRL/IF researchers were initiated through the IISI. These cover 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, Stanford, and the University of Washington is developing a tunable benchmark suite for the design and evaluation of new algorithms for combinatorial auctions. The IISI also sponsored the American Association for Artificial Intelligence (AAAI) Symposium on Uncertainty Within Computation; the 2001 Conference on Empirical Methods in Natural Language Processing (EMNLP 2001); Language Technologies 2001; North American Association for Computational Linguistics (NAACL 2001); School on Statistical Physics, Probability Theory, and Computational Complexity (2002); Workshop on Phase Transition and Algorithmic Complexity at the Institute of Pure and Applied Mathematics (2002); the International Workshop on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems (CP-AI-OR 2002–03); the International Conference on the Principles and Practice of Constraint Programming (CP 2002–03); the Sixth International Conference on Theory and Applications of Satisfiability Texting (SAT 2003); and the Eighteenth International Conference on Artificial Intelligence (IJCAI-03).

To further its research mission, the IISI hosts many short-term visitors, and several scientists who make medium- and long-term visits. Visitors have included researchers from AFRL/IF, AT&T Labs, Hebrew University of Jerusalem, ILOG Corp., Microsoft Research, Stanford University, Technion, University of Barcelona, University of Lisbon, University of Minnesota, Washington University–St. Louis, University of Washington, and York University.

For further information, see *http://www.cis.cornell.edu/iisi.*

Connections with the Cornell Theory Center

The Cornell Theory Center (CTC)*(http://www.tc.cornell.edu)*, directed by CS Professor Thomas Coleman, is Cornell's high-performance computing and interdisciplinary computational-research center, serving more than 150 faculty research groups across the Ithaca campus and at the Weill Medical College in New York City.

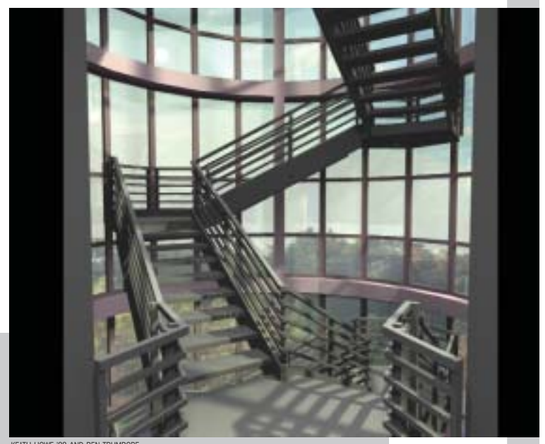

KEITH HOWE '88 AND BEN TRUMBORE

Simulation of a view from the Cornell Theory Center, created four months before the building was completed. The image appeared on the cover of *Cornell Alumni Magazine* in December 1990.

Through a strategic partnership with Microsoft, Dell, and Intel, CTC has pioneered the use of industry-standard computational clusters running Windows™ as a productive large-scale computing environment. CTC's resources, which consist of a cluster complex of more than 1,500 processors, keep Cornell at the forefront of computational science and engineering. New technological advances include the integration of the database software SQLserver into complex engineering applications and the application of .NET and Web services to high-performance computing. CTC operates the first Windows/Dell/Intel–based CAVE 3-D immersive virtual-reality environment, which is used by a variety of projects, including an engineering design course and an architecture course. Students in CEE 479 and M&AE 491 have access to EduCluster, a 16-processor cluster dedicated to student applications.

CTC has three core interdisciplinary research emphases: computational finance, computational biology/genomics, and computational materials.

The computational finance group (*http://www.ctc-manhattan.com/Research/index.asp*) is headed by CTC director Thomas Coleman and includes CS research associate Yuying Li. 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 and is the site of an annual securities derivatives conference.

CTC's Computational Biology Service Unit (CBSU) (*http://cbsu.tc.cornell.edu*), headed by CS professor Ron Elber, applies computational resources and expertise to a variety of applications in the life sciences, ranging from canine genetics to plant breeding to protein-structure modeling. Each summer one or two undergraduates are chosen from a pool of applicants for the CBSU Undergraduate Summer Internship. Through this internship, undergrads onduct research under the guidance of a faculty advisor and in collaboration with CBSU staff. One of the 2003 CBSU Internships was awarded to Keith Jamison, a CS junior.

The Computational Materials Institute at CTC focuses on fracture mechanics and serves as one of the test beds for the adaptive software project led by CS professor Keshav Pingali, who is also a CTC associate director. The adaptive software project is developing software systems that can adapt to changes at the application, algorithmic, and system levels.

CS professor Johannes Gehrke is also involved in interdisciplinary CTC projects. He is applying his datamining expertise to a pilot project involving data acquisition and analysis using the Cornell-operated Arecibo radiotelescope and to a genomics database that tracks pathogens.

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, as well as undergraduate mentors who support teams of high school student developers at remote locations. A number of team members come from Computer Science. SciCentr brings CTC into interaction with research scientists and faculty in the fields of biotechnology, communication, fine arts, theatre arts, music, and architecture. CTC is also engaging undergraduates in development of interactive online lab modules focused on bioinformatics through the BioQUEST Curriculum Library.

For further information, see *http://www.tc.cornell.edu*.

001
11010

Research in Computer Science at Cornell

We believe that the information revolution is transforming universities, because it goes to the heart of what universities are about: the creation and dissemination of knowledge. A few universities will be leading the country and the world into this information-based future.

Artificial Intelligence

Understanding intelligence and creating intelligent agents, the twin goals of artificial intelligence (AI), are two of the final frontiers of modern science. Early pioneers of computer science such as Turing, Von Neumann, and Shannon were captivated by the idea of creating a machine intelligence. Though much progress has been made, computer science and AI are still young fields, and many of the questions and issues considered then are actively being pursued today.

Research in AI at Cornell covers a wide range of topics, including decision theory, information retrieval, knowledge representation, machine learning and data mining, natural-language processing, planning, reasoning under uncertainty, search, and vision. A particular strength of the department is that our research embraces both theory and experiment, with particular emphasis on learning approaches to AI problems. Given the complexity of many of the basic questions in AI, our research often transcends traditional scientific boundaries. We are actively pursuing connections to other disciplines such as bioinformatics, economics, genomics, information science, linguistics, operations research, physics, psychology, and statistics. The department is one of the main participants in the IISI and in two university-wide programs: COGST and CIS.

Computer Architecture and VLSI

Research in architecture and VLSI is part of the Computer Systems Laboratory. Computer-systems research at Cornell 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 system specification and verification.

Human oncogene c-H-ras p21 from CTC's NCRR-funded Parallel Processing Resources for Biomedical Scientists

Computational Biology

The recent completion of the Human Genome Project underlines the need for new computational and theoretical tools in modern biology. The tools are essential for analyzing, understanding, and manipulating the detailed information on life we now have at our disposal.

Problems in computational molecular biology range from understanding sequence data to the analysis of protein shapes, prediction of biological function, study of gene networks, and cell-wide computations.

Cornell has a university-wide plan in the science of genomics; CS is playing a critical role in this initiative. CS researchers are engaged in a broad range of computational-biology projects, from genetic mapping to advanced sequence analysis, fold prediction, structure-comparison algorithms, protein classification, comparative genomics, and long-time simulation of protein molecules.

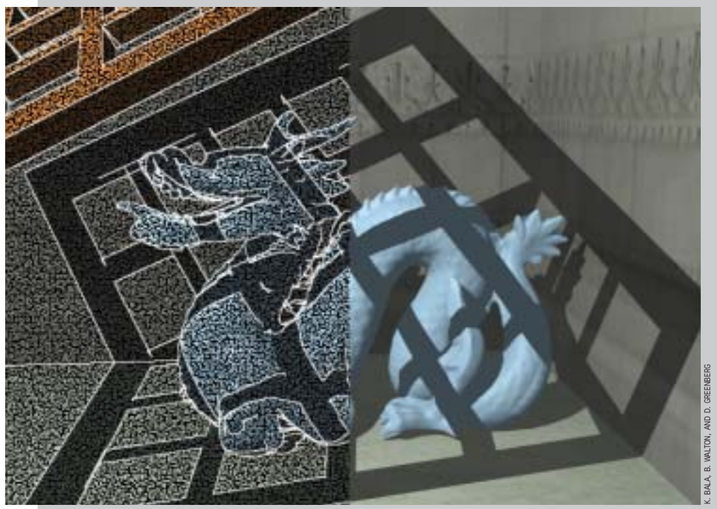

K. BALA, B. WALTON, AND D. GREENBERG

The edge-and-point rendering system (visualization [*left*]; rendered image [*right*])

Database Systems

The Cornell Database Group is exploring issues related to all aspects of data management. Our interests range from developing efficient algorithms for very large data sets to building large-scale systems for new and emerging applications.

In the Cougar project, we are developing database technology for sensor networks. In the Himalaya project, we are exploring new directions in data mining. In the Pepper project, we are developing a query layer for large-scale peer-to-peer systems. In the Quark project, we are building a unified data-management system for both structured and unstructured data. We also collaborate with researchers in related areas such as systems, algorithms, and artificial intelligence.

Languages and Compilation

Cornell has particular strength in programming languages and compilation, with more than eight faculty members and twenty-five graduate students working in the area. Our research ranges from theory, including logics and semantics, to practical engineering issues in verification, optimizing compilers, security, and run-time systems. In addition, there are strong synergies within the languages and compiler groups, and exciting connections with other subdisciplines. For example, Greg Morrisett and Andrew Myers have developed secure programming languages, such as Typed Assembly Language (TAL) and Java Information Flow (Jif), that are used to ensure the safety and security of networked information systems. We also worked closely with the computer industry. For instance, Keshav Pingali's group has recently licensed program optimization tools to Intel for use in their iA-64 compiler product line.

Computer Graphics

Cornell is a leader in the field of computer graphics. Computer graphics is a broad, interdisciplinary field that includes a wide and growing range of applications, from science to communication to entertainment. Research in computer graphics includes algorithms, physics, psychology, computation, computer vision, and architecture, among other fields.

The Program of Computer Graphics (PCG), an interdisciplinary research center with close ties to CS, was one of the pioneering laboratories in computer graphics. Established in 1974, the PCG has made breakthrough contributions in several areas: research topics include reflectance models; physics-based accurate rendering; visual perception for graphics, sketching, and modeling; medical visualization; and digital photography. The state-of-the-art facility includes many tools for advanced research, including a sophisticated light-measurement laboratory, a 128-processor PC cluster, and a high-resolution tiled projection display.

Over the years, the PCG has brought together researchers and students from different disciplines: computer science; physics; mathematics; electrical, structural, and mechanical engineering; architecture; and perception psychology. The recent hires of CS faculty members Bala and Marschner, who are also members of the PCG, have further strengthened our presence in the field.

Operating Systems, Networks, and Distributed Computing

The Operating Systems group at Cornell examines the design and implementation of the fundamental software systems that comprise our computing infrastructure. Our interests span from the very small, including the smart-card systems that fit on a postage stamp–sized die, to the very large, including the wide-area distributed systems that span the globe.

Overall, we are concerned with fundamental questions in systems design. How should our computing infrastructure be structured to address the diverse challenges posed by ubiquitous computing, sensor networks, wide-area distributed computing, and large-scale Web services? What mechanisms and policies are required for a trustable computing infrastructure? What kind of techniques can we use to measure and characterize Web- and Internet-based systems, and how can we apply the lessons learned to the construction of next-generation networked systems?

To answer essential questions like these, we have undertaken many projects on diverse topics, ranging from peer-to-peer systems, operating system services for ad hoc and sensor networks, fault-tolerant communication protocols, application of formal techniques to Web service construction, secure smart-card operating systems, extensible operating systems, intrusion detection, and secure networked service design, among others.

Scientific and Parallel Computing

Scientists and engineers rely more than ever on computer modeling and simulation to buttress their experiments and designs. From improved understanding of the body's circulatory system to the smart design of new medicines, today's scientific and technological advances would be impossible without the combination of powerful computers and the powerful algorithms running on those computers.

The scientific computing group at Cornell develops the algorithms that underlie simulation and optimization. Matrix computations are a recurring theme in our research. The focus is on efficient and robust algorithms with an eye toward modern high-performance parallel and multithreaded architectures.

Security

Cornell is a leader on a broad range of research issues related to computer security. Under the aegis of the Information Assurance Institute, located within CIS, we tackle the fundamental problem of ensuring the security and reliability of our global critical-computing infrastructure.

Many active research projects are aimed at developing a science-and-technology base that enhances information assurance and ensures the trustworthiness of networked information systems. These project areas range from system and network security to reliability and assurance, spanning language-based security, secure online services, advanced type systems for mobile code, static information-flow control, policy specification and enforcement, and proof carrying code.

Overall, the breadth and depth of the projects undertaken at Cornell are a direct result of the well-integrated, diverse, and collegial environment that CS provides. Our work draws its strength from the synergy between the groups working on security, programming languages, operating systems, logic, and formal methods.

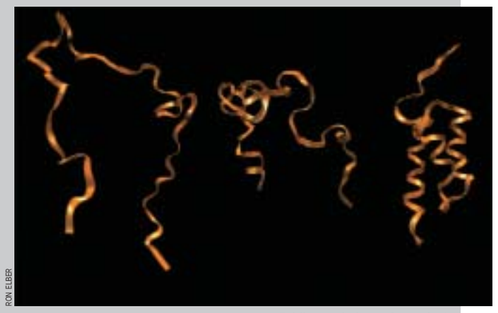

RON ELBER

Folding pathology of protein A

Theory of Computing

The theory of computing is the study of efficient computation, models of computational processes, and their limits. It has emerged over the past few decades as a deep and fundamental scientific discipline. It is a young science, with many central questions still unanswered; and it is a science poised to have considerable impact on current issues in the development of systems and software, the nation's network and communications infrastructure, and the physical and biological sciences. At Cornell, we are proud of our position as a world leader in the ongoing development of theoretical computer science.

Research at Cornell spans all areas of the theory of computing and is responsible for the development of modern computational complexity theory, the foundations of efficient graph algorithms, and the use of applied logic and formal verification for building reliable systems. Our faculty and students are actively involved in areas such as the design of fundamental algorithms, combinatorial optimization, machine learning, computational complexity theory, computational algebra, logic in computer science, computational geometry, and applications to verification, reliable systems, data mining, information science, and the computational sciences.

In addition to its depth in the central areas of theory, Cornell is unique among top research departments in the fluency with which students can interact with faculty members in both theoretical and applied areas and work on problems at the critical juncture of theory and applications.

Corporate Interactions

Gifts and Grants

CS is grateful for the support, including equipment and software, provided by our industrial partners.

Credit Suisse First Boston
$7,500

Google, Inc.
$58,223

Green Hills Software, Inc.
$1,600

Hewlett Packard
$20,000

IBM
$40,000

Intel Foundation
$102,363

Microsoft Corporation
$386,950

PricewaterhouseCoopers, LLP
$1,000

Verizon
$10,000

CIS is also grateful for gifts from the following partners:

Microsoft Corporation
$250,878

McGraw Hill
$4,000

CS faculty members, researchers, and graduate students are conducting leading-edge research in architecture, artificial intelligence, computational biology, databases and digital libraries, languages and compilation, graphics, operating systems, networks and distributed computing, scientific and parallel computing, security, and theory of computing (see pages 22–55 for research summaries). CIS, a university initiative that includes CS, encourages and sponsors interactions with university researchers in interdisciplinary programs including information science and computational biology. Our relationships with corporate partners provide many opportunities for collaboration.

We realize that a true partnership results in mutual satisfaction and gain. Toward this end, we invite our corporate partners to appoint a corporate contact who will work with CIS to build a strategic corporate–CIS partnership, build strong personal relationships on campus, or organize recruiting activities on campus for CS undergraduate and graduate students.

CIS welcomes corporate partners to make unrestricted donations in support of department initiatives, make research grants to individual faculty and researchers, give matching funds to NSF or other granting agencies, create fellowships for graduate students, provide equipment grants, startup funds for new faculty, support for BOOM (Bits on Our Minds, which showcases our student technology work), or course-development grants.

We welcome corporate partners' researchers to the department for long- and short-term visits to work with individual faculty members and research groups. The aforementioned research areas and two institutes, The Information Assurance Institute, and the Intelligent Information Systems Institute, as well as affiliated programs in computational biology, digital arts and graphics, information science, and computational science and engineering, are available to joint researchers on a case-by-case basis.

The Department of Computer Science at Cornell is ranked among the top computer science departments internationally and includes

- 37 faculty members
- 19 full-time research associates
- 110 Ph.D. candidates
- 100 M.Eng. candidates
- 200 undergraduate majors graduating each year

You may address any correspondence regarding corporate interactions to:

Nora Balfour
Department of Computer Science
Cornell University
4130 Upson Hall
Ithaca, NY 14853-7501

E-mail: *corprel@cs.cornell.edu*
Telephone: 607 255-9197
Fax: 607 255-4428
Web site: *http://www.cis.cornell.edu/*

CS faculty members and researchers continued collaborations with the following corporate partners, whose financial contributions support our educational and research missions.

Credit Suisse First Boston and **Microsoft** sponsored **Bits On Our Minds (BOOM)** this past March. The **General Electric Fund** is providing support to identify new programs and approaches to increase the number of women and minorities in computer science. **Green Hills Software** provided support for the ACSU programming contest. **Hewlett Packard** donated equipment to Professor Gün Sirer to facilitate the integration of interactive wireless technology with teaching.

IBM provided the 2002 Faculty Partnership Award for Professor Jayavel Shanmugasundaram.

Intel supported undergraduate teaching labs and provided a fellowship to Ph.D. student Dan Grossman. Intel also supported Professors Gehrke and Shanmugasundaram's scalable-sensor data-management project.

Microsoft supported several faculty and senior researcher projects, including Werner Vogels's distributed systems effort; Professor Johannes Gehrke's work on query caching and routing, and research on a light-weight DBMS and stream processor for sensor devices; and Professor Gün Sirer's research "Assuring the Security of Components in the .NET Framework".

Microsoft also provided a fellowship for Ph.D. candidate Ranveer Chandra, and support for the Information Assurance Institute and the CIS curriculum.

PricewaterhouseCoopers provided support to the CS Undergraduate Computing Association.

Verizon provided support for a CS graduate fellowship.

Programs of the Faculty of Computing and Information Science

Computational Biology

Genomic databases, protein databanks, MRIs of the human brain, and remote-sensing data on landscapes contain unprecedentedly 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 entire 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 computer science, 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 members work primarily in four subareas of computational biology: biomolecular structure and function, 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 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.

Undergraduates can major in computational biology through the new CIS–created undergraduate program of study, which 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 an 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.

Recently Cornell announced a combined graduate program in computational biology with Sloan–Kettering and Rockefeller University. This tri-institutional effort provides three fellowships for Cornell computer science graduate students.

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 multifield program in Computational Molecular Biology. 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.

Through the Cornell Theory Center, two competitive IBM fellowships were granted to undergraduate students doing summer bioinformatic research. The research is a collaboration between the CBSU at the CTC and the Cornell faculty. It exposes the students to high-performance computing and its application to bioinformatics. The CBSU mission is to bridge the gap between molecular biology and mathematical sciences, by helping individual researchers or students, maintaining a computational-biology facility, and by conducting intensive training workshops.

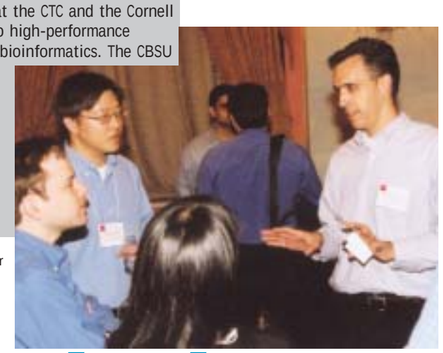

CIS Professor Dan Huttenlocher with alumni in New York City.

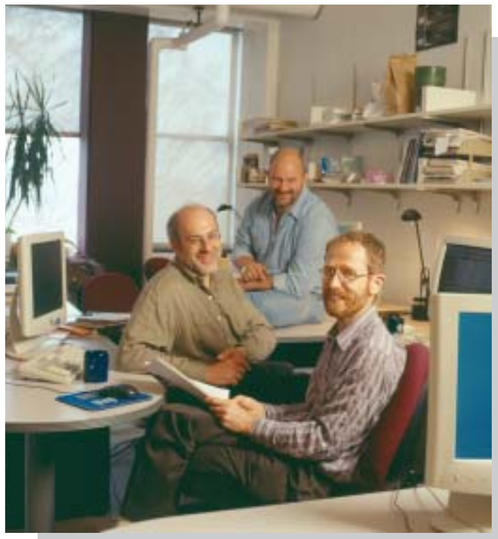

CS Professor Ken Birman [*left*] with researchers Werner Vogels and Robbert van Renesse

Computational Science and Engineering

Many of the faculty members 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.

The CS&E subgroup continues to offer a series of four minicourses, taught by Andrew Pershing:

COM S 401	Introduction to Applied Scientific Computing with MATLAB
COM S 402	Scientific Visualization with MATLAB
COM S 403	Development of Scientific Computing Programs
COM S 404	Survey and Use of Software Libraries for Scientific Computing

Designed primarily for first-year graduate students, these four-week courses provide an efficient introduction to important topics in applied scientific computing. The first two courses focus on the MATLAB programming environment and demonstrate how systems like MATLAB can be used to aid scientific research. The last two courses consider the process of developing scientific software and explore a range of techniques and tools to make this process more efficient. These well-received courses attracted students from across the university.

Digital Arts and Graphics

The digital arts have strong ties to computer graphics at Cornell. With these ties we are building a new academic program. Research in computer graphics includes algorithms, computation, computer architecture, computer vision, physics, and perception psychology.

Cornell's Program of Computer Graphics (PCG), an interdisciplinary research center with close ties to CS, was one of the pioneering laboratories in computer graphics. Established in 1974, the program made breakthrough contributions in radiosity and other aspects of high-quality rendering. Its algorithms for realistic rendering opened the way for computer-aided architectural design, computer-generated motion picture animation, scientific visualization, and more.

Research topics in PCG include reflectance models, physics-based accurate rendering, visual perception for graphics, sketching and modeling, medical visualization, digital photography, and computer animation. The program's modern facility includes many tools for advanced research, including a sophisticated light measurement laboratory, a 128-processor PC cluster, and a high-resolution tiled projection display.

Digital Arts and Graphics (DA&G) involves architects, artists, art historians, city planners, and information scientists. It will explore technical aspects of digital graphics, psychological aspects of vision and perception, the relationship of human senses and robotics, and creating art in a time of digital reproduction.

Information Science

Information Science at Cornell is an interdisciplinary program of CIS 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, cultural, and legal contexts in which information systems function.

Information science has been available since 2002 as an official minor or concentration in all seven of the undergraduate schools or colleges at Cornell. An undergraduate major in information science is currently in the approval process in Arts and Sciences, Engineering, and CALS. We are excited by the enthusiasm with which information science has been received across campus thus far, and look forward to welcoming undergraduates into the information science program this coming year.

Students in the programs obtain an understanding of the core topics of study emerging in this new and rapidly 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.

Annual BOOM expo returns with virtual spiders and velvet switches

By Lissa Harris

Becky Chu, a master's student in computer science, explains her collaborative project, which employs simulations of wolf spiders' courtship dances, to Ron Elber, Cornell professor of computer science, during the BOOM event in Upson Hall, March 5.

Surrounded by intrigued onlookers, Becky Chu, a master's student in computer science, is holding forth on the topic of "spider porn".

No, it's not anime night in the basement of Goldwin Smith. It's the seventh annual BOOM (Bits On Our Minds), the annual expo of student projects hosted by the departments of Computer Science and of Electrical and Computer Engineering. On March 5, 50 displays, ranging from pocket-sized robots to stock market simulators, took over three floors of Upson Hall with a display of digital virtuosity.

Chu's unlikely sounding project is part of a unique collaboration between computer programmers and animal behaviorists. Eileen Hebets, a postdoctoral fellow in neurobiology and behavior, wondered just what it is about the male wolf spider's elaborate courtship dance that females find intriguing. She found that female spiders, being extremely visually oriented creatures, react with interest to video clips of male spiders performing mating dances.

Hebets wanted to be able to edit the video clips—for instance, to systematically change the speed of the dance or the size or color of the male's forelegs—in order to track which characteristics were most important. To help her do that, Chu wrote a piece of software that allows a researcher to construct a virtual male wolf spider, whose behavior and appearance can be manipulated interactively.

"I think of it as kind of like a dating service," Chu joked. Whether the female spiders will be fooled remains to be seen.

Nearby, Jay Ayres '04 and Lin Zhu '03 displayed a few of their small research objects: rectangular blocks just big enough to house a couple of AA batteries and bristling with electronic sensors. The devices, known as "cougar motes," collect information about their environment (such as light, heat and motion) and communicate with each other, forming a kind of loosely organized database that continues to function even if one or more of the motes is destroyed.

"There's just enormous potential for what you can do with this," said Ayres, a computer science major. For example, the motes could be used by the military to scout out areas remotely (in fact, the project receives funding from DARPA, the Defense Advanced Research Projects Agency). Or, in a more benign application, industrial workers could use them to monitor levels of vibration throughout a factory, sensing mechanical breakdowns about to occur.

Although most presenters hailed from computer science or engineering, students from traditionally less byte-crunching parts of campus also were represented. Development sociology major Katrina Becker '03 demonstrated a couple of "on/off switches," part of a project she calls "reflective design," on the relationship between people and technology.

The "switches"—one made of wood and the other covered in plush velvet—don't have any particular function. Their purpose, said Becker, is to study how people react to objects that are obviously technological in nature, and how design and appearance influence people's experience of technology.

By gauging how people respond to their objects, Becker and her collaborators hope to come up with ways to improve the way technology is designed and developed. "The idea is that you get people incorporated into design, to meet real needs rather than just feeding consumption," she said.

Other BOOM '03 projects ranged from a slew of computer games to a program that translates the American Sign Language alphabet into printed letters.

This year's BOOM also, for the first time, was an online exhibit as well as a physical one: Each project had an accompanying display on the World Wide Web. The online display can be accessed through BOOM's Web site: *www.cs.cornell.edu/boom*.

BOOM '03 was supported in part by grants from Microsoft and Credit Suisse First Boston.

Cornell Chronicle, March 13, 2003

New Faculty

Paul Francis

Assistant Professor, CS
francis@cs.cornell.edu
http://www.cs.cornell.edu/people/francis

Paul Francis received his Ph.D. from the University College London (UCL) in 1994. Dr. Francis is one of the industry's foremost scientists in large-scale routing and addressing and internetworking. He has fifteen years of research experience in network routing and addressing, large-scale self-configuring networks, and distributed peer-to-peer search.

Francis has done research at MITRE Corporation, Bellcore, NIT Software Labs, and ACIRI (now ICIR), and was chief scientist at two startups, FastFoward Networks and Tahoe Networks. Dr. Francis' innovations include NAT (Network Address Translation), multicast shared trees (used in PIMSM and CBT), shortcut routing, and landmark routing. He is also the originator of two key IPv6 concepts: the unique host identifier (from Pip) and the use of multiple addresses for multihomed sites.

Dr. Francis' research interests looking forward are in the areas of peer-to-peer applications, overlay networks, network host proximity, Internet scaling, and IP mobility.

Dr. Francis has chaired two IETF working groups, and has published numerous RFCs, U.S. and international patents, and research papers.

SELECTED PUBLICATIONS

"IPNL: A NAT–Extended Internet Architecture". *SIGCOMM 2001* (August, 2001). San Diego. (With R. Gummadi)

"Extending the IP Internet Through Address Reuse", *ACM SIGCOMM Computer Communications Review,* 23(1): 16–33. (As Paul Tsuchiya, with T. Eng)

"The Landmark Hierarchy: A New Hierarchy for Routing in Very Large Networks", *Proceedings SIGCOMM 88 Conference, Stanford, California* (August, 1988). (As Paul Tsuchiya)

Uri Keich

Assistant Professor, CS
keich@cs.cornell.edu

Uri Keich received his Ph.D. in mathematics from the Courant Institute in New York City in 1996, and his M.Sc. in mathematics from Technion in Israel in 1991.

Before coming to CS at Cornell, he was a project scientist at the Department of Computer Science and Engineering of the University of California, San Diego, and assistant professor at the Department of Mathematics of the University of California, Riverside, until 2000. He was also a Von Karman Instructor at the Applied Mathematics Department of the California Institute of Technology.

Keich's research interests include statistical and algorithmic problems that arise in areas of bioinformatics such as motif finding, seed design for similarity search, sequence assembly, and mass spectrometry.

SELECTED PUBLICATIONS

"Designing Seeds for Similarity Search in Genomic DNA". *Proceedings of the Seventh Annual International Conference on Research in Computational Molecular Biology (RECOMB-2003)* (April 2003). Berlin, Germany. (With J. Buhler and Y. Sun)

"Genome-Wide Analysis of Bacterial Promoter Regions". *Proceedings of the Pacific Symposium on Biocomputing (PSB-2003)* (January, 2003). Kaua'i, Hawaii. (With E. Eskin, M. Gelfand, and P. Pevzner)

"Finding Motifs in the Twilight Zone", *Proceedings of the Sixth Annual International Conference on Research in Computational Molecular Biology (RECOMB 2002)* (April 2002). Washington, D.C. ACM Press. (With P. Pevzner)

CS Faculty Honors and Awards

CS faculty members have been recipients of many awards in the past, including

- 1 Air Force Office of Scientific Research Young Investigator
- 3 American Academy of Arts and Sciences Fellows
- 1 Fulbright Scholar
- 5 Guggenheim Fellows
- 1 MacArthur Award winner
- 3 members of the National Academy of Engineering
- 6 National Science Foundation Young Investigators
- 1 Office of Naval Research Young Investigator
- 1 New York State Council for Advancement and Support of Education (CASE) Professor of the Year
- 4 Sloan Fellows
- 2 Turing Award winners

John Hopcroft and Juris Hartmanis

Faculty and Senior Researcher Profiles

Stuart Allen

Research Associate
sfa@cs.cornell.edu
http://www.cs.cornell.edu/info/people/sfa/

Stuart Allen received a bachelor's degree in computer science from the University of New Orleans in 1978, and a Ph.D. in computer science from Cornell in 1987. He has held several positions at Cornell since, and is currently a research associate in CS.

Allen's principal interest is in making computer-manipulable formal data an adjunct to, and ideally a medium for, precise human expression, especially argument. This involves the design, justification, and employment of practical formal systems and notations.

The bulk of his work has been in relation to the PRL project (*http://www.nuprl.org*), which has traditionally focused on constructive theory of types and proof by means of tactics. In addition to theory, application, and explanation of type theory–based practice, he has been interested in formalizing and exploiting conventional mathematical notations, as well as the development of interfaces for user immersion in bodies of formal data.

Most recently, Allen's efforts (as part of the PRL project) have been directed at designing methods for implementing digital collections grounded in formal material, especially proof, emphasizing theoretical neutrality and anticipating the coexistence of material with distinct, possibly conflicting, formal bases, entailing the need for strict yet extensible logical accounting.

William Y. Arms

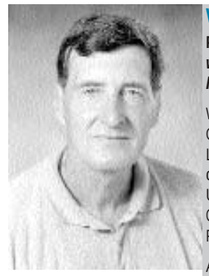

Professor
wya@cs.cornell.edu
http://www.cs.cornell.edu/wya

William Arms received his B.A. degree in mathematics from Oxford University in 1966, and his M.Sc. (Econ.) from the London School of Economics in 1967. He obtained his doctorate (D.Phil.) in operational research from the University of Sussex in 1973. He has been a professor in CS since 1999 and director of the Information Science Program since 2002.

Arms's 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. His book, *Digital Libraries,* was published by the M.I.T. Press in winter 2000. The Cornell Digital Libraries Research Group received a major grant to build the core system for 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. One of Arms's 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.

Professor Arms has recently completed a term as chair of the Association for Computing Machinery (ACM) Publications Board. He is a member of the M.I.T. Press Management Board, and a member of a strategic-planning committee of the American Physical Society.

SELECTED PUBLICATIONS

"Core Services in the Architecture of the National Digital Library for Science Education (NSDL)". *Joint Conference on Digital Libraries* (July, 2002). (With C. Lagoze, S. Gan, D. Hillmann, C. Ingram, D. Krafft, R. Marisa, J. Phipps, J. Saylor, C. Terrizzi, W. Hoehn, D. Millman, J. Allan, S. Guzman-Lara, and T. Kalt) *http://arxiv.org/abs/cs.DL/0201025*.

"A Spectrum of Interoperability. The Site for Science Prototype for the NSDL". *D-Lib Magazine,* 8(1) (January, 2002). (With D. Hillman, C. Lagoze, D. Krafft, R. Marisa, J. Saylor, C. Terrizzi, and H. Van de Sompel) *http://www.dlib.org/dlib/january02/arms/01arms.html*.

"What Are the Alternatives to Peer Review? Quality Control in Scholarly Publishing on the Web". *Journal of Electronic Publishing* 8(1) (August 2002).

Graeme Bailey

Professor
bailey@cs.cornell.edu
http://www.cs.cornell.edu/people/bailey

Graeme Bailey received a bachelor's degree in mathematics in 1973, and a M.Sc. in pure mathematics in 1974, from the University of Birmingham. He obtained a Ph.D. in pure mathematics (low-dimensional topology and combinatorial group theory) in 1977 from the University of Birmingham as well. He has been a professor in CS and an adjunct professor in Cornell's Department of Mathematics since 1988. He is the director of the College of Engineering's Master of Engineering degree program in computer science.

Originally working in low-dimensional topology and combinatorial group theory, through an odd mixture of circumstances Bailey has 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 the Upstate Medical University, in Syracuse, N.Y. This is in the early stages of a program to extend to various pathologies affecting elasticity and aimed towards effective clinical treatments. The group, having made some significant advances in answering questions that had remained unsolved for more than thirty 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 pharmacologists and structural biologists has already yielded some intriguing insights.

Bailey is the recipient of the Kenneth A. Goldman '71 Excellence in Teaching Award 2000, and the Kendall S. Carpenter Memorial Advising Award for 2002.

CS Assistant Professors Steve Marschner and Kavita Bala.

Kavita Bala

Assistant Professor
kb@cs.cornell.edu
http://www.cs.cornell.edu/~kb/

Kavita Bala received her Ph.D. in computer science at the Massachusetts Institute of Technology (M.I.T.). After her doctorate, she worked as a post-doctoral researcher in the Program of Computer Graphics at Cornell. She joined CS as an assistant professor in the fall of 2002.

Bala's research is in the area of computer graphics; her research interests include algorithms and systems for interactive rendering, image-based modeling and rendering, and augmented reality. Increasingly, technology is permitting the acquisition of complex data sets; rendering with these data sets remains a challenge. Bala's research focus is on scalable algorithms for rendering complex scenes with both high fidelity and interactive performance. This research is applicable to both synthetic and augmented-reality rendering. Bala has developed compact, high-dimensional representations and algorithms for interactive rendering of complex dynamic scenes while bounding approximation error. She has also developed hybrid hardware and software algorithms for fast high-fidelity image generation.

SELECTED PUBLICATIONS

"Combining Edges and Points for Interactive High-Quality Rendering". In *Proceedings of SIGGRAPH 2003* (July, 2003). (With B. Walter and D. Greenberg)

"Adaptoive Shadow Maps". In *Proceedings of SIGGRAPH 2001* (August, 2001). (With R. Fernando, S. Fernandez, and D. Greenberg)

"Radiance Interpolants for Accelerated Bounded-error Ray Tracing". *ACM Transactions on Graphics* 18(3) (July, 1999): 213–256. (With J. Dorsey and S. Teller)

Kenneth P. Birman

Professor
ken@cs.cornell.edu
http://www.cs.cornell.edu/ken/

Ken Birman obtained a bachelor's degree in computer science at Columbia University in 1978 and a Ph.D. in computer science at the University of California at Berkeley in 1981. He joined the CS faculty in 1982.

Birman's research is concerned with reliability and security in modern networked environments. In past work on the Isis system, his software became a central part of the New York Stock Exchange and Swiss Stock Exchange (in both settings, Isis runs the core messaging component used to distribute new stock quotes and information about trades reliably and securely), the French air-traffic control system (Isis is used to keep clusters of three to five controller workstations synchronized, and handles failures), the U.S. Navy's Aegis warship's radar system, and other mission-critical computer networks.

Birman's current focus is on a new system called "Astrolabe", which was developed as part of a DARPA–funded Spinglass effort (*http://www.cs.cornell.edu/Info/Projects/Spinglass*). Astrolabe is like a network-wide database in which each computer or component contributes a live tuple. As data changes, 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. A second part of Spinglass is concerned with reliable multicast. Birman's group has developed a scalable multicast protocol that gives probabilistic consistency guarantees, and integrated it with Astrolabe. Underlying both systems is a class of reliable peer-to-peer communication protocols that are extremely scalable and provide probabilistic reliability techniques. The approach permits the development of systems that work as well with ten thousand computers as they do with just ten.

Birman was named an ACM Fellow in 1999 and won the Stephen '57 and Marilyn Miles Excellence in Teaching Award in 2000. He was editor in chief of *ACM Transactions on Computer Systems* from 1993 to 1997, and has served on a number of university committees. Birman is chairman of the Responsible Conduct of Research Committee and is a member of the Founding Committee for the CIS faculty; the Engineering College Policy Committee; and the IP Advisory Council for the Cornell Research Foundation.

SELECTED PUBLICATIONS

"Astrolabe: A Robust and Scalable Technology for Distributed Systems Monitoring, Management, and Data Mining". In *ACM Transactions on Computer Systems* (May, 2003) 21(3). (With R. Van Renesse and W. Vogels)

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

Building Secure and Reliable Network Applications. Manning Publications and Prentice Hall (December, 1996).

"Software for Reliable Networks". *Scientific American* 274(5) (May, 1996): 64–69. (With 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/

Martin Burtscher received his Ph.D. degree in computer science from the University of Colorado at Boulder in 2000 and his B.S./M.S. degree in computer science from the Swiss Federal Institute of Technology (ETH) Zurich in 1996. He is an assistant professor in ECE at Cornell.

His research interests include high-performance microprocessor architecture, instruction-level parallelism, and compiler optimizations.

High-end microprocessors rely on a variety of predictors for good performance. Future CPUs will likely need even more predictors to meet the continuing demand for more and more processing power. Designing, evaluating, and improving such predictors is an important focus of Burtscher's 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.

SELECTED PUBLICATIONS

"Hybrid Load-value Predictors". *IEEE Transactions on Computers* 51(7) (July, 2002). (with B. Zorn)

"Delphi: Prediction-based Page Prefetching to Improve the Performance of Shared Virtual Memory Systems". *International Conference on Parallel and Distributed Processing Techniques and Applications* (June, 2002). (With E. Speight)

"Static Load Classification for Improving the Value Predictability of Data-cache Misses". *Conference on Programming Language Design and Implementation* (June, 2002). (With A. Diwan and M. Hauswirth)

Claire Cardie

Associate Professor
cardie@cs.cornell.edu
http://www.cs.cornell.edu/home/cardie/

Claire Cardie obtained a B.S. in computer science from Yale University in 1982 and an M.S. and Ph.D. in computer science at the University of Massachusetts at Amherst in 1994. She has been a CS faculty member at Cornell since 1994.

Cardie's research is in the areas of natural language processing and machine learning. In particular, her group has focused both on building systems for large-scale natural language processing tasks like information extraction, question-answering, and multidocument summarization, and on developing corpus-based machine learning techniques to address underlying theoretical problems in the syntactic and semantic analysis of natural language.

Cardie is a recipient of a NSF Faculty Early Career Development (CAREER) Award (1996–2000) and was program chair for the Second Conference on Empirical Methods in Natural Language Processing in 1997. She has been secretary of the Association for Computational Linguistics Special Interest Group on Natural Language Learning (1999–2001), and is currently serving a four-year term as secretary of the North American Association for Computational Linguistics.

SELECTED PUBLICATIONS

"Improving Machine Learning Approaches to Coreference Resolution". *Proceedings of the Fortieth Annual Meeting of the Association for Computational Linguistics* (2002). (With V. Ng)

"A Cognitive Bias Approach to Feature Selection and Weighting for Case-based Learners". *Machine Learning* 41(1) (2000): 85–116.

"Error-driven Pruning of Treebank Grammars for Base Noun Phrase Identification". In *Proceedings of the Annual Conference of the Association for Computational Linguistics* (1998): 218–224. (With D. Pierce)

Rich Caruana

Assistant Professor
caruana@cs.cornell.edu
http://www.cs.cornell.edu/~caruana/

Rich Caruana obtained his Ph.D. in computer science from Carnegie Mellon University in 1998. Currently he is an assistant professor in CS, where he does research in machine learning and data mining. His current focus is on ensemble learning, inductive transfer, rank learning, adaptive clustering, and applications of these methods to problems in medical-decision making and protein-folding.

Inductive transfer is a subfield of machine learning that aims to achieve better performance by learning related problems simultaneously—surprisingly, sometimes it is easier to learn 100 problems at the same time than to learn any one of them in isolation. Caruana helped create this subfield by publishing the first paper on multitask learning ten years ago.

Learning rankings is an exciting new area in machine learning that has important applications in information retrieval and medicine. Caruana is developing algorithms that learn rankings for problems in medical-decision making where it may be difficult to assess absolute risk for a patient, but easier to learn to order patients by relative risk. He developed the first machine-learning algorithm specifically designed to learn rankings. This method outperformed a dozen other learning methods in a multi-institutional pneumonia-risk prediction project.

In 2000–01 Caruana led a team of researchers that developed the first automated system for the early detection of bioterrorist releases of anthrax. The system applies data mining to consumer purchases in supermarkets to look for unexplained increases in the sales of products such as analgesics and cough syrup.

Caruana's work in ensemble learning and clustering are new focuses for him. His interest in clustering arose from limitations he discovered when applying traditional clustering methods to the protein-folding problem with colleagues in bioinformatics. The research in ensemble learning arose from a competition in a machine-learning course he teaches at Cornell where students use different learning methods to make accurate predictions for a mystery data set.

A theme that runs through all of Professor Caruana's work is the importance of developing methods that are effective on real-world problems. He likes to mix algorithm development with applications work to insure that the methods are useful.

SELECTED PUBLICATIONS

"Multitask Learning". *Machine Learning* 28 (1997): 41–75.

"Using the Future to `Sort Out' the Present: Rankprop and Multitask Learning for Medical Risk Evaluation". *Advances in Neural Information Processing Systems* 8 (1996). (With S. Baluja and T. Mitchell)

"Early Statistical Detection of Anthrax Outbreaks by Tracking Over-the-counter Medication Sales". *Proceedings of the National Academy of Sciences* 99(8) (April, 2002): 5237–5240. (With A. Goldberg, G. Shmueli, and S. Feinberg)

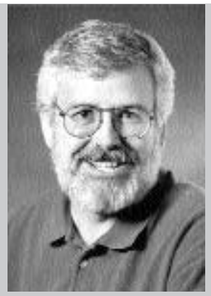

L. Paul Chew

Senior Research Associate
chew@cs.cornell.edu
http://www.cs.cornell.edu/info/people/chew/chew.html

Paul Chew received his Ph.D. in computer science from Purdue University in 1981. He served as a faculty member at Dartmouth College until 1988 when he joined CS at Cornell as a senior research associate.

Chew's primary research interest is in geometric algorithms with an emphasis on practical applications. These practical applications have included placement, motion planning, vision, sensing, mesh generation, molecular matching, and protein shape–comparison. The work on protein shape–comparison has been used as part of the evaluation scheme for CAFASP (Critical Assessment of Fully Automated Structure Prediction), a competition held every two years to evaluate the performance of fully automatic servers for protein-structure prediction. Chew developed *backwards analysis,* a method now widely used for analyzing randomized algorithms. Chew's work on mesh generation has been motivated by the finite-element method, a technique 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. For complex geometries mesh generation can be difficult. Chew has developed methods for automatically generating a high-quality mesh. This work is being used in a large, multidisciplinary project: developing adaptive software for field-driven simulations.

Chew is an associate editor for *Pattern Recognition,* the journal of the Pattern Recognition Society. He is also a member of the steering committee for the International Meshing Roundtable.

SELECTED PUBLICATIONS

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

"Voronoi Diagrams of Lines in 3-space Under Polyhedral Convex Distance Functions". *Journal of Algorithms* 29(2) (1998): 238–255. (With K. Kedem, M. Sharir, B. Tagansky, and E. Welzl)

"Constrained Delaunay Triangulations". *Algorithmica* 4(1) (1989): 97–108.

Thomas F. Coleman

Professor
Director, Cornell Theory Center
Director, CTC–Manhattan
coleman@cs.cornell.edu
http://www.tc.cornell.edu/~coleman/

Thomas F. Coleman obtained his bachelor's degree in mathematics in 1975, and his master's in mathematics in 1976, both from the University of Waterloo. He received a Ph.D. in mathematics from Waterloo as well, in 1979. He is currently a professor of computer science and applied mathematics at Cornell, and also the director of both the CTC, a center for the support of large-scale computational science, and CTC–Manhattan, a computational finance center in New York City.

With colleagues Shirish Chinchalkar, Yohan Kim, Yuying Li, Peter Mansfield, and Arun Verma, Coleman is 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: Jay Henniger, Cristina Patron, Siddharth Alexander, Katharyn Boyle, and Changhong He. In their most recent academic work, "Derivative Portfolio Hedging Based on CVaR", an efficient new way to hedge large portfolios of derivative instruments is proposed.

Coleman's specific interests include the computation of implied volatility surfaces from option prices, hedging techniques, index tracking, portfolio optimization, and the use of parallel computing techniques in computational finance.

Professor Coleman is a member of both the admissions committee and the program committee for the Center for Applied Mathematics. He is the author of two books on computational mathematics, the editor of four proceedings, and has published over sixty journal articles. He was chair of the SIAM Activity Group on Optimization (1998–2001) and serves on the editorial boards of numerous professional journals.

SELECTED PUBLICATIONS

"Reconstructing the Unknown Local Volatility Function". *Journal of Computational Finance* 2(3) (Spring, 1999): 77–102. (With Y. Li and A. Verma)

"The Efficient Computation of Sparse Jacobian Matrices Using Automatic Differentiation". *SIAM Journal on Scientific Computing* 19(4) (1998): 1210–1233. (With A. Verma)

"Large Sparse Numerical Optimization". *Lecture Notes in Computer Science* 165 (May, 1984): 105 p.

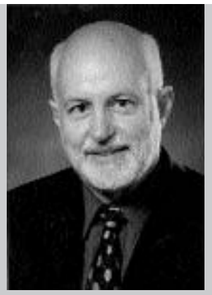

Robert Constable

Professor
Dean for Computing and Information Science
rc@cs.cornell.edu
http://www.cs.cornell.edu/home/rc/

Robert L. Constable is the Dean for Computing and Information Science, and a professor in CS. He obtained his Ph.D. in mathematics from the University of Wisconsin in 1968. He served as CS chair from 1994 to 1999. He was also acting chair from 1993 to 1994.

Constable's research has focused on building a system called 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 his group's case, they 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, a theorem prover, and a formal digital library. Constable's group uses the latest version of Nuprl as the prover.

He is also working with others to build the formal digital library component of the LPE that will allow interactive access to theorems and proofs from Nuprl, MetaPRL, PVS, and other major theorem provers. 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. The group is funded by the Office of Naval Research (ONR) to further develop and explore the concept of a formal digital library of constructive mathematics built around these theorems. Their 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.

Constable is the director of the PRL Project, and a member of the Cognitive Studies executive committee; the applied math policy committee; and the LICS General Committee. He serves as editor for the *Journal of Logic and Computation; Formal Methods in System Design;* and the *Journal of Symbolic Computation.*

SELECTED PUBLICATIONS

"Naive Computational Type Theory" In *Proof and System Reliability,* Proceedings of the International Summer School Marktoberdorf, July 24 to August 5, 2001 (H. Schwichtenberg and R. Steinbruggen, eds): 213–260. NATO Science Series II, Kluwer Academic Publishers, Amsterdam.

"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)

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

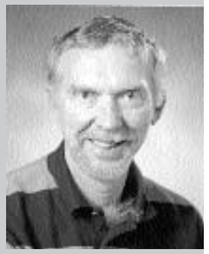

Alan J. Demers

Professor
ademers@cs.cornell.edu

Alan J. Demers received his bachelor's degree in physics from Boston College in 1970. He obtained his Ph.D. in computer science from Princeton University in 1975. Demers was at Cornell from 1975 to 1984. He then worked at Xerox PARC in Palo Alto as a principal scientist, and at Oracle Corporation as an architect, before returning to CS in 2000.

Demers's research concerns aspects of weakly consistent data replication in databases, sensor networks, and distributed systems. With Ken Birman, Robbert van Renesse, Johannes Gehrke, and others, he is studying randomized "gossip protocols". Such protocols are highly fault-tolerant and, when properly designed, extremely scalable as well. The group is studying convergence properties of several flat and hierarchical versions of the basic protocols tailored to specific application requirements.

More specifically, Demers's focus is approximate evaluation of aggregate queries in such a system. He is studying age distributions of gossiped data in order to prove probabilistic bounds on the quality of aggregate query results. Alternatively, the group can use this approach to bound the latency required to probabilistically guarantee a client-specified degree of consistency. In addition, they are considering algorithms for specific problems that frequently occur in high-level protocols: for example, accurate estimation of the size of a group with no initial information.

Finally, they 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.

These techniques are being adapted to achieve energy-efficient routing in ad-hoc sensor networks.

SELECTED PUBLICATIONS

"KELIPS: Building an Efficient and Stable P2P DHT through Increased Memory and Background Overhead". In *Proceedings of the Second International Workshop on Peer-to-Peer Systems (IPTS03)* (February, 2003). (With I. Gupta, K. Birman, P. Linga, and R. Van Renesse)

"Research Issues in Distributed Mining and Monitoring". In *informal proceedings of the NSF workshop on next-generation data mining (NGDM02)* (Baltimore, Maryland: November, 2002). (With J. Gehrke and M. Riedewald)

"Spatial Gossip and Resource Location Protocols". In *Proceedings of the Thirty-third ACM Symposium on Theory of Computing* (July, 2001): 163–172. (With D. Kempe and J. Kleinberg)

Ron Elber

Professor
ron@cs.cornell.edu
http://www.cs.cornell.edu/ron/

Ron Elber obtained a bachelor's degree in chemistry and physics in 1981, and a Ph.D. in theoretical chemistry in 1984 at the Hebrew University of Jerusalem. He was a postdoctoral fellow in theoretical biophysics from 1984 to 1987 at Harvard University. Ron was on the chemistry faculty of the University of Illinois (1987–1992) and on the chemistry and biology faculty at Hebrew University of Jerusalem (1992–1999). Since 1999 he has been on the CS faculty, where he is currently a full professor. Ron is also a faculty member of the Department of Biological Statistics and Computational Biology.

Ron's research is in computational biology and bioinformatics. His group is developing novel tools (MOIL) to simulate dynamics of biological macromolecules. His current research focuses on algorithms to extend the time scales of simulations, and to study complex processes such as the kinetics of protein folding. Ron's techniques for path following and enhanced sampling are in wide use and motivated the development of related algorithms. His bioinformatic investigations focus on protein annotation, using sequence-to-structure matches (LOOPP). LOOPP linked a gene that influences the size of the tomato fruit with a human protein that controls cell growth and may cause cancer.

Ron received the Stein award for his Ph.D. studies (1984), and the Camille and Henry Dreyfus New Faculty Award (1987–1990). He was a University of Illinois Scholar (1991–1992). Ron received the Alon new faculty award (1992–1994), and the Bergman award (1994).

SELECTED PUBLICATIONS

"An Anatomically Detailed Study of the Folding Pathways of Protein A with the Stochastic Difference Equations". *PNAS 99* (2002). (With A. Ghosh and H. Scheraji)

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

"Anharmonic Wavefunctions of Proteins: Quantum Self-consistent Field Calculations of BPTI". *Science* 268 (1995): 1319–1322. (With A. Roitberg, R. Gerber, and M. Ratner)

K-Y. Daisy Fan

Assistant Professor
dfan@cs.cornell.edu
http://www.cs.cornell.edu/~dfan/

Daisy Fan obtained her B.Sc. and M.Sc. degrees in civil engineering at the University of Manitoba in 1994 and 1997, respectively, and her Ph.D. degree in civil and environmental engineering at Cornell in 2002. She is currently an assistant professor in CS. Her research interests include the application of systems-analysis techniques for water-resources and environmental problems. Problems she has investigated include optimal control of multiple-reservoir operation using stochastic dynamic programming and river-basin water-quality management. She teaches CS 100, and with Professor David Schwartz, develops the academic-excellence workshops that are associated with the programming courses.

Fan is the director of the Summer College Explorations in Engineering Seminar for high school students. She actively participates in outreach initiatives, including Cornell's CURIE Academy, which showcases engineering to high school girls.

Fan is a recipient of a Graduate Teaching Assistant Award in CS (2000), a New York State American Water Works Association Russell L. Sutphen Scholarship (2000), and a Cornell School of Civil and Environmental Engineering John E. Perry Teaching Assistant Prize (1999).

SELECTED PUBLICATIONS

"Regression Dynamic Programming for High-dimensional Continuous-state Problems". In preparation for submittal to *Operations Research.* (With C. Shoemaker and D. Ruppert)

"First Programming Course in Engineering: Balancing Tradition and Application". In *Proceedings of the Annual Conference and Exposition of the ASEE* (June, 2003). (With D. Schwartz)

"Stochastic Multiple-reservoir Optimization Using Regression Dynamic Programming". In *Proceedings of World Water and Environmental Resources Congress,* (May, 2001).

Geri Gay

Professor
CIS, joint with Communication
gkg1@cornell.edu
http://www.comm.cornell.edu/faculty/gay.html

Geri Gay is the director of the Human Computer Interaction Group (HCI Group) and a professor in Department of Communication. She received her Ph.D. from Cornell in 1985. The HCI Group is a research-and-development group whose members design and research the use of computer-mediated learning environments. Current research focuses on the use and design of PDAs for communication and collaboration (funded by Intel). Other research examines navigation issues, knowledge management, social network analysis (NSF), knowledge representations, collaborative work and learning (NASA and AT&T Foundation), and activity-centered design of mobile devices.

Professor Gay teaches courses in computer-mediated communication, human–computer interaction, and the social design of communication systems. She was awarded the New York State Chancellor's Award for Excellence in Teaching in 2001.

SELECTED PUBLICATIONS

Activity Centered Design: An Ecological Approach. Boston: M.I.T. Press. [*in press*] (With H. Hembrooke)

"Collaboration in Wireless Learning Networks". *Proceedings of the HICSS Conference* (January, 2002). (With H. Hembrooke)

"E-Graffiti: Evaluating Real-world Use of a Context-Aware System". Interacting with Computers Special Issue on Universal Usability (2001). (With J. Burrell)

CIS Dean Robert L. Constable with Professor Geri Gay.

SLOAN RESEARCH FELLOWSHIP

Johannes Gehrke was a recipient of a 2003 Sloan Research Fellowship.

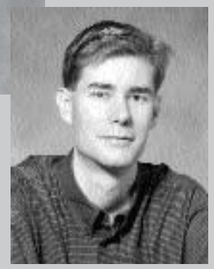

Johannes Gehrke

Assistant Professor
johannes@cs.cornell.edu
http://www.cs.cornell.edu/johannes/

Johannes Gehrke obtained his Ph.D. in computer science at the University of Wisconsin at Madison in 1999, and he has been an assistant professor in CS since then.

Gehrke's research interests are in the areas of data mining, data stream processing, and novel applications of distributed database technology. In his current research, he is working on integrating complex querying capabilities into wireless sensor networks and peer-to-peer networks, and he is developing database techniques for high-speed data streams. His data mining research includes privacy-preserving data mining, theoretical foundations of data mining, and high-performance data mining algorithms, and his group has developed some of the fastest known algorithms for several important data mining tasks.

Gehrke is a recipient of an Alfred P. Sloan Fellowship (2003), a National Science Foundation CAREER Award (2002), an IBM Faculty Award (2000 and 2001), and the Cornell College of Engineering James and Mary Tien Excellence in Teaching Award (2001).

SELECTED PUBLICATIONS

"Approximate Join Processing over Data Streams". *In Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD 2003)* (San Diego, California: June, 2003) (With A. Das and M. Riedewald)

"Limiting Privacy Breaches in Privacy Preserving Data Mining". In *Proceedings of the Twenty-Second ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems* (San Diego, California: June, 2003). (With A. Evfimievski and R. Srikant)

Database Management Systems, third ed. (2002). McGraw Hill. (With R. Ramakrishnan)

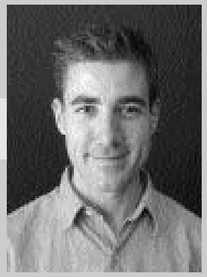

Tarleton Gillespie

Assistant Professor
CIS, joint with Science and Technology Studies
tlg28@cornell.edu

Tarleton Gillespie received his bachelor's degree in English from Amherst College in 1994, his master's in communication in 1997, and his Ph.D. in 2002 from the University of California at San Diego.

His work focuses on the cultural and institutional arrangements surrounding media technologies, considering how power and practice are woven into their use, and the cultural notions of their value. In particular, he is interested in the way that law and technology sometimes do battle, but more often are often brought together to regulate knowledge production.

His research uses recent disputes over copyright and the Internet to analyze the historical contest over the nature of authorship, law, and technology. He is interested in the history of copyright, which he feels has borrowed particular romanticized notions of authorship and traditionally neutral notions of technology—to rationalize and naturalize one system of distribution of creative work, i.e., a corporate-driven commercial market system.

His theoretical purpose is to reject the deterministic tone of most claims about media and cultural expression, and replace them with an understanding of technology as a complex material artifact, but one that may be articulated in ways that seem to support one ideological agenda or another. He finds this argument an important one to make, especially now—precisely because the decisions made today will set the standards by which the Internet is developed and regulated in the future.

SELECTED PUBLICATIONS

"Copyright and Commerce: The DCMA, Trusted Systems, and the Stabilization of Distribution". In *The Information Society* (forthcoming, Fall 2003).

"The Stories Digital Tools Tell". In *New Media: Theses on Convergence, Media, and Digital Reproduction,* (J. Caldwell and A. Everett, eds.) Routledge (February, 2003).

JOHN D. AND CATHERINE T. MacARTHUR FOUNDATION FELLOWSHIP

Paul Ginsparg was a 2002 recipient of a John D. and Catherine T. MacArthur Foundation Fellowship.

Paul Ginsparg

Professor
CIS, joint with Physics
ginsparg@cs.cornell.edu
http://www.physics.cornell.edu/profpages/Ginsparg.htm

Paul Ginsparg received his A.B. in physics from Harvard University in 1977 and his Ph.D. in physics from Cornell in 1981 (Quantum Field Theory, thesis advisor: Kenneth G. Wilson). He was in the Harvard Society of Fellows from 1981–84, and a junior faculty member in the Harvard physics department from 1984–90. From 1990–2001, he was a technical staff member in the theoretical division at the Los Alamos National Laboratory.

Ginsparg came to Cornell in 2001, where he holds a joint appointment with the Department of Physics and the Faculty of CIS. He has been an A. P. Sloane Fellow and a DoE Outstanding Junior Investigator, and has 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 the University of California at Santa Barbara; the Mathematical Science Research Institute at University of California at Berkeley; and at Hebrew University of Jerusalem. In 1991, Ginsparg initiated the "e-print arXiv" as a new form of communications-research infrastructure for physics.

Ginsparg's current research in information science investigates the optimal combination of automated text classification, data mining, machine learning, human–computer interaction, quantum field theory, and related techniques for use in research-communications infrastructure.

SELECTED PUBLICATIONS

"A Remnant of Chiral Symmetry on the Lattice". *Physical Review* D25 (1982): 2649. (With K. G. Wilson)

"2D Gravity + 1D Matter". *Physics Letters* B240 (1990): 333. (With J. Zinn–Justin)

"Winners and Losers in the Global Research Village". In *Proceedings of "Electronic Publishing in Science"*, Sir R. Elliot and D. Shaw, eds. ICSU Press (1996).

Carla P. Gomes

Associate Professor
CIS, joint with Applied Economics and Management
Director, IISI
gomes@cs.cornell.edu
http://www.cs.cornell.edu/gomes

Carla P. Gomes obtained a Ph.D. in computer science in the area of artificial intelligence and operations research from the University of Edinburgh in 1993. She also holds an M.Sc. in applied mathematics from the University of Lisbon.

Gomes's research has covered many areas in artificial intelligence and computer science, including planning and scheduling, integration of CSP and OR techniques for solving combinatorial problems, software agents, and algorithm portfolios.

Her current projects focus on the interplay between problem structure and computational hardness, the use of approximation methods in large-scale constraint-based reasoning systems, and applications of constraint-based reasoning and optimization in multi-agent optimal control, distributed wireless networks, and combinatorial auctions. She was the conference chair of the *Eighth International Conference on Principles and Practice of Constraint Programming* (CP-2002). Gomes is also the director of the Intelligent Information Systems Institute (IISI) at Cornell.

SELECTED PUBLICATIONS

"An Improved Approximation for the Partial Latin Square Extension Problem". in *Proceedings of the ACM–SIAM Symposium on Discrete Algorithms (SODA)*. (2003). (With R. Regis and D. Shmoys)

"Artificial Intelligence and Operations Research: Challenges and Opportunities in Planning and Scheduling". *Journal of Knowledge Engineering Review* 15(1) (2000): 1–10.

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

AAAS FELLOWSHIP

Donald Greenberg was named an AAAS Fellow in 2003.

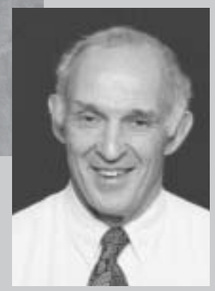

Donald Greenberg

Professor
Member of CIS, 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

Dr. Greenberg received his Ph.D. from Cornell in 1968. He joined the Cornell faculty in 1968, with a joint appointment in the Departments of Architecture and Structural Engineering. His prior education consisted of both the architecture and engineering disciplines at Cornell and Columbia University. From 1960 to 1965, he served as a consulting engineer with Severud Associates, and was involved with the design of numerous building projects, including the St. Louis Arch, New York State Theater of the Dance at Lincoln Center, and Madison Square Garden. Early in his career he taught courses in structural analysis and design, architectural design, shell structures, reinforced concrete, and computer applications in architecture. In 1970–71, he was a guest professor at the ETH in Zurich, Switzerland, and he has been a visiting professor at Yale University.

Professor Greenberg's current research is primarily concerned with physically based image synthesis and with applying graphic techniques to a variety of disciplines. His specialties include color science, parallel processing, and real-time realistic-image generation. His application work includes medical imaging, architectural design, visual perception, digital photography, and computer animation.

Greenberg is a member of Cornell's faculty in the Johnson Graduate School of Management, CS, and the Department of Architecture, and a founding member of the Faculty of CIS. In recent years he has taught courses in computer graphics, computer-aided architectural design, digital photography, and disruptive technologies.

He is the director of the Program of Computer Graphics and was the founding director of the NSF Science and Technology Center for Computer Graphics and Scientific Visualization. More than 300 articles on computer graphics have been published by the Program of Computer Graphics and many of Professor Greenberg's students have been highly recognized in the field, including several who have received the SIGGRAPH Achievement Award and others who have received Hollywood Oscars.

Greenberg received the ACM Steven Coons Award in 1987, the highest honor in the graphics field; the National Computer Graphics Association Academic Award in 1989; the ASCA Creative Research Award in Architecture in 1997; and an honorary doctoral degree from New Jersey Institute of Technology. He is an ACM Fellow and a member of the National Academy of Engineering.

SELECTED PUBLICATIONS

"Visions of Light". *Metropolis Magazine* (by Jonathan Ringen): 138–144, 185–189 (June, 2002)

"Enhancing and Optimizing the Render Cache". In *Proceedings of the Thirteenth Eurographics Workshop on Rendering* (June, 2002). (With B. Walter. G. Drettakis, and O. Deussen)

"Adaptive Shadow Maps". In *Proceedings of SIGGRAPH 2001* (August, 2001). (With R. Fernando, S. Fernandez, and K. Bala)

David Gries

Associate Dean of Engineering for Undergraduate Programs
Professor of Computer Science
Cornell Weiss Presidential Fellow
gries@cs.cornell.edu
http://www.cs.cornell.edu/home/gries/gries.html

Professor Gries's research is aimed at gaining a better understanding of the programming process, with respect to both sequential and concurrent (or parallel) programs. The work requires investigation of theories of program correctness and their application, as well as investigation of other concepts in the semantics of programming languages.

Education is a also a strong interest for Gries, particularly the first few courses in computer science. Under the thesis that logic is the glue that binds together reasoning in all domains, with colleague F. B. Schneider, Gries developed a text, *A Logical Approach to Discrete Math,* which makes a usable "calculational logic" the foundation for almost all the discrete math topics.

Gries is also heavily involved in writing an introductory programming text, based on Java. Earlier, with his son, he developed a "livetext"—a text that comes on a CD and has more than 250 two- to three-minute recorded lectures with synched animation, as well as other innovative features. They are finishing up a paperback text to accompany it.

Gries received the Dr. rer. nat. degree from the Munich Institute of Technology in 1966; a Doctor of Science degree (Honorary) from Miami University in Oxford, Ohio in 1999; and a Doctor of Laws degree (Honorary) from Daniel Webster College in Nashua, New Hampshire in 1996. He won the ACM Karl V. Karlstrom Outstanding Educator Award in 1996, and the Taylor L. Booth Award Education Award, IEEE Computer Society, in 1995.

SELECTED PUBLICATIONS

A Logical Approach to Discrete Math. Springer Verlag, New York. (1993) (With F. Schneider)

ProgramLive. John Wiley & Sons (2000). (With P. Gries). (This livetext comes on a CD and has over 250 recorded lectures, with synched animation.)

Zygmunt J. Haas

Associate Professor
Member of the School of Electrical and Computer Engineering and the Graduate Field of Computer Science
haas@ece.cornell.edu
http://people.ece.cornell.edu/haas/

Zygmunt J. Haas received his Ph.D. degree from Stanford University in 1988 and subsequently joined AT&T Bell Laboratories, where he pursued research on wireless communications, mobility management, fast protocols, optical networks, and optical switching. In August of 1995, he joined the ECE faculty at Cornell.

Haas is an author of numerous technical papers and holds fifteen patents in the fields of high-speed networking, wireless networks, and optical switching. He has organized several workshops, delivered tutorials at major Institute of Electrical and Electronics Engineers (IEEE) and ACM conferences, and serves as editor of several journals and magazines, including the *IEEE Transactions on Networking* and *IEEE Transactions on Wireless Communications.* He has been a guest editor of three IEEE *Journal of Selected Areas on Communications (JSAC*) issues ("Gigabit Networks", "Mobile Computing Networks", and "Ad-Hoc Networks"). Haas is the chair of the IEEE Technical Committee on Personal Communications.

Haas's current interests include: mobile and wireless communication and networks, personal communication service, and high-speed communication and protocols. He heads the Wireless Networks Laboratory (WNL) at Cornell, which performs research in the area of mobility management for wireless networks, ad hoc networking (routing, multicasting, medium access control (MAC), and topology control), security of wireless communications, and cross-layer design of communication protocols. The ad hoc networking technology is the central research area of WNL. In particular, Haas's research group has developed the first hybrid ad hoc routing protocols—the Zone Routing Protocol—which is currently an Internet Engineering Task Force (IETF) draft. The WNL (*http://wnl.ece.cornell.edu*) has also pioneered in its research on ad hoc network security.

Dr. Haas is a recipient of the Michael Tien College of Engineering Teaching award in the years 1997 and 2000.

SELECTED PUBLICATIONS

"On Multicast Flow Control for Heterogeneous Receivers". *ACM/IEEE Transactions on Networking* 10(1) (February, 2002): 86–101. (With R.–H. Gau and B. Krishnamachari)

"The Performance of Query Control Schemes for the Zone Routing Protocol". *ACM/IEEE Transactions on Networking* 9(4) (August, 2001): 427–438. (With M. Pearlman)

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

ACM FELLOW

Joseph Halpern was named an ACM Fellow in 2002.

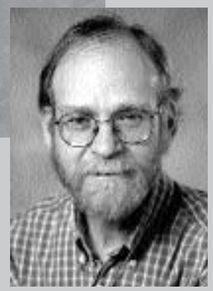

Joseph Halpern

Professor
Codirector, Cognitive Studies Program
halpern@cs.cornell.edu
http://www.cs.cornell.edu/home/halpern/

Joseph Halpern received a B.Sc. in mathematics from the University of Toronto in 1975 and a Ph.D. in mathematics from Harvard in 1981. In between, he spent two years as the head of the mathematics department at Bawku Secondary School in Ghana. After a year as a visiting scientist at M.I.T., he joined the IBM Almaden Research Center in 1982, where he remained until 1996, also serving as a consulting professor at Stanford. In 1996, he joined CS at Cornell.

Halpern's major research interests are in reasoning about knowledge and uncertainty, security, distributed computation, and decision theory. Together with his former student, Yoram Moses, he pioneered the approach of applying reasoning about knowledge to analyzing distributed protocols and multi-agent systems. He has coauthored five patents, a book, *Reasoning About Knowledge*, and over 200 technical publications.

Halpern is a Fellow of the AAAI and the ACM. Among other awards, he received the Godel Prize in 1997, and was a Guggenheim Fellow and a Fulbright Fellow in 2001–02. Two of his papers have won best-paper prizes at International Joint Conferences on Artificial Intelligence (IJCAI). Many of his other papers were invited for special issues of journals. He serves as editor-in-chief of the *Journal of the ACM* and on several other editorial boards.

SELECTED PUBLICATIONS

"Knowledge and Common Knowledge in a Distributed Environment". *Journal of the ACM* 37(3) (1990): 549–587. (With Y. Moses)

"An Analysis of First-order Logics of Probability". *Artificial Intelligence* 46(3) (1990): 311–350.

"Plausibility Measures and Default Reasoning". *Journal of the ACM* 48(4) (2001): 648–685. (With N. Friedman)

GRAND MEDAL OF THE LATVIAN ACADEMY OF SCIENCES *(Lielo Medalu)*

TURING AWARD

Juris Hartmanis received the 2001 Grand Medal of the Latvian Academy of Sciences. He was also a recipient of the ACM Turing Award in 1993.

Juris Hartmanis

Emeritus Walter R. Read Professor of Computer Science and Engineering
jh@cs.cornell.edu

Juris Hartmanis obtained his Ph.D. from the California Institute of Technology in 1955. In 1965, he founded CS and was its first chairman. In 1993 he was awarded the Turing Award. In 2001, Hartmanis was a recipient of the Lielo Medalu, the Grand Medal of the Latvian Academy of Sciences.

Hartmanis is also the founder of the field of computational complexity theory. He believes that 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.

Professor Hartmanis's current research interests focus on understanding the structure of computational complexity classes and exploring how to view computation as construction of complex objects and relate computational complexity to the complexity of constructed objects.

He is a member of the National Academy of Engineering; a foreign member of the Latvian Academy of Sciences; a Fellow of the American Academy of Arts and Sciences, the New York State Academy of Sciences, the American Association for the Advancement of Science, and ACM. He serves as editor of the *Springer-Verlag Lecture Notes in Computer Science; Journal of Computer and Systems Sciences;* and *Fundamenta Informaticae*.

SELECTED PUBLICATIONS

"On the Computational Complexity of Algorithms". *Transactions of the American Mathematical Society* 117(5) (May, 1965): 285–306. (With R. Stearns)

"On Isomorphisms and Density of NP and Other Complete Sets". *SIAM Journal on Computing* 6 (June, 1977): 305–322. (With L. Berman)

"Generalized Kolmogorov Complexity and the Structure of Feasible Computations". In *Proceedings of the Twenty-fourth Annual Symposium on Foundations of Computer Science* (November, 1983): 439–445.

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/

Mark Heinrich is an assistant professor in ECE at Cornell, a cofounder of its Computer Systems Laboratory, and an IISI member.

His research interests include active memory and I/O systems, parallel computer architecture, system-area networks, novel computer architectures, embedded architectures, scalable cache-coherence protocols, multiprocessor design and simulation methodology, and hardware/software codesign.

He received his Ph.D. in electrical engineering from Stanford University under John Hennessy in 1998, where he was a principal designer of the FLASH multiprocessor. He was the author of FlashLite, the system-level simulator of the FLASH machine, as well as four cache-coherence protocols for FLASH. He also developed the first model for evaluating the effect of node-controller occupancy in distributed shared-memory machines.

He received his M.S. degree from Stanford in 1993, and his B.S.E. in electrical engineering and computer science from Duke University in 1991. Heinrich was also the cofounder and chief architect of Flashbase, Inc. an Internet company specializing in automated sweepstakes and database-backed forms and tools for customer acquisition. Flashbase was acquired by DoubleClick Inc. in May 2000.

SELECTED PUBLICATIONS

"Leveraging Cache Coherence in Active Memory Systems". In *Proceedings of the Sixteenth ICS* (June, 2002). (With D. Kim and M. Chaudhuri)

"Active Memory Clusters: Efficient Multiprocessing on Commodity Clusters". In *Proceedings of the Fourth ISHPC, Lecture Notes in Computer Science* 2327 (May, 2002): 78–92. (With E. Speight and M. Chaudhuri)

"FLASH vs. (Simulated) FLASH: Closing the Simulation Loop". In *Proceedings of the Ninth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)* (November, 2000): 49–58. (With J. Gibson, R. Kunz, D. Ofelt, M. Horowitz, and J. Hennessy)

Sheila S. Hemami

Associate Professor
Member of the School of Electrical and Computer Engineering and the Graduate Field of Computer Science
hemami@ece.cornell.edu
http://foulard.ece.cornell.edu/Shemami.html

Sheila S. Hemami received her B.S.E.E (1990) and M.S.E.E. (1992) degrees from the University of Michigan. She obtained her Ph.D. degree from Stanford University in 1994.

Her doctoral work comprised development of real-time, low-complexity lossy signal-processing techniques to provide reconstruction of image and video data lost in transmission over lossy packet networks. During her last year at Stanford, she was a member of the technical staff at Hewlett Packard Laboratories in Palo Alto, California. Upon completing her Ph.D., she joined the ECE faculty at Cornell, where she currently directs the Visual Communications Laboratory.

In 1997, she received a National Science Foundation CAREER Award. In 2000 she received the Eta Kappa Nu C. Holmes MacDonald Outstanding Teaching Award (a national award). In 2002, she was a finalist for the Eta Kappa Nu Outstanding Young Electrical Engineer Award. She is a member of the IEEE, Eta Kappa Nu, and Tau Beta Pi.

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. Dr. Hemami's 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.

SELECTED PUBLICATIONS

"Efficient Sign Coding and Estimation of Zero-quantized Coefficients in Embedded Wavelet Image Codecs". *IEEE Transactions on Image Processing* 12(4)(2003): 420–30. (With D. Deever)

"Universal Multiple Description Scalar Quantization: Analysis and Design". *Data Compression Conference 2003, Snowbird, Utah* (2001). (With C. Tian)

"Contrast-based Quantization and Rate Control for Wavelet-coded Images". In *Proceedings of the IEEE International Conference on Image Processing* (2002). (With D. Chandler)

TURING AWARD

John Hopcroft was a recipient of the ACM Turing Award in 1993.

John Hopcroft

Professor
jeh17@cornell.edu
http://www.cs.cornell.edu/jeh

Professor Hopcroft's research centers on the study of information capture and access. This includes the study of large graphs, spectral analysis of structures, clustering, and queries. He has also been involved in the theoretical aspects of computing, especially analysis of algorithms, formal languages, automata theory, and graph algorithms. He has coauthored four books on formal languages and algorithms with Jeffrey D. Ullman and Alfred V. Aho.

From January 1994 until June 2001, he was the Joseph Silbert Dean of the College of Engineering. He 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, Professor Hopcroft spent three years on the faculty of Princeton University. In 1967, he joined the Cornell faculty, was named professor in 1972 and served as CS chairman from 1987 to 1992. An undergraduate alumnus of Seattle University, Hopcroft was honored with their Doctor of Humanities degree, *Honoris Causa,* in 1990.

SELECTED PUBLICATIONS

Formal Languages and Their Relation to Automata. Addison–Wesley (1969). (With J. Ullman)

"Efficient Planarity Testing". *Journal of the ACM,* 21(4) (October, 1974): 549–568. (With R. Tarjan)

"A Paradigm for Robust Geometric Algorithms". *Algorithmica* 7(4) (1992): 339–380. (With P. Kahn)

Daniel P. Huttenlocher

John P. and Rilla Neafsey Professor of Computing and Information Science and Business
Cornell Weiss Presidential Fellow
CIS, joint with the Johnson Graduate School of Management
dph@cs.cornell.edu
http://www.cs.cornell.edu/~dph/

Dan Huttenlocher received a dual degree in computer science and experimental psychology from the University of Michigan in 1980, and master's and Ph.D. degrees in computer science from M.I.T. in 1984 and 1988, respectively. He has been on the CS faculty since 1988. He currently holds a joint appointment with the Johnson Graduate School of Management at Cornell.

Huttenlocher's research interests are in computer vision, computational geometry, electronic-collaboration tools, financial-trading systems, and the principles of software development. In addition to teaching and research, Dan has considerable experience managing software-development efforts in corporate and academic settings. He is chief technical officer of Intelligent Markets, a leading provider of advanced trading systems. He also spent more than ten years at the Xerox PARC, directing work that led to the ISO JBIG2 image-compression standard, and serving as part of the senior management team.

Huttenlocher has been recognized on several occasions for his teaching and research, including being named a Presidential Young Investigator by the NSF in 1990, the New York State Professor of the Year by CASE in 1993, and a Stephen H. Weiss Fellow by Cornell in 1996. He holds twenty-two U.S. patents and has published more than fifty technical papers, primarily in the areas of computer vision and computational geometry.

SELECTED PUBLICATIONS

"Comparing Images using the Hausdorff Distance". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15(9) (1993): 850–863. (With G. Klanderman and W. Rucklidge)

"The Upper Envelope of Voronoi Surfaces and its Applications". *Discrete and Computational Geometry* 9(3) (1993): 267–291. (With K. Kedem and M. Sharir)

"Recognizing Solid Objects by Alignment with an Image". *International Journal of Computer Vision* 5(2) (1990): 195–212. (With S. Ullman)

Thorsten Joachims

Assistant Professor
tj@cs.cornell.edu
http://www.cs.cornell.edu/people/tj

Thorsten Joachims joined CS as an assistant professor in 2001. Earlier that year, he completed his dissertation, "The Maximum-margin Approach to Learning Text Classifiers: Methods, Theory, and Algorithms" at the Universität Dortmund, Germany, advised by Katharina Morik.

His research interests center on a synthesis of theory and system building in the field of machine learning, with a focus on support-vector machines, text-mining, and machine learning in information access. In particular, Joachims has worked on WebWatcher, an adaptive-browsing assistant for the Web. He has authored the SVM-Light algorithm and software for support-vector learning. His most recent work is on learning from clickthrough data in search engines, and on using unlabeled data for supervised learning in the framework of transduction.

Joachims taught the course "Advanced Topics in Machine Learning" and cotaught "Language Technologies" with Claire Cardie.

Joachims received the dissertation award of the Universität Dortmund. He is member of the editorial board of the *Journal of Machine Learning Research* and *the Journal of Artifical Intelligence Research*. He serves on the program committees of International Conference on Machine Learning (ICML), European Conference on Machine Learning (ECML), Special Interest Group on Information Retrieval (SIGIR), and others.

SELECTED PUBLICATIONS

Learning to Classify Text Using Support Vector Machines, Kluwer (2002).

"Optimizing Search Engines Using Clickthrough Data". *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)* (2002).

"Making Large-scale Support Vector Machine Learning Practical" In *Advances in Kernel Methods Support Vector Learning*, B. Schölkopf, C. Burges, and A. Smola (eds.), M.I.T. Press (1999): 169–184.

Klara Kedem

Professor
kedem@cs.cornell.edu
http://www.cs.cornell.edu/kedem/

Klara Kedem obtained her Ph.D. in computer science at Tel-Aviv University in 1989. She is currently spending the summers as a CS professor at Cornell and during the rest of the year she is a professor in the computer science department at Ben-Gurion University in Israel.

Professor Kedem's research is in computational geometry with applications to robotics, computer vision, and bio-information. She is known for devising the minimum Hausdorff distance for shape matching, a robust method that has had a strong impact and is still being investigated actively.

Recently Professor Kedem and CS collaborators have looked into shape-comparison problems in the life sciences. In computational molecular biology, they have come up with a new metric, the unit-vector root mean square (URMS), to measure substructure resemblance between proteins. This measure has been further applied to the analysis of molecular dynamics. Currently she is working on finding consensus shapes for protein families, and applying string-matching algorithms to protein-shape comparison.

Kedem is on the editorial board of the *Journal of the Pattern Recognition Society*, and serves as guest editor of *Computational Geometry: Theory and Applications.* She won the Mary Upson Visiting Professorship at Cornell for 1997–1998.

SELECTED PUBLICATIONS

"The Upper Envelope of Voronoi Surfaces and its Applications". *Discrete and Computational Geometry* 9 (1993): 267–291. (With D. Huttenlocher and M. Sharir)

"Fast Detection of Geometric Substructure in Proteins". *Journal of Computational Biology* 6(3–4) (1999): 313–325. (With L. Chew, D. Huttenlocher, and J. Kleinberg)

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

Jon Kleinberg

Associate Professor
kleinber@cs.cornell.edu
http://www.cs.cornell.edu/home/kleinber/

Jon Kleinberg received his A.B. in computer science and mathematics from Cornell in 1993 and his Ph.D. in computer science from M.I.T. in 1996. He subsequently spent a year as a visiting scientist at the IBM Almaden Research Center, and is now an associate professor in CS at Cornell.

Kleinberg's research interests are centered around algorithms, particularly those concerned with the structure of networks and information. He focuses on combinatorial and randomized methods in the design of algorithms, with applications to information science, discrete optimization, data mining, and computational biology. His work introduced the notion of network analysis based on hubs and authorities, a framework that has been incorporated into a number of prominent search tools on the Web.

Kleinberg is a recipient of an NSF CAREER Award, an ONR Young Investigator Award, an Alfred P. Sloan Foundation Fellowship, a David and Lucile Packard Foundation Fellowship, and the 2001 National Academy of Sciences Award for Initiatives in Research. He also received the Fiona Ip Li and Donald Li Teaching Award from the Cornell College of Engineering, and the Cornell Association of Computer Science Undergraduates Faculty of the Year Award for 2001–2002.

SELECTED PUBLICATIONS

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

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

"Approximation Algorithms for Classification Problems with Pairwise Relationships: Metric Labeling and Markov Random Fields". *Journal of the ACM (2002)* 49(5): 616–639. (With E. Tardos)

Dexter Kozen

Joseph Newton Pew, Jr. Professor of Engineering
kozen@cs.cornell.edu
http://www.cs.cornell.edu/kozen/

Dexter Kozen received his undergraduate degree in mathematics from Dartmouth College in 1974 and his Ph.D. in computer science from Cornell in 1977. After working as a research staff member at the IBM Thomas J. Watson Research Center for several years, he returned to Ithaca to join the Cornell faculty in 1985.

Kozen's research interests include the design and analysis of algorithms, computation-complexity theory, the complexity of decision problems in logic and algebra, and logics and semantics of programming languages. He is currently involved in a research project involving efficient code certification and its application to malicious firmware. His most recent theoretical project is the development of the theory of Kleene algebra and Kleene algebra with tests, including results on complexity, deductive completeness, expressiveness, and applications to compiler correctness. He developed and taught a new course on this topic in spring 2002. Kozen is the author of three books.

Professor Kozen received the Stephen and Margery Russell Distinguished Teaching Award from Cornell's College of Arts and Sciences in 2001 and was named the Williams College Class of 1960 Scholar in 2000. He is also a recipient of an IBM Outstanding Innovation Award and a former Guggenheim fellow.

SELECTED PUBLICATIONS

"Results on the Propositional mu-calculus". *Theoretical Computer Science* 27 (1983): 333–354.

"Kleene Algebra with Tests". *Transactions on Programming Languages and Systems,* (May, 1997): 427–443.

Dynamic Logic. M.I.T. Press (2000). (With D. Harel and J. Tiuryn)

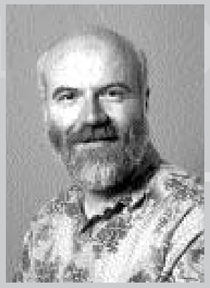

Dean Krafft

Senior Research Associate
Director of Computing Facilities
dean@cs.cornell.edu
http://www.cs.cornell.edu/dean

Dean Krafft received his Ph.D. in computer science from Cornell in 1981. He serves as both a CS researcher and administrator at Cornell. As an administrator, he manages the Computing Facilities Support group and worries about a number of issues including computer security, networking, and building Web services.

On the research side, Krafft is part of the CIS Program in Digital Libraries (*http://www.cis.cornell.edu/infoscience/research/dl/home.html*). As part of that effort, Krafft is a coprincipal investigator on the NSF–funded National Science Digital Library Project at Cornell (*http://www.nsdl.org*). He is currently working specifically on both the NSDL user portal and the central metadata repository. Krafft's 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.

SELECTED PUBLICATIONS

"Core Services in the Architecture of the National Digital Library for Science Education (NSDL)". In *Proceedings of the Second ACM/IEEE–CS Joint Conference on Digital Libraries* (July, 2002). (With Carl Lagoze, et al.) *http://arxiv.org/abs/cs.DL/0201025*

"Dienst: Building a Production Technical Report Server". In *Advances in Digital Libraries* (May, 1995): 211–223. (With J. Davis and C. Lagoze)

"The Challenge of Robotics for Computer Science". *Advances in Robotics, Algorithmic and Geometric Aspects of Robotics* 1 (1986): 7–42. J. Schwartz and C. Yap, eds. Lawrence Erlbaum Associates, Inc. (With J. Hopcroft)

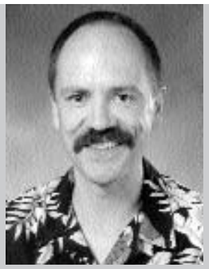

Christoph Kreitz

Senior Research Associate
kreitz@cs.cornell.edu
http://www.cs.cornell.edu/home/kreitz

Christoph Kreitz obtained his Ph.D. in computer science at the FernUniversität Hagen, Germany in 1984. His research has focused on computational models for infinite objects and on the application of automated theorem-proving to the design, verification, and optimization of software systems.

In collaboration with researchers of Robert Constable's Nuprl and Ken Birman's Ensemble groups he has built logic-based tools that automatically improve the code of fault-tolerant communication systems and guarantee that the improvements do not introduce errors. He has also developed techniques for the formal design and verification of adaptive distributed systems. He currently investigates the validation of end-to-end Quality-of-Service behavior of networked systems.

Christoph Kreitz also works on enhancing the automatic reasoning capabilities of tactical theorem provers. Together with his former students from the Technical University of Darmstadt, he has developed and implemented proof-search procedures for classical, intuitionistic, modal, and fragments of linear logic, and algorithms that transform the machine-found proofs into the proof calculus of other systems. His theorem prover has been connected to the interactive proof assistants Nuprl, MetaPRL, and Coq, and is being used to guide the development of proofs in these systems.

SELECTED PUBLICATIONS

"Theory of Representations". *Theoretical Computer Science* 38 (1985): 35–53. (With K. Weihrauch)

"Building Reliable, High-performance Systems from Components". In *Proceedings of the Seventeenth ACM Symposium on Operating System Principles (SOSP '99), Operating Systems Review* 34(5) (December 1999): 80–92. (With X. Liu, R. van Renesse, J. Hickey, M. Hayden, K. Birman, and R. Constable)

"Connection-based Theorem Proving in Classical and Non-classical Logics". *Journal for Universal Computer Science* 5(3) (1999): 88–112. (With J. Otten)

Carl Lagoze

Senior Research Associate
lagoze@cs.cornell.edu
http://www.cs.cornell.edu/lagoze/lagoze.html

Carl Lagoze obtained his master's degree in software engineering from Wang Institute for Graduate Studies in 1987. He is currently a senior research associate in CIS. He is concurrently Director of Technology of the NSF–funded National Science Digital Library (NSDL).

Lagoze's research investigates policies, organization, and architecture of distributed-information spaces. The Web provides the backdrop for the work. The goal is to understand the services and organization that can be built on top of this global information base to increase its functionality, integrity, and ease-of-use. The research is undertaken 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.

Lagoze's research group is recognized for a number of advances in distributed-information systems. These include the Dienst architecture for distributed digital libraries, the FEDORA digital-object model for complex digital content, the ABC event-aware descriptive ontology, and the Open Archives Initiative Protocol for Metadata Harvesting that has been widely adopted as a foundation for information-systems interoperability. His technical leadership of the NSDL provides the opportunity to realize these advances in a major national resource for science and mathematics education.

SELECTED PUBLICATIONS

"Core Services in the Architecture of the National Digital Library for Science Education (NSDL)". *Joint Conference on Digital Libraries* (July, 2002). (With W. Arms, S. Gan, D. Hillmann, C. Ingram, D.Krafft, R. Marisa, J. Phipps, J. Saylor, C. Terrizzi, W. Hoehn, D. Millman, J. Allan, S. Guzman-Lara, and T. Kalt) *http://arxiv.org/abs/cs.DL/0201025*.

"The Open Archives Initiative Protocol for Metadata Harvesting". Protocol Version 2 (June, 2002). (With H. Van de Sompel, M. Nelson, and S. Warner) *http://www.openarchives.org/OAI/2.0/openarchivesprotocol.htm*

"The ABC Ontology and Model". *Journal of Digital Information* 2(2) (November, 2001): 18 p. (With J. Hunter)

Lillian Lee

Associate Professor
llee@cs.cornell.edu
http://www.cs.cornell.edu/home/llee/

Lillian Lee (A.B., Cornell, 1993; Ph.D., Harvard University, 1997) is an assistant professor in CS. Her main research interest is natural language processing, in particular the development of "knowledge-lean" statistical methods that allow computers to automatically learn linguistic and domain knowledge directly from text. A major focus has been the study of distributional similarity and distributional clustering. She and her colleagues have also considered applications ranging from finding word boundaries in streams of Japanese to creating English versions of computer-generated mathematical proofs.

Lee is a recipient of an Alfred P. Sloan Research Fellowship, the Stephen and Marilyn Miles Excellence in Teaching Award, and the James and Mary Tien Excellence in Teaching Award. Her professional activities include serving as the program chair of the 2001 Conference on Empirical Methods in Natural Language Processing, an area chair for the 2001 Annual Meeting of the Association for Computational Linguistics, a member of the editorial boards of the journals *Computational Linguistics* and *Machine Learning*, and a member of the executive board of SIGDAT, the ACL Special Interest Group on Linguistic Data and Corpus-based Approaches.

SELECTED PUBLICATIONS

"Fast Context-free Grammar Parsing Requires Fast Boolean Matrix Multiplication". *Journal of the ACM* 49(1) (2002): 1–15.

"Iterative Residual Rescaling: An Analysis and Generalization of LSI". In *Proceedings of the Twenty-fourth Annual International Conference on Research and Development in Information Retrieval (SIGIR)* (2001): 154–162. (With R. Ando)

"Distributional Clustering of English Words". In *Proceedings of the Thirty-first Annual Meeting of the ACL* (1993): 183–190. (With F. Pereira and N. Tishby)

Yuying Li

Senior Research Associate
yuying@cs.cornell.edu
http://www.cs.cornell.edu/home/yuying/yuying.html

Yuying Li obtained a bachelor's degree in applied mathematics at the Sichuan University in China in 1982, and an Ph.D in computer science at the University of Waterloo in 1988. She has been a research associate in CS since.

Li's research interests include numerical optimization and scientific computation. In addition, she is interested in the application of optimization methods to medical, engineering, and financial problems.

Her current interest has focused on solving problems in financial applications, e.g., volatility estimation, discrete hedging, portfolio compression, and portfolio optimization under different risk measures.

Li is the recipient of the 1993 First Prize of the Sixth Fox Prize Competition in Numerical Analysis, Oxford, England.

SELECTED PUBLICATIONS

"Dynamic Hedging with a Deterministic Local Volatility Function Model". *The Journal of Risk* 4(1) (2001): 64–90.

"An Interior, Trust Region Approach for Nonlinear Minimization Subject to Bounds". *SIAM Journal on Optimization* 6(2) (1996): 418–445. (With T. Coleman)

Large-scale Numerical Optimization. SIAM, Philadelphia (1990). (With T. Coleman)

Hod Lipson

Assistant Professor
CIS, joint with Mechanical and Aerospace Engineering
hl274@cornell.edu
http://www.mae.cornell.edu/lipson

Hod Lipson joined the Faculty of CIS and the faculty of the Sibley School of Mechanical and Aerospace Engineering in 2001 as an assistant professor. Prior to this appointment, he was a postdoctoral researcher at Brandeis University's computer science department, working on evolutionary computation and evolutionary robotics, where he led the Golem Project—creating the first physical artificial life forms. He was also a lecturer in M.I.T.'s mechanical-engineering department, where he taught design and conducted research in design automation.

Professor Lipson's Ph.D. (Technion, 1998) research was on the reconstruction of a three-dimensional object from a single freehand sketch, as a means for human–computer interaction for CAD. Before joining academia, Lipson spent several years as a design engineer in the mechanical, electronic, and software industries, and cofounded two currently-active companies.

Lipson's research interests are in the area of computational synthesis: How do we combine basic building blocks to achieve some high-level functionality? He is interested in understanding the synthesis process of design and emulating it computationally, and he focuses on the ideas of self-organization and self-replication as new paradigms of fully automated design, fabrication, and learning. Primary questions concern automatic discovery of modules, regular and hierarchical composition, and automatic abstraction of functionality. Cornell's computational-synthesis group develops both new theoretical ideas and applies them to various engineering problems, from evolutionary robotics and structures to circuits and game players. Lipson believes that fully automated synthesis holds the key to future competitiveness, and presents a largely unaddressed challenge across engineering, biology, and AI.

Among the awards and honors Lipson has received are: *Time* Magazine's Annual 2001; EXPO 2000 ("Shaping the Future"); and the CIRP International F. W. Taylor Medal in 1997.

SELECTED PUBLICATIONS

"Correlation-based Reconstruction of a 3D Object from a Single Freehand Sketch". *AAAI 2002 Spring Symposium on Sketch Understanding* (2002): 99–104. (With M. Shpitalni)

"Uncontrolled Engineering: A Review of Evolutionary Robotics". *Artificial Life* 7(4) (2001): 419–424.

"Automatic Design and Manufacture of Robotic Lifeforms". *Nature* 406 (2000): 974–978. (With J. Pollack)

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/

Rajit Manohar obtained his B.S. (1994), M.S. (1995) and Ph.D. (1998) degrees in computer science from the California Institute of Technology. He is currently an assistant professor in ECE, and a member of the graduate fields of computer science and applied mathematics. The focus of his research effort is on the design of efficient computation structures. His group is currently working on the following:

The SNAP project develops novel energy-efficient clockless architectures for sensor network applications. The NoC project (joint with Professor Lang Tong) has demonstrated that asynchronous circuit techniques applied to wireless network modeling can provide three-orders-of-magnitude increase in simulation speed, compared to traditional approaches.

Professor Manohar received the NSF CAREER Award (2000–2004), the Cornell IEEE Teacher of the Year Award (2001), the College of Engineering Sonny Yau Excellence in Teaching Award (2001), and the Tau Beta Pi and Cornell Society of Engineers Excellence in Teaching Award (2000).

SELECTED PUBLICATIONS

"SNAP: A Sensor-Network Asynchronous Processor". In *Proceedings of the Nonth International Symposium on Asynchronous Curcuits and Systems* (May 2003). (With C. Kelly and V. Ekanayake)

"Network on a Chip: Modeling Wireless Networks with Asynchronous VLSI". *IEEE Communications Magazine* 39(11) (November, 2001). (With Clinton Kelly IV)

"Slack Elasticity in Concurrent Computing". In *Proceedings of the Fourth International Conference on the Mathematics of Program Construction, Lecture Notes in Computer Science* 1422 (June, 1998): 272–285. Springer-Verlag. (With A. Martin)

Steve Marschner

Assistant Professor, CS
srm@cs.cornell.edu
http://www.cs.cornell.edu~srm/

Steve Marschner obtained his B.S. degree in mathematics and computer science from Brown University in 1993 and his Ph.D. from Cornell in 1998. He held research positions at Hewlett–Packard Labs, Microsoft Research, and Stanford University before joining the CS faculty in 2002.

Marschner's research interests are in the field of computer graphics, focusing on realistic rendering, especially models for light reflection and scattering, and high-resolution geometric modeling. Recent projects include a new model to efficiently simulate translucent materials, which has been widely implemented by the film-effects industry; ongoing work on processing very high resolution geometric data for the Digital Michelangelo Project at Stanford; and an experimental and theoretical investigation into the scattering of light from human hair. The overall goal of his work is to use measurements to capture the complexity of real objects and understand the subtleties of real materials, thereby increasing the richness and realism of computer generated images .

SELECTED PUBLICATIONS

"Light Scattering from Human Hair Fibers". ACM Transactions on Graphics 22: 3. *Proceedings of SIGGRAPH 2003 (2003)*. (With H. Jensen, M. Cammarano, and P. Hanrahan)

"A Practical Model for Subsurface Light Transport". In *Proceedings of SIGGRAPH 2001* (August, 2001). (With H. Jensen, M. Levoy, and P. Hanrahan)

"Image-based BRDF Measurement Including Human Skin". In *Proceedings of 10th Eurographics Workshop on Rendering* (June, 1999): 139–152. (With S. Westin, E. Lafortune, K. Torrance, and D. Greenberg)

Jeanna Neefe Matthews

Assistant Professor
jnm@cs.cornell.edu
http://www.cs.cornell.edu/info/people/jnm/

Jeanna Matthews obtained her Ph.D. in computer science at the University of California at Berkeley in 2000. She is currently an assistant professor in CS. Matthews' research lies in the areas of operating systems, storage systems, and networks. She is actively involved in several projects aimed at integrating hands-on exposure to research results into computer science courses.

SELECTED PUBLICATIONS

"Improving the Performance of Log-structured File Systems with Adaptive Methods". In *Proceedings of the Sixteenth ACM Symposium on Operating System Principles* (October, 1997): 238–251. (With D. Roselli, A. Costello, R. Wang, and T. Anderson)

"Serverless Network File Systems". In *Proceedings of the ACM Symposium on Computer Systems* (February, 1996). (With T. Anderson, M. Dahlin, D. Patterson, D. Roselli, and R. Wang)

"Serverless Network File Systems". (Award paper) *Proceedings of the Fifteenth ACM Symposium on Operating System Principles* (December, 1995): 109–126. (With T. Anderson, M. Dahlin, D. Patterson, D. Roselli, and R. Wang)

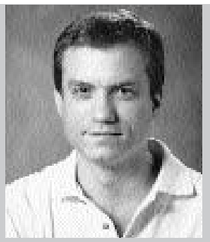

Greg Morrisett

Associate Professor
jgm@cs.cornell.edu
http://www.cs.cornell.edu/home/jgm

Greg Morrisett obtained his Ph.D. in computer science from Carnegie Mellon University in 1995. He is currently an associate professor in CS at Cornell.

Morrisett's research focuses on programming-language design, implementation, and semantics. He is particularly interested in the emerging area of language-based security. He is best known for the development of TAL and Certifying Compilation. These are important mechanisms that can be used to automatically verify an important class of safety properties for machine code. More recently, Morrisett has concentrated on type systems for legacy software. His Cyclone project provides type safety for C code without sacrificing control over data structures, calling conventions, or memory management. Other projects include work on run-time code specialization, type-safe reflection, type-based alias analysis, region-based memory management, and in-lined reference monitors.

Morrisett is an editor for the *Journal of Functional Programming,* and an associate editor for *ACM Transactions on Programming Languages and Systems.* In 2000, he was given a Presidential Early Career Award for Scientists and Engineers. He is also a recipient of a Sloan Foundation Fellowship, an NSF CAREER Award, and the Allen Newell Medal of Research Excellence.

SELECTED PUBLICATIONS

"Cyclone: A Safe Dialect of C". In *Usenix Annual Technical Conference* (June, 2002): 275–288. (With T. Jim, D. Grossman, M. Hicks, J. Cheney, and Y. Wang)

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

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

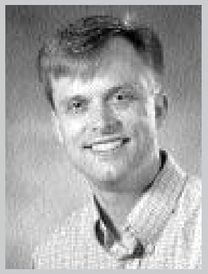

Andrew Myers

Assistant Professor
andru@cs.cornell.edu
http://www.cs.cornell.edu/andru

Andrew Myers received his Ph.D. in computer science from M.I.T. in 1999. He is currently an assistant professor in CS. Myers is particularly interested in using language-level information to improve security guarantees, performance, and transparency for distributed systems and mobile code.

A current focus is on the protection of confidential data, a problem that is gaining importance in our connected world. Methods are needed for building practical systems while guaranteeing that they enforce strong security properties. Myers has developed novel and efficient static-analysis techniques to identify and control privacy violations in complex programs. These techniques have been employed in the Jif compiler and run-time system for writing secure programs. Jif has been applied to distributed systems containing untrusted components, and to systems in which security requirements change dynamically.

Myers received a NSF CAREER award in 2001, and the Alfred P. Sloan Research Fellowship and the Excellence in Teaching Award from the College of Engineering in 2002.

SELECTED PUBLICATIONS

"Using Replication and Partitioning to Build Secure Distributed Systems". In *IEEE Symposium on Security and Privacy* (May, 2003). (With L. Zheng, S. Chong, and S. Zdancewic)

"Language-based Information-flow Security". In *IEEE Journal on Selected Areas in Communications*, special issue on Formal Methods for Security (January, 2003) 21(12): 5–19. (With A. Sabelfeld)

"Secure Program Partitioning". *ACM Transactions on Computing Systems (TOCS)* 20(3)(August, 2002): 283–328. (With S. Zdancewic, L. Zheng, and N. Nystrom)

Anil Nerode

Goldwin Smith Professor of Mathematics
Member of the Graduate Field of Computer Science
anil@math.cornell.edu
http://www.math.cornell.edu/~anil/

Anil Nerode obtained his Ph.D. in mathematics, under Saunders MacLane, from the University of Chicago in 1956. He was a NSF postdoctoral fellow with Kurt Godel, at the Institute for Advanced Study, from 1957–58; visiting assistant professor with Alfred Tarski at the University of California at Berkeley from 1958–59; was brought to Cornell by J. Barkley Rosser in 1959; appointed professor in 1965; and named Goldwin Smith Professor in 1990. He served as chair of the Department of Mathematics from 1982–87, and was director of the Mathematical Sciences Institute from 1987–1996. He also served as director of the Center for Foundations of Intelligent Systems from 1996–2001.

Nerode's research areas include mathematical logic, computability theory, recursive mathematics, nonstandard logics, nonmonotonic logics, AI, applied mathematics, control theory, hybrid systems, and complex system design.

SELECTED PUBLICATIONS

"Control Synthesis in Hybrid Systems with Finsler Dynamics" *Houston Journal of Mathematics 28(2)* (2002): 352–375. (With Kohn Wolf and Vladimir Brayman)

[in preparation] Constructive Logics and Lambda Calculi: 500 p. (With G. Odifreddi)

Automata Theory and Its Applications: Birkhauser (2001). 480 p. (With B. Khoussainov)

Keshav Pingali

Professor
pingali@cs.cornell.edu
http://www.iss.cs.cornell.edu/

Keshav Pingali obtained a bachelor's degree in electrical engineering at the Indian Institute of Technology (I.I.T.), Kanpur in 1978, and an Sc.D. in computer science at M.I.T. in 1986. Since 1986, he has been on the CS faculty where he is currently a full professor. Pingali is also an ECE faculty member.

Pingali's research has focused on programming languages and compiler technology for program understanding, restructuring, and optimization. His group is known for its contributions to memory-hierarchy optimization; some of these have been patented. Algorithms and tools developed by his projects are used in many commercial products such as Intel's IA-64 compiler, SGI's MIPSPro compiler, and Hewlett–Packard's PA-RISC compiler. In his current research, he is investigating language-based fault-tolerance, and highly adaptive software systems for large-scale simulations.

Among other awards, Pingali has won the President's Gold Medal at I.I.T., Kanpur (1978), IBM Faculty Development award (1986–87), NSF Presidential Young Investigator award (1989–94), Ip-Lee teaching award of the College of Engineering at Cornell (1997), and the Russell teaching award of the College of Arts and Sciences at Cornell (1998). In 2000, he was a visiting professor at I.I.T., Kanpur where he held the Rama Rao Chaired Professorship.

SELECTED PUBLICATIONS

"Algorithms for Computing the Static Single Assignment Form". *Journal of the ACM* (2003): 375–425. (With G. Bilardi)

"A Comparison of Empirical and Model-driven Optimization". *ACM Programming Language Design and Implementation (PLDI)* (June, 2003): 52–66. (With K. Yotov et al.)

"Fractal Symbolic Analysis". *International Conference on Supercomputing* (June, 2001): 38–49. (With N. Mateev and V. Menon)

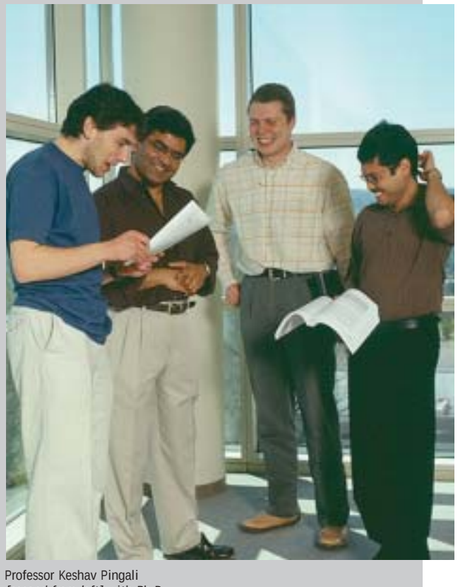

Professor Keshav Pingali *[second from left]* with Ph.D. students Greg Bronevetsky *[left]*, Kamen Yotov *[second from right]*, and Rohit Fernandes *[right]*.

Robbert van Renesse

Senior Research Associate
rvr@cs.cornell.edu
http://www.cs.cornell.edu/home/rvr/

Robbert van Renesse received his M.Sc. in mathematics and computer science from the Vrije Universiteit in 1985, under the supervision of Andrew S. Tanenbaum, with the honorary addendum cum laude. He obtained his Ph.D. in computer science from the Vrije Universiteit in 1989, also under the supervision of Professor Tanenbaum.

His research focus is in large-scale, self-organizing network protocols and distributed applications. Currently, he is involved in four projects. First, the Astrolabe system is a peer-to-peer implementation of a DNS–like directory service that supports on-the-fly aggregation of resource information. It incorporates epidemic algorithms to ensure robustness and efficiency, and is used among others to build scalable multicast protocols. This is joint work with Ken Birman and Werner Vogels. Second, with Lidong Zhou, he is developing an implementation of IPv6 as an overlay network, using distributed hash tables and Astrolabe. Third, the MediaNet system is a distributed multimedia facility that supports user-specified adaptation. This is joint work with Mike Hicks, Bob Constable, Mark Bickford, and Christoph Kreitz. Fourth, with Fred Schneider and Visiting Professor Dag Johansen, he is investigating techniques for filtering high volume–event streams.

In addition to his current research, van Renesse is a technical advisor for Fast Search and Transfer, ASA, a company that develops search engines.

SELECTED PUBLICATIONS

"Astrolabe: A Robust and Scalable Technology for Distributed Systems Monitoring, Management, and Data Mining". In *ACM Transactions on Computer Systems* (May, 2003) 21(3). (With K. Birman and W. Vogels)

"User-specified Adaptive Scheduling in a Streaming Media Network". *Proceedings of Open ARCH 03* (April, 2003). (With M. Hicks and A. Nagarjan)

"COCA: A Secure Distributed Online Certification Authority ". In *ACM Transactions on Computer Systems* (November, 2002). (With L. Zhou and F. Schneider)

Mats Rooth

Professor
CIS, joint with Linguistics
rooth@cornell.edu
http://www.people.cornell.edu/pages/mr249/

Mats Rooth's research is concerned with theories and applications in linguistics and computational linguistics, which combine theoretical–linguistic formalisms, knowledge, and problem statements with numerical modeling and parameter-estimation techniques. Using current methodology, it is possible to create approximately complete grammars of human languages, and using parsing algorithms and the grammars, to map sentences to representations that represent their syntax and meaning. However, sentences of human languages are very ambiguous, to the extent that it would be possible know everything about the syntax of a language, without having any operative means of identifying the intended syntax and meaning of the sentences that people use. This problem is addressed by numerical models that put weights on possible representations. Numerical models and optimization algorithms also allow linguistic information (in particular, syntactic and semantic properties of individual words) to be learned from large data samples.

Rooth, whose appointment is joint with the Department of Linguistics and the Faculty of CIS, also works on the semantics of natural language, using logical methods and formalisms. He developed an approach to the meaning of intonation, which is known as alternative semantics. Currently, he is working on interactions between the grammar of ellipsis and the grammar of intonation.

Rooth has a B.S. degree in mathematics from M.I.T., and a Ph.D. in linguistics from the University of Massachusetts at Amherst. Before joining the Cornell faculty, he was chair of theoretical computational linguistics at the University of Stuttgart, and member of the technical staff at AT&T Bell Laboratories.

SELECTED PUBLICATIONS

"A Theory of Focus Interpretation". *Natural Language Semantics* 1 (1992): 75–116.

"Inside–outside Estimation of a Lexicalized PCFG for German". In *Proceedings of the Thirty-seventh Annual Meeting of the Association for Computational Linguistics* (1999). (With F. Beil, G. Carroll, D. Prescher, and S. Riezler)

"Parse Forest Computation of Expected Governors". In *Proceedings of the Thirty-ninth Annual Meeting of the Association for Computational Linguistics* (2001): 458–465. (With H. Schmid)

Radu Rugina

Assistant Professor
rugina@cs.cornell.edu
http://www.cs.cornell.edu/people/rugina/

Radu Rugina received a bachelor's degree in computer science from the University Politechnica of Bucharest in 1996, and a Ph.D. degree in computer science from the University of California at Santa Barbara in 2001. Between 1997 and 2001 he was a visiting scholar at the Laboratory for Computer Science at M.I.T.

His research interests lie in the area of programming languages and compiler support for program understanding, maintaining, and debugging; for checking safety properties of programs; and for program transformations and optimizations. Rugina has developed program analysis techniques capable of analyzing memory accesses in recursive and multithreaded programs that heavily manipulate pointers. Concrete applications of these analyses include automatic parallelization of sophisticated divide-and-conquer problems, static detection of array-bounds violations, and data-race detection in multithreaded programs that use pointers and pointer arithmetic.

In his current research, he is investigating program analysis approaches to improve software reliability and security, by automating the process of checking the properties that guarantee program safety and functionality.

SELECTED PUBLICATIONS

"Static Analysis of Accessed Regions in Recursive Data Structures". In *Proceedings of the Tenth International Static Analysis Symposium* (June, 2003). (With S. Chong)

"Pointer Analysis for Structured Parallel Programs". In *ACM Transactions on Programming Languages and Systems* 25(1) (January, 2003). (With M. Rinard)

"Symbolic Bounds Analysis of Pointers, Array Indices, and Accessed Memory Regions". In *Proceedings of the ACM SIGPLAN 2000 Conference on Programming Languages Design and Implementation* (June, 2000). (With M. Rinard)

HONORARY DOCTORATE

Fred Scheider was awarded an honorary doctorate degree by the University of Newcastle-upon-Tyne in 2003.

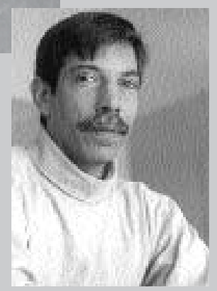

Fred B. Schneider

Professor
Director, Information Assurance Institute
Chief Scientist, Griffiss Institute
fbs@cs.cornell.edu
http://www.cs.cornell.edu/fbs/

Fred B. Schneider has studied concurrent and distributed systems since joining Cornell's faculty in 1978. His early work concerned programming methodology and formal methods. He is known for formalizing "safety" and "liveness" properties as well as for developing methods to reason about concurrent and distributed programs. His work in fault-tolerant distributed systems led to now well-known protocols and structures (including the "failstop processor" abstraction, a seminal survey on the state machine approach, hypervisor-based fault tolerance, and various protocols used in today's air-traffic–control systems).

Most recently, Schneider's attention has turned to questions related to computer security:

- exploiting insights from formal methods and programming languages as a basis for relocating trust and enforcing application-specific security policies; and
- the design of systems and protocols to support both fault-tolerance and security in distributed systems.

Both of these efforts have led to practical new tools. For example, Schneider and collaborators are currently building a third-generation inlined reference monitor suite (targeted to Microsoft's CLR) to better understand practical problems with enforcing fine-grained security policies through rewriting object-code. Work also continues with collaborators on the COCA (Cornell On-line Certification Authority) project, with attention now focused on implementing secure and scalable publish/subscribe protocols.

Schneider has been professor-at-large at the University of Trömso (Norway) since 1996, is a fellow of ACM and AAAS, and received an honorary doctorate in May 2003 from the University of Newcastle-upon-Tyne. He is associate editor-in-chief for IEEE Security and Privacy and serves on the editorial boards of several other journals.

SELECTED PUBLICATIONS

On Concurrent Programming. Springer-Verlag, N.Y. (1997): 473 p.

Trust in Cyberspace. (ed.) National Academy Press (December, 1998): 331 p.

"Enforceable Security Policies". *ACM Transactions on Information and System Security* 3(1) (February, 2000): 30–50.

David I. Schwartz

Assistant Professor
dis@cs.cornell.edu
http://www.cs.cornell.edu/dis/

David I. Schwartz obtained his Ph.D. in civil engineering at the State University of New York at Buffalo in 1999. He is currently an assistant professor in CS.

Schwartz's research and interests involve educational technology, the support of undergraduate research, textbook writing, and graduate-student development. He continues to work on developing a multidisciplinarian curriculum for computer game–design courses that incorporates technical and artistic aesthetics of computer-game design in a collaborative environment. Students from diverse backgrounds in engineering, computer science, fine art, and music formed teams that developed and implemented computer games and associated game-development tools. The project encourages women and underrepresented minorities to enter the field of computer science.

SELECTED PUBLICATIONS

Introduction to Mapl.e (2nd ed.) Prentice Hall: New Jersey (2003).

Introduction to UNIX. Prentice Hall: New Jersey (1999).

"A Constraint-based Approach for Qualitative Matrix Structural Analysis". *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* 9 (1995): 23–36. (With S. Chen)

AAAS FELLOWSHIP

Bart Selman was the recipient of an AAAS Fellowship in 2003.

Bart Selman

Associate Professor
selman@cs.cornell.edu
http://www.cs.cornell.edu/home/selman/

Bart Selman obtained a Ph.D. in computer science from the University of Toronto in 1991. Currently an associate professor in CS, he spent the previous six years at AT&T Bell Laboratories in the principles of artificial intelligence–research department. His research has covered many areas in artificial intelligence and computer science, including tractable inference, knowledge representation, stochastic search methods, theory approximation, knowledge compilation, planning, default reasoning, and the connections between computer science and statistical physics (phase-transition phenomena). His current projects focus on planning, multi-agent systems, and the integration of learning and reasoning techniques.

Bart Selman has received an NSF CAREER Award (1998–2002) and an Alfred P. Sloan Research Fellowship (1999–2001). He has received four best paper awards at the American and Canadian national artificial-intelligence conferences, and at the International Conference on Knowledge Representation.

SELECTED PUBLICATIONS

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

"Determining Computational Complexity from Characteristic Phase Transitions". *Nature* 400(8) (1999): 133–137. (With R. Monasson, R. Zecchina, S. Kirkpatrick, and L. Troyansky)

"Knowledge Compilation and Theory Approximation". *Journal of the ACM* 43(2) (1996): 193–224. (With H. Kautz)

NSF CAREER AWARD

Phoebe Sengers was a recipient of a 2003 NSF CAREER Award.

Phoebe Sengers

Assistant Professor
CIS, joint with Science and Technology Studies
sengers@cs.cornell.edu
http://www.cs.cornell.edu/people/sengers/

Phoebe Sengers received her Ph.D. in artificial intelligence and cultural theory in 1998 from Carnegie Mellon University. She was a Fulbright Scholar at the Center for Art and Media Technology (ZKM) in Karlsruhe, Germany, and spent two years as a research scientist at the German National Research Center for Information Technology (GMD). She joined the Faculty of CIS in October, 2001, and has a joint appointment with the Department of Science and Technology Studies.

Sengers works in human–computer interaction, especially problems that bridge cultural issues and technology design. She develops culturally embedded systems; i.e., new kinds of interactive technology that respond to and encourage critical reflection on the place of technology in culture. Her current research, funded by a five-year NSF CAREER Award, explores everyday computing, or interactive media devices for non-work contexts, and draws on techniques from computer science, cultural analysis, design, and the arts. She uses insights from analysis of consumer culture to rethink the work-based assumptions underlying technologies for the home, developing both new application areas for everyday computing, including systems to support personal reflection, and new techniques for designing systems, including the use of self-experiemnt in design and new forms of evaluation for open-ended systems. She works on the National Research Council's Committee on Information Technology and Creativity, which develops policy suggestions for interdisciplinary research in information technology and the arts, humanities, and other creative areas.

SELECTED PUBLICATIONS

"The Enigmatics of Affect". *Conference on Designing Interactive Systems* (June, 2002). (With R. Liesendahl, W. Magar, C. Seibert, B. Müller, T. Joachims, W. Geng, P. Mårtensson, and K. Höök)

"Practices for Machine Culture: A Case Study of Integrating Artificial Intelligence and Cultural Theory". *Surfaces* 8 (1999).

"Designing Comprehensible Agents". In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence* (August, 1999): 1227–1232.

CIS Professor Phoebe Sengers discusses switches with CIS Professor Paul Ginsparg.

Jayavel Shanmugasundaram

Assistant Professor
jai@cs.cornell.edu
http://www.cs.cornell.edu/people/jai

Jayavel Shanmugasundaram obtained his Ph.D. degree in computer science from the University of Wisconsin at Madison in 2001. He is currently an assistant professor in CS.

Shanmugasundaram's research interests include Internet data management, information retrieval, and query processing in emerging system architectures. His research group is currently working on three projects. The Quark project aims to unify the database and information retrieval worlds by developing a next-generation data-management system for handling both structured and unstructured data. The Deep Glue project develops a platform for integrating and querying Internet-attached databases, also referred to as the "deep web". The Pepper project (joint with Johannes Gehrke) develops highly robust indexing and query-processing strategies for evaluating complex queries over large-scale, distributed peer-to-peer systems.

Shanmugasundaram's research ideas have been incorporated in commercial data-management products, and have resulted in several patents.

Shanmugasundaram is a recipient of an IBM Faculty Partnership Award.

SELECTED PUBLICATIONS

"XRANK: Ranked Keyword Search Over XML Documents". In *ACM SIGMOD Conference on Management of Data* (2003). (With L. Guo, F. Shao, and C. Botev)

"Querying XML Views of Relational Data". In *Proceedings of the Conference on Very Large Data Bases* (2001): 261–270. (With J. Kiernan, E. Shekita, C. Fan, and J. Funderburk)

"Relational Databases for Querying XML Documents: Limitations and Opportunities". In *Proceedings of the Conference on Very Large Data Bases* (1999): 302–314. (With K. Tufte, C. Zhang, G. He, D. DeWitt, and J. Naughton)

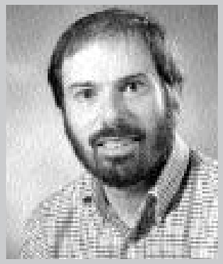

David B. Shmoys

Professor
Member of the School of Operations Research and Industrial Engineering and the Graduate Field of Computer Science
shmoys@cs.cornell.edu
http://www.cs.cornell.edu/home/shmoys/shmoys.html

David Shmoys obtained his Ph.D. in computer science at the University of California at Berkeley in 1984. He has faculty appointments in both CS and the School of Operations Research and Industrial Engineering. Shmoys' research has focused on the design and analysis of efficient algorithms for discrete optimization problems.

His work has highlighted the central role that linear programming plays in the design of approximation algorithms for NP–hard problems. In particular, he is known for his results on scheduling and clustering problems, including the first constant-performance guarantees for several problems central to the literature, including the k-center and k-median problems, the generalized-assignment problem, as well as scheduling problems in which the aim is to minimize the average job-completion time. Furthermore, his work on polynomial-time approximation schemes for scheduling problems introduced techniques that have subsequently been applied to a variety of other settings. His current work includes the application of discrete optimization techniques to several issues in computational biology.

Professor Shmoys is a Fellow of the ACM, and is the recipient of a National Science Foundation Presidential Young Investigator's Award and the Cornell College of Engineering Sonny Yau Excellence in Teaching Award (twice). He is currently editor-in-chief of *SIAM Journal on Discrete Mathematics,* and on the editorial board of several other journals, including *SIAM Journal on Computing,* and *Mathematics of Operations Research.*

SELECTED PUBLICATIONS

"Using Dual Approximation Algorithms for Scheduling Problems: Theoretical and Practical Results". *Journal of the ACM* 34(1) (1987): 144–162. (With D. Hochbaum)

"Fast Approximation Algorithms for Fractional Packing and Covering Problems". *Mathematics of Operations Research* 20: 257–301 (1995). (With S. Plotkin and E. Tardos).

"Scheduling to Minimize the Average Completion Time: On-line and Off-line Approximation Algorithms". *Mathematics of Operations Research* 22 (1997): 513–544. (With L. Hall, A. Schulz, and J. Wein)

Emin Gün Sirer

Assistant Professor
egs@cornell.edu
http://www.cs.cornell.edu/people/egs/

Emin Gün Sirer obtained his Ph.D. from the University of Washington in 2002. His research interests span operating systems, networking, distributed systems, and ubiquitous computing, with recent emphasis on ad hoc networks, peer-to-peer systems, and secure network services.

Professor Sirer is currently involved in four projects. The MagnetOS project investigates operating-system support for ad hoc and sensor networks. Specifically, Sirer's group is designing and building a new operating system that improves the longevity and reliability of ad hoc networks through energy-aware, adaptive object migration, multipath route selection, and novel hybrid routing protocols. The CliqueNet project is a peer-to-peer, self-organizing system for anonymous communication that provides an information-theoretic guarantee of privacy. The Pepper system, with Professors Gehrke and Shanmugasundaram, examines peer-to-peer data-location techniques for ubiquitous computers. Finally, the WebGuard project investigates how to build secure Web services.

Sirer's past work on the SPIN and Kimera projects examined novel operating-system architectures. The SPIN kernel developed the techniques for safely extending operating systems with application-specific code. The Kimera system introduced a new virtual-machine architecture that enables Java systems of drastically higher manageability, security, and performance, while reducing their resource requirements. The techniques developed in the Kimera project have been adopted throughout the industry, including Hewlett–Packard, Microsoft, Sun, and Schlumberger Inc.

SELECTED PUBLICATIONS

"SHARP: A Hybrid Adaptive Routing Protocol for Mobile Ad Hoc Networks". In *Proceedings of the Symposium on Mobile Ad Hoc Networking* (June, 2003). (With V. Ramasubramanian and Z. Haas)

"Path-set Selection in Mobile Ad Hoc Networks". In *Proceedings of the Third ACM International Symposium on Mobile Ad Hoc Networking and Computing* (2002). (With P. Papadimitratos and Z. Haas)

"On the Need for System-level Support for Ad hoc and Sensor Networks". *ACM Operating Systems Review* 36(2) (April, 2002): 1–5. (With R. Barr, J. Bicket, D. Dantas, B. Du, T. Kim, and B. Zhou)

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/

Evan Speight obtained his Ph.D. in electrical engineering at Rice University in 1998. Speight is an assistant professor in ECE, and is a field member of CS. His research interests include distributed and parallel computing, computer architecture, and affinity-directed mobility in mobile computing environments.

Speight's current projects include Active Memory Clusters (joint work with Professor Mark Heinrich), which seeks to leverage the increased functionality of a programmable memory controller to provide hardware-distributed shared-memory performance from commodity clusters. The Delphi project (joint work with Professor Martin Burtscher) explores the benefits of utilizing value-prediction techniques borrowed from the architectural community in improving the performance of cluster-based shared-memory multiprocessors. The Tern project represents work on examining the possible performance and fault-tolerant benefits of thread migration between hosts in an MPI parallel-runtime environment. Finally, the Bifrost project (joint work with Professor John Bennett at the University of Colorado at Boulder and sponsored by Microsoft) provides a framework for mobile computing that relies on "affinity" to automatically direct application and data throughout a wide geographic region for optimal user access.

SELECTED PUBLICATIONS

"Delphi: Prediction-based Page Prefetching to Improve the Performance of Shared Virtual Memory Systems". In *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications* (June, 2002). (With M. Burtscher)

"Active Memory Clusters: Efficient Multiprocessing on Commodity Clusters". In *Proceedings of the Fourth International Symposium on High Performance Computing, Lecture Notes in Computer Science* 2327 (May, 2002): 78–92 (With M. Heinrich and M. Chaudhuri)

"Using Multicast and Multithreading to Reduce Communication in Software DSM Systems. In *Proceedings of the Fourth Symposium on HPCA* (1998): 312–323. (With J. Bennett)

Paul Stodghill

Research Associate
stodghil@cs.cornell.edu
http://www.cs.cornell.edu/stodghil/

Paul Stodghill obtained his bachelor's degree in mathematics and computer science from Dickinson College in 1988. He obtained his Ph.D. in computer science from Cornell in 1997. Since 1997, he has been a post-doctoral research associate and research associate in CS.

With the deployment of high-bandwidth networks, computational science is entering a new era of distributed and collaborative computing. Stodhill's research interests focus on supporting this effort. For example, he has worked closely with a number of computational scientists to develop novel, high-performance distributed scientific applications. Currently, he is developing fault-tolerant support for parallel applications and infrastructure for deploying scientific simulations as Web services. He is also helping to develop model-based and empirical optimization techniques that allow codes to be migrated between platforms without loss of performance.

SELECTED PUBLICATIONS

"Computational Science Simulations Based on Web Services". In *International Conference on Computational Science 2003*. (With L. Chew, et al.)

"Automated Application-level Checkpointing of MPI Programs". In *Principles and Practices of Parallel Programming (PPOPP)* (July, 2003). (With C. Bronevetsky, D. Marques, and K. Pingali)

"A Comparison of Empiricxal and Model-driven Optimization" In *Programming Languages Design and Implementation (PLDI)* (2003). (With K. Yotov)

Éva Tardos

Professor
eva@cs.cornell.edu
http://www.cs.cornell.edu/home/eva/eva.html

Éva Tardos received her Ph.D. at the Eötvös University in Budapest, Hungary in 1984. After teaching at Eötvös and at M.I.T., she joined Cornell in 1989. She is currently a full professor in CS. She is a member of the American Academy of Arts and Sciences, an ACM Fellow, was a Guggenheim Fellow, a David and Lucille Packard Fellow in Science and Engineering, a Sloan Fellow; a Presidential Young Investigator; and has received the Fulkerson Prize in 1988 (awarded jointly by the American Mathematical Society and the Mathematical Programming Society for a paper in discrete mathematics). She is the editor of several journals.

Tardos's research interest focuses on the design and analysis of efficient algorithms for combinatorial-optimization problems on graphs or networks. Such problems arise in many applications such as vision, and the design, maintenance, and management of communication networks. She is mostly interested in fast combinatorial algorithms that provide provably optimal or close-to-optimal results. She is most known for her work on network-flow algorithms, approximation algorithms for network flows, cut, and clustering problems. Her recent work focuses on algorithmic game theory, an emerging new area of designing systems and algorithms for selfish users.

SELECTED PUBLICATIONS

"How Bad is Selfish Routing?" *Journal of the ACM*, 49(2) (2002): 236–259. (With T. Roughgarden)

"Approximation Algorithms for Classification Problems with Pair-wise Relationships: Metric Labeling and Markov Random Fields". In *Journal of the ACM* 49(5) (2002): 616–639. (With J. Kleinberg)

"A Strongly Polynomial Minimum Cost Circulation Algorithm". *Combinatorica* 5(3) (1985): 247–255.

Professor Éva Tardos with a student.

Tim Teitelbaum

Associate Professor
tt@cs.cornell.edu
http://www.cs.cornell.edu/info/people/tt/tim_teitelbaum.html

Tim Teitelbaum received a B.S. in mathematics from the Massachusetts Institute of Technology in 1964, and his Ph.D. in computer science from Carnegie Mellon University in 1975.

His 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. He is working to improve the performance and functionality of generic dependence-graph technology, and also exploring the use of the technology in various application domains, including software development, maintenance and reengineering of legacy code, test-data generation, security-assurance and safety-assurance inspection, and semantic interference checking in configuration-management systems.

Teitelbaum's earlier work on programming environments and incremental computation resulted in the Cornell Program Synthesizer and the Synthesizer Generator, two of the earliest systems to have demonstrated the viability of integrated language-based programming environments and syntax-directed editors.

He is a cofounder and chairman of GrammaTech, Inc., and a panelist for the National Science Foundation.

SELECTED PUBLICATIONS

"The Cornell Program Synthesizer: A Syntax-directed Programming Environment". *Communications of the ACM* 24(9) (September, 1981): 563–573. (With T. Reps)

"Incremental Context-dependent Analysis for Language-based Editors". *ACM Transactions on Programming Languages and Systems (TOPLAS)* 5(3) (July, 1983): 49–477 (With T. Reps and A. Demers)

"Systematic Derivation of Incremental Programs". *Science of Computer Programming* 24(1) (1995): 1–39. (With Y. Liu)

Charles Van Loan

Professor and Chair, CS
cv@cs.cornell.edu
http://www.cs.cornell.edu/cv/

Charles Van Loan received his Ph.D. in mathematics from the University of Michigan in 1973. After being a postdoctoral research fellow at the University of Manchester, he joined CS as an assistant professor in 1975.

Professor Van Loan works in the matrix-computation field, specializing in least-squares and eigenvalue problems that arise in control engineering and signal processing. Block-matrix computations are a current interest with a special emphasis on novel algorithms that exploit Kronecker-product structure. Kronecker products are increasingly important because of the role that they play in fast transforms and various multilinear applications. He is currently focusing on low-rank approximations of high-dimensional tensors using the singular value decomposition.

Professor Van Loan is the author of five textbooks: *Matrix Computations* (with G. H. Golub), *Computational Frameworks for the Fast Fourier Transform, Introduction to Scientific Computation—A Matrix Vector Approach Using Matlab, Introduction to Computational Science and Mathematics,* and *Handbook for Matrix Computations* (with T. Coleman).

SELECTED PUBLICATIONS

"An Analysis of the Total Least Squares Problem". *SIAM Journal on Numerical Analysis* 17(1980): 883–893. (With G. Golub)

"The WY Representation for Products of Householder Transformations". *SIAM Journal of Scientific and Statistical Computing* 8 (1987): s2–s13 (1987). (With C. Bischof)

"The Ubiquitous Kronecker Product". *Journal of Computational and Applied Mathematics* 123 (2000): 85–100.

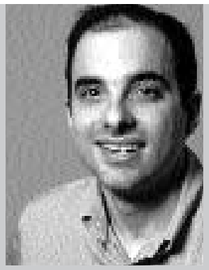

Stephen A. Vavasis

Professor
vavasis@cs.cornell.edu
http://www.cs.cornell.edu/home/vavasis/

Stephen Vavasis received his Ph.D. in computer science from Stanford University in 1989. Since then, he has taught at Cornell where he is currently a full professor. He has also held summer and sabbatical research positions at Sandia National Laboratories, RIACS, Xerox PARC, Bell Labs, and Argonne National Laboratory.

Vavasis's research centers on scientific computing. He is known for bridging the gap between theory and practice in numerical algorithms. His contributions include the first provably good mesh generator for three-dimensional finite-element analysis (with former student S. Mitchell, now at Sandia), the first interior-point method for linear programming whose polynomial running time does not depend on the objective function or constraint right-hand side (with Y. Ye of Stanford), and guaranteed-quality geometric mesh partitioning (with Miller, Teng, and Thurston). More recently (with former student G. Jonsson, now at deCode), he has developed the first resultant-based algorithm for solving polynomial equations with provable accuracy guarantees.

Since coming to Cornell, Vavasis has received a Presidential Young Investigator Award (1990–1995), a Guggenheim Fellowship (1996–1997), and an Ip-Lee Teaching Award (1999).

SELECTED PUBLICATIONS

"A Primal-dual Interior Point Method Whose Running Time Depends Only on the Constraint Matrix". *Mathematical Programming* 74 (1996): 79–120. (With Y. Ye)

"Geometric Separators for Finite–Element Meshes". *SIAM Journal of Scientific Computing* 19 (1998): 364–386. (With G. Miller, S.–H. Teng, and W. Thurston)

"Quality Mesh Generation in Higher Dimensions". *SIAM Journal of Computing* 29 (2000): 1334–1370. (With S. Mitchell)

Werner Vogels

Research Associate
vogels@cs.cornell.edu
http://www.cs.cornell.edu/vogels/

Werner Vogels obtained his Ingenieur Hogere Informatica degree in 1989 from the Haagse Hogeschool in The Hague, The Netherlands. After a number of years as a researcher in various European Esprit projects, he joined CS in 1994, where he is now a research associate.

Vogels' research interest is in communication technologies for scalable distributed systems, with a focus on the interaction between applications, operating systems and network protocols, and in the design of high-performance run-time systems for advanced distributed operations on cluster-computing systems. He is a principal investigator in the Spinglass project, where, in collaboration with Ken Birman and Robbert van Renesse, he works on the development of a new generation of high-scalable reliable network protocols based on the principles of epidemic information dissemination. He also leads the Galaxy project, which focuses on the distributed-systems needs of enterprise cluster-computing systems, in particular providing practical solutions to the scalability problems that arise in these systems.

SELECTED PUBLICATIONS

"An Overview of the Galaxy Management Framework for Scalable Enterprise Cluster Computing". In *Proceedings of the IEEE International Conference on Cluster Computing: Cluster–2000* (December, 2000). (With D. Dumitriu)

"File System Usage in Windows NT 4.0". In *Proceedings of the Seventeenth ACM Symposium on Operating Systems Principles* (December, 1999).

"U-Net: A User-level Network Interface for Parallel and Distributed Computing". In *Proceedings of the Fifteenth ACM Symposium on Operating Systems Principles* (December, 1995). (With T. von Eicken, A. Basu, and V. Buch)

Golan Yona

Assistant Professor
golan@cs.cornell.edu
http://www.cs.cornell.edu/golan/

Golan Yona obtained a bachelor's degree (honor program) in physics and mathematics, and Ph.D. in computer science at the Hebrew University of Jerusalem in 1999. He was a Burroughs–Welcome postdoctoral fellow in computational molecular biology at Stanford University from 1998 until 2000.

Yona's research focuses on computational molecular biology, with an emphasis on developing tools and methodologies for large-scale analysis of the protein universe. The goal of his research is to explore high-order organization and obtain a global view of the protein space. The global view is expected to yield valuable insights about the nature and function of new genes and can lead to the discovery of global principles in the protein space.

Yona's research is rooted in two different disciplines, computer science and molecular biology, and 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, 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. His study so far resulted in two large databases that are being used by biologists to study new genes, ProtoMap (*http://protomap.cornell.edu*) and BioSpace (*http://biospace.cornell.edu*).

Professor Yona is recipient of a National Science Foundation CAREER Award (2002).

SELECTED PUBLICATIONS

"Towards a Complete Map of the Protein Space based on a Unified Sequence and Structure Analysis of all Known Proteins". In *Proceedings of ISMB 2000:* 395–406, AAAI Press. (With M. Levitt)

"ProtoMap: Automatic Classification of Protein Sequences, a Hierarchy of Protein Families, and Local Maps of the Protein Space". *Proteins: Structure, Function and Genetics* 37 (1999): 360–378. (With N. Linial and M. Linial)

"Global Self Organization of All Known Protein Sequences Reveals Inherent Biological Signatures". *Journal of Molecular Biology* 268 (1997): 539–556 (With M. Linial, N, Linial, and N. Tishby)

Ramin Zabih

Associate Professor
rdz@cornell.edu
http://www.cs.cornell.edu/rdz/

Ramin Zabih received undergraduate degrees in computer science and in mathematics from the Massachusetts Institute of Technology, and a Ph.D. in computer science from Stanford University in 1994. He joined CS in 1994, and was promoted to associate professor in 2001. In 2001, he was also given a joint appointment in the Department of Radiology at Cornell's Joan and Sanford I. Weill Medical College.

Zabih's research interests are in computer vision and its applications, especially in medical imaging. He is best known for the work his group has done in applying combinatorial-optimization methods, such as graph cuts, to computer-vision problems. He is currently supervising several Ph.D. students who are working on applying such methods to the automated analysis of magnetic resonance imagery. He has also done extensive consulting for Microsoft, where his work had a major impact on Internet Explorer.

He received the Abraham Wong teaching award from the College of Engineering in 1995. In 2002 he received the best paper award at the European Conference in Computer Vision.

SELECTED PUBLICATIONS

"Fast Approximate Energy Minimization via Graph Cuts". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23(11) (November, 2001): 122–1239. (With Y. Boykov and O. Veksler)

"Color-spatial Indexing and Applications". *International Journal of Computer Vision* 35(3) (1999): 245–268. (With J. Huang, S. Kumar, M. Mitra, and W.-J. Zhu)

"Non-parametric Local Transforms for Computing Visual Correspondence". In *Proceedings of the Third European Conference on Computer Vision* 94 (May, 1994): 151–158. (With J. Woodfill)

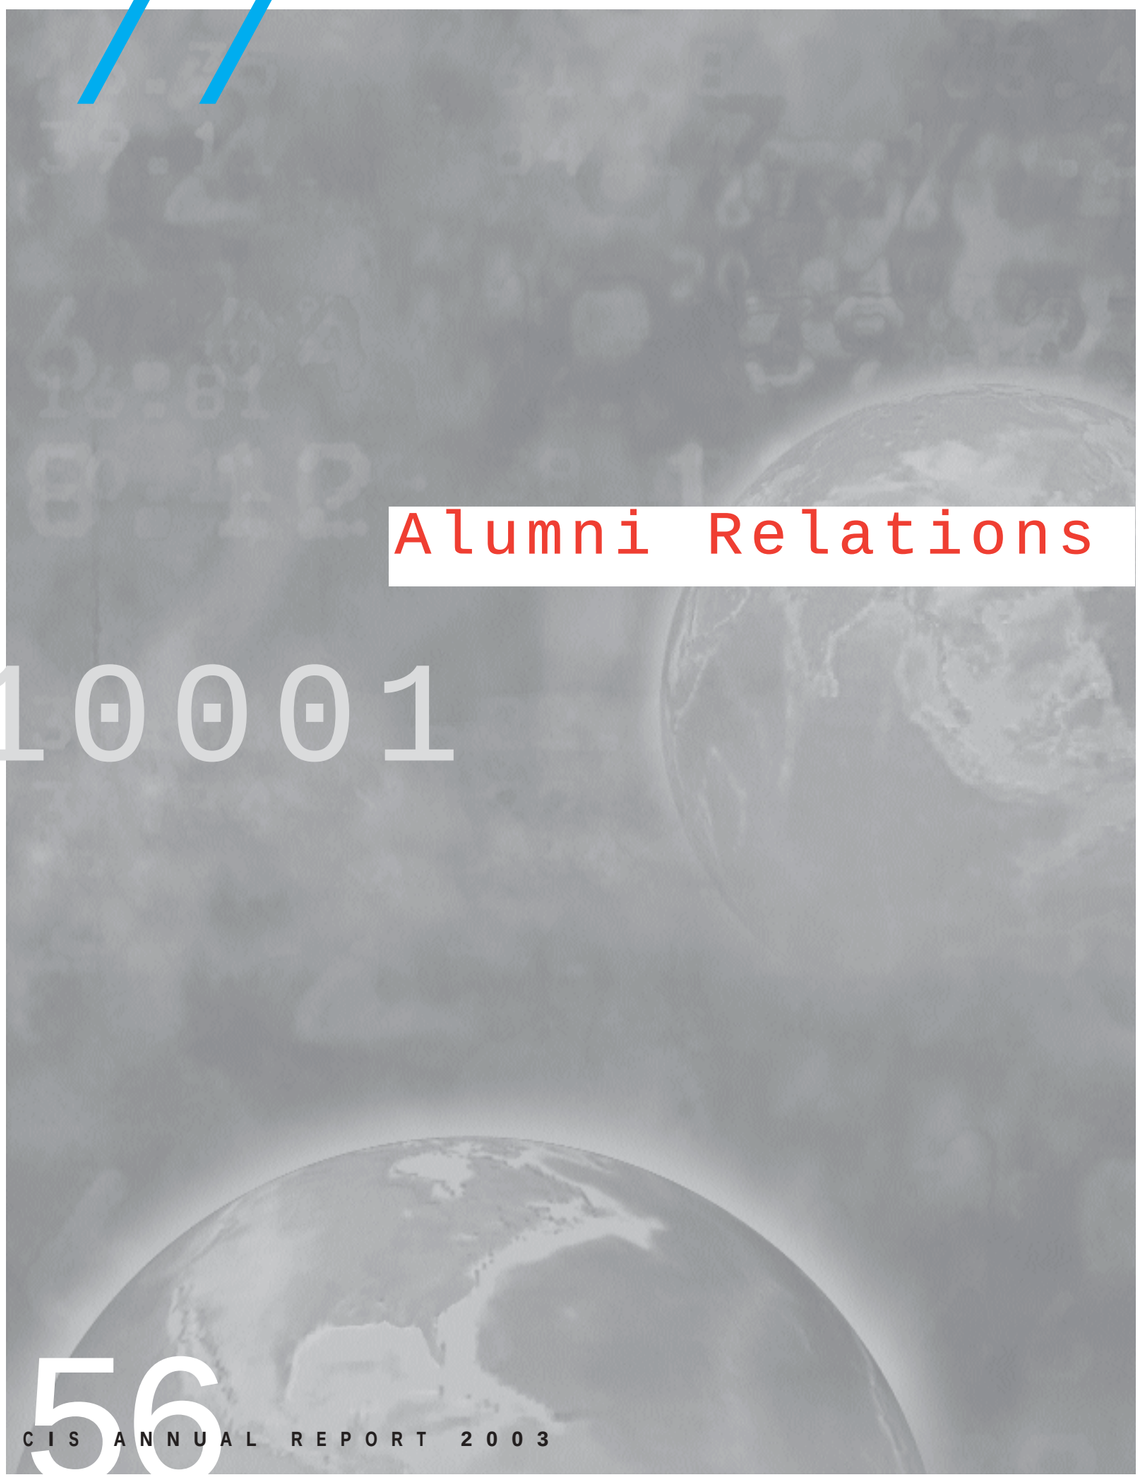

Alumni Relations

"An Evening with the Faculty" was the theme of a large gathering of alumni and friends on May 12, 2003 at the Plaza in New York City. More than eighty people came to hear professors Robert Constable, Dan Huttenlocher, and Ramin Zabih brief some of Cornell's most active alumni on the progress of computing and information science. Topics included the new Information Science majors in the colleges of Engineering, Arts and Sciences, and Agriculture and Life Sciences; the expansion of the digital arts and graphics program; the growing success of the digital libraries initiative; collaborative efforts in biology and medicine; and the establishment of several key security institutes.

Robert Constable

John Belizaire '94

The Faculty of Computing and Information Science was pleased to welcome the involvement of business leader and educational visionary **Narayana Murthy** this year as an advisor for Cornell's efforts in computing and information science. Mr. Murthy recently joined the Cornell University Board of Trustees, and is the CEO and founder of Infosys, a worldwide leader in software development. A campus talk by Murthy is planned for the CIS Distinguished Lecture Series in fall 2004.

Carl Bass '78, Executive Vice President at Autodesk, has arranged for a generous academic site license for several top design and graphics software packages at Cornell. Autodesk is the leading mechanical and architectural design software developer in the world and in recent years has come out with a line of high-end digital content creation tools for game and film animation design.

CS alumni, **John Belizaire '94 M.Eng. '95** and **Guilherme Hoefel '02** have been actively reaching out to minority high school students on behalf of Cornell. Belizaire, a partner in NEXT STAGE, LLC, who resides in Manhattan, has an interest in establishing programs and scholarships for talented students from the northeast. Guilherme works for Qualcomm in San Diego (**Irwin Jacobs '54** is the founder and CEO of Qualcomm) and has been working with Cornell to connect with minority students in his area.

In January 2002, Belizaire also participated as an industry panelist at an annual Cornell Silicon Valley event celebrating entrepreneurship. Belizaire and fellow CS classmate **Julian Pelenur '94 M.Eng. '95** made a huge splash when they sold their start-up company, The Theory Center Inc., to BEA Systems for $100 million in 1999.

CS/Fine Arts alum **Rama Hoetzlein '01** joined forces with Professors David Schwartz (CS) and David Borden (Music) to establish a ground-breaking educational project in computer game design. The project involved more than 45 students from a variety of fields working collaboratively to develop alternative game formats in a multidisciplinary setting. The main purpose of the project, beyond teaching game design principles, was to encourage technical and liberal arts students to develop interesting alternatives to the glut of violent games in the market place.

In December of 2002 a panel of women leaders addressed a crowd of more than 75 students, faculty, and alumni at an event called "Perspectives on Women in Computer Science". Ph.D. student and Cornell alumna **Vicky Weissmann '96 M.Eng. '99** served effectively as panel moderator. Association of Computer Science Undergraduates President, **Priyanka Nishar '03** was one of five panelists discussing issues ranging from myths about women in computing-related fields to effective programs for mentorship and outreach. Other alumni attending this event were: **Jordan Erenrich '02 M.Eng. '03,** Professor **Daisy Fan, Ph.D. '03, Dan Jenkins '82,** Professor **Dexter Kozen, Ph.D. '77,** and Professor **Lillian Lee '93.**

Cornell had another successful year competing in the ACM Programming Championships thanks to **Nikola Valerjev '96** who once again secured generous sponsorship from **Green Hills Software.** A Cornell team won first place in the northeast regional competition and moved on to Beverly Hills, California to compete in the world finals. Green Hills has sponsored Cornell teams for several years.

Vanitha (Badari) Milberg '92 and husband **Justin Milberg '88,** made a generous gift of $50,000 to establish scholarships for students in the Colleges of Arts and Sciences, and Engineering. Vanitha is a former ACSU Vice President.

This year's **Degenfelder Family Scholarship** was split three ways in 2002. New honoree **Craig Lowe '04** joined previous winners **Ben Mathew '03** and **Vlad Muste '04** to split the $5,000 award which was created to recognize students who are 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.

Stephan Paternot '96, who made boom/bust history taking a dot com public in the late 1990s, is now the star of his own movie. Stephan teamed up with Linda Hamilton (of Terminator movie fame) and super model Shannyn Sossamon, to create a short film called "Wholey Moses". The film has been selected as a finalist in the 2003 USA Film Festival.

This year's Jonathan E. Marx Senior Prizes were presented to **Robin Ghi-Hao Lim '03** and **Praveen Sethupathy '03** as part of the Computer Science Graduation Ceremony on Hoy Field on May 25, 2002. **Jonathan E. Marx '85** was a CS major who died in a skiing accident shortly after his graduation in 1985. The Marx family established the Marx Senior Prizes to recognize students who have most demonstrated a positive spirit among their classmates, held significant leadership roles, and have been of service in the community. The Marx family also established a teaching award in the name of Jonathan's father, the late **Alan S. Marx, J.D. '61.** The Alan Marx Memorial Prize for Excellence Supporting Undergraduate Education was awarded to **Benjamin T. Mathew '03**, recognizing his efforts as a consultant and teaching assistant for CS 100.

The Computer Science Prize for academic excellence, given by the CS faculty, was awarded to **Omar Habib Khan '03**, in recognition of his academic and research achievements.

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.*

Courses

CS Student Wins Computing Research Association's 2003 Outstanding Undergraduate Research Award

Omar Khan received his Bachelor's degree in Computer Science in May 2003.

Omar has done significant research in data analysis techniques. He has addressed a wide variety of problems at both the theoretical and implementation levels. Omar's work involves attempting to cluster all documents in NEC Corporation's CiteSeer collection and determining how the clustering changes with time. Omar posed fundamental questions about the nature of structures found by clustering algorithms. He contributed to the development and implementation of a sophisticated clustering technique that he then validated using several independent methods. His range of skills includes theoretical analysis, careful experimentation, and explanation of results. Additionally, he obtained research results in stochastic search and in sensor fusion. Omar and his advisors are now writing papers that will disseminate his work.

Omar ranked first in his class of nearly 700 students at Cornell in his freshman and sophomore years. He has been a teaching assistant and a course consultant at Cornell. He has also been a research assistant at Cornell, a summer research intern at McGill University, a summer research intern at Xerox PARC, and a student researcher and project leader at the Cornell Theory Center. Omar has won national recognition in mock trial competitions. At Cornell University, he was awarded the 2002–03 Frank and Rosa Rhodes Scholarship and has been named to the Dean's List in every semester of his undergraduate studies. He has participated in a variety of outreach activities with the Cornell Theory Center.

Below are titles of the courses offered

by the Department of Computer Science. For more details, see *http://www.cs.cornell.edu/CUCS/courses_degreeprogs/.*

Undergraduate Courses

COM S 099	Fundamental Programming Concepts
COM S 100	Introduction to Computer Programming
COM S 101	Introduction to Cognitive Science
COM S 113	Introduction to C
COM S 114	Unix Tools
COM S 130	Introductory Web Programming
COM S 172	Computation, Information, and Intelligence
COM S 191	Media Arts Studio I
COM S 201	Cognitive Science in Context Laboratory
COM S 211	Computers and Programming
COM S 212	Java Practicum
COM S 213	C++ Programming
COM S 214	Advanced Unix Programming and Tools
COM S 230	Intermediate Web Design
COM S 280	Discrete Structures
COM S 312	Data Structures and Functional Programming
COM S 314	Computer Organization
COM S 321	Numerical Methods in Computational Molecular Biology
COM S 322	Introduction to Scientific Computation
COM S 324	Computational Linguistics
COM S 330	Applied Database Systems
COM S 381	Introduction to Theory of Computing
COM S 392	Topics in High-level Vision
COM S 401	Introduction to Applied Scientific Computing with MATLAB
COM S 402	Scientific Visualization with MATLAB
COM S 403	Development of Scientific Computing Programs
COM S 404	Survey and Use of Software Libraries for Scientific Computing
COM S 409	Data Structures and Algorithms for Computational Science
COM S 411	Programming Languages and Logics
COM S 412	Introduction to Compilers
COM S 413	Practicum in Compilers
COM S 414	Systems Programming and Operating Systems
COM S 415	Practicum in Operating Systems
COM S 417	Introduction to Computer Graphics
COM S 418	Practicum in Computer Graphics
COM S 421	Numerical Analysis
COM S 430	Information Discovery
COM S 432	Introduction to Database Systems
COM S 433	Practicum in Database Systems
COM S 472	Foundations of Artificial Intelligence
COM S 473	Practicum in Artificial Intelligence
COM S 474	Introduction to Natural Language Processing
COM S 478	Machine Learning
COM S 481	Introduction to Theory of Computing
COM S 482	Introduction to Analysis of Algorithms
COM S 483	Quantum Computation
COM S 486	Applied Logic
COM S 490	Independent Reading and Research

Graduate Courses

COM S 501	Software Engineering
COM S 502	Web Information Systems
COM S 504	Applied Systems Engineering I
COM S 505	Applied Systems Engineering II
COM S 513	System Security
COM S 514	Intermediate Computer Systems
COM S 515	Practicum in Systems
COM S 517	Advanced Rendering
COM S 518	Special Topics: Computer Animation
COM S 519	Computer Networks
COM S 522	Computational Tools and Methods for Finance
COM S 572	Heuristic Methods for Optimization
COM S 574	Language Technologies
COM S 576	Decision Theory I
COM S 577	Decision Theory II
COM S 578	Empirical Methods in Machine Learning and Data Mining
COM S 601	System Concepts
COM S 611	Advanced Programming Languages
COM S 612	Compiler Design for High-performance Architectures
COM S 613	Concurrent Programming
COM S 614	Advanced Systems
COM S 615	Adaptive Systems
COM S 621	Matrix Computations
COM S 622	Numerical Optimization and Nonlinear Algebraic Equations
COM S 624	Numerical Solutions of Differential Equations
COM S 626	Computational Molecular Biology
COM S 632	Advanced Database Systems
COM S 664	Machine Vision
COM S 671	Introduction to Automated Reasoning
COM S 672	Advanced Artificial Intelligence
COM S 674	Natural Language Processing
COM S 676	Reasoning About Knowledge
COM S 677	Reasoning About Uncertainty
COM S 678	Advanced Topics in Machine Learning
COM S 681	Analysis of Algorithms
COM S 682	Theory of Computing
COM S 683	Advanced Design and Analysis of Algorithms
COM S 684	Approximation and Network Algorithms
COM S 685	The Structure of Information Networks
COM S 686	Logics of Programs
COM S 709	Computer Science Colloquium
COM S 711	Seminar in Advanced Programming Languages
COM S 713	Seminar in Systems and Methodology
COM S 715	Seminar in Programming Refinement Logics
COM S 717	Topics in Parallel Architectures
COM S 719	Seminar in Programming Languages
COM S 721	Topics in Numerical Analysis
COM S 726	Problems and Perspectives in Computational Molecular Biology
COM S 732	Seminar in Database Systems
COM S 750	Evolutionary Computation and Design Automation
COM S 751	Media Research and Critical Design
COM S 752	Seminar on Scholarly Information Architecture
COM S 754	Systems Research Seminar
COM S 772	Seminar in Artificial Intelligence
COM S 775	Seminar in Natural Language Understanding
COM S 786	Introduction to Kleene Algebra
COM S 789	Seminar in Theory of Algorithms and Computing
COM S 790	Special Investigations in Computer Science
COM S 890	Special Investigations in Computer Science
COM S 990	Special Investigations in Computer Science

Research Grants

10001

FUNDED RESEARCH Computing and Information Science/Computer Science

Investigator	*Sponsor*	*Award (dollars)*	*Title*
Arms	NSF	799,085	The National Science Digital Library (NSDL) Central System
Birman	AFOSR	427,608	A Testbed for Highly-scalable Mission-critical Information Systems (DURIP)
Birman	DARPA/AFRL	288,000	Scalable Data Redundancy for Network-centric Military Applications
Birman/Constable	DARPA/AFRL	3,839,383	Spinglass Adaptive Probabilistic Tools for Advanced Networks
Birman/Gehrke/Demers	AFOSR	4,000,000	Scalable Technology for a New Generation of Collaboration Applications
Cardie	NSF/POWRE	68,695	Integrating Natural Language Processing and Information Retrieval for Intelligent Text-processing
Cardie	DARPA/ONR	948,169	Rapidly Portable Translingual Information Extraction and Interactive Multidocument Summarization
Cardie	MITRE Corp	61,743	ARDA–NRRC Workshop
Cardie	NSF	500,000	Reducing the Corpus Annotation Bottleneck for Natural Language Learning
Constable	NSF	283,975	Educational Innovation: Creating and Evaluating Formal Courseware for Mathematics and Computing
Constable	NSF	20,800	U.S.-Germany Cooperative Research: Enhancing Proof Assistant Systems
Constable	DARPA/AF	2,004,156	An Open Logical Programming Environment: A Practical Framework for Sharing Formal Models
Constable	ONR	1,938,148	Building Interactive Digital Libraries of Formal Algorithmic Knowledge
Constable	NSF	300,000	Innovative Programming Technology for Embedded Systems
Department	NSF	1,331,298	CISE Research Infrastructure: A Next Generation Computing and Communications Substrate
Elber	NSF	465,742	Kinetics of Ion Channels by Atomically Detailed Computer Simulations
Elber	NIH	943,467	Long Time Dynamics of Biomolecules
Gehrke	AFRL	735,000	Flexible Decision Support in Device-saturated Environments—SenseIT
Gehrke	NSF	235,000	Interactive and Online Data Mining
Gehrke	NSF	210,000	Scalable Decision Tree Construction
Gehrke	Lockheed Martin	25,000	Intelligent Data Cleansing Technologies
Gehrke	NSF	340,000	CAREER: Towards Sensor Database Systems
Gehrke	Sloan	40,000	Sloan Research Fellowship
Gehrke/Demers	KDD thru NSF	760,000	Distributed Mining and Monitoring
Gehrke	Microsoft	35,000	Monitoring the Tsunami—A Light-Weight DBMS and Stream Processor for Sensor Devices
Gehrke/Sirer/Shanmugasundaram	Microsoft	70,878	Query Caching and Routing for Mobile Ad Hoc Clients in the .Net Framework
Gehrke/Sirer/Shanmugasundaram/ Demers/Birman	NSF	907,320	ITR: Massively Convergent Distributed Computing
Ginsparg	NSF	958,798	E-Print Archive
Gomes	AFRL	3,100,000	Intelligent Information Systems Institute
Halpern	ONR	433,962	Semantic Consistency in Information Exchange
Halpern	ONR	526,058	Software Quality and Infrastructure Protection for Diffuse Computing
Halpern	AFOSR	322,788	Formulating and Reasoning About Security Policies
Halpern	ONR	354,763	A Logical Foundation for Reasoning About Security
Halpern	NSF	300,000	Towards Improved Logics For Reasoning About Security
Hopcroft/Selman	NSF	30,000	ITR: Emerging Communities in Large Linked Networks: Theory Meets Practice
Joachims	Google	58,223	Learning Metrics and Learning Ranking Functions, and Transductive Text Classification
Kleinberg	Packard	625,000	Algorithmic Methods for Networks
Kleinberg/Lee/Cardie/Selman	NSF ITR	450,000	The Construction and Analysis of Information Networks
Kozen	NSF	210,000	Kleene Algebra
Kreitz	NSF	285,000	Proof Automation in Constructive Type Theory
Lagoze	Digital Library Federation	177,467	Open Archives Initiative
Lagoze	NSF	255,700	Metadata for Resource Discovery of Multimedia Digital Objects Harmony
Lagoze	NSF	291,650	Integrating and Navigating Eprint Archives through Citation-linking
Lagoze	Univ of Virginia/Mellon Found.	495,325	The Open Source FEDORA Repository Development Project
Lee	Sloan	40,000	Sloan Research Fellowship
Morrisett	Sloan	35,000	Sloan Research Fellowship
Morrisett	NSF	205,000	CAREER: Design, Applications, and Foundations of Safe Low Level Program Languages
Morrisett	AFOSR/PECASE	1,000,000	Next Generation Systems Languages
Myers	NSF	349,999	CAREER: Practical Language-Based End-to-end Security
Myers	Sloan	40,000	Sloan Research Fellowship
Myers	NSF	330,000	End-to-end Integrity and Confidentiality for Distributed Systems
Pingali	NSF	190,993	MATLAB Extensions and Compiler Techniques for High Performance Computing

FUNDED RESEARCH Computing and Information Science/Computer Science

Investigator	*Sponsor*	*Award (dollars)*	*Title*
Pingali	NSF	548,314	Synthesis of Block-recursive Codes for Deep Memory Hierarchies
Pingali	NSF	236,826	A Framework for Developing Complex Applications on High-End Petaflop-class Machines
Pingali	NSF	590,000	ITR/SY: A New Framework for Program Optimization
Schneider	AFOSR	2,050,000	AFRL/Cornell Information Assurance Institute
Schneider	AFOSR	592,657	CIPIAF for Information Assurance Institute
Schneider	AFOSR	4,138,325	AFRL/Cornell Information Assurance Institute
Schneider	Microsoft	150,000	Information Assurance Institute
Schneider	Intel	73,000	Research on Language-based Security
Schneider/Myers	DARPA/AF	2,709,784	Containment and Integrity for Mobile Code
Schneider/Morrisett	AFOSR	844,408	Language-based Security for Extensible Systems
Schneider/Morrisett	AFOSR	471,107	Trust in Security-policy Enforcement Mechanisms
Schneider/Morrisett/Kozen/Myers	ONR	4,247,977	Language-based Security for Malicious Mobile Code
Selman	Sloan	35,000	Sloan Research Fellowship
Selman	NSF	300,000	CAREER: Compute Intensive Methods for AI
Selman/Gomes	AFRL	550,000	Principled Analysis and Synthesis of Agent Systems Using Tools from Statistical Physics
Selman/Gomes	DARPA/AF	1,426,881	Controlling Computational Cost: Structure, Phase Transitions, and Randomization
Sengers	NSF	500,000	CAREER: Using Cultural Theory to Design Everyday Computing
Sirer	Microsoft	105,589	Assuring the Security of Components in the .Net Framework
Tardos	NSF	249,559	Algorithmic Issues in Communication Networks
Tardos	DARPA/ONR	256,212	Efficient Algorithms for Transportation in Dynamic Networks
Tardos	ONR	1,176,548	Algorithmic Issues in Network Design and in Information Access
Tardos/Zabih	NSF	300,000	ITR/SY: Combinatorial Optimization Algorithms for Information Access
Van Loan	NSF	247,874	New Applications and Algorithms that Involve the Kronecker Product
Vogels	Microsoft	20,000	High-Performance Computing using SSCLI
Vogels	Microsoft	15,000	CollabNet
Vogels/Gehrke/Shanmugasundaram	Microsoft	175,000	Distributed Systems Support for the Global Real Time Enterprise
Yona	Bio & Life Sciences	48,658	Global Self-organization of the Protein Space: Towards a Map of the Protein Space
Yona	NSF	1,103,917	CAREER: Global Self-organization of all Known Proteins—Toward a Complete Map of the Protein Space

TOTAL EXPENDITURES for Fiscal Year 2002–2003: $14,600,000

Collaborative Research at Cornell

Investigator	*Sponsor*	*Award (dollars)*	*Title*
Arms/Krafft/Lagoze/Eng Library	NSF	1,637,500	Collaborative Project: Core Integration of the National SMETE Digital Library
Arms/Krafft/Lagoze/Eng Library	NSF	8,745,453	Collaborative Project: Core Integration—Leading NSDL Toward Long-term Success
Coleman/TC	TG Information Network	840,000	Financial Engineering and Tools
Coleman/TC	NYS	1,200,000	Theory Center Operational Support FY2003
Coleman/TC	KYNEX	20,000	Numerical Pricing of Bonds
Coleman/TC	WOTN	390,000	Sublicense Agreement
Coleman/TC	Aximetric Inc.	30,000	Aximetric Inc. Collaboration Agreement
Coleman/TC	SKG	157,750	Development Agreement
Department/TC/CURIE/ Summer College	GE	200,000	Program Continuum for Attracting and Retaining Women to/in CS Studies for Information Technology Careers
Elber/Kleinberg/Chew/Kedem/MGB	NSF	899,000	Multiscale Hierarchical Analysis of Protein Structure and Dynamics
Elber/BSCB	anonymous	361,113	Two-track Program in Computational Biology and Medicine (CBM) as a Part of the Tri-institutional Research Program
Lagoze/Comm/Olin Library	NSF	2,425,899	Security and Reliability in Component-based Digital Libraries
Pingali/TC	NSF	1,500,000	CISE Research Infrastructure: A Two-tier Computation and Visualization Facility for Multiscale Problems
Pingali/Vavasis/Chew/TC/Phy	NSF ITR	5,035,425	Adaptive Software for Field-driven Simulations
Selman/Gomes/ECE	DARPA	109,736	Configuring Wireless Transmission and Decentralized Data Processing for Generic Sensor Networks
Selman/Gomes/M&AE	AFOSR MURI	258,567	Cooperative Control in Uncertain Adversarial Environments
Sirer/Gehrke/Demers/ECE	NSF	410,000	The Ad Hoc Classroom: Integrating Emerging Wireless Communications and Networking Technologies into Mainstream Computer Science and Electrical Engineering Curricula
Yona/Bio Sciences	NSF	1,000,000	CRCNS: Modeling Pathfinding and Target Recognition in the Olfactory System

Submitted Grant Proposals

Investigator	*Sponsor*	*Amount (dollars)*	*Title*
Birman/Myers	NSF sub of U Virginia	2,852,000	ITR: Center for Securing Computation in Physical World Applications
Constable	DARPA sub of Boeing	380,000	Model-driven Generation of Verifiable Distributed Real-time Embedded Systems
Constable	NSF	499,892	Enabling Large Scale Coherency Among Mathematical Texts in the NSDL
Elber	ACS-PRF	80,000	Molecular Dynamic Simulations of EPR Spectra in Proteins
Elber	Sloan	120,000	Computational Studies of the R to T transition in Hemoglobin
Elber/Joachims	NIH	1,692,032	Optimization of Folding and Threading Proteins
Gehrke	NSF	363,757	Processing Set-valued Query Results over Data Streams
Gehrke	NSF	500,000	SENSORS: Data-driven Sensor Networks
Gehrke/Shanmugasundaram	NSF	2,221,695	ITR: On Queries and Privacy in P2P Data Sharing Systems
Joachims	NSF	400,000	CAREER: Improving Information Access by Learning from User Interactions
Lee/Kleinberg	NSF	449,897	Graph-Based Approaches to Text Processing
Pingali	DARPA sub of IBM	510,000	PERCS: Phase II
Shanmugasundaram	NSF	406,750	CAREER: Towards Unifying Database Systems and Information Retrieval Systems
Shanmugasundaram/Gehrke	AFRL	41,000	ORIS: Peer to Peer Object Repository with Integrated Security
Tardos	NSF	150,000	Approximation Algorithms and Applications in Network Games
Warner/Lagoze	NSF	300,000	Collaborative Project: Extension of the OAI–PMH Metadata Harvesting Framework to Peer-to-peer Networks
Yona	NSF	3,952,416	ITR: A Comprehensive Protein and DNA Characterization, Classification, and Management Database System: From Genes to Families, Pathways, and Organisms

Submitted Collaborative Research at Cornell

Investigator	*Sponsor*	*Amount (dollars)*	*Title*
Arms/Kleinberg/Ginsparg/Physics	NSF	14,988,311	ITR: Untangling the Web
Cardie/Human Development	NSF	3,688,726	ITR: Addressing Child Language Data Annotation Challenges in a Virtual Linguistic Lab Using Weakly Supervised Natural Language Learning
Caruana/TC	NSF	379,999	Multi-algorithm Parallel Optimization of Costly Functions
Coleman/TC	TG Information Network	675,000	Financial Engineering and Tools
Coleman/TC	USDA	232,180	Computation Agriculture Initiative
Guckenheimer/Math/Physics/ CCB/MBG/Biomedical/ Molecular Med/M&AE/ Sociology/TAM/ECE	NSF	3,338,800	IGERT: Program in Nonlinear Systems
Rooth/Linguistics	NSF	1,229,736	Detection of Linguistic Events Using Statistical Parse Forest Algorithms and Treebank-aligned Feature Constraint Grammars
Sirer/Tardos/ECE	NSF	3,797,452	ITR: An Integrated Approach to Designing Very Large Scale Networks
Tardos/Kleinberg/Huttenlocher/ Halpern/ORIE	NSF	3,659,186	ITR: Networks of Strategic Agents: Theory and Algorithms

Cornell University is an equal-opportunity affirmative-action educator and employer.

Communication and Marketing Services

The following photographs were supplied by University Photo:
Robert Barker, pages 30, 45, 49;
Charles Harrington, pages 2, 64;
Frank DiMeo, pages 4–5;
Nicola Kountoupes, pages 18, 21, 22, 24, 52.
All other images were supplied by CIS.

Printed on recycled paper.
1M 8/2003 MP