Computing and Information Science at Cornell

Education and Research

2004

CIS ANNUAL REPORT 2004

CORNELL

Editor: Una Moneypenny
Publication Coordinators: Una Moneypenny and Nora Balfour

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

Annual Report 2004

Computing and Information Science at Cornell

CONTENTS

CORNELL

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

Those of us working in CIS experienced the sense of steady acceleration throughout the year.

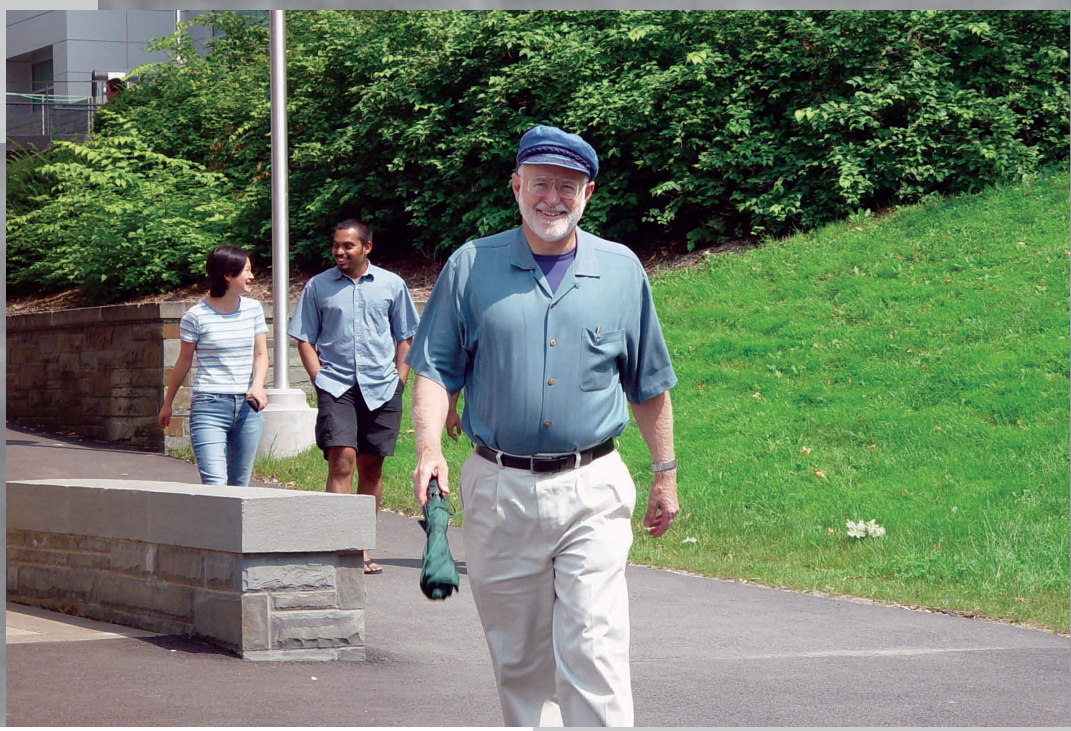

Robert L. Constable

Message from the Dean of the Faculty of Computing and Information Science

It started with a bang when the university trustees appointed Professor Keshav Pingali to a CIS endowed chair, called the India Chair, provided by an anonymous gift to support CS and to strengthen Cornell's ties to India. Then came the inauguration on October 16 of Cornell's eleventh president, Jeffrey Sean Lehman. One of President Lehman's first steps in working with CIS was to invite Bill Gates to visit Cornell—this happened on February 25, 2004, and this report includes photos from the event. It was a very exciting visit because Bill Gates engaged us in detail on a broad range of issues including the future of the Internet, software reliability, program verification, bioinformatics, and the role of CIS in the university. He expressed considerable interest in CIS, especially the fact that the new majors we have created in computational biology and in information science have made up for the loss of enrollment in the CS major.

Later in February, President Lehman and I represented Cornell at the ground-breaking ceremony for the Bridging the Rift (BTR) Research Center that straddles the border between Israel and Jordan. BTR represents a major effort to use science as a common meeting ground to advance the cause of peace—peace through cooperation. The research initiative for this center is called the Library of Life, a joint effort with Stanford University. CIS is responsible for the bioinformatics portion of this major research initiative, planned and led by Ron Elber and Steve Tanksley. It is described in more detail in this report. Still later in March, Cornell trustee Narayana Murthy discussed the idea of a new building for CS with President Lehman, and subsequently President Lehman asked CIS for details on its idea for a new building and a possible "information campus". We now have a process in place to come up with three design options. Dean of Engineering Kent Fuchs is a significant partner in our plans because all of the building options enable Engineering to take advantage of the space vacated by CS and other CIS units that use space adjacent to Engineering, for example the Program of Computer Graphics (PCG) and the Cornell Theory Center (CTC). Moreover, our first choice would include the School of Operations Research and Industrial Engineering (OR&IE) as an occupant of the building, thus freeing up even more space for Engineering.

In May, CIS co-sponsored "Graphics Week" with the PCG and the Department of Architecture. Don Greenburg organized this exciting event, which brought to campus Ed Catmull, president of Pixar, to speak on "Crisis in Production", or "the stuff that we don't normally talk about". Catmull was joined by his colleague, Rob Cook, vice president for research and development at Pixar, who kicked off the event with "A Behind-the-scenes Tour of Moviemaking". Marc Levoy, professor at Stanford University, discussed "The Digital Michelangelo Project", focusing on the intersection of art and computers. George Joblove, senior vice president for technology at Sony Imageworks, and Douglas Kay, chairman, Mondo Media, rounded out the week's events with two talks on "Digital Imagery in Entertainment". All speakers reinforced Levoy's focus on the crossroads where art and computing meet.

In late May, we joined forces with the CTC and the Division of Alumni Affairs and Development to establish a joint CIS office for development, and we entered into an agreement with the CTC to tie our programs in Computational Science and Engineering (CS&E) more tightly to CTC capabilities, which include the massive Windows NT cluster computers. In addition, the NSF awarded Cornell a new IGERT grant to support CS&E under the leadership of John Guckenheimer and Steve Strogatz.

Also in late May, President Lehman and the academic deans met with Bill Clinton, who gave the Convocation Address to the graduating seniors. His visit prompted us to think about the role of universities in American life and in the body politic.

It has been a very good year for CIS, and I was pleased to accept the offer by the president and provost of a second five-year term as dean of the Faculty of Computing and Information Science.

One pleasant official duty as dean was to reappoint Charles Van Loan to chair the CS department for two more years; this unprecedented step of extending a CS chair's appointment recognizes the high regard in which Charlie is held by his colleagues. Charlie will take a well-earned sabbatical at the end of his term, two years hence.

The Chair Search Committee consisted of Claire Cardie, Juris Hartmanis (chair), John Kleinberg, and Johannes Gehrke. They did an outstanding job of taking the pulse of the department and drawing the faculty into thinking about the future of the department.

Other key appointments included Bill Arms and Claire Cardie as co-directors of the Information Science Program, Dan Huttenlocher as the director of graduate studies for that program, and Juris Hartmanis as the chair of the CIS Council. David Shalloway continues to lead our successful computational biology major, and he is leading the effort to create a field with this name as well.

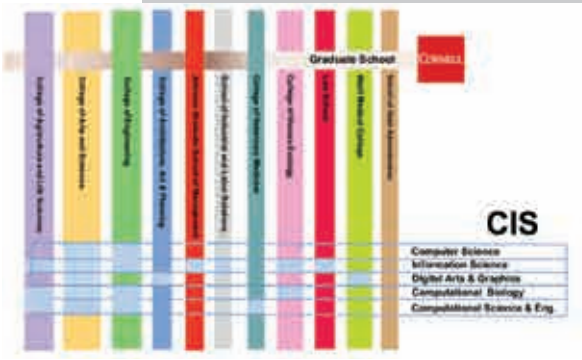

We devoted a considerable amount of our recruiting energy this year to helping the Department of Biological Statistics and Computational Biology and the Department of Statistical Sciences. The CIS faculty has been strongly advocating that we play a role in helping the university establish more strength in core statistics because it is so important to our efforts in artificial intelligence, computational biology, and information science. We are introducing ideas from computational statistics into courses serving these subareas, and we feel that perhaps more than most units, we need to support fundamental strength in core statistics.

Multidisciplinary research is 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.

Charles Van Loan

Message from the Chair of the Department of Computer Science

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 10 percent. The number of outside departments 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.

Let me briefly mention a few of the honors and awards won by CS faculty members during the past year. Steve Marschner, together with co-authors Henrik Wann Jensen and Pat Hanrahan, received a Technical Achievement Award from the Academy of Motion Picture Arts and Sciences. Steve and Rich Caruana received NSF Faculty Early Career and Development (CAREER) awards. Johannes Gehrke was a co-recipient of the Provost's 2004 Award for Distinguished Scholarship. Lillian Lee received Best Paper Award at the joint Human Language Technology Conference and Annual Meeting of the North American Chapter of the Association for Computational Linguistics (HLT–NAACL) (with Regina Barzilay). Dexter Kozen became a fellow of the Association for Computing Machinery (ACM). The department has two new endowed professorships: Keshav Pingali is the India Professor of Computer Science and John Hopcroft is the IBM Professor of Engineering and Applied Mathematics—a pair of well-deserved honors.

CS's faculty members continue to have leadership roles within the university making possible new collaborations and initiatives. David Gries continues as the associate dean for undergraduate education in the College of Engineering. We have joint appointments with the Johnson Graduate School of Management (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 Information Technology Research (ITR) grant concerned with high-performance code generation for scientific and engineering applications (Keshav Pingali, Steve Vavasis, and Paul Chew). We have CS leadership in the CTC (Tom Coleman), the Financial Industry Solutions Center (FISC) (Tom Coleman), the Intelligent Information Systems Institute (IISI) (Carla Gomes), the Information Assurance Institute (IAI) (Fred Schneider), the PCG (Don Greenberg), and the National Science Digital Library (NSDL) (Bill Arms.) Ron Elber was named director of the Library of Life, which is part of the new Bridging the Rift Center to be located on the border between Jordan and Israel. These multidisciplinary adventures are supported by the department's commitment to collegiality and core CS research.

Saddle up. It's Big Sky Country!

CS's faculty members continue to have leadership roles within the university, making possible new collaborations and initiatives.

CAROL MINTON MORRIS

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 Computer Systems Laboratory.

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

National Science Digital Library (NSDL)

Cornell and Columbia Universities share the distinction of heading up the National Science Digital Library, a project that links the resources of more than 100 smaller Internet resources.

Funded by the National Science Foundation, the electronic library aims to bring teachers and students a wealth of ideas and information in mathematics, science, engineering, and technology. Heading the project is an expert on digital libraries and electronic publishing, who also directs CIS's Information Science program.

For ten years before it won this assignment, Cornell's digital libraries group had been researching architectures, protocols, services, and policies of distributed information that facilitate the creation, management, accessibility, and longevity of distributed information.

Digital library research at Cornell is rooted in the practical problems of large-scale electronic publishing, Web information systems, scholarly communication, and the long-term preservation of digital information. One important focus is interoperability. How can coherent information services be created from sources around the world that are managed by independent organizations, with very different goals and technologies? Cornell's particular interests are in simplicity, lowering barriers to collaboration, and finding automatic ways to carry out labor-intensive tasks.

The group's achievements in building coherent services from heterogeneous, independently managed digital libraries include the Open Archive Initiative Protocol for Metadata Harvesting, which enables technically inexperienced groups to share information, and the FEDORA mechanisms, which enable more sophisticated users to work with digital library content.

For more 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. 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

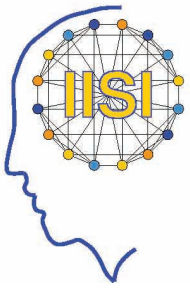

The mission of the IISI, founded in December of 2000, is threefold: to perform and stimulate research in computational and data-intensive methods for intelligent decision-making systems; to foster collaborations within the scientific community; and to play a leadership role in the research and dissemination 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 Carlos Ansotegui (encodings and solvers for combinatorial problems using propositional logic [SAT], many-valued SAT, or constraints [CSP]); Krishna Athreya (branching processes, Markov chains, mathematical statistics, and applications to computer science); Claire Cardie (natural-language understanding and machine learning); Rich Caruana (machine learning, data mining, and bioinformatics); Raffaello D'Andrea (dynamics and control); Carmel Domshlak (modeling and reasoning about preferences and uncertainty, combinatrial search and optimization, and 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); 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); David Schwartz (computer-game design); Meinolf Sellmann (OR and constraint programming [CP] for hard combinatorial problems); Bart Selman (knowledge representation, complexity, and multi-agent systems); Phoebe Sengers (intelligent systems in human and social content and human–computer interaction); David Shmoys (algorithms for large-scale discrete optimization); Chris Shoemaker (large-scale optimization and modeling); 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. Over the past year, the IISI has hosted or sponsored several conferences and workshops: the Tenth International Conference on Principles and Practice of Constraint Programming; the Seventh International Conference on Theory and Applications of Satisfiability Testing (SAT '04); the 2004 HLT–NAACL; the International Conference on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems (CP–AI–OR '04); the IISI Workshop on Strategic Research Directions in AI; and the AFRL/IISI Workshop on Mixed Initiative Decision Making.

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, Albert-Ludwigs-Universität Freiburg, AT&T Labs, Ben-Gurion University, Carnegie Mellon University, Hebrew University of Jerusalem, ILOG Corporation, Microsoft Research, New York University, Rutgers University, Stanford University, Technion, University of Alberta, University of Barcelona, University of British Columbia, University of Lisbon, University of Massachusetts, University of Minnesota, University of Pennsylvania, University of Dallas, University of Washington, York University, and Washington University at St. Louis.

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

Light Work

Prof brings realism to computer graphics

On Valentine's Day, two weeks before hundreds of couture-clad celebrities watched the latest *Lord of the Rings* film win a record-tying eleven Oscars, a more modest ceremony was held at the Ritz Carlton in Pasadena. Hosted by "Alias" star Jennifer Garner, the Academy of Motion Picture Arts and Sciences' Scientific and Technical Awards honored the people who make movie magic possible through advances in such fields as digital effects, robotics, and sound production. Some in the film community call them the "supernerds"—and they mean it affectionately. "These technologies," says Richard Edlund, chairman of the academy's Scientific and Technical Awards Committee, "are part of the creative force behind moviemaking."

Among the big brains in black tie: Steve Marschner, PhD '98. A Cornell computer graphics professor, Marschner and two colleagues from Stanford were honored for advances that have made digitally rendered skin much more realistic. Their work has been used in several high-profile pictures—most notably in the character of Gollum, the computer-generated creature who accompanies the hobbits on their quest to destroy the Ring of Power in the *Lord of the Rings* trilogy. The technology has also been used in such films as *Harry Potter and the Chamber of Secrets* (for Dobby the House Elf, Harry's high-strung servant), *Terminator 3: Rise of the Machines* (for the lovely-but-lethal T-X robot), and *The Matrix Reloaded* (for superimposing actors' heads onto stuntmen's bodies). "This is one of the holy grails of computer graphics," says Edlund. "One of the difficulties of creating lifelike characters in the computer world is the problem that skin is not opaque. If you render a faithfully scanned or created character and the skin is opaque, it doesn't look real."

The issue Marschner and his colleagues addressed, called subsurface scattering, describes the way light both penetrates into and reflects off skin and a wide array of other substances—including marble, cloth, wax, milk, snow, and ocean water. Earlier models had treated such substances as purely reflective rather than translucent, a technique that made computer renderings appear false to the human eye. A digital depiction of a marble statue, for example, would make it look like plaster. "Without accounting for translucency, it's very hard to set things up so the skin looks like skin," says the thirty-two-year-old Marschner. "It looks too hard, or too smooth, or too shiny. But once you account for translucency, then suddenly you can get another level of realism."

The work grew out of a Stanford research effort known as the Digital Michelangelo Project, a marriage of art and science in which extremely high-resolution scans were made of several statues, including the *David*. One of the project's questions was whether, in the context of such detailed scans, marble's translucency could cause problems with the data collection. The answer turned out to be no, but it sparked the team's interest in subsurface scattering. Taking principles of physics into account, they created a simple model to describe the behavior of light in relation to translucent materials, then translated it into a method of rendering such materials via computer graphics. The researchers didn't create a new kind of software, but rather a tool for effects masters to use in their own rendering systems. "When they give these awards, frequently they go to the studios and the people who did the implementing for the films," computer science professor Don Greenberg '55, PhD '68, says of the academy nod. "But enough studios were using this that they looked deeper and went back to the source, which is the research."

Like many of his contemporaries, Marschner can date his fascination with computer graphics and special effects to 1977 and the release of *Star Wars*. As a child in suburban Chicago, he read everything he could get his hands on about how the filmmakers used intricately designed models and motion-controlled cameras to create the movie's spaceship battles. "I was completely fascinated by *Star Wars,* and really excited about it," recalls Marschner. "I had a video camera and played with trying to make little models myself, although that never turns out quite as well as you expect it to."

Marschner's work on modeling skin—as well as projects on rendering hair and other materials—has applications beyond motion pictures and video games. There are many other fields in which it's vital to reproduce reality in exacting detail: hospitals need accurate simulations for training surgeons, for example, and online retailers want to show consumers the exact drape and luster of that silk blouse. "One of the great things about computer graphics is that it's an interdisciplinary area at the boundary between computer science and other things—physics, animation, optics, the dynamics of how things move," Marschner says. "I get to learn about new things, all the way from science to art, as part of my work."

At the Valentine's Day ceremony, Marschner and colleagues Henrik Wann Jensen and Pat Hanrahan received certificates, not Oscar statuettes. Still, the honor puts him in the fraternity of Cornell-educated computer graphics gurus—including George Joblove '76, MS '79; Doug Kay '76, MS '79; Roy Hall, March '83; and Rob Cook, March '82—whom the Academy has honored for contributions to the field. "It's very gratifying to see the stuff being used, especially in films that I really enjoy," Marschner says. "Watching *Lord of the Rings,* even as a computer graphics person, my attention was essentially focused on the performance of Gollum, rather than on the details of the rendering. Usually I'm always looking at all the little flaws. But that one was done so well that it's the performance you're seeing rather than the technology."

CS Professor Steve Marschner *[right]* with colleague Pat Hanrahan at the awards ceremony

BETH SAULNIER
(reprinted from *Cornell Alumni Magazine,* May/June 2004: 38–39)

The computing and information revolution is transforming society. The Cornell Computer Science Ph.D. program is a leader in this transformation, producing cutting-edge research in many important areas.

Research in Computer Science at Cornell

[left to right] CS Professors Uri Keich, Bart Selman, Paul Francis, Dexter Kozen, David Gries (also associate dean for undergraduate programs in engineering), Robert Constable (also CIS dean), and Éva Tardos.

Computer Architecture and Very Large Scale Integration (VLSI)

Experts in computer hardware work closely with members of the School of Electrical and Computer Engineering, forming interdisciplinary research teams. Graduate students are admitted to either graduate field. They usually enroll in ECE if their primary interest is in computer architecture, multiprocessor design, VLSI, computer-aided design (CAD) and circuit design; they enroll in CS if their interest is in compilers, operating systems, and programming environments. Students are not rigidly classified: ECE students can have a thesis advisor in CS and vice-versa. Research is conducted in the CSL and 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.

For more information, see *http://www.cs.cornell.edu/Research/Architecture/* and *http://www.csl.cornell.edu*.

Artificial Intelligence

Research in AI at Cornell is both theoretical and experimental, and covers a range of topics, including decision theory, knowledge representation, machine learning and data mining, natural-language processing, planning, reasoning under uncertainty, search, and computer vision. Natural-language processing, the ability to interact with machines using plain English, is also a research area of special interest to researchers in information science. Computer vision encompasses both archival storage of images and the function of robotic eyes. Research in AI often involves such other disciplines as biology, economics, linguistics, medicine, operations research, physics, and psychology. A particular strength is in computer-intensive approaches to AI problems. Professors who work in AI are key participants in the university-wide Cognitive Studies Program (CSP) and in the IISI.

For more information, see *http://www.cs.cornell.edu/Research/ai/* and *http://www.cis.cornell.edu/iisi/*.

Computational Biology

Computational tools are being developed to analyze, understand, and manipulate newly available, richly-detailed information on the basic structures of life. Without advanced computing, biological databases with gigabytes of information would defy attempts to recognize patterns in them. Cornell has a university-wide initiative in the science of genomics, in which CS and others are taking part. CIS oversees interdisciplinary teaching and research in many aspects of computational biology. CIS initiated a novel graduate program in Computational Biology and Medicine in collaboration with Rockefeller University and Sloan–Kettering Research Institute. The Tri-institutional Program in Computational Biology and Medicine, now in its second year, builds on the strengths in the medical sciences in the New York City campuses and on the strengths in computer science at the Ithaca campus (for more information, see *http://www.cs.cornell.edu/grad/cbm/*).

Within CS, the focus is on computational molecular biology. This includes pattern matching, protein classification, protein-fold prediction, long-time simulation of protein dynamics and function, and structure-comparison algorithms. CS is leading a new project, "The Library of Life", in collaboration with the Bridging the Rift (BTR) Center, which will integrate (and create) data from classical and molecular biology of all living systems. The BTR Center also aims to promote peace in the Middle East and is located on the border between Israel and Jordan. The project will start with the creation of the Library of the Desert.

For more details, see *http://www.news.cornell.edu/features/BTR/*.

Another multiyear project seeks to create a database of the world's approximately six million proteins. (For more information, see *http://protomap.cornell.edu/* and *http://biospace.cornell.edu/*.)

For more information on computational biology at Cornell, see *http://www.cs.cornell.edu/Research/compbio/* and *http://www.cis.cornell.edu/cb/*.

Database Systems and Digital Libraries

Data-management topics range from efficient algorithms for very large data sets to large-scale systems for new and emerging applications. While traditional databases stored information numerically, today's scientific, journalistic, and other databases often store pictures and related nontextual data. Summarizing this information for efficient storage presents a big challenge. Researchers in database systems collaborate with computer-science colleagues who work on theory, artificial intelligence, and operating systems. The Himalaya Project is exploring new directions in data mining. The Quark Project is building a system for querying semistructured Internet sources, and is related to the XML Data-management Project, which is part of Interactive Digital Libraries of Formal Algorithmic Knowledge. The Pepper Project is building a query layer for large-scale peer-to-peer systems. The Cougar Project is developing database technology for sensor networks.

For more information, see *http://www.cs.cornell.edu/Research/datab/*.

With the excellence of our faculty and students, and their drive to discover and collaborate, our leadership will continue to grow.

Programming Languages and Compilation

Research in programming languages and compilation ranges from theory (including logic and semantics) to practical engineering issues (including verification, optimizing compilers, security, and run-time systems). Strong synergies exist with other subfields, for example in the development of secure programming languages, such as Typed Assembly Language (TAL) and Java Information Flow (Jif). Faculty members in this subfield work closely with the computer industry. One research group recently licensed program optimization tools to Intel for use in their iA-64 compiler product line. On-going research projects include Proof/Program Refinement Logic (PRL); several language-based security projects: Cyclone, a safe dialect of C; TAL (Typed Assembly Language); Jif; ECC (Efficient Code Certification); Polyglot, an extensible compiler framework for Java; JMatch, iterable pattern matching for Java; and Intelligent Systems Software.

For more information, see *http://www.cs.cornell.edu/Research/lang*.

Computer Graphics

Cornell is a leader in the field of computer graphics, a broad, interdisciplinary field that includes a wide and growing range of applications from science to communication to entertainment.

Research in computer graphics involves algorithms, physics, psychology, computation, computer vision, and architecture, among other fields. At Cornell, research and teaching in computer graphics are centered in CS and the closely affiliated PCG, one of the world's leading computer-graphics laboratories and a dominant force in the international computer-graphics community for more than thirty years. The PCG is particularly famous for its work in realistic rendering, simulating environments that are physically accurate and perceptually indistinguishable from real-world scenes.

The interests of the computer-graphics group are broadly centered on the topic of high-quality rendering. Current research thrusts focus on the interrelated topics of improving the models of light scattering that underlie realism, deepening our understanding of how human viewers perceive computer-generated images, and developing algorithms for high-quality rendering at interactive rates. Other areas of interest include image-based modeling and texturing, architectural modeling, animation, graphics-hardware programming, and digital photography.

The PCG's state-of-the-art facility includes many tools for advanced research, including a sophisticated light-measurement laboratory with unique capabilities for directional light measurement, a large PC cluster, and a high-resolution tiled projection display.

For more information, see *http://www.cs.cornell.edu/Research/graphics/*.

[left to right] CS Professors Thorsten Joachims, Dexter Kozen, Bart Selman, and Éva Tardos

Operating Systems, Networks, and Distributed Computing

The systems group at Cornell is concerned with the design and implementation of the fundamental software systems that constitute a computing infrastructure. Our interests span from the very small, such as the smart-card systems that fit on a postage stamp–sized die, to the very large, like the wide-area distributed systems that span the globe. Past research by the group in fault-tolerance, distributed communication, extensible systems, and Internet networking is widely cited and used. Concrete software artifacts developed at Cornell are used by the New York Stock Exchange and the French air traffic control system, and are also deployed in commercial operating systems.

Over the last year:

The Spinglass group (led by Ken Birman), which focuses on scalable group communication, has developed a series of protocols and systems for scalable information sharing, distributed monitoring, data mining, and control.

Paul Francis, who recently joined CS, has initiated the Firebreak project to add a ring of protection to the Internet Protocol that combats distributed denial-of-service attacks.

The Jif group (led by Andrew Myers) has shown that secure distributed systems can be constructed by automatically introducing encryption, partitioning, and replication to satisfy explicit policies for data confidentiality and integrity.

Robbert van Renesse has worked on scalable peer-to-peer systems, and has developed self-organizing protocols for high-throughput consistent data replication, aggregation of distributed information, failure detection, overlay multicast, and publish/subscribe.

Fred Schneider has led a group attempting to understand how fault-tolerance can be combined with attack-tolerance and has investigated threshold cryptography with proactive secret sharing in the Cornell Single Sign-On (CorSSO) and Cornell Online Certification Authority (COCA) authentication services.

The Beehive project (led by E. Gün Sirer) has introduced a novel replication technique that enables fast (O(1)) lookups in distributed, fault-tolerant, scalable hash-tables, and used this technique to build an alternative to the legacy Domain Name Service, a critical part of the cyber infrastructure used to translate Internet names to addresses.

For more information, see *http://www.cs.cornell.edu/Research/Systems/*.

Scientific and Parallel Computing

Scientists and engineers increasingly rely on computer models and simulation to buttress their experiments and designs. From the prevention of metal fatigue in engines to the design of new medicines, scientific and technological advances rely on powerful computers and the algorithms that run them. The scientific computing group develops the algorithms that underlie simulation and optimization. Matrix computations are a recurring theme in the research. The focus is on efficient and robust algorithms with an eye toward modern high-performance parallel and multithreaded architectures.

For more information, see *http://www.cs.cornell.edu/Research/scientif/*.

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 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.

For more information, see *http://www.cs.cornell.edu/Research/Security/*.

Theory of Computing

Cornell is a world leader in the theory of computing with fundamental research that spans the breadth of the field. Our faculty members and students have contributed to such areas as the design and analysis of algorithms, combinatorial optimization, computational complexity, computational algebra, program logic and semantics, automated deduction, and computational geometry. Cornell is also at the forefront in applying theoretical ideas to practical problems in areas such as machine learning, data mining, software and hardware verification, reliable systems, computer vision and medical imaging, computer graphics, programming languages and compilers, information science, and the computational sciences. Yet, while the theory of computing is by now recognized as a deep and foundational scientific discipline with broad impact, many of its central questions remain unanswered. As we enter the information age, continued progress on the fundamental theoretical questions will be essential for heightened understanding of complex systems and software, the nation's network and information infrastructure, and the physical and biological sciences.

For further information, see *http://www.cs.cornell.edu/Research/theory/*.

Common Cause

Computer vision scientists help radiologists see the big picture

Over the past four years, electrical engineer Anthony Reeves and computer scientist Ramin Zabih have each made the commute between Ithaca and New York City more than 100 times. But the two have a lot more in common than the thousands of miles each has logged traveling between the University's main campus and its Medical College in Manhattan. Both researchers are experts in computer vision, creating mathematical algorithms to analyze digitized images.

CS Professor Ramin Zabih

And through collaborations with clinicians in Weill Cornell Medical College's Department of Radiology, each has used his expertise to enhance the ability of physicians to accurately diagnose their patients.

Radiology is one of the most computerized fields in medicine, yet its practice lags far behind the sophisticated technology available. While the resolution and sheer number of images radiologists can obtain have increased dramatically, assessment has been less formalized. "The standard for measuring lung nodules today is that the radiologist puts calipers on a two-dimensional image, looks at the largest extent, and says, 'That's the size,'" says Reeves, who works in the field of lung cancer detection and diagnosis. "The technology is producing so many more images in so much more detail that the concept of having a human look at them when a machine has so many more advantages is simply impractical. It's a no-brainer." But first the images have to be transformed into numbers. "Once you have something in quantitative form," says Zabih, "then you can look for patterns."

The task of computer vision researchers is deceptively complex: it involves figuring out what the human mind does instantly and intuitively, and translating that process into mathematical language. "If you ask a person, is it easy or hard to do calculus, they say it's hard," says Zabih. "If you ask a person, is it easy or hard to count the number of people in a room, they say it's easy. But for a computer, the opposite is true." Ultimately, whether a computer vision researcher analyzes photographs of crowds, Computed Tomography (CT) scans, or Magnetic Resonance (MR) images matters little; the same strategies apply. In each of their projects, Zabih and Reeves are building an arsenal of what Zabih calls "power tools," algorithms that translate a radiologist's medical knowledge into a computer program to automate such tasks as maximizing the quality of an image or even helping to diagnose such conditions as lung cancer, aneurysms, and breast cancer.

In 1997, Reeves teamed up with Weill Cornell radiologists Claudia Henschke and David Yankelevitz to develop algorithms for analyzing CT scans of the lung. Henschke, the project's principal investigator, has long recognized the need for computer-aided techniques in the field. After earning a doctorate in mathematical statistics in 1969, Henschke consulted on clinical trials for the Veteran's Administration and the National Academy of Science. Then she decided to go to medical school. It was 1977, and CT, a technique that was the first to code diagnostic images into numbers, was just being introduced. "My whole idea for going into radiology was because I was a statistician and a computer programmer," says Henschke, now division chief of Weill Cornell's Chest Imaging. "I wanted to do things with CT numbers, and what is now known as CT image analysis."

With the advent of helical CT scanning in the early 1990s, she and Yankelevitz set out to demonstrate its superiority to traditional chest X-ray screening methods. For help, they contacted the Engineering college on the Ithaca campus, where they found Reeves. Their research focuses on automating the detection, measurement, and diagnosis of pre-cancerous nodules. "In our early days," says Reeves, "we pioneered the concept of trying to measure exactly where is the lesion, where is the blood vessel or chest wall attached to it, where is the joining point between the two. Now this is pretty well accepted as the way you would measure these nodules on the computer, and most CT manufacturers now have products that follow that strategy." What made the difference was Reeves's ability to see the problem from an engineer's perspective. "When we got together, my first reaction was, this is a three-dimensional problem," says Reeves. As a result, the algorithm the trio developed treated a CT scan not as a flat image but as a three-dimensional object consisting of a set of two-dimensional images. The result was a method for measuring the size of lung nodules that far surpassed anything available at the time.

The ability to accurately measure the size of tumors holds the potential to improve the cure rates of patients with lung cancer. "To date," says Reeves, "the best predictor of malignancy is rate of growth."

A. KATSNELSON
(reprinted from *Cornell Alumni Magazine,* May/June 2004)

Paving the Way

Scientists Lay the Groundwork for Cross-campus Collaborations

In recent years, Cornell has made fostering cross-campus research ventures a priority, one that ranks high on the agenda of President Jeffrey Lehman. "There's a growing interest in collaborations, not just between scientific disciplines, but also between science and the social sciences, between science and the humanities, public health, and outreach," says Vice Provost Lisa Staiano-Coico, PhD '81, who holds a joint appointment at both the Medical College and the University. "The more we highlight these collaborations, the more we can leverage our expertise and our strengths."

Seventy faculty members currently participate in twenty-one cross-campus collaborations, spanning areas of research from robotics and biochemistry to the social aspects of aging, but very few have been clinical in nature. Computer vision experts Anthony Reeves and Ramin Zabih have more than a decade of combined experience working with Medical College professors, and the challenges have been myriad. From harnessing technology to span the distance between researchers to funding graduate students and making joint applications for grants, Reeves and Zabih have consistently had to blaze a trail. "There was nothing there to make it easy," says Reeves. "There was no mechanism, at least when we started seven years ago, for these interactions."

A November 2003 study led by Weill Cornell cardiothoracic surgeon Nasser Altorki found that even minute differences in tumor size have a measurable effect on patient survival, and underscored the need for early detection. Yet nodules are usually detected only after they have grown past the point of easy surgical removal, leading to the historically low cure rate for the disease. With accurate nodule measurement, high-risk patients can be screened periodically, allowing radiologists to analyze the growth of a lesion over time and thus assess its malignancy. "Before, I would look at a scan and say, 'Well, this looks a little bit bigger,' and have to make the decision of whether to do a biopsy—an invasive procedure," says Henschke. "But this technique, as it gets better and better, will make us more confident in making that recommendation." Ultimately, the team hopes their algorithm might even supplant the need for biopsy altogether.

Unlike Reeves, who was sought out by clinicians, Zabih relied on serendipity to yield his collaborations. In April 2000, the day after the computer science department in Ithaca voted to grant him tenure, Zabih phoned Weill Cornell's Department of Radiology and requested permission to spend his sabbatical there. "I figured maybe at the end of the year they'd be interested enough in what I was doing to make some kind of collaboration."

Early on, Zabih met Martin Prince, director of Weill Cornell's MRI service. Despite MRI's spectacular resolution, it is highly prone to motion artifacts, caused both by quirks in the imaging technology and the natural movements of the human body. "Even though it's been around for fifteen or twenty years as an imaging modality, MRI is still relatively new,"

Continued on page 16

Particularly in bridging the gulf between clinical medicine and basic research, collaborative research depends heavily on proximity. "I'm a technical problem solver," says graduate student Ashish Raj, a member of Zabih's team. "And I can only solve a problem when I know about it. But to do that we need to interact more with the medical community."

Once theoreticians and clinicians have found one another, they must also overcome profound differences between how they approach research questions. "In computer vision there generally isn't one really good answer," says graduate student Amy Gale, another member of Zabih's research team. "If you offer a computer algorithm a picture of a hand and say, 'How many objects are in this?'—well, is there one hand or are there five fingers?" Computer vision scientists, she says, might conclude that a problem can't be solved and call it a day. "But there's no real point to saying, 'This is not a solvable problem' in medicine. You have to come up with something. If there's no perfect answer, you have to come up with a very good answer." Zabih, who has held a joint appointment between Ithaca and the Medical College since 2001, puts it another way: "Academics care about impressing other academics. Doctors couldn't care less about anything that doesn't work."

What has undoubtedly helped these collaborations bloom is the common technical backgrounds the New Yorkers and the Ithacans share. "Mathematics helps bridge the gap," says radiologist Claudia Henschke of her work with Reeves. MR specialist Martin Prince, Zabih's collaborator, earned a doctorate in engineering at MIT before going to medical school.

A more logistical challenge in cross-campus research has been getting Ithaca-based graduate students to New York City. While Reeves has developed an extensive Web-based system that has made it feasible for his students to connect with the radiologists electronically, the situation, says Henschke, is not ideal. "Coming and seeing the medical environment and the process gives them a different insight in terms of what is really valuable, how things need to be presented to the patient. What would help a lot is to have a fund so that students can regularly come down and spend a week or maybe a month working with us."

Zabih's students work for a significant amount of time in the city, where they have found themselves in the strange position of having to register as visiting students at a branch of their own university. Gale, for example, alternates semesters, spending the fall in Ithaca and the spring in New York. "I've never really been clear on what my relationship to the system is," she says. "I can read notices on elevator walls, but I'm not really in the loop." While she managed to secure housing in a Medical College dorm last year and was covered by the University's health plan in Ithaca, she wasn't eligible to be treated at the college's student health center.

The hurdles encountered by Zabih's students revealed an administrative gap that might not otherwise have been obvious. "We realized that the success of this collaboration depended heavily on graduate student and faculty exchange, and so we sought to eliminate some of the associated administrative burdens," says Staiano-Coico. This January, the Medical College announced the Cornell University Graduate Student Synergy (CUGSS) Program, which gives students who must spend at least a semester working at Weill Cornell essentially the same status as students in the Graduate School of Medical Science. The students now have access to Weill Cornell's graduate housing pool and student health services, as well as cultural events. They are also able to register and earn credit for courses. Similar plans are in the works for reciprocal exchanges of New York–based graduate students to the Ithaca campus. Staiano-Coico reports that initial feedback from faculty has been positive. "It makes their life easier," she says. "Now they have somebody to contact; they have a formalized process so that it's not reinventing the wheel every single time they want a student to come down."

says Prince. "And the tools that you need in order to figure things out haven't become commercially available." With Prince, Zabih has developed an algorithm that automatically corrects for motion artifacts in MR angiography, a non-invasive technique for visualizing blood vessels, and provides radiologists with significantly superior images. "For a very specific task," says Zabih, "it's actually as good as an expert radiologist."

Zabih has also partnered with Dr. Ruth Rosenblatt, director of Women's Imaging at Weill Cornell, to develop strategies to encode the biological properties of human tissue, thereby distinguishing potential tumors from surrounding flesh. After entering the MR scanner, patients are injected with a tiny amount of a harmless gadolinium-based compound that acts as a contrast agent as it passes through the body. While normal tissue releases the gadolinium evenly over time, tumors have "leaky capillaries," which cause the gadolinium to diffuse out of the tissue in a rush. By measuring the rate at which the contrast agent passes through the tissue, a radiologist could potentially differentiate benign from malignant tissue with the sweep of a computer mouse. "It's like having a spell-checker that recognizes a certain pattern," says Rosenblatt.

One of the biggest challenges facing Zabih and his Medical College collaborators has been the 240-mile distance between them. Reeves and Henschke's lung cancer group has minimized the distance by relying on the Internet, but for Zabih's student Ashish Raj, the miles are not virtual, but mind-numbingly real. Like Bill Murray's character in the movie *Groundhog Day,* who must repeat the events of one day in his life until he gets it right, Raj travels between Ithaca and New York perfecting his image correction algorithm, spending about one week each month in the city. At the Medical College, he works with Prince to identify scans where motion artifacts have occurred, then returns with the images to his home base in Ithaca. After refining the algorithm to improve the images, he goes back to the city, sits down with Prince again, and asks the radiologists to rate a "double-blind" assortment of corrected and uncorrected images. "Every time he gives a score I ask him, 'Why did you give it this score and not that score?' or 'Why did you think this image was better?'" says Raj. "The things he tells me are pretty much common sense: you want more contrast, more detail, less background. It's just that it really does help to sit down with him when he points out what he's looking for in each case."

Ultimately it is the interaction between the clinical and the theoretical that drives the researchers' insights. "We're trying to be scientists," says Rosenblatt, "but we're also clinical. By bringing two different points of view together, one will help the other." Such collaborations really work, says Zabih, when problems are tackled from both ends. A researcher can design a solution, but a clinician has to test it. " They'll push back on you," says Zabih. "They'll say, 'Hmm, it does a pretty good job here, but in the following circumstances it doesn't work. Do you have any ideas?' Then you can say, 'Let me work on it.'"

Speed of Light

Group Goes Online to Bridge Distance

One of the widest gulfs between Cornell's Ithaca campus and the Medical College is the 240 miles that separate the two. "I can't just go across the street and be in the medical community," says electrical engineering Professor Anthony Reeves. "I can't just attend a meeting. Logistics have to come into it." But having surmounted the initial challenge of finding one another, Reeves and Weill Cornell radiologists Claudia Henschke and David Yankelevitz have transformed distance from a hindrance into one of the driving elements of their work.

On the long wall of the conference room of Weill Cornell's Lung Cancer Screening Program hangs a map of the world, marked with nearly three dozen colored flags pinpointing locations across the United States and cities around the world—in France and Finland, China and the Philippines. Each represents a clinical trial site in the International Early Lung Cancer Action Project, I-ELCAP. Along with NY-ELCAP, a parallel program running in a network of medical centers throughout New York City, I-ELCAP investigates the benefits of CT scanning for lung cancer detection.

At one of their first meetings, Reeves watched Henschke filling out forms by hand. "We could put them on the Web," he told her. What developed is the group's online management system, a databank that allows the team to match specific scans with information from the clinical trials. So far, it contains more than 20,000 baseline CT images and as many repeat scans, which are compared to investigate lung health and lung cancer development in at-risk patients. Using the ELCAP database, Henschke and her team published a groundbreaking study in the July 1999 *Lancet* demonstrating that death rates from lung cancer could potentially be reduced from almost 80 percent to 20 percent if at-risk patients (smokers and former smokers) were regularly screened with CT scans.

"One of the great things about the ELCAP project is that we're facilitating the pooling of data," says Reeves. "By gathering information, we're able to address many more questions than we would be able to do with these studies individually." The system also guides the entire protocol by ensuring that procedures are as consistent as possible at trial sites around the world. By maintaining a fully standardized protocol in real time, says Henschke, "we've brought about a whole new paradigm on how studies can be done in the future."

Line in the Sand

Science Center Launched at Israeli–Jordanian Border

In early March, delegations from Cornell and Stanford, including presidents Jeffrey Lehman '77 and John Hennessy, traveled to a desolate 150-acre patch of land straddling the Israeli–Jordanian border. Here in the Arava desert, known as the Wadi Araba on the Jordanian side, they broke ground for the Bridging the Rift Center, a multi-million-dollar research facility intended to bring together American, Israeli, and Jordanian scientists to research the genetic code of all living things.

"This has the potential to be one of the biggest science projects of all time," says Ron Elber, an Israeli professor of computer science at Cornell and the center's director. "It could be bigger than putting a man on the moon." Politicians say the venture could also boost Israeli–Jordanian relations with the kind of cooperation they envisioned when the two countries signed a peace treaty a decade ago, and create an oasis of intellectual exchange in an otherwise politically volatile region. Ra'annan Gissin, a spokesperson for Israeli Prime Minister Ariel Sharon, calls the center a "major breakthrough" in grass-roots relations between Israel and Jordan, "and potentially, between Israel and the Arab world."

BTR takes its name from two sources: the physical separation caused by the Jordan Rift Valley that divides the two countries and the ideological discord between them. Inspiration for the center came in 2000 from Israeli businessman Mati Kochavi, who dreamed of a scientific endeavor to bring Jews and Arabs together. Kochavi, chairman of a holding company that invests in technology and energy companies, established the nonprofit in 2002 and has since raised "multi-millions" from private donors.

The three-part groundbreaking celebration in early March—Jordan's King Abdullah II and Israel's Sharon conducted separate ceremonies in Amman and Jerusalem to appease hard-liners in both countries—is a sign of how complex the start-up process has been. The project was kept secret until February, and even put on hold several times, out of fear that publicity about a joint Israeli–Jordanian initiative might create a backlash. Each country had to create new laws to allow the donation of land. "We are grateful that the governments of Israel and Jordan have taken the first steps to show how this collaboration can evolve," says Lehman.

In addition to cooperation from both governments, Kochavi wanted to involve top-tier U.S. schools, to help draw scientists from around the world and create a prestigious educational opportunity for Israelis and Jordanians. He was interested in agriculture; King Abdullah was keen on information technology. So Kochavi turned to Cornell and Stanford because of their strengths in these fields.

Then the brainstorming began: what would the center actually do? Kochavi wanted a field with growth potential, and one that would attract important scientists and provide fodder for spin-off industries. Information technology wasn't just the king's preference—"it was something easy to place anywhere and has global implications," says Cornell plant sciences professor Steven Tanksley, whose expertise fit well with Kochavi's goals for the center. Then Tanksley suggested what was to become the facility's centerpiece, an idea he called the "Library of Life"—a computer databank of genetic information on everything from humans and animals to plants and microbes. The center, say Cornell faculty, will surpass in importance GenBank, the database operated by the National Institutes of Health and part of an international collection of DNA information. That's because the Library of Life will not only record genetic codes but also incorporate digitized images and global positioning data to analyze how genes interact with and adapt to their environments and how species co-evolved, allowing scientists to make predictions at the genetic level.

Cornell President Jeffrey Lehman talks with Jordan's King Abdullah II.

CS Professor Klara Kedem *[left]* and CIS Dean Robert Constable *[right]* at King Abdullah's palace

Through BTR, scheduled to open in three to five years, Cornell and Stanford will offer doctoral degrees to Israelis, Jordanians, and others who will pursue course-work on the U.S. campuses and conduct their fieldwork at the center. Tanksley and other Cornell and Stanford experts in genetics, biology, ecology, and computer science will collaborate with Jordanian scientists from the University of Jordan in Amman, the University of King Hussein, and the University of Albalka; Israel is contributing scientists from Tel Aviv University, Hebrew University of Jerusalem, and the Weizmann Institute of Science. Jordan's Ministry of Education and Ministry of Higher Education are also involved. The schools do not have formal affiliations with the center because, says Kochavi, "we wanted neutral U.S. universities" to be the key players. Plans call for the center to accommodate about 150 people initially and eventually up to 1,000.

Now the academic committee, led by Stanford biology professor Marc Feldman, is divvying up duties: Stanford researchers will likely collect data while Cornell scientists will identify sampling needs and develop computer modeling. Organizers hope that their research will eventually foster spin-off enterprises, creating tangible medical and economic benefits such as new treatments for genetic diseases.

The Arava site fifty miles south of the Dead Sea was chosen for its relative safety and quiet, far from Israeli–Palestinian clashes on the West Bank and Gaza Strip. Additionally, says Feldman, the area is "perfect for researching what happens in extreme environments—places with high salinity and high temperatures." For Elber, the center's location has symbolic significance, too. "The Middle East is the place of the birth of civilization," he says. "Perhaps this could help the Middle East become the center of civilization again."

TAMAR MORAD

(reprinted from *Cornell Alumni Magazine,* May/June 2004: 30–31)

CIS Dean Robert Constable *[left]* shares a light moment with Bill Gates.

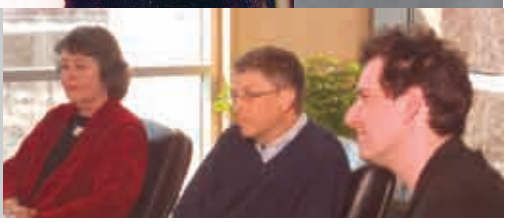

CIS Professors Geri Gay *[left]* and Hod Lipson *[right]* at a round-table discussion with Bill Gates.

Corporate Interactions

We realize that a true partnership results in mutual satisfaction and gain.

Gifts and Grants

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

Association for Behavior Analysis $5,500
Green Hills Software, Inc. $3,800
IBM $40,000
Intel Foundation $96,239
Microsoft Corporation $544,478
Sun Microsystems, Inc. $97,915
Verizon $57,500

CIS is grateful for gifts from the following partners:

Bloomberg L.P. $8,000
Bridging the Rift $40,000
Creative Realities, Inc. $1,004
Credit Suisse First Boston $7,000
Intel Foundation $92,060
Microsoft Corporation $316,467

PRINTER: silhouette this image- soften edge slightly

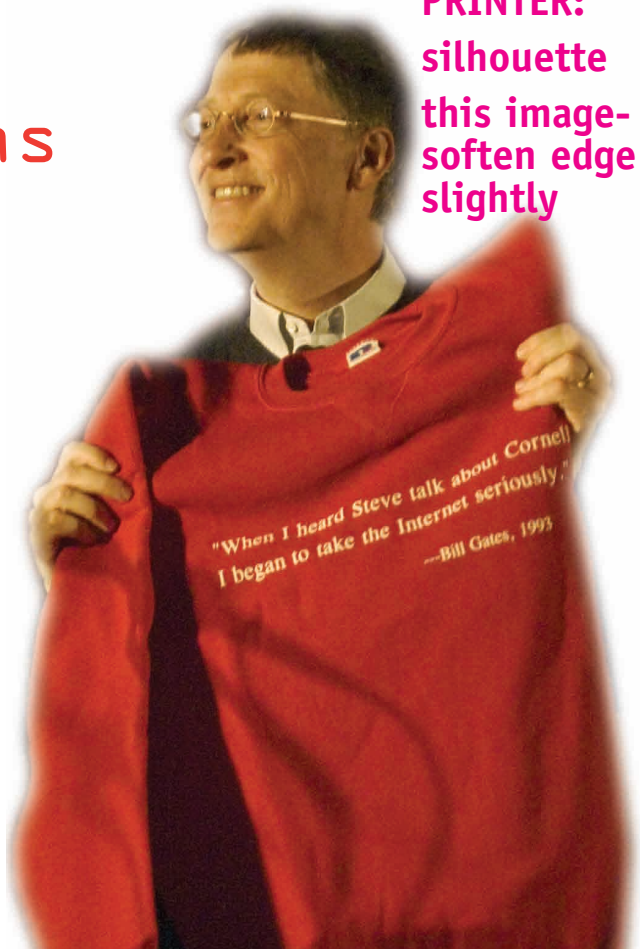

CS faculty members, researchers, and graduate students are conducting leading-edge research in many broad areas, including architecture, artificial intelligence, computational biology, databases and digital libraries, languages and compilation, multimedia and graphics, operating systems, networks and distributed computing, scientific and parallel computing, security, and theory of computing (see pages 10–12 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 partnering and 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, and organize recruiting activities on campus for CS undergraduate and graduate students.

CIS welcomes unrestricted donations by corporate partners in support of department initiatives, research grants to individual faculty members and researchers, matching funds to NSF or other granting agencies, creation of fellowships for graduate students, equipment grants, startup funds for new faculty members, 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 IAI and the IISI, 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.

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 sites: *http://www.cis.cornell.edu* and *http://www.cs.cornell.edu*.

CS at Cornell is ranked among the top six computer-science departments internationally, and includes

44	faculty members
2	Turing Award winners
3	members of the National Academy of Engineering
5	Guggenheim Fellows
6	faculty members who have received honorary doctoral degrees
6	National Science Foundation Young Investigators
1	Office of Naval Research Young Investigator
1	Air Force Office of Scientific Research Young Investigator
3	American Academy of Arts and Sciences Fellows
7	Sloan Fellows
1	Fulbright Scholar
1	New York State Council for Advancement and Support of Education (CASE) Professor of the Year
19	full-time research associates
110	Ph.D. candidates
100	M.Eng. candidates
200	undergraduate majors graduating each year

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

Creative Realities, Inc. provided support for research on culturally-embedded computing.

Credit Suisse First Boston and **Bloomberg L.P.** co-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** continued its support in funding two teams that participated in the New York regional ACM contest. The lead team went on to participate in the World Championship Competition.

IBM provided a 2003 Faculty Partnership Award for Professor Jayavel Shanmugasundaram.

Intel supported the computer graphics instructional lab, provided a fellowship for Ph.D. student Nathaniel Nystrom, and funded Professor Fred B. Schneider's language-based security project.

Microsoft supported Professor Johannes Gehrke's research, Professor David Schwartz's course on gaming, provided a fellowship for Ph.D. student Ranveer Chandra, and donated software to the department. They also supported Werner Vogel's research on "Distributed Systems Support for the Global Real-time Enterprise" and continued support for Professor Gün Sirer's research "Assuring the Security of Components in the .NET Framework". They also provided significant support for research being done by the Information Assurance Institute under the direction of Professor Fred B. Schneider.

Sun Microsystems, Inc. awarded Professor Golan Yona an Academic Excellence Grant and provided support for Professor Fred Schneider in his research on "Language-based Security for Malicious Mobile Code".

Verizon provided support for several Verizon Fellowships, as well as continued funding for the M.Eng. program to enhance efforts in faculty hiring.

The Association for Behavior Analysis provided support for Zoheb Sait's studies in the M.Eng. program.

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

The Faculty of Computing and Information Science (CIS) is Cornell's central home for computing and information research and education.

Programs of the Faculty of Computing and Information Science

Computational Biology

Computation has been become essential to biological research. Genomic databases, protein-structure databanks, MRIs of the human brain, and remote-sensing data on landscapes contain unprecedented amounts of detailed information that are transforming almost all of biology. Complex patterns, structures, and interactions raise fundamental and fascinating questions that can only be addressed using computational methods, and computational biologists are creating new tools to analyze these new sources of data.

More than fifty Cornell professors in six university colleges, including the Weill Medical College in New York City, are involved in computational biology. They tend to be quantitatively-oriented biologists, biophysicists, doctors and medical researchers, mechanical engineers, and others, and they are looking at everything from the genetic differences between humans and chimpanzees, to how protein chains "fold" into three-dimensional structures, to how hearts pump.

Problems investigated by computational biologists span a wide spectrum, including:

- the genetics of disease susceptibility
- comparing entire DNA genomes to uncover the secrets of evolution
- using protein structures to design new therapeutic drugs
- mathematical modeling of cellular signaling networks
- predicting how ecosystems will respond to climate change
- designing recovery plans for endangered species
- animal behavior and ecology
- cognitive psychology and neurobiology
- modeling the spread of diseases

Research at Cornell in these areas usually involves interdepartmental collaborations that take advantage of Cornell's great breadth in the biological sciences and strength in computer science, mathematics, and the physical sciences. For example:

Researchers in the Departments of Computer Science, Molecular Biology and Genetics, and Chemistry and Chemical Biology focus on macromolecular biology. Predicting the structures and functions of proteins is a primary aim of research supported by the National Science Foundation. The group is strong, well organized, and has ties to the Computational Biology Service Unit in the Cornell Theory Center.

The group studying evolutionary genomics includes researchers from the Departments of Biological Statistics and Computational Biology, Molecular Biology and Genetics, and Ecology and Evolutionary Biology, and is one of the strongest of its kind in the country. This group, like the macromolecular biology group, has strong ties to Cornell's Genomics Initiative, which has computational and statistical genomics as one of its major thrust areas.

Studies of nonlinear systems bring together ecologists from the Departments of Ecology and Evolutionary Biology and Natural Resources, neurobiologists from the Department of Neurobiology and Behavior, and mathematicians from the Center for Applied Mathematics for increasingly sophisticated analyses of some the highly complex coupled systems that arise in biology.

Intracellular signaling networks are studied by a collaboration involving biologists, physicists and engineers from the Departments of Chemistry and Chemical Biology, Physics, and Chemical and Biomolecular Engineering, and the Cornell Theory Center.

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

Computational Science and Engineering

Many of the faculty members in engineering and the sciences engage in research that is computationally driven. 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 Dr. Andrew J. 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 attract students from across the university.

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

Spanning the entire campus, CIS serves to bring together experts in computing with researchers and scholars in a variety of disciplines.

Digital Arts and Graphics (DA&G)

Computer graphics is a rapidly evolving field that has had significant impact on most scientific, artistic, and engineering fields, and graphics is today the most common and efficient means of man–machine communication. At Cornell, research and teaching in computer graphics are centered in CS and the closely affiliated PCG, one of the world's leading computer-graphics laboratories and a dominant force in the international computer-graphics community for more than thirty years.

The field of computer graphics relates to the broader area of digital arts and graphics, a new CIS program area that involves computer-graphics researchers, architects, artists, art historians, city planners, and information scientists. This program will explore the interaction of graphics and arts by considering the technical aspects of digital graphics, psychological aspects of vision and perception, and the creation of art in a time of digital reproduction.

Research in graphics requires a multidisciplinary team with knowledge in algorithms, systems, numerical simulation, machine vision, software and hardware engineering, physics, optics, and perception psychology. The PCG is particularly famous for its work on realistic rendering—simulating environments that are physically accurate and perceptually indistinguishable from real-world scenes. Current graphics research at Cornell includes realistic interactive rendering, advanced material modeling, human visual perception in graphics, modeling complex scenes, image-based modeling and rendering, animation, and display technology.

For more information, see *http://www.graphics.cornell.edu/*.

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, 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 has been approved 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 in the coming year.

Students in the program will 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.

For more information, see *http://www.infosci.cornell.edu*.

CS Professor Gün Sirer *[left]* with a student at BOOM '04.

CIS Professor Geri Gay *[right]* looks on as a student describes her project at BOOM '04.

Bits on Our Minds (BOOM)

BOOM

BOOM explodes with student invention

With soccer-playing robots downstairs and computers that can play chess upstairs, this year's eighth annual BOOM (Bits On Our Minds) exhibition looked like something out of "The Jetsons."

Some 120 presenters with a total of 64 projects crowded three floors in Cornell's Upson Hall on March 3 to partake in the annual expo hosted by the Department of Computer Science and the School of Electrical and Computer Engineering.

"We have this expo every year for two reasons," said Emin Gün Sirer, assistant professor of computer science and faculty coordinator for BOOM. "We want to reach out to undecided majors and to people who are not in college yet to show them the opportunities computer science holds. We also do it as a teach-in, to show colleagues what the cutting-edge research is."

Irene Chung's project is an example. Chung '04, college scholar, displayed her Web concept managing tool, a program that allows the user to generate many different Website styles for the same information. Chung already has sold her program to Production IG, a prominent animation firm in Japan. Production IG used Chung's technology to configure the animation in the recent feature film "Kill Bill."

"I have a lot of work to still do with this program," Chung said. "But I think what is going on right now with the project is awesome, and I'm glad a company likes it."

The RoboCup team drew a big crowd. The team's soccer-playing robots operate on artificial intelligence programs that team members write. "Every year we build the project from the ground up," said graduate student Nathan Pagel, M. Eng. "It is great to see the robots through to completion." Cornell's team "rules the world in robotic soccer," Sirer said, referring to the university's success in competition.

The theme of robotics was common to many projects, including a robot built by Ithaca High School students for the FIRST ("For Inspiration and Recognition of Science and Technology") Robotics competition. Cornell undergraduates work as mentors to the Ithaca High students. This year's team will compete in the second round competition in Toronto in April.

Not all projects came from the computer science school of thought. Lindsay Lyman-Clarke, a graduate student in textiles and apparel, displayed her project, which uses body scan data to design clothing sizing systems. She uses lasers to acquire body scan data, in conjunction with computer patterning programs, to design clothing.

Above: CIS Dean Robert Constable and CS Professor Gün Sirer talk with Kathy Okun, wife of Cornell president Jeffrey Lehman at BOOM.

BOOM was sponsored by Bloomberg and Credit Suisse First Boston, which now has sponsored BOOM for three years in a row.

"We see BOOM as a way to tap into the talent here at Cornell," said Carolyne Phillips of Credit Suisse First Boston.

Other BOOM projects ranged from computer animations to a self-guided submarine to types of computer games.

"It is wonderful to see what students are doing," said Kathy Okun, wife of President Jeffrey Lehman. "There are lots of interesting things going on in computer science, and it's so great that students have the opportunity to do these things."

RACHEL EINSCHLAG '04

(reprinted from *Cornell Chronicle*, March 11, 2004)

Alumni and External Relations

In the fall of 2003, Cornell President **Jeffrey Lehman '77** was inaugurated as the eleventh president of Cornell University. More than 5,000 people filled Barton Hall, which had been brilliantly transformed into a stately venue, complete with massive red columns, Cornell banners, and acres of red carpet to honor the faculty, visiting dignitaries, and many special guests from around the world.

Lehman, a self-confessed technophile, has enthusiastically embraced the importance of computing and information science to Cornell's future. Following his inauguration, he visited a broad range of high-tech companies, meeting alumni and asking for feedback on host of issues important to Cornell.

At **Microsoft** in Redmond, Washington, more than forty enthusiastic alumni sat down for lunch with the new president and talked with him about Cornell's overall mission and his own interests. The alums at Microsoft asked President Lehman to outline his main challenges as a new president, to share his understanding of current endowment levels for student financial aid, to talk about his personal passions, and to comment on the role corporations can play in supporting academia and research. Lehman talked about his interest in Cornell's unique history in the world and its reputation for providing a more adventurous, path-breaking form of leadership.

The CS ACM Programming Contest Team jumps for joy before heading off to the World Championship in Prague. The team was sponsored by Green Hills Software.

Looking forward, he wanted to know what ideals alums thought Cornell should aspire to and vowed to bring high-profile speakers to campus to further engage students and faculty members in this important discussion. Less than three months later, Microsoft's own **Bill Gates** answered the call and arrived at Cornell to give a talk.

On the day of Lehman's inauguration in Ithaca, a special lecture was given by **Narayana Murthy,** founder and CEO of **Infosys,** a worldwide leader in software development. Murthy, a newly appointed trustee of the university, gave an inspired speech, calling on Cornellians to apply their talents to solving global problems. He spoke of the power of what he called "old-fashioned leadership," describing this as direction based on a vision of higher principles. He shared his belief that average people can be motivated to become an extraordinary force for positive change and urged Cornellians to carry their values to every corner of the world. He also encouraged fellow corporate leaders to earn the trust of society, to relate more effectively to the universal needs of people, and to strive to lead simple lives.

Speaking of socially responsible leaders, we are pleased to announce that former ACSU president **Jordan Erenrich '02, M.Eng. '03,** has agreed to serve as the chair of the **Computer Science Alumni Association**. Jordan plans to create an alumni listserv and is thinking about other ways to enable computing and information science graduates of Cornell to network with each other. If you have an interest in reaching out to other alums, contact Jordan at *jordan.erenrich@citigroup.com*. The alumni breakfast is a great place to get reacquainted and discover opportunities for outreach. At the last Cornell Reunion, a record nineteen alums joined CS chair Charles Van Loan for breakfast.

On another important front, the newly designated **"Game Design Initiative at Cornell University"** (GDIAC) is thriving under the leadership of CS alums **Mohan Rajagopalan M.S. '02** and **Rama Hoetzlein '01,** and Professor **David I. Schwartz**. GDIAC provides a great example of how alumni can stay connected; without Rama and Mohan, GDIAC would not exist as it does today. GDIAC is introducing its first official course in computer-game design, CIS 300 Introduction to Computer Game Design, which will be showcased at the end of each semester in a public open house. You can download student games and projects, learn about our courses, and get the latest news about game R&D at *http://www.cs.cornell.edu/projects/game/*.

In October of 2003, CS faculty hosted the first talk of a newly developed Distinguished Career Lecture Series. The inaugural talk, "Magnificent Pathways", was given by **Barbara Liskov,** Ph.D., a senior faculty member in computer science at M.I.T. Dr. Liskov shared the details of her distinguished career with an audience of fifty enthralled students, faculty members, and alums. One of the highlights of her talk was hearing that Princeton had refused to consider her for a faculty position because she was a woman—thankfully times have changed!

In November, a panel of women leaders addressed a group of about thirty-five students, faculty members, and alums at an event titled "Opportunities in Computer Science". Ph.D. student **Alexa Sharp** was engaging as the panel moderator, encouraging attendees to think of effective ways to communicate with students about computing-related studies and professions. Cornell alumna **Aleta Ricciardi '84, M.S. '89, Ph.D. '93** and CS major **Radha Narayan '05** were two of four panelists discussing the various opportunities available in research, graduate study, and careers for women in computing-related fields, and effective programs for mentorship and outreach. Also on the panel were CS Professor **Kavita Bala** and **Veronica Vazquez '96,** the dean of students of the Cascadilla School in Ithaca. Other alums attending this event were CS Professor **Daisy Fan, Ph.D. '03, Dan Jenkins '82,** and Ph.D. candidate **Vicky Weissman '96, M.Eng. '99**.

Away from campus, Cornell students had another successful year competing in the Association for Computing Machinery (ACM) Programming Championships, thanks to **Nikola Valerjev '96,** who once again secured generous sponsorship from **Green Hills Software, Inc.** Worldwide, 3,150 teams competed, with only seventy-three advancing to the finals. The Cornell team won first place in the northeast regional competition and moved on to Prague, Czech Republic to compete in world finals, where they were awarded an honorable mention. Green Hills has sponsored Cornell teams for several years.

This year's **Degenfelder Family Scholarship** was awarded to **Paul Shafer '05.** This $5,000 award recognizes students who are working at the interface between computer science and biology. **Joseph R. Degenfelder '60** and his wife **Dr. Pauline Degenfelder '61** worked with CS Professor Ron Elber to establish an endowment for this scholarship.

The **Jonathan E. Marx Senior Prizes** were presented to **Omar Ahmed Nayeem** and **Ramona Pousti** as part of the CS graduation ceremony in the Statler Auditorium on May 30, 2004. **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 **Aaron Justin Fink.** The Computer Science Prize for Academic Excellence was awarded to **Matthew Wachs** and **Asher Walkover.**

For more information about alumni or external relations in CIS or CS, please contact Dan Jenkins at *jenkins@cs.cornell.edu*.

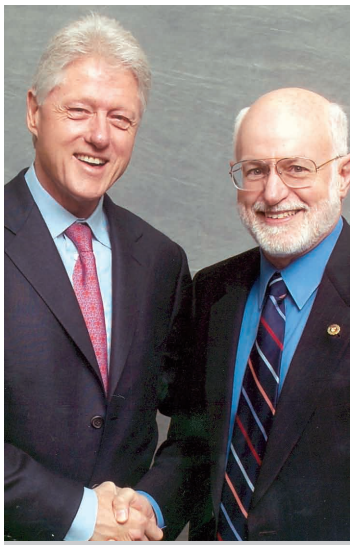

Bill Clinton visited the Cornell campus in May of 2004 to give the Convocation Address to graduating seniors. He is pictured here with CIS Dean Robert Constable.

Courses

CIS encourages the development of new courses for undergraduates and graduate students in the areas of computational biology, information science, digital arts and graphics, and computational science and engineering.

More than a dozen departments now cross-list courses with computer science.

CIS Courses Offered in 2003–04

CIS 130	Introductory Web Programming (also COM S 130)
CIS 191	Media Arts Studio I (also ART 391, THETR 391)
CIS 230	Intermediate Web Design (also COM S 230)
CIS 330	Applied Database Systems (also COM S 330)
CIS 387	The Automatic Lifestyle: Consumer Culture and Technology (also ST&S 387)
CIS 401	Introduction to Applied Scientific Computing with MATLAB
CIS 402	Scientific Visualization with MATLAB
CIS 403	Development of Scientific Computing Programs
CIS 404	Survey and Use of Software Libraries for Scientific Computing
CIS 409	Data Structures and Algorithms for Computational Science
CIS 430	Information Discovery (also COM S 430)
CIS 431	Web Information Systems (also COM S 431, formerly CIS/COM S 502)
CIS 490	Independent Reading and Research
CIS 515	Culture, Law, and Politics of the Internet
CIS 518	Computer Animation (also ART 372)
CIS 630	Representing and Accessing Digital Information (also COM S 630)

CS Courses Offered in 2003–04

COM S 099	Fundamental Programming Concepts
COM S 100j	Introduction to Computer Programming
COM S 100m	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	Web Design and Programming
COM S 165	Computing in the Arts
COM S 167	Visual Imaging in the Electronic Age
COM S 172	Computation, Information, and Intelligence
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 215	Introduction to C#
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 400	Science of Programming
COM S 401	Applied Scientific Computing with MATLAB
COM S 402	Scientific Visualization with MATLAB
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 421	Numerical Analysis
COM S 426	Introduction to Computational Biology
COM S 427	Practicum in Computational Biology
COM S 428	Intro to Computational Biology
COM S 430	Information Discovery
COM S 431	Web Information Systems
COM S 432	Introduction to Database Systems
COM S 433	Practicum in Database Systems
COM S 465	Computer Graphics I
COM S 467	Computer Graphics II
COM S 468	Computer Graphics Practicum
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 480	Introduction to Cryptology
COM S 481	Introduction to Theory of Computing
COM S 482	Introduction to Analysis of Algorithms
COM S 483	Quantum Information Processing
COM S 486	Applied Logic (also MATH 486)
COM S 501	Software Engineering
COM S 504	Applied Systems Engineering
COM S 505	Applied Systems Engineering II
COM S 513	System Security
COM S 514	Intermediate Computer Systems
COM S 519	Computer Networks
COM S 522	Computational Tools and Methods for Finance
COM S 565	Computer Animation
COM S 572	Heuristic Methods for Optimization
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 611	Advanced Programming Languages
COM S 612	Compiler Design for High-Performance Architectures
COM S 614	Advanced Systems
COM S 615	Peer-to-Peer Systems
COM S 621	Matrix Computations
COM S 622	Numerical Optimization and Nonlinear Algebraic Equations
COM S 624	Numerical Solution of Differential Equations
COM S 626	Computational Molecular Biology
COM S 627	Computational Biology: The Machine Learning Approach
COM S 630	Representing and Accessing Digital Inform
COM S 632	Advanced Database Systems
COM S 664	Machine Vision
COM S 665	Advanced Rendering
COM S 667	Physically Based Rendering
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	Algorithmic Game Theory
COM S 686	Logics of Programs
COM S 709	Department Colloquium
COM S 711	Seminar in Advanced Programming Languages
COM S 715	Seminar on PRL
COM S 717	Programming for Fault Tolerance
COM S 718	Topics in Computer Graphics
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 754	Systems Seminar
COM S 772	Seminar in Artificial Intelligence
COM S 775	Seminar on Natural Language Understanding
COM S 778	Topics in Machine Learning
COM S 786	Introduction to Kleene Algebra
COM S 789	Seminar in Theory of Algorithms and Computing

CS Undergraduate, Eugene Lee, Places First in Intel Student Research Contest

CS Professor Kavita Bala with CS undergraduate Eugene Lee

Eugene Lee won first place in the Intel Student Research Contest, which took place in Santa Clara, California, for his project, "Hardware Acceleration of the Edge-and-point System for Interactive Rendering". The Intel Student Research Contest was created to "stimulate inventiveness by challenging students in the sciences and engineering disciplines to explore frontiers of future computing".

Eugene worked with CS Professor Kavita Bala and her research group on a hardware-acceleration implementation of an interactive, high-quality rendering system—the edge-and-point rendering system.

The results of Bala's research on the edge-and-point rendering system co-authored with Bruce Walter and Program of Computer Graphics Director Donald Greenberg appeared at the ACM SIGGRAPH 2003 Conference. Computing realistic shading that captures the complex lighting of the real world is extremely expensive and therefore is traditionally associated with slow, offline applications. The edge-and-point rendering system makes it possible to achieve this high-quality rendering at interactive rates. The key insight is that the human visual system is sensitive to sharp changes in shading; this insight enables efficient rendering and display.

Lee's research uses the graphics processor on a desktop computer to implement the edge-and-point rendering system on modern graphics processors. This hybrid rendering system implements key edge-respecting image-reconstruction kernels on the graphics processor.

Lee is a native of Los Angeles, and an accomplished photographer. He interned with Microsoft in both Korea and Washington, and intends to work for Microsoft after graduation.

Funded Research Grants

Investigator	*Sponsor*	*Award*	*Title*
Birman	AFOSR	427,608	A Testbed for Highly-scalable Mission-critical Information Systems (DURIP)
Birman	DARPA/AFRL	395,000	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
Birman/Gehrke/Francis	DARPA/AFRL	1,098,321	QuickSilver: Middleware for Scalable Self-regenerative Systems
Cardie	MITRE Corp	61,743	ARDA NRRC Workshop
Cardie	NSF	500,000	Reducing the Corpus Annotation Bottleneck for Natural Language Learning
Cardie	ARDA sub of U Pittsburgh	463,368	Opinions in Question Answering
Caruana	NSF	507,000	CAREER: Meta Clustering: Improving User Efficiency on Real-world Clustering Applications
Caruana/Joachims	NSF	270,000	Optimizing Classification Models to Application-specific Performance Metrics
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	NSF	460,000	Enabling Large-scale Coherency Among Mathematical Texts in the NSDL
Constable	ONR	1,938,148	Building Interactive Digital Libraries of Formal Algorithmic Knowledge
Constable	NSF	300,000	Innovative Programming Technology for Embedded Systems
Constable	AF sub of ATC–NY	29,000	SCorES, A Logical Programming Environment for Distributed Systems
Constable	Boeing (McDonnell–Douglas)	161,857	Model-driven Generation of Verifiable Distributed Real-time 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
Elber	NIH	1,113,170	Long Time Dynamics of Biomolecules
Elber	ACS–PRF	80,000	Molecular Dynamic Simulations of EPR Spectra in Proteins
Elber	BTR	40,000	Library of Life
Francis	NSF	496,421	SP: Very Fine-grained Proximity Addressing
Gehrke	NSF	235,000	Interactive and Online Data Mining
Gehrke	NSF	210,000	Scalable Decision Tree Construction
Gehrke	NSF	340,000	CAREER: Towards Sensor Database Systems
Gehrke	NSF	500,000	SENSORS: Data-driven Sensor Networks
Gehrke	Sloan	40,000	Sloan Research Fellowship
Gehrke/Demers	KD–D thru NSF	1,193,000	Distributed Mining and Monitoring
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
Gomes	AFRL	5,000,000	Intelligent Information Systems Institute
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	300,000	ITR: Emerging Communities in Large Linked Networks: Theory Meets Practice
Joachims	NSF	400,000	CAREER: Improving Information Access by Learning from User Interactions
Kleinberg	Packard	625,000	Algorithmic Methods for Networks
Kleinberg/Lee/Cardie/Selman	NSF	450,000	ITR: Construction and Analysis of Information Networks

Investigator	*Sponsor*	*Award*	*Title*
Kozen	NSF	210,000	Kleene Algebra
Kreitz	NSF	285,000	Proof Automation in Constructive Type Theory
Lagoze	NSF	255,700	Metadata for Resource Discovery of Multimedia Digital Objects Harmony
Lagoze	U Virginia/Mellon Found	495,325	The Open Source FEDORA Repository Development Project
Lee/Kleinberg	NSF	449,897	Graph-based Approaches to Text Processing
Lee	Sloan	40,000	Sloan Research Fellowship
Marschner	NSF	400,000	CAREER: Modeling the Properties and Appearance of Materials
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	DARPA sub of IBM	255,000	PERCS: Phase II
Pingali/Stodghill	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
Pingali/Rugina	NSF	850,000	A System for Semi-automatic Application-level Checkpointing of Parallel Programs
Schneider	AFOSR	592,657	CIPIAF for Information Assurance Institute
Schneider	AFOSR	4,138,325	AFRL/Cornell Information Assurance Institute
Schneider	Microsoft	300,000	Information Assurance Institute
Schneider	Sun	20,000	Language-based Security for Malicious Mobile Code
Schneider	Intel	36,500	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
Schwartz	Microsoft	25,000	Game Project
Selman	Sloan	35,000	Sloan Research Fellowship
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
Shanmugasundaram	NSF	406,750	CAREER: Towards Unifying Database Systems and Information Retrieval Systems
Shanmugasundaram	IBM	40,000	2003 Faculty Award
Shanmugasundaram/Gehrke	AFRL	41,000	ORIS: Peer-to-peer Object Repository with Integrated Security
Sirer	Microsoft	75,000	The Ad Hoc Classroom: Integrating Emerging Wireless Communications and Networking Technologies into Mainstream Computer Science and Electrical Engineering Curricula
Tardos	NSF	150,000	Approximation Algorithms and Applications in Network Games
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	35,000	High-performance XML/SOAP Processing for Pub/Sub Engines
Vogels	Microsoft	20,000	CollabNet
Vogels/Gehrke/ Shanmugasundaram	Microsoft	175,000	Distributed Systems Support for the Global Real Time Enterprise
Yona	NSF	1,103,917	CAREER: Global Self-organization of All Known Proteins—Toward a Complete Map of the Protein Space
Yona	GNS	13,171	Predicting Protein–Protein Interactions

Total Expenditures for Fiscal Year 2003–2004: $15,063,346

Funded Collaborative Research at Cornell

Investigator	*Sponsor*	*Award*	*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/Comm	NSF	8,745,453	Collaborative Project: Core Integration—Leading NSDL Toward Long-term Success
Caruana/CTC	NSF	379,999	Multi-algorithm Parallel Optimization of Costly Functions
Coleman/CTC	TG Information Network	1,065,000	Financial Engineering and Tools
Coleman/CTC	EPA	28,440	Design, Setup, and Testing of High Performance Computing Cluster to Support Water Modeling Needs at EPA
Coleman/CTC	Aximetric Inc.	30,000	Aximetric Inc. Collaboration Agreement
Coleman/CTC	USDA	232,180	Computation Agriculture Initiative
Demers/Huttenlocher/ Kleinberg/Marschner /Astronomy	NSF	1,799,183	Petabyte Storage Services For Data-driven Science
Department/CTC/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/MBG	NSF	899,000	Multiscale Hierarchical Analysis of Protein Structure and Dynamics
Elber/Joachims	NIH	1,043,276	Optimization of Folding and Threading Proteins
Elber/BSCB	anonymous	361,113	Two-track Program in Computational Biology and Medicine as a Part of the Tri-institutional Research Program
Guckenheimer/Math/ Physics/CCB/MBG/ Biomedical/Mole Med/ Eng/Soc	NSF	3,371,200	IGERT: Program in Nonlinear Systems
Lagoze/Comm/Olin Library	NSF	2,425,899	Security and Reliability in Component-based Digital Libraries
Pingali/CTC	NSF	1,500,000	CISE Research Infrastructure: A Two-tier Computation and Visualization Facility for Multiscale Problems
Pingali/Vavasis/Chew/CTC/Phy	NSF ITR	5,035,425	Adaptive Software for Field-driven Simulations
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
Tardos/Kleinberg/ Huttenlocher/Halpern/ OR&IE/Econ	NSF	2,468,677	ITR: Networks of Strategic Agents: Theory and Algorithms
Vavasis/CEE	NSF	500,000	MSPA–MCS: Automatic Geometric Simplification
Yona/Bio Sci	NSF	1,000,000	CRCNS: Modeling Pathfinding and Target Recognition in the Olfactory System

Submitted Grant Proposals

Investigator	*Sponsor*	*Award*	*Title*
Constable	NSF	1,298,766	Designing for Accountability
Demers/Gehrke/Riedewald	NSF	945,634	A Formal Approach to Data Stream Management Systems
Elber	NIH	500,000	A Computer Cluster for Computational Biology
Elber	NIH sub of U Pittsburgh	58,305	Computational Prediction of Biomolecular Dynamics
Francis	Cisco	99,930	Solving BGP Scaling Through Tunneled Virtual Super-prefixes
Francis	NSF	658,599	NeTS NR: Firebreak: An IP Perimeter Defense Architecture

Investigator	*Sponsor*	*Award*	*Title*
Francis	Telcordia	175,491	Network and Service Architecture for DTN
Gehrke	DOE	300,000	Analyzing Spatio-temporal Simulation Data
Gehrke	NSF	886,849	ITR: (NHS)—(dmc): On Privacy-preserving Data Mining
Ginsparg	NSF	796,395	Classification, Analysis and Navigation Tools for Physics Research Communication
Halpern	ONR sub of U Penn	189,000	Trustworthy Infrastructure, Mechanisms, and Experimentation for Diffuse Computing
Halpern	AFOSR	419,853	Reasoning About Authorization and Security
Keich	NSF	803,058	SEI: Collaborative Research: Designing Algorithmic and Statistical Tools for Computational Discovery of Transcription Factor Binding Sites
Lagoze/Payette	U Virginia/Mellon	711,763	The Open Source FEDORA Repository Development Project: Phase II
Lagoze/Warner	NSF	2,400,000	II: Pathways
Marschner	NSF	324,629	Collaborative Research: Efficient Representation, Acquisition, and Rendering with Measured Reflectance in Computer Graphics
Myers	ONR	300,000	Using Information Flow To Construct Trustworthy Distributed Systems
Myers/Birman/Schneider	NSF	2,233,524	Integrating Security and Fault Tolerance in Distributed Systems
Rugina	NSF	430,838	Region-based Memory Management for Pointer-based Programs
Selman	DARPA sub of SRI	299,516	Architectures for Cognitive Information Processing
Tardos	ONR	534,544	Network Games and Approximation Algorithms

Submitted Collaborative Research at Cornell

Investigator	*Sponsor*	*Award*	*Title*
Birman/Schneider/Sirer/ECE	NSF	3,200,000	Team for Research in Ubiquitous Secure Technology (TRUST)
Birman/Sirer/ECE	NSF	2,206,999	NeTS–NOSS: Ultra Low-power Self-configuring Wireless
Caruana/ECE	NSF	542,378	Machine Learning Approaches to Understanding and Adapting Computer Architecture and System Behavior and Performance
Caruana/Riedewald/Lab of Ornith	NSF	2,701,622	ITR—(ASE+EVS)—(dmc+sim): Tracking Environmental Change through the Data Resources of the Bird-monitoring Community
Chew/Vavasis/M&AE	NSF	1,467,585	ITR—(ASE)—(sim+dmc): Collaborative Research: Algorithms for Large-scale Simulations of Turbulent Combustion
Coleman/CTC	USDA	207,085	Computation Agriculture Initiative
Elber / Gehrke/Plant Science	Library of Life	12,196,109	Bridging the Rift
Elber/Kedem/Yona/CCB	NIH	3,358,894	BioComUnity: NCBC for Systome Research
Marschner/Bala/PCG	NSF	3,577,319	ITR—(ASE+NHS)—(sim): The Visual Turing Test: Realistic Image Synthesis for Future Displays
Pingali/Gomes/CTC	NSF	3,972,702	ITR—(ASE)—(dmc+sim+int): Agile Anomaly Detection and Dynamic Evaluation of Responses
Sirer/ECE	NSF	3,130,996	IGERT: Mathematics, Physics, and Systems for Large Scale Sensor Networks
Tardos/OR&IE	NSF	1,909,194	ITR: Collaborative Research—(ASE+EVS+NHS)—(soc+dmc+int)—Fairness Issues in Heterogeneous Systems: an Algorithmic Approach
Tardos/Cardie/Huttenlocher/ Econ/OR&IE/Comm	NSF	2,722,795	IGERT: Doctoral Training in Information Science
Yona/multiple depts	NSF	425,638	Acquisition of Instrumentation for Establishing a Microarray Core

Connections with the Cornell Theory Center

CTC Director and CS Professor Thomas Coleman *[right]* with CTC Associate Director and CS Professor Keshav Pingali and CS Professor Johannes Gehrke

The 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 campus and at the Weill Medical College in New York City.

Through a strategic partnership with Microsoft and Intel, CTC has pioneered the use of industry-standard computational resources running Windows™ to create 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 SQL Server into complex engineering applications and the application of .NET and Web services to high-performance computing. CTC operates the first multiwall Windows/Intel computer-aided virtual environment (CAVE), providing a three-dimensional, stereo-immersive, virtual-reality environment for viewing scientific, engineering, architectural, and art applications. The CAVE enhances a variety of research projects and courses at both Cornell and Syracuse University, including an architectural course and an engineering design course that is taught in conjunction with the NASA/New York State/AT&T–sponsored Advanced Interactive Discovery Environment for the Engineering Education Project.

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 senior 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 lifesciences, ranging from canine genetics and 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 conduct 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.

The Computational Materials Institute at CTC focuses on fracture mechanics and serves as one of the test beds for the Adaptive Software Project (ASP) led by CS Professor Keshav Pingali, who is also a CTC associate director. ASP is developing software systems that can adapt to changes at the application, algorithmic, and systemic levels.

CS Professor Johannes Gehrke is also involved in interdisciplinary CTC projects. He is applying his data-mining expertise to a pilot project involving data acquisition and analysis by 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 educational outreach program, 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 CS. Through SciCentr, CTC shares research conducted by Cornell scientists and faculty members in the fields of biotechnology, communication, fine arts, theatre arts, music, and architecture, by using current technology that engages students. CTC is also involving undergraduates in the development of interactive online lab modules focused on bioinformatics through the BioQUEST Curriculum Library.

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

New Faculty

Rachel Prentice

Assistant Professor

CIS, joint with Science and Technology Studies

http://www.sts.cornell.edu/

Rachel Prentice is an assistant professor of the social, political, and ethical implications of information technology in the Department of Science and Technology Studies at Cornell. She obtained her A.B. in comparative literature from Columbia College, Columbia University, in 1987, and her Ph.D. in science, technology, and society from M.I.T. in 2004. Her dissertation, "Bodies of Information: Reinventing Bodies and Practice in Medical Education," is an ethnography about groups of physicians, engineers, and computer experts building computer applications and simulations for teaching anatomy and surgery. She is interested in how interdisciplinary groups are creating new methods for representing and interacting with bodies. Before graduate school, she worked as a newspaper reporter in Washington state, New Mexico, and Rome, Italy, focusing on science, environment, and government issues.

Prentice will teach "From Surgery to Simulation", and "Computers: From Babbage to Gates" this fall. She will also teach "Qualitative Research Methods for Studying Science", and "Exploring Cyberworlds: Thinking With and About Machines" in spring 2005.

PUBLICATIONS

Review of "Figurations: Child, Bodies, Worlds," by C. Castaneda. *Technology and Culture,* January 2004.

AWARDS AND HONORS

Hugh Hampton Young Fellowship (2003–04)

Jacob K. Javits Fellowship, U.S. Department of Education (1998–2002)

Siegel Prize for Best Paper on Science and Technology Studies, for "Calculating Women, Calculating Machines: The Rise of Scientific and Technical Computation in England, 1920–1945" Science, Technology, and Society Program, M.I.T. (1998)

Ida M. Green Fellowship, M.I.T. (1997–98)

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 University in 1987.

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.

PUBLICATIONS

"Abstract identifiers, intertextual reference and a computational basis for record-keeping". *First Monday 9(2) (February '04). URL: http://firstmonday.org/issues/issue9_2/allen/index.html.*

William Y. Arms

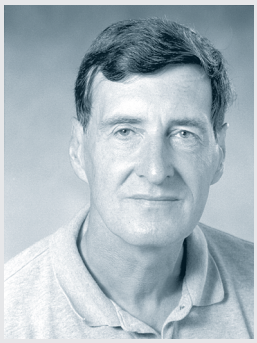

Professor

wya@cs.cornell.edu

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

William Arms received his B.A. in mathematics from Oxford University (Balliol College) 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 CS professor since 1999.

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. He is the principal investigator of a major grant to build the core system for the NSF's new digital library for science, mathematics, engineering, and technology education. This is one of the largest and most heterogeneous digital libraries yet attempted.

Professor Arms is co-director of Cornell's new Information Science Program.

PUBLICATIONS

"Mixed Content and Mixed Metadata: Information Discovery in a Messy World". In *Metadata in Practice* (D. Hillmann and E. Westbrooks, eds., to be published by ALA Editions in 2004). (With C. Arms)

"A Case Study in Metadata Harvesting: the NSDL". *Library Hi Tech* 21(2) (2003). (With N. Dushay, D. Fulker, and C. Lagoze).

LECTURES

"Free Access to Information Today. Who Benefits? What are the Risks? Who Pays?". Keynote address, Association of Learned and Professional Society Publishers, London (April 4, 2003).

PROFESSIONAL ACTIVITIES

Member, Management Board, M.I.T. Press.

Series Editor, M.I.T. Press Series on Digital Libraries and Electronic Publishing.

Member, National Research Council, Public–Private Partnerships in the Provision of Weather and Climate Service.

Member, Program Committee, Joint Conference on Digital Libraries.

UNIVERSITY ACTIVITIES

Co-director, Information Science Program.

Director, eCornell.

Member, Cornell University Library Board.

Member, Faculty Senate.

Graeme Bailey

Professor

bailey@cs.cornell.edu

http://www.cs.cornell.edu/~bailey

Originally working in low-dimensional topology and combinatorial group theory, through an odd mixture of circumstances Dr. 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, Syracuse, New York. 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.

He is also becoming more actively involved again in the areas of digital music, exploiting his twin areas of professional expertise in an area he first worked at half a lifetime ago.

PROFESSIONAL ACTIVITIES

Member, Selection Committees: Rhodes, Marshall, Churchill, and Fulbright Fellowships.
Member, Board of Directors of Engineers for a Sustainable World.
Member, Advisory Board, Cornell University Research in Engineering (CURIE).

UNIVERSITY ACTIVITIES

Adjunct Professor, Mathematics.
Director, Computer Science Master of Engineering Program.
Faculty Advisor, Judo Club; Cornell Lunatics.
Member, Baccalaureate Review Committee.
Member, Cornell University Emergency Medical Service.
Member, Faculty Committee for Residence Life.
Member, Faculty Committee on Music.
Member, General Committee of the Graduate School.
Member, Health Careers Advisory Board Committee.
Member, Master of Engineering Committee.
Member, Transition Committee, West Campus Housing Initiative.
Risley Faculty Fellow.

AWARDS AND HONORS

ACSU Faculty of the Year, 2004.
Kendall S. Carpenter Memorial Advising Award, 2002.
Kenneth A. Goldman '71 Excellence in Teaching Award, 2000.
ACSU Faculty of the Year, 2000.

Kavita Bala

Assistant Professor

kb@cs.cornell.edu

http://www.cs.cornell.edu/~kb/

Kavita Bala received her Ph.D. in computer science at M.I.T. After her doctorate, she worked as a post-doctoral researcher in the Program of Computer Graphics at Cornell University. 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 image synthesis; feature-based rendering and texturing; and image-based modeling and rendering. With technology permitting the acquisition of increasingly complex data sets, rendering with these data sets remains a challenge. Bala's research focus is on scalable algorithms for high-fidelity, interactive image synthesis of complex synthetic and augmented-reality scenes.

PUBLICATIONS

"Feature-based textures". In *Fifteenth Eurographics Symposium on Rendering* (June '04). (With G. Ramanarayanan and B. Walter).

"Combining edges and points for interactive high-quality rendering". In *Computer Graphics Proceedings, Annual Conference Series, ACM Special Interest Group on Graphics and Interactive Techniques (SIGGRAPH)*, J. Hodgins, ed., (July '03). (With B. Walter, and D. Greenberg).

Advanced Global Illumination. A. K. Peters Ltd., 2003. (With P. Dutré and P. Bekaert).

"Detail synthesis for image-based texturing". In *ACM SIGGRAPH 2003, Symposium on Interactive 3D Graphics* 171–176. (With R. Ismert and D. Greenberg).

"Adaptive Shadow Maps". In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH* (E. Fuime, ed.) (August 2001) (With R. Fernando, S. Fernandez, and D. Greenberg).

LECTURES

"Edge-and-point rendering and texturing". University of Washington, graphics mini symposium (March 29, 2004).

"Reusing Shading for Interactive Global Illumination". Game Developers Conference, San Jose, California (March 25, 2004).

PROFESSIONAL ACTIVITIES

Co-chair, Eurographic Symposium on Rendering, 2005.

Member, Papers Program Committee, Graphics Interface 2005.

Member, Papers Program Committee, SIGGRAPH 2003–2004.

Member, Program Committee, Eurographics Symposium on Rendering, 2003–2004.

Member, Program Committee, Pacific Graphics 2004.

Member, Program Committee, Eurographics Symposium on Point-based Computer Graphics 2004.

Member, Papers Program Committee, SIGGRAPH 2003.

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 research focuses on the development of Quicksilver, a platform supporting massively scalable distributed applications (see *http://www.cs.cornell.edu/projects/quicksilver.htm*). For example, Quicksilver will be used to build several kinds of publish–subscribe functionality. Whereas conventional publish–subscribe systems get fragile and slow with even a thousand users, Quicksilver-based solutions should be able to support tens of thousands of publishers and subscribers and hundreds of thousands of publication topics. Key to the approach is the use of new kinds of scalable distributed computing infrastructure tools that employ probabilistically reliable epidemic (gossip) protocols between peers to achieve stability even under stress that can cripple more standard protocols. Birman's group has developed several of these components, notably Bimodal Multicast, Astrolabe, and Kelips, and the hope is to unite them with virtual synchrony group communication protocols within an easily used but powerful network. The work is supported by DARPA under their new Self Regenerative Systems program (SRS).

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, including the Responsible Conduct of Research Committee; the council of CIS faculty; the Engineering College Policy Committee; and the Internet Protocol (IP) Advisory Council for the Cornell Research Foundation.

PUBLICATIONS

Reliable Distributed Computing: Technologies, Web Services, and Applications. [forthcoming (target: November 2004)]. Springer-Verlag.

"Kache: Peer-to-peer Caching Using Kelips". [submitted to *ACM Transactions on Information Systems (TOIS)*, June, 2004]. (With P. Linga and I. Gupta).

"Adding High Availability and Autonomic Behavior to Web Services". In *Proceedings of the Twenty-sixth Annual International Conference on Software Engineering (ICSE 2004)*. May 23–28, 2004. Edinburgh. Scotland. (With R. van Renesse and W. Vogels).

"Astrolabe: A Robust and Scalable Technology for Distributed System Monitoring, Management, and Data Mining". *ACM Transactions on Computer Systems* 21(2): 164–206. May, 2003. (With R. van Renesse and W. Vogels).

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

"The Process Group Approach to Reliable Distributed Computing". *Communications of the ACM* 36(12): 37–53. December, 1993.

PROFESSIONAL ACTIVITIES

Member, Program Committee, 2003 ACM Symposium on Operating Systems Principles.

UNIVERSITY ACTIVITIES

Chairman, Engineering College Facilities Strategic Planning Committee (2003).

Chairman, College of Engineering Facilities Task Force (2003).

AWARDS AND HONORS

Stephen and Marilyn Miles Excellence in Teaching Award, 2000.

ACM Fellow, 1998.

Martin Burtscher

Assistant Professor of the School of Electrical and Computer Engineering and Member of 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 (central processing units) will likely need even more predictors to meet the continuing demand for ever-faster processors. Designing, evaluating, and improving such predictors as well as the corresponding infrastructure is an important focus of Burtscher's research.

Ongoing projects include adding compiler and software support to aid and simplify the prediction hardware, devising means to make microprocessors more energy efficient without compromising their performance, designing adaptive and self-optimizing hardware, studying superspeculative execution cores, devising novel compression approaches for program traces and messages in parallel MPI programs, and developing high-speed processor simulators.

PUBLICATIONS

"VPC3: A Fast and Effective Trace-compression Algorithm". *Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems 167–176 (SIGMETRICS 2004)* New York, New York (2004).

"Compressing Extended Program Traces Using Value Predictors". *Proceedings of the 2003 International Conference on Parallel Architectures and Compilation Techniques, Institute of Electrical and Electronics Engineers (IEEE) Computer Society* 159–169. New Orleans, Louisiana (2003). (With M. Jeeradit).

LECTURES

VPC3: A Fast and Effective Trace-compression Algorithm. Joint International Conference of the Special Interest Group on Measurement and Modeling of Computer Systems (SIGMETRICS 2004). New York, New York.

Low-power Load-value Prediction. Intel Corporation. Santa Clara, California.

Energy-efficient Load-value Predictors. Transmeta Corporation. Santa Clara, California.

Compressing Extended Program Traces Using Value Predictors. International Conference on Parallel Architectures and Compilation Techniques. New Orleans, Louisiana.

Techniques for Improving the Predictability of Critical Loads. Intel Corporation. Santa Clara, California.

PROFESSIONAL ACTIVITIES

Cochair, Second Value-prediction and Value-based Optimization Workshop.

Guest Editor, *Journal of Instruction-level Parallelism*.

Reviewer, *ACM Transactions on Architecture and Code Optimization, IEEE Transactions on Parallel and Distributed Systems, International Symposium on High-performance Computer Architecture, IEEE Computer, Annual International Symposium on Microarchitecture,* and *International Conference on Parallel Architectures and Compilation Techniques.*

UNIVERSITY ACTIVITIES

Member, Computer Advisory Committee, ECE.

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 syntactic and semantic analysis of natural language.

Cardie is a recipient of a 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 just completed her second term as secretary of the North American Association for Computational Linguistics (2000–2003).

PUBLICATIONS

"Evaluating an Opinion Annotation Scheme Using a New Multiperspective Question and Answer Corpus". *American Association for Artificial Intelligence (AAAI) Spring Symposium on Exploring Attitude and Affect in Text,* AAAI Press (2004). (With V. Stoyanov, J. Wiebe, and D. Litman).

"Weakly Supervised Natural Language Learning Without Redundant Views". *Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT–NAACL 2003)* 173–180. (With V. Ng).

"Bootstrapping Coreference Classifiers with Multiple Machine Learning Algorithms". *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing (EMNLP 2003),* Association for Computational Linguistics (2003). (With V. Ng).

"Combining Low-level and Summary Representations of Opinions for Multiperspective Question Answering". *2003 AAAI Spring Symposium on New Directions in Question Answering* 20–27, AAAI Press (2003). (With J. Wiebe, T. Wilson, and D. Litman).

"Recognizing and Organizing Opinions Expressed in the World Press". *2003 AAAI Spring Symposium on New Directions in Question Answering* 12–19, AAAI Press (2003). (With J. Wiebe, E. Breck, C. Buckley, P. Davis, B. Fraser, D. Litman, D. Pierce, E. Riloff, T. Wilson, D. Day, and M. Maybury).

LECTURES

"Machine Learning for Noun Phrase Coreference Resolution".

Syracuse University, School of Information Studies (June 2004).

IBM Watson Research Center (June 2004).

General Motors, Detroit, Michigan (June 2004).

University of Washington Department of Computer Science Colloquium, (January 2004).

Xerox Research and Development Center, Webster, New York (January 2004).

S.U.N.Y. at Buffalo Department of Computer Science Colloquium (December 2003).

University of Pennsylvania Distinguished Lecture Series (December 2003).

"Noun Phrase Coreference Resolution". Cornell Information Science Seminar (October 2003).

PROFESSIONAL ACTIVITIES

Secretary, North American Association for Computational Linguistics.

Action Editor, *Journal of Machine Learning Research.*

Member, Editorial Board, *Journal of Artificial Intelligence Research.*

Member, Editorial Board, *Machine Learning.*

Member, Executive Board, Special Interest Group of the Association for Computational Linguistics for Linguistic Data and Corpus-based Approaches to Natural Language Processing (NLP).

Member, Nominating Committee, North American Association for Computational Linguistics.

Member, Organizing Committee, AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications.

Member and Area Chair, Program Committee, Forty-second Annual Meeting of the Association for Computational Linguistics (ACL).

Member, Program Committee, Eighth International Conference on Computational Natural Language Learning.

Member, Program Committee, EMNLP.

Member, Program Committee, Tenth ACM Special Interest Group on Knowledge Discovery and Data Mining (SIGKDD) International Conference.

Member, Program Committee, ACL–2004 Workshop on Reference Resolution and its Applications.

Reviewer, *Computational Linguistics*.

Reviewer, *Cognitive Systems*

UNIVERSITY ACTIVITIES

Co-director, Information Science Program.

Director, Undergraduate Studies, Information Science (Colleges of Arts and Sciences, and Agriculture and Life Sciences).

Co-director, Undergraduate Studies, Information Science, Systems, and Technology.

Member, Computing and Information Science Council.

Member, Chair Search Committee, Computer Science.

Board Member, Cornell Presidential Research Scholars.

Member, Provost's Advisory Group of Women in Science and Engineering.

Member, College Scholar Advisory Board.

Member, Independent Major Advisory Board.

College Scholar Advisor.

AWARDS AND HONORS

CAREER award (1996–2000).

Lilly Teaching Fellow, Cornell University (1996–1997).

Ralph S. Watts College of Engineering Excellence in Teaching Award, Cornell University (1996).

Graduate Fellow, University of Massachusetts (1993–1994).

Massachusetts Regents Fellowship, University of Massachusetts (1992–1993).

Best Written Paper Award, Ninth National Conference on Artificial Intelligence, Honorable Mention (1991).

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, adaptive clustering, learning rankings, and applications of these learning methods to problems in medical decision-making and bioinformatics.

Inductive transfer is a subfield of machine learning where better performance is achieved by learning many related problems at the same time—it is easier to learn 100 related problems together than to learn any one of them in isolation. Caruana helped create this subfield of machine learning by publishing the first paper on multitask learning more than ten years ago.

Learning rankings is an exciting area in machine learning where the goal is not to predict a classification or value for an item, but to predict an ordering for a set of items. Caruana is developing algorithms that learn rankings for problems in medical decision-making where it is difficult to assess absolute risk for any given patient, but easier to learn to order patients by relative risk. One method he developed to learn rankings outperformed a dozen other machine learning methods in a large multi-institutional pneumonia-risk prediction project.

In 2000–2001 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 cough syrup that may signal the onset of symptoms from a recent attack. Because consumers tend to self-medicate using easily available products such as cough syrup and throat lozenges before consulting physicians, the system can detect the onset of flu-like symptoms twenty-four to forty-eight hours before these can be detected by visits to hospitals and doctors offices.

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 a protein-folding problem. Caruana recently received a CAREER award to pursue this research in clustering. The research in ensemble learning arose from an empirical comparison of machine-learning methods he and students were performing where an ensemble of different learning methods outperformed ensembles of any individual learning methods. The ensemble selection method they are developing may be the first high-performance machine-learning method that can be optimized to almost any performance metric.

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 he develops are useful in practice.

PUBLICATIONS

"Learning from Libraries of Models with Ensemble Selection". *2004 International Conference on Machine Learning (ICML)*. (With A. Niculescu, G. Crew, and A. Ksikes).

"Evaluating the C-section Rate of Different Physician Practices: Using Machine Learning to Model Standard Practice". In *Proceedings of the American Medical Informatics Association Conference* (November 2003). (With S. Niculescu, B. Rao, and C. Simms).

LECTURES

Invited Speaker, Williams College, Williamstown, Massachusetts (April 2004)

Invited Speaker, University of Edinburgh, Scotland (May 2003).

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

Invited Speaker, University of Skovde, Sweden (December 1994).

PROFESSIONAL ACTIVITIES

Co-chair, KDD–CUP for the 2004 ACM SIGKDD International Conference.

Area Chair, Advances in Neural Information Processing Systems (NIPS) (December 2003).

Area Chair, ICML (July 2003).

Member, Program Committee for International Conference on Data Mining (ICDM) (2003).

Member, Program Committee for the 2003 ACM SIGKDD International Conference.

Member, Program Committee, American Medical Informatics Association (AMIA) Conference (2003).

Co-chair of Workshops, Advances in Neural Information Processing (1999, 2000).

AWARDS AND HONORS

CAREER award, May 2004.

Invited speaker, Williams College, April 2004.

Nomination for Best Paper: "Evaluating the C-section Rate of Different Physician Practices: Using Machine Learning to Model Standard Practice". *AMIA Conference* (November 2003). (With S. Niculescu, B. Rao, and C. Simms).

Award for Outstanding Contribution to Research, Philips Labs, 1988.

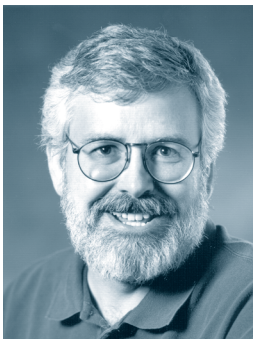

L. Paul Chew

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, shape comparison, 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 the *Journal of the Pattern Recognition Society*. He is also a member of the steering committee for the International Meshing Roundtable.

PUBLICATIONS

"Finding the Consensus Shape for a Protein Family". *Algorithmica* 38(1) (2004). (With K. Kedem).

"Computational Science Simulations Based on Web Services". *International Conference on Computational Science (2003)*. (With N. Chrisochoides, S. Gopalsamy, G. Heber, T. Ingraffea, E. Luke, J. Neto, K. Pingali, A. Shih, B. Soni, P. Stodghill, D. Thompson, S. Vavasis, and P. Wawrzynek).

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 consulting center in New York City.

With colleagues Shirish Chinchalkar, Yuying Li, Peter Mansfield, and Cristina Patron, 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, Dimitry Leuchenkov, 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 more than sixty journal articles. He was chair of the Society for Industrial and Applied Mathematics (SIAM) Activity Group on Optimization (1998–2001) and is on the editorial board of numerous professional journals.

PUBLICATIONS

"Hedging Guarantees in Variable Annuities (Under Both Market and Interest Rate Risks)". *Fourteenth Annual Derivative Securities and Risk Management Conference,* New York (April 23–24, 2004). (With Y. Li and M. Patron)

"Derivative Portfolio Hedging Based on CvAR". In *Risk Measures for the Twenty-first Century* (G. Szegö, ed.) 339–363, Wiley (2004). (with S. Alexander and Y. Li)

"An Object-oriented Framework for Valuing Shout Options on High-performance Computer Architectures". *Journal of Economic Dynamics and Control* 27(6): 1133–1161 (2003). (With H. Windcliff, K. Vetzall, P. Forsyth, and A. Verma)

"Discrete Hedging Under Piecewise Linear Risk Management". *Journal of Risk* 5: 39–65 (spring 2003). (With M, Patron and Y. Li)

LECTURES

Financial Engineering on Computational Clusters, University of Hong Kong, Hong Kong (May 14, 2004).

Microsoft–TGINet Seminar on Finance, Tokyo (May 12, 2004).

Clusterworld Conference and Expo 2004, San José, California (April 2004).

Fast Portfolio Calculations, on a Cluster, using Web Services. Hedge Funds World, New York, New York (March 29, 2004).

New Directions for the Efficient and Accurate Computation of Value-of-risk for a Portfolio of Complex Derivatives, Global Association of Risk Professionals 2004, New York (February 23–26, 2004).

Cluster Computing for Risk Management and Financial Engineering, FIST Global, Seoul, Korea (December 10, 2003).

Cluster Computing for Risk Management and Financial Engineering, Professional Risk Manager's International Association, Paris (March, 2003).

PROFESSIONAL ACTIVITIES

Director, CTC–Manhattan (formerly Financial Industry Solutions Center).

Member, International Visiting Committee, Computational Finance Engineering Program, National University of Singapore (2004–).

Member, High-performance Computing Advisory Committee to the Council on Competitiveness (2004–2007).

Co-organizer, The Fourteenth Derivatives Securities Conference, New York, New York (April 26–27, 2004).

Member, Editorial Board, *Computational Optimization and Applications* (1992–).

Member, Editorial Board, *Applied Mathematics Letters* (1989–).

Referee and/or reviewer for numerous professional journals.

UNIVERSITY ACTIVITIES

Director, Cornell Theory Center.

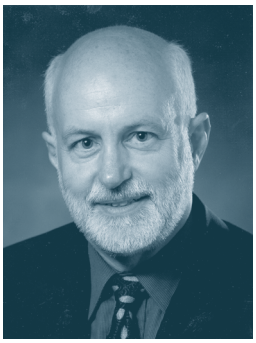

Robert Constable

Professor

Dean of the Faculty of Computing and Information Science

rc@cs.cornell.edu, cis-dean@cornell.edu

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

Robert Constable is the dean of the CIS Faculty 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 a formal digital library that will allow interactive access to theorems and proofs from Nuprl, MetaPRL, HOL, PVS and other major theorem provers. The Library includes more than 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; a member of the Cognitive Studies executive committee; the applied math policy committee; and the general committee for the IEEE Conference on Logic in Computer Science (LICS). He serves as editor for *Journal of Logic and Computation* and *Formal Methods in System Design*.

PUBLICATIONS

Expressing and Implementing the Computational Content Implicit in Smullyan's Account of Boolean Valuations (2003). (With S. Allen and M. Fluet).

"Information-intensive Proof Technology and Computation". *Springer-Verlag Lecture Notes for the Marktoberdorf NATO Summer School* (2004).

"MetaPRL—A Modular Logical Environment". In *Proceedings of the Sixteenth International Conference on Theorem Proving in Higher Order Logics. Lecture Notes in Computer Science* 2758: 287–303 (D. Basin and B. Wolff, eds.), Springer-Verlag (2003). (With J. Hickey, A. Nogin, B. Aydemir, E. Barzilay, and L. Lorigo).

"Practical Reflection in Nuprl". In *Eighteenth Annual IEEE Symposium on Logic in Computer Science* (P. Kolaitis, ed.), Ottawa, Canada (June 22–25, 2003). (With E. Barzilay and S. Allen).

"Recent Results in Type Theory and Their Relationship to Automath". In *Thirty-five Years of Automating Mathematics,* (F. Kamareddine, ed.) 1–11, Kluwer Academic Publishers (2003).

PROFESSIONAL ACTIVITIES

Member, Advisory Council, Computer Science Department, Princeton University.

Chairman, Advisory Board, Computer Science Department, University of Chicago.

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

Director, NATO Summer School, Marktoberdorf, Germany.

Member, General Committee, LICS.

Member, Computing Research Association Board.

UNIVERSITY ACTIVITIES

Dean, Computing and Information Science.

Member, Applied Math Policy Committee and Cognitive Studies Executive Committee.

AWARDS AND HONORS

ACM Fellow, 1995.

Guggenheim Fellow, 1991.

Cornell Outstanding Educator Award, 1989.

Alan J. Demers

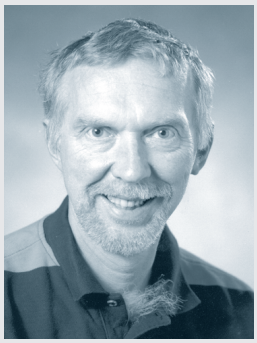

Professor

ademers@cs.cornell.edu

http://www.cs.cornell.edu/annual_report/00-01/bios.htm#demers

Alan J. Demers received his bachelor's degree in physics from Boston College in 1970. He obtained his Ph.D. degree in computer science from Princeton University in 1975. A CS professor, he has been at Cornell since 1982. Prior to his arrival at Cornell, Demers worked for the Oracle Corporation as an architect, and at Xerox Parc in Palo Alto as a principal scientist.

Demers' research concerns aspects of databases and distributed systems.

As part of the Cougar Project, he is studying "sensor database systems" using novel forms of in-network processing to do energy-efficient query processing over sensor networks.

With Johannes Gehrke, Mirek Riedewald, and others, he is studying stream-processing databases. The group is studying efficient query evaluation and multi-query optimization for extensions of publish–subscribe systems. It is also considering the formal underpinnings of such systems, developing an algebra of events that can describe parameterization and aggregation.

With Ken Birman, Johannes Gehrke, Robbert van Renesse, 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 variants of the basic protocol tailored to specific application requirements. It is also investigating algorithms that maintain approximate group membership views with probabilistic guarantees against partitioning under "churn".

PUBLICATIONS

"The Architecture of the Cornell Knowledge Broker". In *Proceedings of the Second Symposium on Intelligence and Security Informatics (ISI 2004),* Tucson, Arizona (June 2004). (With J. Gehrke and M. Riedewald).

"The Cougar Project: A Work-in-progress Report". In *SIGMOD Record* 34(4) (December 2003). (With J. Gehrke, R. Rajaraman, N. Trigoni, and Y. Yao).

"Energy-efficient Data Management for Sensor Networks: A Work-in-progress Report". *Second IEEE Upstate New York Workshop on Sensor Networks,* Syracuse, New York (October 2003). (With J. Gehrke, R. Rajaraman, N. Trigoni, and Y. Yao).

Shimon Edelman

Professor of the Department of Psychology and Member of the Graduate Fields of Psychology, Computer Science, and Information Science

se37@cornell.edu

http://kybele.psych.cornell.edu/~edelman/

Shimon Edelman received his Ph.D. in computer science from the Department of Applied Mathematics and Computer Science at the Weizmann Institute, Rehovot, Israel, in 1988. He taught and conducted research at M.I.T., the Weizmann Institute, and at the University of Sussex (U.K.), before assuming his current position at Cornell.

Over the years, Edelman has taught courses and led seminars that matched his wide range of interests in cognitive sciences. His research in the past dealt with motor control, reading, perceptual learning, visual recognition and categorization, natural language processing and computational linguistics.

Edelman works on developing mathematical solutions to the problems at hand, on experimental assessment of these solutions as models of human cognition, and on bringing the theoretical understanding of cognition to bear on data from neurobiological studies of the brain. Publications stemming from this research appeared in journals ranging from computational (such as the *IEEE Transactions on Pattern Analysis and Machine Intelligence*) to empirical (*Neuron*), spanning the fields of vision (*International Journal of Computer Vision*; *Vision Research*) and language (*Journal of Computational Linguistics; Journal of Linguistics*). Edelman's book, *Representation and Recognition in Vision,* was published by M.I.T. Press in 1999.

PUBLICATIONS

"Bridging Computational, Formal and Psycholinguistic Approaches to Language". *Proceedings of the Twenty-sixth Cognitive Science Society Conference,* Chicago, Illinois. (August 2004). (With Z. Solan, D. Horn, and E. Ruppin).

"Unsupervised Context-sensitive Language Acquisition from a Large Corpus". *Proceedings of the 2003 NIPS Conference* 15. (L. Saul, ed.), M.I.T. Press (2004). (With Z. Solan, D. Horn, and E. Ruppin).

"Unsupervised Statistical Learning in Vision: Computational Principles, Biological Evidence". *Extended abstract distributed to the participants of the European Conference on Machine Learning Workshop on Statistical Learning in Computer Vision,* Prague (May 2004). (With N. Intrator).

"Unsupervised Context-sensitive Language Acquisition from Large, Untagged Corpora". *AAAI Spring Symposium on Language Acquisition,* Stanford, California. (March 2004). (With Z. Solan, E. Ruppin, and D. Horn).

"Metric Category Spaces of Biological Motion". *Vision Sciences Society* (May 2003). (With M. Giese and I. Thornton).

"Unsupervised Efficient Learning and Representation of Language Structure". *Proceedings of the Twenty-fifth Cognitive Science Society Conference,* Boston, Massachusetts. (July 2003). (With Z. Solan, D. Horn, and E. Ruppin).

"A New Vision of Language" [extended abstract]. *Proceedings of the Twenty-fifth Cognitive Science Society Conference,* Boston, Massachusetts. (July 2003).

"Automatic Acquisition and Efficient Representation of Syntactic Structures". *Proceedings of the 2002 NIPS Conference*. 15 (S. Thrun, ed.), M.I.T. Press (2003). (With Z. Solan, E. Ruppin, and D. Horn).

"Towards Structural Systematicity in Distributed, Statically Bound Visual Representations". *Cognitive Science* 27: 73–110 (2003). (With N. Intrator).

LECTURES

On What It Could Mean To See, and What Could Be Done About It, Computation and Neural Systems Program Colloquium, California Institute of Technology (March 2003).

A Vision of Language, NSF Workshop on Integrated Cognitive Science, Arlington, Virginia (October 2–3, 2003).

Unsupervised Acquisition of Context-sensitive Recursive Structure from Language-like Data, Biology Colloquium, C.U.N.Y. (December 2003).

Rich Syntax from a Raw Corpus: Unsupervised Does It, Conference on Neural Information Processing Systems Workshop on Syntax, Semantics, and Statistics, Whistler, British Columbia (December 2003).

Computational Principles for Unsupervised Learning in Vision (and in Language Acquisition), Engineering Colloquium, Brown University (March 2004).

Computational Principles for Unsupervised Learning in Vision, Special Psychology Colloquium, Stanford University (March 2004).

Unsupervised Statistical Learning in Vision: Computational Principles, Biological Evidence, ECCV–2004 Workshop on Statistical Learning in Computer Vision, Prague (May 2004).

Object Recognition and Categorization: Some Lessons from Psychophysics, Neurobiology, and Computer Vision, IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2004 Workshop on Generic Object Recognition, Washington, D.C. (June 2004).

PROFESSIONAL ACTIVITIES

Associate Editor, *Cognitive Science*.

Associate Editor, *Behavioral and Brain Sciences*.

UNIVERSITY ACTIVITIES

College of Arts and Sciences admissions, Spring 2004.

Director, Cornell Cognitive Studies Program.

Co-organizer, Language Universals (Cornell Cognitive Studies Spring 2004 symposium).

AWARDS AND HONORS

Levinson Prize in Mathematics, 1996.

Holder of the Sir Charles Clore Career Development Chair, 1994–1998.

Yigal Alon Fellowship, 1992–1995.

Koret Foundation Postdoctoral Fellowship, 1990–1992.

Chaim Weizmann Postdoctoral Fellowship, 1988–1990.

Grants for participation in the Cold Spring Harbor course on Computational Neuroscience, from the Aharon Katzir Fund, Weizmann Institute; and the Cold Spring Harbor Laboratory, 1987.

Dean's Award for Achievement, Feinberg Graduate School, Weizmann Institute, 1985.

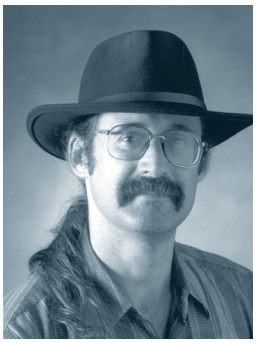

Ron Elber

Professor

Director, Ithaca campus of the Tri-institutional Program in Computational Biology; Director, Computational Biology Service Unit; Director, Library of Life Project

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 1984–1987 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 (1992–1999). Since 1999 he has been on the CS faculty at Cornell where he is currently a full professor.

Ron's research is in computational biology and bioinformatics. His group is developing novel tools (like 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.

PUBLICATIONS

"Computing Time Scales from Reaction Coordinates by Milestoning". *Journal of Chemical Physics* 120: 10880–10889. (2004) (With T. Faradjian).

"Large-scale Linear Programming Techniques for the Design of Protein Folding Potentials". *Mathematical Programming* [in press]. (With M. Wagner and J. Meller).

"The Evolutionary Capacity of Protein Structures". *Proceedings of the Eighth Annual International Conference on Research in Computational Molecular Biology (RECOMB 2004)* (2004). (With L. Meyerguz, D. Kempe, and J. Kleinberg).

"Computational Analysis of Sequence Selection Mechanisms". *Structure* 12: 547–557 (2004). (With L. Meyerguz, C. Grasso, and J. Kleinberg).

"Enriching the Sequence Substitution Matrix by Structural Information". *Proteins, Structure, Function, and Genetics* 54: 41–48 (2004). (With O. Teodorescu, T. Galor, and J. Pillardy).

"Atomically Detailed Simulations of Helix Formation with the Stochastic Difference Equation". *Biophysical Journal* 85: 2919–2939 (2003). (With A. Cárdenas).

"Kinetics of Cytochrome C Folding: Atomically Detailed Simulations". *Proteins, Structure, Function, and Genetics* 51: 245–257(2003). (With A. Cárdenas).

"The Dominant Interaction Between Peptide and Urea is Electrostatic in Nature: A Molecular Dynamics Simulation Study". *Biopolymers* 68: 359–369 (2003). (With D. Tobi and D. Thirumalai).

"Ion Permeation through the Gramicidin Channel: Atomically Detailed Modeling by the Stochastic Difference Equation". *Proteins, Structure, Function, and Genetics* 50: 63–80 (2003). (With K. Siva).

"The Stochastic Difference Equation as a Tool to Compute Long-time Dynamics". Chapter in *Bridging the Time Scale Gap* (P. Nielaba, M. Mareschal, and G. Ciccotti, eds.), Springer Verlag, Berlin, 335–363 (2002). (With A. Ghosh and A. Cárdenas).

"Bridging the Gap between Reaction Pathways, Long-time Dynamics, and Calculation of Rates". *Advances in Chemical Physics* 126: 93–129 (2003). (With A. Ghosh, A. Cárdenas, and H. Stern).

LECTURES

"Computing Rates by Milestoning". Applied Math Seminar, New York University (April 2004).

"The Evolutionary Capacity of Protein Structures".

Princeton, Chemical Engineering (May 2004).

NATO School on Soft-matter Physics, Edinburgh (April 2004).

Bioinformatic Program, Ann Arbor, Michigan (March 2004).

"Long-time Dynamics and Protein Folding".

Biophysical Society, Baltimore, Maryland (February 2004)

NIH, Chemical Physics Laboratory, (November 2003).

European Conference on Activated Processes, Paris (October 2003).

Energy Landscapes Conference, Telluride (August 2003).

"The Temperature of Evolution", Bioinformatic Conference, Buffalo, New York (June 2003).

PROFESSIONAL ACTIVITIES

Member, Editorial Board: *Biophysical Journal; Theoretical Chemistry Accounts; Computer Physics Communication*.

UNIVERSITY ACTIVITIES

Director, Ithaca campus of the Tri-institutional Program in Computational Biology; The Computational Biology Service Unit; Library of Life Project

Member, Life Science Advisory Council.

AWARDS AND HONORS

The Bergman Award, 1994.

The Alon New Faculty Award, 1992–1994.

University of Illinois Scholar, 1991–1992.

The Camille and Henry Dreyfus New Faculty Award, 1987–1990.

The Stein Award for Ph.D. Studies, 1984.

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, Canada, 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 COM S 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.

PUBLICATIONS

"First Programming Course in Engineering: Balancing Tradition and Application". [forthcoming] *Computers in Education Journal* (July–September 2004). (With D. Schwartz).

"Penalized Regression Dynamic Programming (PR–DP) for High-dimensional Continuous-state Stochastic Control Problems with Application to Reservoir Operation". [in review] *Operations Research*. (With C. Shoemaker and D. Ruppert).

PROFESSIONAL ACTIVITIES

American Society for Engineering Education Annual Conference and Exposition, Salt Lake City, Utah (June 2004).

SIGCSE Technical Symposium, Norfolk, Virginia (March 2004).

UNIVERSITY ACTIVITIES

Reviewer for Center for the Environment (Cornell) Graduate Research Grant Program.

College of Arts and Sciences Admissions Committee (spring 2004).

Member, Faculty Leadership Committee, CURIE Academy (plans the one-week outreach program for high school girls).

FIRST Robotics Club, Cornell.

Faculty Advisor for the robotics club that helps high school and middle school students build and program robots to compete in the annual FIRST Robotics and FIRST Lego League competitions.

Faculty panels and hosting events for engineering admission:

Society of Women Engineers Hosting Day (spring)

Minority Hosting Weekend (fall)

Faculty Fellow: Fellow in one of the first-year student residence halls—developed academic and social programs for residents.

AWARDS AND HONORS

Cornell Society of Engineers Achievement Award (2004).

Graduate Teaching Assistant Award (CS, Cornell, 2000–2001).

New York State American Water Works Association Russell Sutphen Scholarship (2000).

John E. Perry Teaching Assistant Prize (Civil and Environmental Engineering, Cornell) (1999).

Natural Sciences and Engineering Research Council (NSERC Canada) Postgraduate Scholarship (1994–96).

Paul Francis

Associate Professor

francis@cs.cornell.edu

http://www.cs.cornell.edu/people/francis

Paul Francis received his Ph.D. from University College of 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 the AT&T Center for Internet Research at the International Computer Science Institute, and was chief scientist at two startups, FastForward Networks and Tahoe Networks. Dr. Francis's innovations include NAT (Network Address Translation), multicast shared trees (used in protocol independent multicast sparse mode 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's research interests looking forward are in the areas of peer-to-peer applications, overlay networks, network host proximity, Internet scaling, and DDoS protection.

Dr. Francis has chaired two Internet Engineering Task Force (IETF) working groups, and has published numerous requests for comments, U.S. and international patents, and research papers.

PUBLICATIONS

"Is the Internet Going NUTSS?" IEEE Computing Magazine (November–December 2003).

LECTURES

"NAT and IP v6, We Meet at Last". North American Network Operators Group 30, Miami (February 2004).

COM S 212 Lecture on IP networks and sockets (November 2003).

OR&IE 480 Two lectures on IP networks (September 2003).

"NUTSS: the DeFacto Next-generation Internet Architecture". IAI@Rome Summer Seminar Series (August 2003).

CS Lecture on domain name systems and CDNs (August 2003).

PROFESSIONAL ACTIVITIES

Reviewer, *Computer Networks Journal* (Elsevier Science Publisher)

AWARDS AND HONORS

NSF award (Directorate for Computer and Information Science and Engineering (CISE), Computer and Network Systems), "Very Fine-grained Proximity Addressing" (September 2003).

Eric Friedman

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

friedman@orie.cornell.edu

http://www.orie.cornell.edu/~friedman/

Eric Friedman received an A.B. in physics from Princeton (1985) and a M.A. in physics (1987) and M.S. and Ph.D. (1993) in operations research from the University of California at Berkeley. He was on the faculty at Duke (decision sciences) and Rutgers (economics) before joining the faculty at OR&IE at Cornell in 2001.

Eric's research is at the intersection of game theory, computer science, and operations research. Current projects include: constructing reputation systems for peer-to-peer networks, designing fair and efficient Web-serving algorithms, and allocating bandwidth in heterogeneous wireless systems. He is also interested in self-organized critical systems, learning in games, and the geometric structure of cost-allocation methods.

PUBLICATIONS

"Paths and Consistency in Additive Cost Sharing". *International Journal of Game Theory* [forthcoming] (2004).

"The Behavior of Coupled Automata". *Physical Review E* (2003).

"Large Scale Synchrony, Global Interdependence, and Contagion". *Quantitative Finance* 3(4): 296–305 (August 2003). (With S. Johnson and A. Landsberg).

"Asynchronous Learning with Limited Information: An Experimental Analysis". *Games and Economic Behavior* (2003). (With M. Shor, S. Shenker, and B. Sopher).

"Pricing WiFi at Starbucks—Issues in Online Mechanism Design". *Fourth ACM Conference on Electronic Commerce (EC 2003)* (2003). (With D. Parkes).

"Strong Monotonicity in Surplus Sharing". *Economic Theory* (2003).

"Propping and Tunnelling". *Journal of Comparative Economics* (2003). (With S. Johnson and T. Mitton).

"Fairness and Efficiency in Processor Sharing Protocols to Minimize Sojourn Times". *SIGMETRICS* 229–237 (2003). (With S. Henderson).

"Asynchronous Learning in Decentralized Environments: A Game Theoretic Approach". In *Collectives and the Design of Complex Systems* (K. Turner and D. Wolpert, eds.) Springer-Verlag (2003).

LECTURES

Fairness and Stability in Allocating Wireless Bandwidth: Caltech, Rochester, IBM Watson—2003; Dagstuhl, University of California at Berkeley—2004.

Cost Allocation: CAM Cornell: 2003.

UNIVERSITY ACTIVITIES

Program committee: Second workshop on the Economics of Peer-to-peer Systems.

Geri Gay

Professor

CIS, joint with Communication

gkg1@cornell.edu

http://www.comm.cornell.edu/People_Professors.html#gkg1

Geri Gay is the director of the Human–Computer Interaction Group (HCI Group) and a professor in the 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, mental models and metaphors (NSF), knowledge representations, collaborative work and learning (NASA and the AT&T Foundation), and system design of interactive computing systems.

Professor Gay teaches courses in computer-mediated communication, human–computer interaction, and the social design of communication systems.

PUBLICATIONS

Activity-centered Design: An Ecological Approach to Designing Smart Tools and Usable Systems. Cambridge, Massachusetts: M.I.T. Press (2004). (With H. Hembrooke).

"Usability, Learning, and Subjective Experience: User Evaluation of K–MODDL in an Undergraduate Class". *Joint Conference on Digital Libraries* [accepted] (2004). (With B. Pan, J. Saylor, H. Hembrooke, and D. Henderson).

"Eye-tracking Analysis of User Behavior in WWW Search". In *Proceedings of Twenty-eighth Annual ACM Conference on Research and Development in Information Retrieval (SIGIR 2004)*, Sheffield, U.K. (2004). (With L. Granka and T. Joachims).

"The Determinants of Web Page Viewing Behavior: An Eye-tracking Study". In *Proceedings of Eye Tracking Research and Applications* (S. N. Spencer, ed.), New York, New York: *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*. (2004). (With B. Pan, H. Hembrooke, G. Granka, M. Feusner, and J. Newman).

"Culturally Embedded Computing". *IEEE Pervasive Computing, Special Issue on Art, Design, and Entertainment in Pervasive Environments* (2004). (With P. Sengers, J. Kaye, K. Boehner, J. Fairbank, E. Medynskiy, and S. Wyche).

"The Lecture and the Laptop: Multitasking in Wireless Learning Environments". *Journal of Computing in Higher Education* 15(1): 46–65 (2003). (With H. Hembrooke).

"MetaTest: Evaluation of Metadata from Generation to Use". *Proceedings of the Joint Conference on Digital Libraries* (2003). (With L. Liddy, E. Allen, H. Hembrooke, T. Finneran, and L. Granka).

"User-recalled Occurrences of Usability Errors: Implications on the User Experience". *Computer–Human Interaction (CHI) 2003 Extended Abstracts, on Human Factors in Computing Systems* 434–436 (2003). (With H. Mentis).

"Technology Acceptance and Social Networking in Distance Learning". *Educational Technology and Society* 6(2): 50–61 (2003). Available at *http://ifets.ieee.org/periodical/6-2/6.html*. (With J. Lee, H. Cho, B. Davidson, and A. Ingraffea).

LECTURES

Invited Talk, Affective Presence Seminar—Hillsboro, Oregon (January 20–21, 2004).

Invited Speaker, Digital Media and Communication, held in Philadelphia, Pennsylvania under the sponsorship of the Annenberg Public Policy Center at the University of Pennsylvania, (October 31–November 1, 2003).

Presentation, Wireless Computing in Museums. Centre for Industrial and Medical Informatics Museums, Europe Meeting. Edinburgh, Scotland (November 2002).

Invited Speaker, Intel Symposium and Workshop, Understanding Visitor Expectations and Museums as Mobile Computing Environments: Hand-helds in the Museum Landscape. Hillsboro, Oregon (May 2002).

PROFESSIONAL ACTIVITIES

Chair, Department of Communication (2004–present).

Director of Graduate Studies for Information Science Ph.D. Program (2003–2004)

UNIVERSITY ACTIVITIES

Member, CIS Council Committee (2004–present).

Member, CIS Building Committee (2003–present).

Advisor, CALS Technology Committee (2003–present).

Member, Museum Faculty Advisory Committee, Herbert F. Johnson Museum of Art, Cornell University (2001).

Board of Directors. Boyce Thompson Institute, Cornell University (2001).

Member, Founder's Board, Faculty of Computing and Information Science, (1999–present).

Member, Computing and Information Sciences Task Force, Cornell University (1999).

Member, Distance Learning Committee, College of Agriculture and Life Sciences, Cornell University (1997–1999).

AWARDS AND HONORS

Stephen H. Weiss Presidential Fellowship Teaching Award, Cornell University, 2004.

Chancellor's Award for Excellence in Teaching, S.U.N.Y., 2001.

Merrill Presidential Scholar Faculty, 2000.

Innovative Teaching Award, College of Agriculture and Life Sciences, 2000.

Distinguished Teaching Award of the Cornell Chapter of Sigma Delta, the Honor Society of Agriculture and Life Sciences, Human Ecology, and Veterinary Medicine, 1996.

Johannes Gehrke

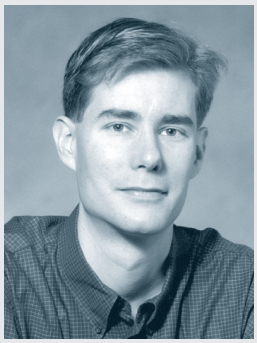

Assistant Professor

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.

Johannes's research interests are in the areas of data mining and database systems. With Professor Al Demers, he is working on distributed data management for wireless sensor networks. With Professor Jayavel Shanmugasundaram, he is building a peer-to-peer database system that scales to thousands of nodes. He is also interested in techniques for processing high-speed data streams and in data privacy. His data-mining research includes privacy-preserving data mining, theoretical foundations of data mining, and applications of data mining to problems in the sciences. His group has developed some of the fastest known algorithms for several important data-mining tasks.

PUBLICATIONS

"Querying Peer-to-peer Networks Using P–trees". In *Proceedings of the Seventh International Workshop on the Web and Databases (WebDB 2004),* Paris, France (June 2004). (With A. Crainiceanu, P. Linga, J. Gehrke, and J. Shanmugasundaram).

"Approximation Techniques for Spatial Data". In *Proceedings of the ACM International Conference on Management of Data (SIGMOD).* Paris, France (June 2004). (With A. Das and M. Riedewald).

"The Architecture of the Cornell Knowledge Broker". In *Proceedings of the ISI,* Tucson, Arizona (June 2004). (With A. Demers and M. Riedewald).

"Privacy Preserving Mining of Association Rules". *Information Systems* 29(4): 343–364 (June 2004). Special Issue with Best Papers from KDD 2002. (With A. Evfimievski, R. Srikant, and R. Agrawal).

"P–Tree: A P2P Index for Resource Discovery Applications". In *Proceedings of the Thirteenth International World Wide Web Conference,* New York, New York (May 2004). Poster paper. (With A. Crainiceanu, P. Linga, and J. Shanmugasundaram).

"A Storage and Indexing Framework for P2P Systems". In *Proceedings of the Thirteenth International World Wide Web Conference,* New York, New York (May 2004). Poster paper. (With A. Crainiceanu, P. Linga, A. Machanavajjhala, and J. Shanmugasundaram).

"Sketch-based Multiquery Processing Over Data Streams". In *Proceedings of the Ninth International Conference on Extending Database Technology,* Heraklion–Crete, Greece (March 2004). (With A. Dobra, M. Garofalakis, and R. Rastogi).

"DualMiner: A Dual-pruning Algorithm for Itemsets with Constraints". *Data Mining and Knowledge Discovery* 7(3): 241–272. *Special Issue on "Selected Papers from the Eighth ACM SIGKDD International Conference—Part I.* (With C. Bucila, D. Kifer, and W. White).

"The Cougar Project: A Work-in-progress Report". *Sigmod Record* 32(4) (December 2003). (With A. Demers, R. Rajaraman, N. Trigoni, and Y. Yao).

"MAFIA: A Performance Study of Mining Maximal Frequent Itemsets". In *Workshop on Frequent Itemset Mining Implementations (FIMI 2003).* Melbourne, Florida. (November 2003). (With D. Burdick, M. Calimlim, J. Flannick, and T. Yiu).

"Energy-efficient Data Management for Sensor Networks: A Work-in-progress Report". In *Second IEEE Upstate New York Workshop on Sensor Networks.* Syracuse, New York. (October 2003). (With A. Demers, R. Rajaraman, N. Trigoni, and Y. Yao).

"Computing Aggregate Information Using Gossip". In *Proceedings of the Forty-fourth Annual IEEE Symposium on Foundations of Computer Science (FOCS 2003),* Cambridge, Massachusetts (October 2003). (With D. Kempe and A. Dobra).

"Decision Tree Construction". In *Handbook of Data Mining* (N. Ye, ed.). Lawrence Erlbaum Associates (2003).

LECTURES

Privacy Breaches in Privacy Preserving Data Mining. Computer Science Colloquium, North Carolina State University. Raleigh, North Carolina (April 2004).

Energy-efficient Data Management in Sensor Networks. NSF–RPI Workshop on Pervasive Computing and Networking. Troy, New York (April 2004).

Towards a Theory of Constraint-based Itemset Mining. Keynote at the Workshop on Inductive Databases and Constraint Based Mining. Freiburg, Germany (March 2004).

Privacy-preserving Data Mining. KD–D Review Meeting. Ithaca, New York (February 2004).

Distributed Mining and Monitoring. KD–D Principal Investigators Meeting. Washington, D.C. (November 2003).

Privacy Breaches in Privacy-preserving Data Mining. Computer Science Colloquium, Department of Computer Science, University of Buffalo. Buffalo, New York (November 2003).

RPH–Trees: An Adaptive, Scalable, High-performance Publish/Subscribe System. Air Force Rome Labs Science Advisory Board. Rome, New York (November 2003).

Data Streams and Data Mining. Lawrence Livermore National Laboratory. Livermore, California (November 2003).

Mining Streaming Data *and* Distributed Sensor Databases. Cornell Information Assurance Institute, Ithaca, New York (October 2003).

Distributed Mining and Monitoring. Annual Multidisciplinary University Research Institute (MURI) Review Meeting. Ithaca, New York (October 2003).

Processing Data Streams. University of Karlsruhe. Karlsruhe, Germany (September 2003).

Results of the KDD Cup 2003. Presentation at the 2003 ACM SIGKDD International Conference. Washington, D.C. (August 2003). (With J. Kleinberg)

An Overview of Database Research at Cornell. Microsoft Research, Data Mining and Exploration Group. Seattle, Washington (July 2003).

PROFESSIONAL ACTIVITIES

Committees:

Member, ACM SIGKDD Curriculum Committee.

Member, Workshop Steering Committee, Workshop on GeoSensor Networks. Portland, Maine (October 2003).

Editorial Boards:

Member, Editorial Board, *Journal of Privacy Technology* (2004–present).

Member, Editorial Board, *Machine Learning* (2003–present).

Action Editor, *Data Mining and Knowledge Discovery* (2003–present).

Associate Editor, *Bulletin of the Technical Committee on Data Engineering,* IEEE Computer Society (2002–2004).

Member, Editorial Board, *Journal of Database Management* (2000–present).

Editorial Activities:

Guest Editor, *Data Engineering Bulletin.* Special Issue on Privacy and Security (March 2004).

Program Chairmanships

Area Chair, Twentieth ICML, Washington, D.C. (August 2003).

Co-chair, KDD–Cup, Ninth ACM SIGKDD International Conference, Washington, D.C. (August 2003).

Program Committees:

Member, Program Committee, Twenty-third ACM SIGMOD–SIGACT–SIGART (Special Interest Group on Artificial Intelligence) Symposium on Principles of Database Systems. Paris, France (June 2004).

Member, Program Committee, Second NSF/NIJ Symposium on Intelligence and Security Informatics (ISI 2004). Tucson, Arizona. (June 2004).

Member, Program Committee, Ninth ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery. Paris, France (June 2004).

Member, Program Committee, Thirteenth International World Wide Web Conference (WWW 2004). New York, New York (May 2004).

Member, Program Committee, Fourth SIAM International Conference on Data Mining (SDM 2003). Orlando, Florida. (April 2004).

Member, Program Committee, Twenty-fourth International Conference on Distributed Computing Systems. Tokyo, Japan (March 2004).

Member, Program Committee, Twentieth IEEE International Conference on Data Engineering (ICDE). Boston, Massachusetts (March 2004).

Member, Program Committee, Fifth International Conference on Mobile Data Management. San Jose, California (January 2004).

Member, Program Committee, ICDM. Melbourne, Florida (December 2003).

Member, Program Committee, Ninth ACM SIGKDD International Conference. Washington, D.C. (August 2003).

Member, Program Committee, Fifteenth International Conference on Scientific and Statistical Database Management. Boston, Massachusetts (July 2003).

UNIVERSITY ACTIVITIES

Member, Department Chair Search Committee, Department of Computer Science (2004).

Faculty Mentor, Japanese Graduate Student Association (January 2003–present).

AWARDS AND HONORS

Cornell University Provost's Award for Distinguished Scholarship, 2004.

Alfred P. Sloan Foundation Fellowship, 2003.

NSF CAREER award, 2002.

Cornell College of Engineering James and Mary Tien Excellence in Teaching Award, 2001.

IBM Faculty Development Award, 2000 and 2001.

Tarleton Gillespie

Visiting Assistant Professor

CIS, joint with Communication and Science and Technology Studies

tlg28@cornell.edu

http://www.sts.cornell.edu/viewprofile.php?ProfileID=5

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

His past research used recent disputes over copyright and the Internet to analyze the historical contest over the nature of authorship, law, and technology. He uses these cases to examine the cultural and institutional arrangements surrounding media and Internet 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 battle, but more often are often brought together to regulate knowledge production.

He is continuing work on his first book, which will address this "legal turn to technology". What are the implications when, rather than legislating individual behavior, or mandating design for some broad public benefit, lawmakers legislate technology so that it specifically regulates individual behavior? Cases include controversies around copy protection and Internet filtering, DVD encryption and the DMCA, and the most recent example, the FCC's mandated "broadcast flag" for digital television—a controversy certain to impact the ongoing public discussion of this broader shift, and the co-production of law, digital media, and commercial institutions. Using these cases as examples, he is investigating not only this shift in American legal, political, and commercial strategy, and its consequences for legal doctrine, political subjectivity, and technological innovation; but also the shifting relationships between the corporate and governing institutions it represents and depends on, and the impact of these institutional shifts for digital culture.

PUBLICATIONS

"Manufacturing a Principle: 'End-to-end' in the Design of the Internet." [submitted] *Social Studies of Science*.

"Copyright and Commerce: The DMCA, Trusted Systems, and the Stabilization of Distribution." *The Information Society* 20(4) (June 2004).

LECTURES

"Manufacturing a Principle: 'End-to-end' in the Design of the Internet". Presented to the Department of Science and Technology Studies, Cornell University (January 2004).

"The Symbolic Shape of Media Technologies". Presented at the Conference of the Society for the Social Studies of Science (4S), Atlanta, Georgia (October 2003).

Panel:

"The Making of Language and Metaphor in Technoscientific Discourse". Organized for the 4S Conference, Atlanta, Georgia (October 2003).

Panel and comment:

"Building Digital Stuff". Co-organized (With P. Sengers) for "Connecting STS: The Academy, The polity, and the World". Cornell University (September 2003).

PROFESSIONAL ACTIVITIES

Editorial Advisor, *Social Studies of Science*.

AWARDS AND HONORS

Humanities Council Research Grant, Society for the Humanities, Cornell University (2003).

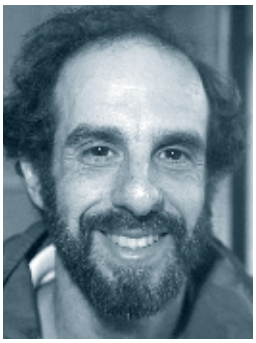

Paul Ginsparg

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 in the Department of Physics and the Faculty of CIS. He has been an A. P. Sloane Fellow and a Department of Energy 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.

PUBLICATIONS

"Can Peer Review Be Better Focused?". *Science and Technology Libraries* 22(3/4): 5–18 (2003).

"Scholarly Information Architecture, 1989–2015". *Data Science Journal* 3(4): 29–37 (February 2004).

"e-Print ArXiv Project". *T–Division Sixtieth Anniversary Book* (F. Harlow and J. Sprouse, eds.), Los Alamos National Laboratory (2003).

"Mapping Subsets of Scholarly Information". *Proceedings of the National Academy of Sciences* 101: 5236–5240 (April 6, 2004). (With P. Houle, T. Joachims, and J.-H. Sul).

"Overview of the 2003 KDD Cup". ACM *SIGKDD Explorations* 5 (2): 149–151 (December 2003). (With J. Gherke and J. Kleinberg).

"Scholarly Information Network". *Complex Networks, Lecture Notes in Physics* (E. Ben-Naim, H. Frauenfelder, and Z. Toroczkai, eds.), Springer Verlag (2004).

LECTURES

Colloquium, Google, Mountain View, California.

Dibner/Sloan History of Recent Science and Technology Conference, M.I.T., Cambridge, Massachusetts.

Colloquium, Department of Biochemistry, Harvard Medical School, Cambridge, Massachusetts.

Vollmer Fries Lecture, Rensselaer Polytechnic Institute, Troy, New York.

NAS Sackler Colloquium, Mapping Knowledge Domains, University of California at Irvine.

Keynote Talk, "Networks: Structure, Dynamics and Function" conference, Santa Fe, New Mexico.

Colloquium, Aspen Center for Physics, Aspen, Colorado.

Seminar, MacArthur Fellows Meeting, Racine, Wisconsin.

Colloquium, Society of Physics Students, Cornell University, Ithaca, New York.

Seminar, Students of Science Reporting, Cornell University, Ithaca, New York.

PROFESSIONAL ACTIVITIES

Member, Advisory Board, Public Library of Science.

Member, National Advisory Board, NIH PubMedCentral.

Member, Publications Oversight Committee, American Physical Society.

Member, Advisory Board, French Centre National Recherche Scientifique "Centre Communication Scientifique Directe".

Member, Advisory Board, "Journal Club for Condensed Matter Physics".

UNIVERSITY ACTIVITIES

Member, Cornell University Faculty Advisory Board on Information Technology.

Member, Cornell University Library Board.

Member, Cornell Computing and Information Science Council.

Member, Cornell Information Science Graduate Field Committee.

AWARDS AND HONORS

John D. and Catherine T. MacArthur Foundation Fellow, 2002.

Fellow of the American Physical Society, 2000.

Lingua Franca "Tech 20" Award, 1999.

P.A.M. (Physics Astronomy Math) award from the Special Libraries Association, 1998.

Distinguished Performance Award, Los Alamos National Laboratory, 1992.

DOE Outstanding Junior Investigator, 1986–91.

A. D. White Fellow, Cornell University, 1977–81.

A. P. Sloan Foundation Fellow, 1986–90.

Harvard Society of Fellows, 1981–84.

Carla P. Gomes

Associate Professor

CIS, joint with Applied Economics and Management

Director, Intelligent Information Systems Institute (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 a M.Sc. in applied mathematics from the University of Lisbon.

Gomes is currently the director of IISI at Cornell. Her 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)*.

PUBLICATIONS

"The Challenge of Generating Spatially Balanced Scientific Experiment Designs". *Proceedings of the Sixth International Conference on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems (CP–AI–OR '04),* Nice, France (2004). (With M. Sellmann, C. van Es, and H. van Es).

"Approximations and Randomization to Boost CSP Techniques". *Annals of Operations Research* (2004). (With D. Shmoys).

"Regular–SAT: A Many-valued Approach for Solving Combinatorial Problems". *Discrete Applied Mathematics,* Elsevier (2004). (With R. Bejar, A. Cabiscol, C. Fernandez, and F. Manya).

"Sensor Networks and Distributed CSP: Communication, Computation and Complexity". *Artificial Intelligence Journal* (2004). (With R. Bejar, C. Domshlak, C. Fernandez, B. Krishnamachari, B. Selman, and M. Valls).

"Complete Randomized Backtrack Search". *Constraint and Integer Programming: Toward a Unified Methodology* (M. Milano, ed.), Kluwer, 233–283 (2003).

"Pareto-like Distributions in Random Binary CSP". *Frontiers in Artificial Intelligence and Applications—Artificial Intelligence Research and Development,* IOS Press 100: 451–461, ISSN 0922-6389 (2003). (With C. Bessiere, C. Fernandez, and M. Valls).

"Backdoors To Typical Case Complexity". *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI 2003)* (2003). (With R. Williams and B. Selman).

LECTURES

Heavy-tailed Behavior in Computation. Institute for Mathematics and its Applications (IMA). (2004).

Exact Randomized Search Methods. Institute for Operations Research and the Management Sciences. Banff, Alberta (2004).

Backdoors in Combinatorial Problems. Canadian Operational Research Society, Banff, Alberta. (2004).

Randomization, Structure, and Complexity in Combinatorial Optimization. Institute for Pure and Applied Mathematics, University of California at Los Angeles. (2003).

PROFESSIONAL ACTIVITIES

AAAI Executive Council (2002–2005).

Editorial Board:

Journal of Artificial Intelligence Research

Journal of Knowledge Engineering Review

Journal of Satisfiability

International Journal on Artificial Intelligence Tools

Program Committees:

Sixth International Conference on the Integration of AI and OR in CP for Combinatorial Optimization (CP–AI–OR) (2004).

International Conference on Automated Planning and Scheduling 2003–04.

Seventh International SAT Conference (2004).

Co-chair, AFRL/IISI Workshop on Mixed Initiative Decision Making, Ithaca, New York (2003). (With H. Kautz and C. Domshlak).

Ninth International Conference on the Principles and Practice of Constraint Programming, CP 2003.

Sixth International SAT Conference (2003).

UNIVERSITY ACTIVITIES

Director, Intelligent Information Systems Institute (IISI).

Donald Greenberg

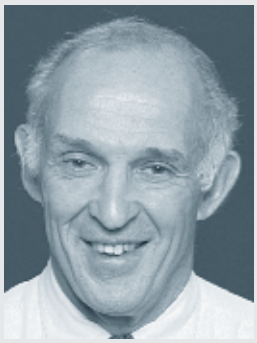

Jacob Gould Schurman Professor

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

Since 1966, Dr. Greenberg has been researching and teaching in the field of computer graphics. During the last fifteen years, he has been primarily concerned with research advancing the state-of-the-art in computer graphics and with utilizing these techniques as they may be applied to a variety of disciplines. His specialties include real-time realistic-image generation, geometric modeling, color science, and computer animation. He presently teaches the computer graphics courses in Computer Science, computer-aided design in Architecture, computer animation in Art and CIS, and technology strategy in the Johnson Graduate School of Management.

Working with the General Electric Visual Simulation Laboratory, he produced a sophisticated computer graphics movie, "Cornell in Perspective", as early as 1971. He is the author of hundreds of articles on computer graphics (including two published in *Scientific American,* May 1974 and February, 1991, both of which have been highly publicized); and he has lectured extensively on the uses of computer graphics techniques in research applications.

He was the founding Director of the National Science Foundation Science and Technology Center for Computer Graphics and Scientific Visualization. He has been the Director of the Program of Computer Graphics for thirty-one years and was the originator and former Director of the Computer-aided Design Instructional Facility at Cornell University.

In 1987, he was awarded the prestigious ACM SIGGRAPH Steven Coons Award for outstanding creative contributions to computer graphics.

Of his two hundred plus graduate students, many have gone on to become leaders in the fields of computer graphics, computer animation, and computer-aided design for architecture. Five have now won Hollywood "Oscars".

Dr. Greenberg joined the faculty of Cornell 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 University 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. He has taught courses in structural analysis and design, architectural design, shell structures, reinforced concrete, and computer applications in architecture.

PUBLICATIONS

"Combining Edges and Points for Interactive High-quality Rendering". *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH* (July 2003). (With K. Bala and B. Walter).

"Detail Synthesis for Image-based Texturing". *Interactive 3D Graphics (I3D)* (April, 2003). (With R. Ismert and K. Bala).

LECTURES

Cornell Silicon Valley, Cornell Entrepreneur Network, and Cornell Entrepreneurship and Personal Enterprise Program New Year Entrepreneurship Event (January 2004).

"Research Challenges in Computer Graphics," Computer Science and Telecommunications Board, Planning Meeting on Fundamental Research Challenges in Computer Graphics (December 2003).

" Design Environments of the Future: Learning form the Past," New York Chapter, American Institute of Architects (October 2003).

"Computer Graphics". Beginning with Children School (June 2003).

PROFESSIONAL ACTIVITIES

Founding Director, National Science and Technology Center for Computer Graphics and Scientific Visualization.

Technical Advisory Board, Intel Corporation.

UNIVERSITY ACTIVITIES

Professor, Architecture, Art, Computer Science, and Johnson Graduate School of Management.
Member, Fields of Civil Engineering and History of Science and Technology.

AWARDS AND HONORS

Fellow, American Association for the Advancement of Science, 2002.
Honorary Doctoral Degree, New Jersey Institute of Technology, 1999.
Architectural Schools Computing Association Creative Research Award in Architecture, 1997.
Fellow, ACM, 1994.
Founding Fellow, American Institute of Medical and Biological Engineering, 1992.
Member, National Academy of Engineering, 1991.
NCGA Academic Award (highest educational award given by the National Computer Graphics Association), 1989.
ACM SIGGRAPH Steven A. Coons Award for Outstanding Creative Contributions to Computer Graphics, 1987.

David Gries

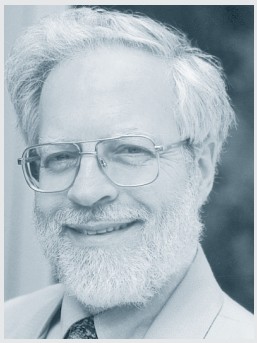

Associate Dean of Engineering for Undergraduate Programs

Professor of Computer Science

Cornell Weiss Presidential Fellow

gries@cs.cornell.edu

http://www.cs.cornell.edu/gries/

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 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, Gries and colleague F. B. Schneider wrote a text, *A Logical Approach to Discrete Math,* which makes a usable "calculational logic" the foundation for almost all the discrete math topics.

Earlier, with his son, Gries developed a "livetext"—a text that comes on a CD and has over 250 two-to-three–minute recorded lectures with synched animation, as well as other innovative features. A paper text to accompany it, *Multimedia Approach to Programming Using Java,* will appear in fall 2004.

Gries received the Dr. rer. nat. degree from the Munich Institute of Technology in 1966. He chaired the Computing Research Association during the time it opened an office in Washington and began representing the research community.

LECTURES

The Mathematics of Programming and Why We Should Teach It. Consortium for Computing Sciences in Colleges in the Northeast 2004, Union College, New York (April 24, 2004).

Fundamentals of Parallelism Workshop on Concurrency. Sri Sathya Sai Institute of Higher Learning. India. (July 4–6, 2003).

PROFESSIONAL ACTIVITIES

Member, Steering Group, American Association for the Advancement of Science (AAAS) Section T on Information, Computing, and Communication (1999–2003).

Member, IEEE Piore Award Committee.

Editorial Board, *Acta Informatica, Information Processing Letters.*

Main Editor, *Acta Informatica* (1982–present).

Series Editor, *Springer Verlag Texts and Monographs in Computer Science* (1973–present).

Managing Editor, *Information Processing Letters* (1972–2003).

UNIVERSITY ACTIVITIES

Associate Dean for Undergraduate Programs.

AWARDS AND HONORS

Doctor of Science (Honorary Degree), Oxford University, Miami, Ohio, 1999.

Doctor of Laws (Honorary Degree), Daniel Webster College, Nashua, New Hampshire, 1996.

ACM Karl V. Karlstrom Outstanding Educator Award, 1996.

Cornell University Computer Science Department Faculty of the Year, 1995–96. (ACSU)

Weiss Presidential Fellow (for contributions to undergraduate education). Cornell University, 1995.

Taylor L. Booth Award Education Award, IEEE Computer Society, 1994.

ACM Fellow (Charter member: among the first group to be inducted), 1994.

ACM SIGCSE Award for Outstanding Contributions to Computer Science Education, 1991.

Computing Research Association Award for Service to the Computing Community, 1991.

Fellow of the AAAS, 1990.

Chosen by a Cornell Merrill Presidential Scholar (Thomas Yan) as the faculty member who had the most positive influence on his education at Cornell, 1990.

Clarke Award for Excellence in Undergraduate Teaching; College of Arts and Science, Cornell University, 1986–87.

Education Award, American Federation of Information Processing Societies, 1986.

Guggenheim Fellowship, 1983–84.

ACM Programming Systems and Languages Paper Award, 1977 (With S. Owicki).

Zygmunt J. Haas

Professor of the School of Electrical and Computer Engineering, and Member of the Graduate Fields of Computer Science and the Center for Applied Mathematics

haas@ece.cornell.edu

http://people.ece.cornell.edu/haas/

Zygmunt J. Haas received his Ph.D. 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 1995, he joined the faculty of the School of Electrical and Computer Engineering at Cornell University.

Dr. 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 IEEE and ACM conferences, and serves as editor of several journals and magazines, including the *ACM/IEEE Transactions on Networking, the IEEE Transactions on Wireless Communications,* and the Kluwer's journal on *Wireless Networks*. He has been a guest editor of three issues of the IEEE's *Journal of Selected Areas on Communication* ("Gigabit Networks", "Mobile Computing Networks", and "Ad-Hoc Networks"). Dr. Haas is a senior member of IEEE and the Chair of the IEEE Technical Committee on Personal Communications. He is an IEEE/COMSOC Distinguished Speaker.

Professor Haas's current interests include: mobile and wireless communication and networks, performance evaluation of communication systems, and biologically-inspired systems and networks. He heads the Wireless Networks Laboratory (WNL) (*http://wnl.ece.cornell.edu*) 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 protocol—the Zone Routing Protocol—which is currently an IETF draft. The WNL 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, 2000, and 2003.

PUBLICATIONS

"On the Scalability of Wireless Networks with Omnidirectional Antennas". [accepted] *Wireless Communications and Mobile Computing*. Wiley and Sons (2004). (With O. Arpacioglu).

"Independent Zone Routing: An Adaptive Hybrid Routing Framework for Ad Hoc Wireless Networks". [accepted] *ACM/IEEE Transactions on Networking*. (With P. Samar and M. Pearlman).

"Implementation of Virtual Factory Communication Systems Using Manufacturing Message Specification". In *The Handbook of Industrial Information Technology,* CRC Press (March 2004). (With D.-S. Kim and W. Kwon).

"Concurrent Search of Mobile Users in Cellular Networks". *ACM/IEEE Transactions on Networking* 12(1) (February 2004). (With R.-H. Gau).

"On Optimizing the Backoff Interval for Random Access Schemes". *IEEE Transactions on Communications* 51(12): 2081–2090 (December 2003). (With J. Deng).

"Efficient Computations for Evaluating Extended Stochastic Petri Nets Using Algebraic Operations". *International Journal of Control, Automation and Systems* 1(4) (December 2003). (With D.-S. Kim, H. Moon, and W. Kwon).

"Performance Evaluation of Modified IEEE 802.11 MAC for Multichannel Multihop Ad-Hoc Network". *Journal of Interconnection Networks (JOIN) Special Issue on Advanced Information Networking: Architectures and Algorithms* 4(3): 345–359 (2003). (With J. Li, M. Sheng, and Y. Chen).

Conference Papers:

"Impact of Concurrent Transmissions on Downstream Throughput in Multihop Cellular Networks". In *Proceedings of the 2004 IEEE International Conference on Communications (ICC 2004),* Paris, France (June 20–24, 2004) (With J. Cho).

"On the Scalability and Capacity of Wireless Networks with Omnidirectional Antennas". *Third International Symposium on Information Processing in Sensor Networks (IPSN 2004),* Berkeley, California (April 27–28, 2004). (With O. Arpacioglu).

"Analyzing Split Channel Medium Access Control Schemes with ALOHA Reservations". *Second International Conference on Ad-Hoc Networks and Wireless,* Montreal, Canada (October 8 –10, 2003). (With J. Deng and Y. Han).

"A Novel Packet Scheduling in an Enhanced Joint CDMA/NC–PRMA Protocol for Wireless Multimedia Communications". IEEE Vehicular Technology Conference (*VTC 2003*), Orlando, Florida (October 4–9, 2003). (With S. Lee, A. Ahmad, and K. Kim).

"Secure Data Transmission in Mobile Ad Hoc Networks". *ACM Workshop on Wireless Security (WiSe 2003),* San Diego, California (September 19, 2003). (With P. Papadimitratos).

Book Chapters:

"A Sensor Network for Biological Data Acquisition". *Handbook on Sensor Networks* (M. Ilyas, ed.) CRC Press (2004). (With T. Small, A. Purgue, and K. Fristrup).

"Secure Mobile Ad Hoc Networks". *Ad Hoc Wireless Networking,* Kluwer Academic Publisher (2004). (With P. Papadimitratos).

"Hybrid Routing: The Pursuit of an Adaptable and Scalable Routing Framework for Ad Hoc Networks". *Ad Hoc Wireless Networking,* Kluwer Academic Publisher (2004). (With P. Samar and M. Pearlman).

LECTURES

"Scalability of Wireless Ad-Hoc Networks". Keynote speech, NSF International Workshop on Theoretical and Algorithmic Aspects of Sensor, Ad Hoc Wireless, and Peer-to-peer Networks, Radisson Bahia Mar Beach Resort, Fort Lauderdale, Florida (February 20–21, 2004).

"The Analysis of Power-controlled MAC Layer for Wireless Ad-Hoc Networks". VTC 2003, Orlando, Florida (October 4–9, 2003). (With H. Inaltekin).

"Throughput Enhancement by the Multihop Relaying in Cellular Radio Networks with Non-uniform Traffic Distribution". VTC 2003, Orlando, Florida (October 4–9, 2003). (With J. Cho).

PROFESSIONAL ACTIVITIES

Editorial Boards:

IEEE Transactions on Wireless Communications.

ACM/IEEE Transactions on Networking.

IEEE Communications Magazine.

Elsevier Ad Hoc Networks Journal.

Wireless Communications and Mobile Computing, Journal, John Wiley and Sons.

Journal of High Speed Networks.

ACM/Kluwer Wireless Networks.

Committees:

General Chair, ACM Mobicom 2004, Philadelphia, Pennsylvania (September 26–October 1, 2004).

Chair, IEEE TCPC Award Committee (September 2003).

Chair, IEEE Ithaca Section (January 1, 2003–December 31, 2003).

Member, Steering Committee of ACM MobiCom.

Membership in Technical Program Committees:

Wireless Communications and Networking Conference 2004 (WCNC 2004), Atlanta, Georgia (March 21–25, 2004).

The 2003 ACM Workshop on Security of Ad Hoc and Sensor Networks, George Mason University, Fairfax, Virginia (October 31, 2003).

IPSN 2004, Berkeley, California (April 27–28, 2004).

The Fifth IEEE Conference on Mobile and Wireless Communications Networks (IEEE MWCN 2003), Singapore (October 27–29, 2003).

Fall IEEE VTC, Orlando, Florida (October 4–9, 2003).

The Ninth Annual International Conference on Mobile Computing and Networking, San Diego, California (September 14–19, 2003).

The Fourteenth IEEE International Symposium on Personal, Indoor, and Mobile Radio Communications, Beijing, China (September 7–10, 2003).

The Third Workshop on Applications and Services in Wireless Networks (ASWN 2003), Berne, Switzerland (July 2–4, 2003).

UNIVERSITY ACTIVITIES

Elected Member, ECE Policy Committee.

Member, ECE Recruiting Committee.

Member, ECE Ad Hoc Tenure Promotion Committee.

Member, Ad Hoc Tenure Promotion Committee, College of Engineering.

AWARDS AND HONORS

Distinguished Lecturer, IEEE Communications Society (2004–2005).

Michael Tien '72 Award, Cornell College of Engineering Excellence in Teaching Award for 2002–2003 (November 2003).

Joseph Halpern

Professor

Director of Graduate Studies

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; two books, *Reasoning About Knowledge* and *Reasoning About Uncertainty;* and more than 200 technical publications.

PUBLICATIONS

Reasoning About Uncertainty, M.I.T. Press (2003).

Reasoning About Knowledge, M.I.T. Press (2003).). (Paperback edition; originally published in 1995.) (With R. Fagin, Y. Moses, and M. Vardi).

"Common Knowledge Revisited". In *Knowledge Contributors* (V. Henricks, K. Jorgensen, and S. Pedersen, eds.), Kluwer, 2003, 87–104. (With R. Fagin, Y. Moses, and M. Vardi). [This is a reprint of an article originally published in *Annals of Pure and Applied Logic* in 1999.]

"A Computer Scientist Looks at Game Theory". *Games and Economic Behavior* 45(1): 114–132 (2003).

"A Logical Reconstruction of SPKI". *Journal of Computer Security* 11(4): 581–614 (2003). (With R. van der Meyden).

"Representation Dependence in Probabilistic Inference". *Journal of Artificial Intelligence Research* 319–356 (2004). (With D. Koller).

"Using First-order Logic to Reason About Policies". *Proceedings of the Sixteenth IEEE Computer Security Foundations Workshop* 187–201, (July 2003). (With V. Weissman).

"Anonymity and Information Hiding in Multiagent Systems". *Proceedings of the Sixteenth IEEE Computer Security Foundations Workshop* 75–88 (July 2003). (With K. O'Neill).

"Great Expectations. Part I: On the Customizability of Generalized Expected Utility". *Proceedings of the IJCAI* 291–296 (August 2003). (With F. Chu).

"Great Expectations. Part II: Generalized Expected Utility as a Universal Decision Rule". *Proceedings of the IJCAI* 297–302 (August 2003). (With F. Chu).

"Responsibility and Blame: A Structural-model Approach". *Proceedings of the IJCAI* 147–153 (August 2003). (With H. Chockler).

"A Logic for Reasoning About Evidence". *Proceedings of the Nineteenth Conference on Uncertainty in AI* 297–304 (August, 2003). (With R. Pucella).

"Rational Secret Sharing and Multiparty Computation". *Proceedings of the Thirty-sixth ACM Symposium on the Theory of Computing* (June 2004). (With V. Teague).

"Sleeping Beauty Reconsidered: Conditioning and Reflection in Asynchronous Systems". *Ninth International Conference on Principles of Knowledge Representation and Reasoning (KR 2004)* (June 2004).

"Intransitivity and Vagueness". *Ninth KR* (June 2004).

LECTURES

Characterizing the Common Prior Assumption, Fifteenth Italian Meeting on Game Theory and Applications, Urbino, Italy (July 2003).

Rational Secret Sharing and Multiparty Computation, Cowles Foundation Workshop on Complexity in Economic Theory, Yale University (September 2003).

Reasoning About Uncertainty in Multiagent Systems, The International Symposium on Modern Computing, Ames, Iowa (October 2003).

A Decision-theoretic Approach to the Design, Analysis, and Specification of Systems, University of Iowa (October 2003).

Reasoning About Uncertainty in Multiagent Systems, Brown Bag Lunch (September 2003).

Updating Probabilities, AI Seminar (January 2004).

Rational Secret Sharing and Multiparty Computation, Workshop on Complexity in Game Theory, Theory Seminar (April 2004).

A Data-acquisition Model for Learning and Cognitive Development and Its Implications for Autism, Cognitive Science Colloquium (May 2004).

Great Expectations. Part I: On the Customizability of Generalized Expected Utility, IJCAI, Acapulco, Mexico (August 2003).

Great Expectations. Part II: Generalized Expected Utility as a Universal Decision Rule, IJCAI, Acapulco, Mexico (August 2003).

Responsibility and Blame: A Structural-model Approach, IJCAI, Acapulco, Mexico (August 2003).

A Logic for Reasoning about Evidence, Nineteenth Conference on Uncertainty in AI, Acapulco, Mexico (August 2003).

Sleeping Beauty Reconsidered: Conditioning and Reflection in Asynchronous Systems, Ninth KR, Whistler, British Columbia (June 2004).

Intransitivity and Vagueness, Ninth KR, Whistler, British Columbia (June 2004).

Using First-order Logic to Reason about Policies, Rome Air Force Lab (August 2003).

PROFESSIONAL ACTIVITIES

Fellow, American Association of Artificial Intelligence.

Fellow, ACM.

Consulting Editor, *Chicago Journal of Computer Science*.

Editorial boards, *Artificial Intelligence Journal and Journal of Logic and Computation*.

Coordinator, CoRR (Computing Research Repository).

Member, Advisory Board, LICS.

President, Board of Directors, Corporation for Theoretical Aspects of Reasoning About Knowledge.

Program Co-chair, Twentieth Conference on Uncertainty in AI.

AWARDS AND HONORS

Fellow of the ACM (2002).

Fulbright Fellow (2001–02).

Guggenheim Fellow (2001–02).

Milner Lecturer at Edinburgh University (May 2000).

Editor-in-chief, *Journal of the ACM* (1997–2003).

Gödel Prize for outstanding paper in the area of theoretical computer science for "Knowledge and Common Knowledge in a Distributed Environment" (1997).

Fellow of the American Association of Artificial Intelligence (1993).

First Plateau Invention Achievement Award, IBM (1992).

Best Paper Awards at IJCAI 1985 and IJCAI 1989.

Outstanding Innovation Awards from IBM:

for work on clock synchronization (1988).

for work on reasoning about knowledge (1987).

Fourteen papers invited to special issues of journals devoted to best papers in conferences.

Juris Hartmanis

Emeritus Walter R. Read Professor of Computer Science and Engineering

Turing Award Winner

jh@cs.cornell.edu

http://www.cs.cornell.edu/annual_report/00-01/bios.htm#hartmanis

Juris Hartmanis obtained his Ph.D. from the California Institute of Technology in 1955. In 1965, he founded CS and was its first chairman.

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 are in computational complexity, structure and management of research organizations. His main focus has been 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.

PUBLICATIONS

"Separation of Complexity Classes". *Journal of the ACM, Fiftieth Anniversary Special Issue* 50(1): 58–62 (January 2003).

PROFESSIONAL ACTIVITIES

Editor, *Springer-Verlag Lecture Notes in Computer Science; Fundamenta Informatica; Journal of Computer and System Sciences.*

Member, Science Steering Committee, Santa Fe Institute for Complex Systems.

Member, Science Board, Santa Fe Institute for Complex Systems.

Member, Advisory Board, Foundations of Computer Science.

Member, Advisory Board, *European Association for Theoretical Computer Science (EATCS) Monographs on Theoretical Computer Science.*

UNIVERSITY ACTIVITIES

Member, CIS Council.

AWARDS AND HONORS

Grand Medal, Latvian Academy of Science, 2001.

CRA Distinguished Service Award, 2000.

Doctor of Science (honoris causa), University of Missouri, Kansas City, 1999.

B. Bolzano Gold Medal of the Academy of Science, Czech Republic, 1995.

Doctor of Science (honoris causa), University of Dortmund, Germany, 1995.

Senior U.S. Scientist Humboldt Award, Max Plank Institute, Saarbruecken, Germany, 1993–94.

ACM Turing Award (shared with R. E. Stearns, 1993).

Fellow, American Academy of Arts and Sciences, 1992.

Member, Latvian Academy of Science, 1990.

Member, National Academy of Engineering, 1989.

Fellow, American Association for the Advancement of Science, 1981.

Mark Heinrich

Assistant Professor of the School of Electrical and Computer Engineering and Member of the Graduate Field of Computer Science

heinrich@csl.cornell.edu

http://www.csl.cornell.edu/~heinrich/

Mark Heinrich, an assistant professor in ECE at Cornell in 2004, a co-founder of its Computer Systems Laboratory, and an IISI member, is currently an associate professor in the School of Computer Science at the University of Central Florida.

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 co-design.

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 co-founder 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.

PUBLICATIONS

"Exploring Virtual Network Selection Algorithms in DSM Cache Coherence Protocols". [accepted] *IEEE Transactions on Parallel and Distributed Systems (TPDS)* (August 2004). (With M. Chaudhuri).

"Architectural Extensions for Executing Coherence Protocols on Multithreaded Processors with Integrated Memory Controllers". [accepted] *Proceedings of the Thirty-first International Symposium on Computer Architecture* (June 2004). (With M. Chaudhuri).

"Architectural Support for Uniprocessor and Multiprocessor Active Memory Systems". *IEEE Transactions on Computers,* 53(3): 288–307, March 2004. (With D. Kim, M. Chaudhuri, and E. Speight).

"The Impact of Negative Acknowledgments in Shared Memory Scientific Applications". *IEEE TPDS* 15(2): 134–150 (February 2004). (With M. Chaudhuri).

"Exploiting Active CMP–based Devices in System Area Networks". *Third Workshop on System Area Networks,* held in conjunction with the International Symposium on High-performance Computer Architecture (HPCA) (February 2004). (With M. Hao).

"Latency, Occupancy, and Bandwidth in DSM Multiprocessors: A Performance Evaluation". *IEEE Transactions on Computers* 52(7): 862–880 (July 2003). (With M. Chaudhuri, C. Holt, and others).

UNIVERSITY ACTIVITIES

Member, Intelligent Information Systems Institute.

Member, College of Engineering Teaching Awards Committee.

Project Coordinator and Committee Member, CURIE Academy for Women in Engineering.

AWARDS AND HONORS

CAREER award, 2000–2004, "Flexible Architectures for Data-intensive Computing".

Sheila S. Hemami

Associate Professor of the School of Electrical and Computer Engineering and Member of 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) from the University of Michigan, and her M.S.E.E. (1992) and Ph.D. (1994) degrees from Stanford University.

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.

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.

She is a member of the IEEE, Eta Kappa Nu, and Tau Beta Pi.

PUBLICATIONS

"An Embedded Image Coding System Based on Tarp Filter with Classification". *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing,* Montreal, Quebec (May, 2004). (With C. Tian).

"Distortion Optimized Multiple Channel Image Transmission Under Delay Constraints". *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing,* Montreal, Quebec, Canada (May 2004). (With W. Xu).

"Scalable Image Embeddings From Arbitrary Wavelet-based Perceptual Models". *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing,* Montreal, Quebec (May 2004). (With M. Gaubatz).

"Local Entropy Estimation for Low-rate Wavelet Image Coding". *Proceedings of Thirty-eighth Annual Conference on Information Sciences and Systems,* Baltimore, Maryland (March 2004). (With M. Gaubatz).

"Efficient Bit Allocation for Dependent Video Coding". *Proceedings of the Data Compression Conference 2004,* Snowbird, Utah (March 2004). (With Y. Sermadevi).

"Sequential Design of Multiple Description Scalar Quantizers". *Proceedings of Data Compression Conference 2004,* Snowbird, Utah (March 2004). (With C. Tian).

"A Metric for Continuous Quality Evaluation of Compressed Video with Severe Distortions". *Signal Processing: Image Communication* 19(2):133–146 (February 2004). (With M. Masry).

"MINMAX Rate Control with a Perceived Distortion Metric". *Proceedings of the International Society for Optical Engineering—Visual Communications and Image Processing,* San Jose, California (January 2004). (With Y. Sermadevi and M. Masry).

"Robust Adaptive Transmission of Images and Video over Multiple Channels". *Signal Processing: Image Communication* 18(10): 981–1000 (November 2003). (With W. Xu).

"Quantifying the Visual Quality of Wavelet-compressed Images Based on Local Contrast, Visual Masking, and Global Precedence". *Asilomar Conference on Signals, Systems, and Computers,* Pacific Grove, California (November 2003). (With D. Chandler and M. Masry).

"Digital Watermarking Using Local Contrast-based Texture Masking". *Asilomar Conference on Signals, Systems, and Computers,* Pacific Grove, California (November 2003). (With M. Masry and D. Chandler).

"Linear Approximations for Rate Control in Video Coding". *IEEE International Conference on Image Processing 2003,* Barcelona, Spain (September 2003). (With Y. Sermadevi).

"Frame Rate Preferences in Low Bit Rate Video". *IEEE International Conference on Image Processing 2003,* Barcelona, Spain (September 2003). (With G. Yadavalli and M. Masry).

"Effects of Natural Images on the Detectability of Simple and Compound-wavelet Sub-band Quantization Distortions". *Journal of the Optical Society of America: A* 20(7) (July 2003). (With D. Chandler).

AWARDS AND HONORS

CAREER award, 1997.

Eta Kappa Nu C. Holmes MacDonald Outstanding Teaching Award (a national award), 2002.

Finalist, Eta Kappa Nu Outstanding Young Electrical Engineer Award, 2002.

John Hopcroft

Professor

Turing Award Winner

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 a Doctor of Humanities Degree (Honoris Causa), in 1990.

PUBLICATIONS

"Natural Communities in Large Linked Networks". Poster paper *Ninth ACM SIGKDD International Conference,* Washington D.C., (August 24–27, 2003). (With O. Khan, B. Kulis, and B. Selman).

"Tracking Evolving Communities in Large Linked Networks". [accepted] Arthur M. Sackler Symposium, *Proceedings of the National Academy of Sciences,* Irvine California. (With O. Khan, B. Kulis, and B. Selman).

LECTURES

"The Future of Computer Science". University of North Texas (April 5, 2002).

PROFESSIONAL ACTIVITIES

Member, Board on Mathematical Science and Applications, National Academy of Sciences.

Member, Science Advisory Board, Packard Foundation.

Member, Financial Management Committee, Society of Applied Mathematics.

Trustee, Boyce Thompson Institute.

Editor and member of Executive Committee, *Algorithmica.*

Editor, *International Journal of Computational Geometry and Applications.*

Editor, *Journal of Computer and System Sciences.*

Associate Editor, *Information Sciences.*

Council Delegate, American Association for the Advancement of Science.

Chair, International Advisory Committee on Informatics and Engineering, National College of Ireland.

UNIVERSITY ACTIVITIES

Chair, Provost's Committee on Intellectual Property.

AWARDS AND HONORS

Fellow, Association for Computing Machinery, 1994.

Member, National Academy of Engineering, 1989.

Fellow, American Academy of Arts and Sciences, 1987.

Fellow, American Association for the Advancement of Science, 1987.

Fellow, Institute of Electrical and Electronics Engineers, 1987.

Recipient, ACM Turing Award, 1986.

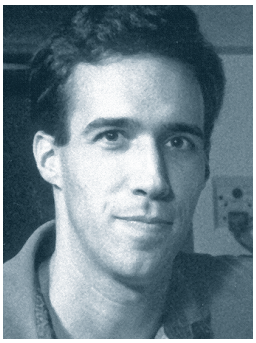

Daniel Huttenlocher

John P. and Rilla Neafsey Professor of Computing, 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 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 has been 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 CASE Professor of the Year in 1993, and a Stephen H. Weiss Fellow by Cornell in 1996. He holds twenty-four U.S. patents and has published more than fifty technical papers, primarily in the areas of computer vision and computational geometry.

PUBLICATIONS

"A Unified Spatiotemporal Articulated Model for Tracking, *Proceedings of the IEEE Computer Vision and Pattern Recognition Conference (2004)*. (With X. Lan).

"Efficient Belief Propagation for Early Vision". *Proceedings of the IEEE Computer Vision and Pattern Recognition Conference (2004)*. (With P. Felzenszwalb).

"Traffic-based Feedback on the Web". *Proceedings of the National Academy of Sciences* 6 (January 2004). (With J. Aizen, J. Kleinberg, and T. Novak).

"Fast Algorithms for Large-state-space HMMs with Applications to Web Usage Analysis:. [accepted] *Advances in Neural Information Processing Systems (NIPS)* 16 (2003). (With P. Felzenszwalb and J. Kleinberg).

LECTURES

Fast Belief Propagation for Low-level Vision Microsoft Research, Redmond, Washington (May 2004).

Distance Transforms for Matching Images, University of Washington, Computer Science Department (May 2004).

The Computer Science of Computer Vision, M.I.T. Computer Science Artificial Intelligence Laboratory Distinguished Lecture Series (April 2004).

Information Technology as a Business Multiplier, IBM Professional Leadership Technical Exchange (March 2004).

Pictorial Structures for Object Recognition, University of Oxford, Robotics Seminar, (January 2004).

Tutorial on Fast Algorithms for Matching, International Conference on Computer Vision (ICCV) (October 2003). (With P. Torr).

PROFESSIONAL ACTIVITIES

Member, Program Committee, IEEE Conference on Computer Vision and Pattern Recognition.

Member, Program Committee, ICCV.

Member, National Academy of Sciences Study "Assessing the Feasibility, Accuracy, and Technical Capability of a National Ballistics Database".

UNIVERSITY ACTIVITIES

Director of Graduate Studies, Information Science.

Member, CS Operations Review Committee.

Member, CS Recruiting Committee.

Member, Technology Management Planning Group, Johnson Graduate School of Management.

Member, Steering Committee, Library and Related Information Systems.

Member, CIS Council.

Member, CIS Building Committee.

Member, Information Science Undergraduate Major Committee.

Member, CIS DA&G Committee.

Member, Cornell–Queens M.B.A. Program Advisory Group.

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 with the title "The Maximum-margin Approach to Learning Text Classifiers: Methods, Theory, and Algorithms" at the Universität Dortmund, Germany, advised by Professor Katharina Morik.

Joachims's 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, he has worked on WebWatcher, an adaptive browsing assistant for the Web. He has developed and authored the SVM–Light algorithm and software for inductive and transductive support-vector learning. His most recent work is on learning from clickthrough data in search engines, and on discriminative training for predicting complex multivariate objects. Joachims taught the course "Representing and Accessing Digital Information", and the "Introduction to Machine Learning".

PUBLICATIONS

"Support Vector Machine Learning for Interdependent and Structured Output Spaces". *Proceedings of the ICML* (2004). (With I. Tsochantaridis, T. Hofmann, and Y. Altun).

"Eye-tracking Analysis of User Behavior in WWW–Search". Poster abstract, *Proceedings of Twenty-eighth Annual ACM Conference on Research and Development in Information Retrieval (SIGIR 2004)*. (With L. Granka and G. Gay).

"Mapping Subsets of Scholarly Information". *Proceedings of the National Academy of Sciences* 10: 1073. (2004). (With P. Ginsparg, P. Houle, and J.-H. Sul).

"Learning a Distance Metric from Relative Comparisons". *Proceedings of the 2003 NIPS Conference* 15. (With M. Schultz).

"Transductive Learning via Spectral Graph Partitioning". *Proceedings of the ICML* (2003).

LECTURES

Optimizing Search Engines Using Clickthrough Data

Brown University, Department of Computer Science (September 2003).

University of Illinois at Urbana–Champaign, Department of Computer Science (September 2003).

Learning to Predict Complex Outputs

Keynote Talk, Belgium/Netherlands Conference on Machine Learning (January 2004).

Cornell University, Weill Medical School (February 2004).

Cornell University, Information Science (March 2004).

Google Inc. (March 2004).

Transductive Learning via Spectral Graph Partitioning

Cornell University, AI Seminar (October 2003).

PROFESSIONAL ACTIVITIES

Member, Editorial Board, *Journal of Machine Learning Research*.

Member, Editorial Board, *Journal of Artificial Intelligence Research*.

PC/Reviewing: ICML, European Conference on Machine Learning (ECML), NIPS, Special Interest Group on Information Retrieval (SIGIR), HLT, World Wide Web Conference (WWW).

Workshop Co-organizer, *Implicit Measures of User Interests and Preferences SIGIR Workshop (2003)*. (With S. Dumais, K. Bharat, and A. Weigend).

UNIVERSITY ACTIVITIES

Member, Ph.D. Admissions Committee.

AWARDS AND HONORS

CAREER award, Improving Information Access using Implicit Feedback, 2003.

Dissertation Award, Universität Dortmund, 2002.

Rotary Stipend to Study Abroad, 1994.

Studienstiftung des Deutschen Volkes, 1991.

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 is the chair of 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 life-science disciplines. 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 *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 University for 1997–1998.

PUBLICATIONS

"Algorithms for Placing and Connecting Facilities, and Their Comparative Analysis". In *Proceedings of the Nineteenth European Workshop of Computational Geometry* (2003) 151–154, Bonn, Germany. (With I. Rabaev and N. Sokolovski).

LECTURES

"Shape-matching Algorithms for Proteins: Pairwise Matching and Consensus Shapes".
Computational Biology Seminar, Dartmouth University, Hanover, New Hampshire,
Theory Seminar, Department of Computer Science, Cornell University.
Triangle Biophysics Symposium, Duke University, Chapel Hill, North Carolina.

PROFESSIONAL ACTIVITIES

Member, Editorial Board, *Journal of the Pattern Recognition Society*.

Member, Steering Committee, Centers of Excellence in Computational Geometry, The Israeli Academy of Science.

Uri Keich

Assistant Professor

keich@cs.cornell.edu

http://www.cs.cornell.edu/~keich

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 at San Diego, and assistant professor at the Department of Mathematics of the University of California at 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, and sequence assembly.

PUBLICATIONS

"Checking for Base-calling Errors in Repeats". [submitted] (With D. Zhi, P. Pevzner, S. Heber, and H. Tang).

"A Faster Reliable Algorithm to Estimate the P-value of the Multinomial llr Statistic". [submitted] (With N. Nagarajan).

"Designing Seeds for Similarity Search in Genomic DNA". *Journal of Computer and System Sciences (JCS) Special Issue on Bioinformatics.* [in press]. (With J. Buhler and Y. Sun).

"Efficiently Computing the P-value of the Entropy Score". *Journal of Computational Biology* (In press).

"On Spaced Seeds for Similarity Search". *Discrete Applied Mathematics* 138: 253–263 (2004), (With M. Li, B. Ma, and J. Tromp).

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

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

LECTURES

Department of Statistics Colloquium, University of Toronto (2003)

Tri-institutional Research Seminar in Computational Biology at the Weill Medical College, Cornell University (2003).

PROFESSIONAL ACTIVITIES

Reviewer, *Journal of Computer and System Sciences, Proceedings of the National Academy of Sciences, Bioinformatics.*

Conference Reviewer, RECOMB (2004).

Member, Program Committee, First Annual RECOMB Satellite Workshop on Regulatory Genomics (March 2004).

UNIVERSITY ACTIVITIES

Courses:

Discrete Structures. Spring 2004, (with J. Halpern).

Problems and Perspective in Computational Molecular Biology. Fall 2003 and spring 2004 (with R. Elber).

Committees:

Member, University-wide Statistics Search Committee.

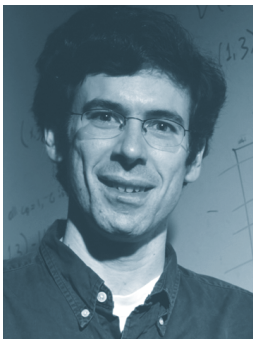

Jon Kleinberg

Associate Professor

kleinber@cs.cornell.edu

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

Jon Kleinberg received 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 World Wide Web.

PUBLICATIONS

"Segmentation problems". *Journal of the ACM* 51(2): 263–280 (2004). (With C. Papadimitriou and P. Raghavan).

"Traffic-based Feedback on the Web". *Proceedings of the National Academy of Sciences* 101 (Suppl. 1): 5254–5260 (2004). (With J. Aizen, D. Huttenlocher, and A. Novak).

"Using Mixture Models for Collaborative Filtering". *Proceedings of the Thirty-sixth ACM Symposium on the Theory of Computing* (2004). (With M. Sandler).

"Network Failure Detection and Graph Connectivity". *Proceedings of the Fifteenth Annual ACM–SIAM Symposium on Discrete Algorithms* (2004). (With M. Sandler and A. Slivkins).

"The Evolutionary Capacity of Protein Structures". *Proceedings of the International ACM RECOMB Conference* (2004). (With L. Meyerguz, D. Kempe, and R. Elber).

"Overview of the 2003 KDD Cup". *ACM SIGKDD Explorations* (2004). (With P. Ginsparg and J. Gehrke).

"The Small-world Phenomenon and Decentralized Search". [an essay as part of Math Awareness Month 2004, appearing in] *SIAM News* 37(3) (April 2004).

"Analysing the Scientific Literature in its Online Context". *Nature Web Focus on Access to the Literature* (April 2004).

"Bursty and Hierarchical Structure in Streams". *Data Mining and Knowledge Discovery* 7(4): 373–397 (2003).

"Maximizing the Spread of Influence through a Social Network. *Proceedings of the Ninth ACM SIGKDD International Conference* (2003). (With D. Kempe and E. Tardos).

"Fast Algorithms for Large-state-space HMMs with Applications to Web Usage Analysis". *Advances in Neural Information Processing Systems (NIPS)* 16 (2003). (With P. Felzenszwalb and D. Huttenlocher).

"The Link Prediction Problem for Social Networks". *Proceedings of the Twelfth International Conference on Information and Knowledge Management* (2003). (With D. Liben-Nowell).

"A Graph-theoretic Approach to Comparing and Integrating Genetic, Physical and Sequence-based Maps". *Genetics* 165(2003). (With I. Yap, D. Schneider, D. Matthews, S. Cartinhour, and S. McCouch).

LECTURES

"Cascading Behavior and Bursty Dynamics in Computational Models of Social Networks". Invited plenary talk at Twentieth Conference on Uncertainty in Artificial Intelligence (July 2004).

"Complex Networks, Search Algorithms, and the Evolution of the Web". Invited plenary talk at SIAM Conference on Discrete Mathematics (June 2004).

"Cascading Behavior and Bursty Dynamics in Computational Models of Social Networks". Princeton Computer Science Department Colloquium (April 2004).

"Network Failure Detection and Graph Connectivity". Institute for Advanced Study (April 2004).

"Information Flow in Social Networks". AAAS Annual Meeting (February 2004).

"Overview of the 2003 KDD Cup". KDD Cup Presentation at Ninth ACM SIGKDD International Conference (August 2003).

SELECTED PRESS APPEARANCES:

"Where to Start to Launch the Butterfly Effect" *ACM Tech News* (February 18, 2004).

"Is Google the Bling-Bling of the Netted Crowd?" *Wall Street Journal* (May 3, 2004).

"How to Influence People". *New Scientist* (November 15, 2003).

"Small-world Networks". Interview on Talk of the Nation: Science Friday, National Public Radio (August 8, 2003).

"Google is Most Popular Search Site, But Others Sometimes Do it Better". *Wall Street Journal* (August 18, 2003)."The Write-stuff Software". *Discover* (June 2003).

PROFESSIONAL ACTIVITIES

National Academy's Computer Science and Telecommunications Board Study on Fundamentals of Computer Science (2000–2004).

Member, Editorial Board, Springer book series in *Information Science and Statistics*; Springer *Lecture Notes in Computer Science*; *Journal of Internet Mathematics*.

Deputy vice-chair, Search Track, International World Wide Web Conference (2004).

Member, Program Committee, ACM SIGKDD International Conference, 2004.

UNIVERSITY ACTIVITIES

Member, Founders of the CIS Council.

AWARDS AND HONORS

Best Research Paper Award, ACM SIGKDD International Conference, 2003.

Cornell Association of Computer Science Undergraduates, Faculty of the Year Award, 2002.

IBM Outstanding Innovation Award, 2002.

National Academy of Sciences Award for Initiatives in Research, 2001.

Fiona Ip Li '78 and Donald Li '75 Excellence in Teaching Award, College of Engineering, Cornell University, 2000.

Best Paper Award, ACM Symposium on Principles of Database Systems, 2000.

David and Lucile Packard Foundation Fellowship, 1999.

Office of Naval Research Young Investigator Award, 1999.

Alfred P. Sloan Research Fellowship, 1997.

NSF Faculty Early Career Development Award, 1997.

Machtey Award for Best Student Paper, IEEE Symposium on Foundations of Computer Science, 1996.

George M. Sprowls Ph.D. Dissertation Prize, M.I.T. Department of Electrical Engineering and Computer Science, 1996.

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 from Dartmouth College in mathematics 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 CS faculty in 1985.

Dexter'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 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.

PUBLICATIONS

"Some Results in Dynamic Model Theory". *Science of Computer Programming* 51(1–2): 3–22 (May 2004). Special Issue: *Mathematics of Program Construction.* (E. Boiten and B. Moller, eds.).

"Computational Inductive Definability". *Annals of Pure and Applied Logic* 126(1–3): 139–148 (April 2004). Special issue: *Provinces of Logic Determined. Essays in the Memory of Alfred Tarski. Parts I, II and II* (Z. Adamowicz, S. Artemov, D. Niwinski, E. Orlowska, A. Romanowska, and J. Wolenski, eds.).

"Automata on Guarded Strings and Applications". *Matematica Contemporanea* 24: 17–139 (2003).

"Kleene Algebra with Tests and the Static Analysis of Programs". *Technical Report TR2003–1915, Computing and Information Science, Cornell University* (November 2003).

"On the Representation of Kleene Algebras with Tests". *Technical Report TR2003–1910, Computing and Information Science, Cornell University* (September 2003).

"KAT–ML: An Interactive Theorem Prover for Kleene Algebra with Tests". In *Proceedings of the Fourth International Workshop on the Implementation of Logics* (B. Konev and R. Schmidt, eds.) 2–12. University of Manchester (September 2003). (With K. Aboul-Hosn).

"Substructural Logic and Partial Correctness". *Transactions on Computational Logic* 4(3): 355–378 (July 2003). (With J. Tiuryn).

LECTURES

Tenth International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (September 2003).

Workshop on Logic and Computation, Nelson, Arizona (January 2004).

Latin American Theoretical Informatics (April 2004).

PROFESSIONAL ACTIVITIES

Program committees 2003–2004:

Mathematical Foundations of Computer Science (2004).

Workshop on Logic, Language, Information, and Computation (WoLLIC) (2003).

Chair, International Conference on the Mathematics of Program Construction (2004).

Seventh International Seminar on the Relational Methods in Computer Science (May 2003).

Second International Workshop Applications of Kleene Algebra (May 2003).

Editorial Boards 2003–2004:

Theoretical Computer Science, special issue for WoLLIC (2004).

Journal of Logical Methods in Computer Science (2003–present).

Theory of Computing Systems (2001–present).

Journal of Relational Methods in Computer Science (2000–present).

External Committees and Advisory Boards (2003–2004):

Centre for Basic Research in Computer Science (BRICS), Aarhus University.

Chair (2003), Gödel Prize Committee, ACM (2000–2003).

IEEE Symposium on Logic in Computer Science (1999–present).

UNIVERSITY ACTIVITIES

Freshman advisor, College of Arts and Sciences (2003–2004).
Chair, Academic Integrity Hearing Board, College of Engineering.
Faculty Advisor, Cornell Men's Rugby Club.
Faculty Advisor, Cornell Women's Rugby Club.
Faculty Advisor, Johnson Graduate School of Management Rugby Club.
Faculty Advisor, Absolute A Cappella Student Chorus.

AWARDS AND HONORS

Fellow, Association of Computing Machinery, 2003.
Research Prize, Polish Ministry of Education, 2001.
Stephen and Margery Russell Distinguished Teaching Award, College of Arts and Sciences, Cornell, 2001.
Class of 1960 Scholar, Williams College, 2000.
Faculty of the Year, Association of Computer Science Undergraduates, Cornell, 1994.
Research Prize, Polish Ministry of Education, 1993.
Fellow, John Simon Guggenheim Foundation, 1991.
Outstanding Innovation Award, IBM Corporation, 1980.
John G. Kemeny Prize in Computing. Dartmouth College, 1974.

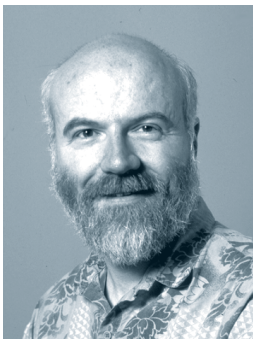

Dean Krafft

Senior Research Associate

Director of Computing Facilities

dean@cs.cornell.edu

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

Dean Krafft received his Ph.D. in computer science from Cornell University in 1981. He serves as both a CS researcher and administrator at Cornell. As an administrator, he manages the Computing Facilities Support group and focuses on a number of issues including computer security, networking, and Web services. As part of the Nomad research project, he led a campus-wide pilot of 802.11b wireless LAN technology. This service, called RedRover, is now supported by Cornell Information Technologies as part of the basic campus networking infrastructure.

On the research side, Krafft is part of the Digital Libraries Research group in Computing and Information Science. In particular, he is a co–principal investigator on the NSF–funded National Science Digital Library Project (*http://nsdl.org*) at Cornell. He shares responsibility for developing key components of the Core Integration Technology for the library, and he is team leader for the overall infrastructure-development effort. 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.

UNIVERSITY ACTIVITIES

Member, IT Policy Advisory Group.

Christoph Kreitz

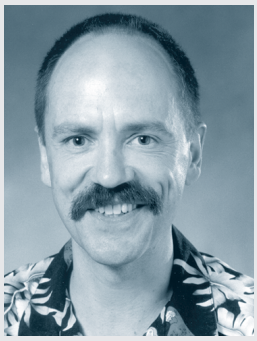

Senior Research Associate

kreitz@cs.cornell.edu

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

Christoph Kreitz is a senior research associate at Cornell University and a professor of computer science at the University of Potsdam, Germany. He 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 theorem-proving environments. Together with his former students and colleagues from Germany, 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 JProver 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.

PUBLICATIONS

"A Nuprl-PVS Connection: Integrating Libraries of Formal Mathematics". *Technical Report, Computing and Information Science, Cornell University* (2003).

A Matrix Characterization for Multiplicative Exponential Linear Logic, *Journal of Automated Reasoning* 32(2) (2004). (With H. Mantel).

"Building Reliable, High-performance Networks with the Nuprl Proof Development System". *Journal of Functional Programming* 14(1) (2004).

"MetaPRL—A Modular Logical Environment". *International Conference on Theorem Proving in Higher-order Logics* (2003). (With J. Hickey, A. Nogin, R. Constable, B. Aydemir, E. Barzilay, Y. Bryukhov, R. Eaton, A. Kopylov, V. Krupski, L. Lorigo, S. Schmitt, C. Witty, and X. Yu).

PROFESSIONAL ACTIVITIES

Member, Program Committee, IJCAR 2004 Workshop on Computer-supported Mathematical Theory Development (2004).

Member, Program Committee, Fourth International Workshop on Logical Frameworks and Meta-languages (LFM '04).

Member, Program Committee, NASA Intelligent Systems/Automated Reasoning NASA Research Announcement Technical Review Panel (2003).

Member, Program Committee, IJCAI 2003 Workshop on Agents and Automated Reasoning (2003).

Member, Program Committee, International Conference on Automated Reasoning with Analytic Tableaux and Related Methods (2003).

UNIVERSITY ACTIVITIES

Head, Faculty Search Committee, Department of Computer Science, University of Potsdam.

Member, Development and Planning Committee. University of Potsdam.

Member, University Elections Committee, University of Potsdam.

AWARDS AND HONORS

Prize for Excellence in Teaching and Research, Adolf-Messer Stiftung, Darmstadt University of Technology, Germany, 1995.

DFG Computer Science Postdoctoral Fellowship, Deutsche Forschungsgemeinschaft, Germany, 1985–1986.

Prize for Excellence as Young Investigator, University of Hagen, Germany, 1985.

Prize for Excellence on Master's Studies (Springorum-Gedenkünze der Rheinisch-Westfälische Technische Hochschule Aachen), Aachen, Germany, 1982.

Graduate Fellow, Studienstiftung des Deutschen Volkes, Germany, 1980–1981.

Carl Lagoze

Senior Research Associate

lagoze@cs.cornell.edu

http://www.cs.cornell.edu/lagoze/

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 the Faculty of Computing and Information Science at Cornell University.

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 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, and the Open Archives Initiative Protocol for Metadata Harvesting that has been widely adopted as a foundation for information systems interoperability. His scientist role in the NSF–funded NSDL project provides the opportunity to realize these advances in a major national resource for science and mathematics education.

PUBLICATIONS

"Towards a Policy Language for Humans and Computers". *ECDL 2004,* Bath, U.K. (2004).

"Rethinking Scholarly Communication: Building the System That Scholars Deserve". *D-Lib Magazine* (September, 2004).

"Bridging the Past and Future: Scholarly Communication in the Twenty-first Century". *International Symposium on Digital Libraries and Knowledge Communities in Networked Information Society (DLKC 2004)*. Tsukuba, Japan (2004).

"The Making of the Open Archives Initiative Protocol for Metadata Harvesting". *Library Hi Tech* 21(2) (2003). (With H. van de Sompel).

"Towards a Core Ontology for Information Integration". *Journal of Digital Information* 4(1) (2003). (With M. Doerr and J. Hunter).

LECTURES

"Back to the Past: Rediscovering Digital Libraries as Interactive Spaces". DLKC 2004, Tsukuba, Japan (2004).

"Technology, Scalability, Metadata (or not)". NSDL Annual Meeting, Washington D.C. (2003).

"Open Archives Initiative: Where We Are, Where We Are Going" Fourth Open Archives Forum Workshop, Bristol, U.K. (2003).

"Towards the Post–DL Age". NSF Post–DC Futures Workshop, Chatham, Massachusetts (2003).

"National Science Digital Library". Common Solutions Group, Austin, Texas (2003).

PROFESSIONAL ACTIVITIES

Member, Advisory Board, Computational Linguistics for Metadata Building.

Member, Advisory Committee, NSF National Virtual Observatory.

Member, Executive Team, NSF NSDL; Open Archives Initiative.

Guest Editor, *IEEE Internet Computing*.

Member, Program Committee, Eighth European Conference on Research and Advanced Technology for Digital Libraries; Joint Conference on Digital Libraries (2004).

Member, Advisory Board, Los Alamos National Laboratory Library.

AWARDS AND HONORS

Library and Information Technology Association Frederick G. Kilgour Award (2004).

Lillian Lee

Associate Professor

llee@cs.cornell.edu

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

Lillian Lee (A.B. Cornell University 1993; Ph.D. Harvard University 1997) is an associate professor in CS. Her main research interest is natural language processing, in particular the development of "knowledge-lean" statistical methods allowing 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 text to creating English versions of computer-generated mathematical proofs to determining document sentiment polarity (e.g., whether a review is "thumbs up" or "thumbs down").

PUBLICATIONS

"Catching the Drift: Probabilistic Content Models, with Applications to Generation and Summarization". *Proceedings of HLT–NAACL 2004*. (With R. Barzilay).

"Corpus Structure, Language Models, and Ad-hoc Information Retrieval". *Proceedings of Twenty-eighth Annual ACM Conference on Research and Development in Information Retrieval (SIGIR 2004)*. (With O. Kurland).

"A Sentimental Education: Sentimental Analysis Using Subjectivity Summarization". *Proceedings of the Association for Computational Linguistics* (2004). (With B. Pang).

"A Matter of Opinion: Sentiment Analysis and Business Intelligence [position paper]. *IBM "Architecture of On Demand Business" Faculty Summit* (2004).

LECTURES

"What is the Matter? Explorations in Text Categorization". Invited Talk at the Twentieth Conference on Uncertainty in Artificial Intelligence. Banff, Canada (2004).

"Knowledge-lean Approaches to Natural Language Processing". Princeton University Computer Science Department Colloquium (2004).

"A Matter of Opinion: Sentiment Analysis and Business Intelligence". IBM "Architecture of On-demand Business" Faculty Summit, IBM Watson (2004).

PROFESSIONAL ACTIVITIES

Secretary , NAACL (2004–2005).

Review Panel, Forty-second Annual Meeting of the ACL (2004).

Journal Referee, *Proceedings of the National Academy of Sciences* (2003–2004).

Member, Advisory Board, *Machine Learning Journal* special issue on Learning in Speech and Language Technologies (2003).

Member, NSF Review Panel (2003).

Journal Referee, *IEEE Transactions on Knowledge and Data Engineering (TKDE)* (2003).

Member, Organizing Committee, ACL Special Interest Group on Linguistic Data and Corpus-based Approaches to NLP (2001–present).

Member, Editorial Board, *Machine Learning Journal* (2001–2004).

Member, Editorial Board, *Computational Linguistics*.

UNIVERSITY ACTIVITIES

Member, Steering Committee, Cognitive Studies (2001–present).

SELECTED PRESS APPEARANCES

"Software Paraphrases Sentences". *Technology Research News* (December 3/10, 2003).

"Get Me Rewrite!" "Hold On, I'll Pass You to the Computer". *New York Times* (December 25, 2003); *ACM Tech News* (December 29, 2003).

AWARDS AND HONORS

Best Paper Award, HLT–NAACL 2004 (With R. Barzilay).

Alfred P. Sloan Research Fellowship (2002–2004).

James and Mary Tien Excellence in Teaching Award (2002).

Stephen and Marilyn Miles Excellence in Teaching Award (1999).

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 a 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.

PUBLICATIONS

"Derivative Portfolio Hedging Based on CvaR". *New Risk Measures in Investment and Regulation* (G. Szegö, ed.) Wiley (2003). (With Alexander and Coleman).

"Discrete Hedging Under Piecewise Linear Risk Minimization". *The Journal of Risk* 5:3 9–65 (2003). (With Patron and Coleman).

AWARDS AND HONORS

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

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 M.I.T. and Brandeis University, in the area of design automation and robotics. Lipson received his Ph.D. in Mechanical Engineering (1988) from the Technion—Israel Institute of Technology, in the field of computer-aided design.

Research in Lipson's group focuses on developing doman-independent computational methods for open-ended design automation. These methods are largely inspired by biological processes such as evolution, co-evolution, and swarm intelligence, and are applied to a variety of systems from mechanics and robotics to software and control systems. Lipson's research introduced the first physical robotic systems whose entire morphology and control were synthesized entirely automatically.

Current research focuses on pushing the limits of what these algorithms can design, by exploring the role of modularity, regularity, and hierarchy in synthesis.

PUBLICATIONS

"Networks, Dynamic and Modularity". *Physical Review Letters* 22(18). (2004). (With E. Variano and J. McCoy).

"Automated Damage Diagnosis and Recovery for Remote Robotics". *IEEE International Conference on Robotics and Automation* 3545–3550. (2004). (With J. Bongard).

"How to Draw a Straight Line Using a GP: Benchmarking Evolutionary Design Against Nineteenth-century Kinematic Synthesis". *Proceedings of Genetic and Evolutionary Computation Conference*. (2004).

"Generative Encodings for the Automated Design of Modular Physical Robots". *IEEE Transactions on Robotics and Automation* 19(4): 703–719. (2003). (With G. Hornby and J. Pollack).

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

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

LECTURES

"Modularity, Regularity and Hierarchy in Evolved Engineering Systems", Invited Plenary Speaker, Meeting of the National Academies, Irvine, California, (November 15 2003).

"Innovation in evolutionary processes", Invited Speaker, Founders Workshop, Santa Fe Institute for Complex Systems, Santa Fe, New Mexico (January 15, 2004).

PROFESSIONAL ACTIVITIES

Co-chair, Genetic and Evolutionary Computation Conference (2005).

Chair, AAAI Symposium on Computational Synthesis, Stanford California (March 24–26, 2003).

AWARDS AND HONORS

Silver Medal for Human-competitive Automated Invention, GECCO 2004 (*TIME Magazine's* annual 2001 "most important events of the year" selection).

Shaping the Future Award, EXPO 2000.

Rajit Manohar

Associate Professor of the School of Electrical and Computer Engineering and Member of the Graduate Field of Computer Science

rajit@csl.cornell.edu

http://vlsi.cornell.edu/~rajit/

Rajit Manohar received his B.S. (1994), M.S. (1995), and Ph.D. (1998) in computer science from the California Institute of Technology. He has been a member of the Cornell faculty since 1998 where he cofounded its Computer Systems Laboratory. His group conducts research on efficient asynchronous computation structures.

Research in Professor Manohar's group combines formal methods, algorithms for design automation, and the VLSI design of asynchronous circuits and systems. His group introduced a new class of high-performance asynchronous programmable logic structures, the first ultralow-power processor for sensor networks, and an asynchronous event-based architecture for discrete-event simulation. His work introduced the notion of slack elasticity and projection, which provides the theoretical foundation for fine-grained pipelining in asynchronous circuits.

PUBLICATIONS

"Fault Detection and Isolation Techniques for Quasi Delay-insensitive Circuits". *Proceedings of the International Conference on Dependable Systems and Networks* (July 2004). (With C. LaFrieda).

"Nonuniform Access Asynchronous Register Files. *Proceedings of the Tenth International Symposium on Asynchronous Circuits and Systems* (April 2004). (With D. Fang).

"Static Tokens: Using Dataflow to Automate Concurrent Pipeline Synthesis". *Proceedings of the Tenth International Symposium on Asynchronous Circuits and Systems* (April 2004). (With J. Teifel).

"Highly Pipelined Asynchronous FPGAs". *Proceedings of the Twelfth ACM International Symposium on Field-programmable Gate Arrays,* Monterey, California (February 2004). (With J. Teifel).

"An Event-synchronization Protocol for Parallel Simulation of Large-scale Wireless Networks". *Seventh IEEE International Symposium on Distributed Simulation and Real Time Applications* (October 2003). (With C. Kelly).

"Programmable Asynchronous Pipeline Arrays". *Proceedings of the Thirteenth International Conference on Field Programmable Logic and Applications,* 345–354, Lisbon, Portugal (September 2003). (With J. Teifel).

"SNAP: A Sensor Network Asynchronous Processor". *Proceedings of the Ninth International Symposium on Asynchronous Circuits and Systems* 24–33, Vancouver, British Columbia (May 2003). (With C. Kelly and V. Ekanayake).

LECTURES

"Simulation and Design of Sensor Networks with Asynchronous VLSI" Invited Speaker, IEEE Computer Society Symposium on VLSI, Lafayette, Louisiana (February 2004).

"How Asynchronous Should We Be?" Invited Panelist, IEEE Computer Society Symposium on VLSI, Lafayette, Louisiana (February 2004).

"Ultra Low Power Asynchronous VLSI." Invited Talk, DARPA Workshop on Ultralow-power Technologies, M.I.T. (January 2004).

"Asynchronous Event-processing." Seminar, Analog VLSI and Biological Systems, M.I.T. (November 2003).

"SNAP: A Sensor-network Asynchronous Processor." Electrical and Systems Engineering Colloquium, University of Pennsylvania (October 2003).

"Designing an Efficient Sensor Network Processor." VLSI Seminar Series, Cornell University (September 2003).

PROFESSIONAL ACTIVITIES

Program Committees:

IEEE/ACM Symposium on Asynchronous Circuits and Systems (May 2004).

Second IEEE Update New York Workshop on Sensor Networks (October 2003).

AWARDS AND HONORS

CAREER award.

Cornell IEEE Teacher of the Year Award.

College of Engineering Sonny Yau Excellence in Teaching Award.

Tau Beta Pi and Cornell Society of Engineers Excellence in Teaching Award.

Steve Marschner

Assistant Professor

srm@cs.cornell.edu

http://www.cs.cornell.edu/~srm

Steve Marschner obtained his Sc.B. 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 is in computer graphics, particularly realistic rendering, and focuses on modeling the optics of materials. His research aims to develop new material models, grounded in physics and supported by measurements of real materials that can represent the complex natural materials we encounter every day. These models will allow computer-graphics practitioners to achieve goals, such as creating realistic virtual human actors, that are currently hampered by the limitations of existing material models. Marschner is also exploring related areas of computer graphics, from image-based modeling and 3D scanning to visual perception and mechanics.

Recent projects include an investigation of light scattering from human hair and an efficient method for rendering translucent materials that has been widely implemented by the film-effects industry and resulted in an Academy Award for Technical Achievement.

PUBLICATIONS

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

PROFESSIONAL ACTIVITIES

Member, Papers Committee, ACM SIGGRAPH 2002 and 2003.

Member, Papers Committee, CVPR 2003.

AWARDS AND HONORS

CAREER Award.

Academy Award for Technical Achievement, 2003.

José F. Martínez

Assistant Professor of the School of Electrical and Computer Engineering and Member of the Graduate Field of Computer Science

martinez@csl.cornell.edu

http://csl.cornell.edu/~martinez/

José Martínez graduated in computer science and engineering in 1996 from the Universidad Politécnica de Valencia, and earned M.S. (1999) and Ph.D. (2002) degrees in computer science from the University of Illinois at Urbana–Champaign. A two-time recipient of the Spanish government's prestigious National Award for Academic Excellence, he was a Bank of Spain graduate fellow for four years.

Martínez's recent work on speculative synchronization was featured in IEEE Micro's 2003 Top Picks from Microarchitecture Conferences.

His research interests include, but are not limited to, parallel architectures, microarchitecture, reconfigurable hardware, and hardware-software interaction. He is a member of the ACM and the IEEE Computer Society.

PUBLICATIONS

"The Thrifty Barrier: Energy-efficient Synchronization in Shared-memory Multiprocessors". In *International Symposium on High-performance Computer Architecture,* Madrid, Spain (February 2004). (With J. Li and M. Huang).

"Speculative Synchronization: Programmability and Performance for Parallel Codes". In *IEEE Micro Special Issue: Top Picks from Microarchitecture Conferences* (November–December 2003). (With J. Torrellas).

"A Case for Resource-conscious Out-of-order Processors". In *IEEE Technical Committee on Computer Architecture Letters* 2 (October 2003). (With A. Cristal, J. Llosa, and M. Valero).

LECTURES

Energy-aware Mechanisms in Multiprocessors. Los Alamos National Laboratory (April 2004).

The Thrifty Barrier: Energy-aware Synchronization in Shared-memory Multiprocessors. Opening session of the International Symposium on High-performance Computer Architecture, Madrid, Spain (February 2004).

Checkpointed Early Resource Recycling. IBM T. J. Watson Research Center (December 2003).

Speculative Resource Management in Out-of-order Processors. Computer Science Colloquium, Cornell University (November 2003).

PROFESSIONAL ACTIVITIES

Member, Program Committee, 2004 International Symposium on High-performance Computer Architecture (HPCA).

Member, Program Committee, 2004 International Conference on Parallel Architectures and Compilation Techniques (PACT).

UNIVERSITY ACTIVITIES

Member, Graduate Committee, School of Electrical and Computer Engineering.

Member, CURIE Academy for Women in Engineering.

AWARDS AND HONORS

National Award for Academic Excellence, Ministry of Education, Spain (1995 and 1997).

Bank of Spain Graduate Fellow (1995–1999)

Inducted by the Phi Kappa Phi Honor Society (1999)

IEEE Micro's Top Picks from Microarchitecture Conferences, on "The Most Industry-relevant and Significant Papers of the Year in Computer Architecture" (2003).

Jeanna Neefe Matthews

Adjunct 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 adjunct 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.

PUBLICATIONS

"Xen and the Art of Repeated Research". *2004 USENIX Annual Technical Conference FREENIX Track* (June 2004). (With B. Clark, T. Deshane, E. Dow, S. Evanchik, M. Finlayson, and J. Herne).

"Deciding Layers in a Multi-Layer User Interface". *Eighth European Research Consortium for Informatics (ERCIM) User Interface Four (UI4) All Workshops* (June 2004). (With B. Clark).

"Common Sense Computing". *Public Radio Exchange* (June 2004).

"The Case for Repeated Research in Operating systems". *Operating System Review* 38(2). (April 2004).

PROFESSIONAL ACTIVITIES

Secretary–Treasurer, Special Interest Group on Operatings Systems (SIGOPS) (2003–present).

Organizing Chair (Webcast and Video-archiving), ACM Symposium on Operating Systems (2003).

AWARDS AND HONORS

IBM Linux Challenge Winners from class projects (COM S 414 Operating Systems: Cornell University, 2002).

IBM Linux Challenge Overall Institutional Prize (2001).

IBM Linux Challenge winners from class projects (CS 644 Advanced Operating Systems: Clarkson University, 2001).

Intel Foundation Graduate Fellowship Award, 1999–2000.

Cal VIEW Fellow Award, for excellence in teaching, 1998–1999.

Award Paper, ACM Symposium on Operating System Principles, 1995.

National Science Foundation Gradaute Research Fellowship, 1994–1998.

Sally A. McKee

Assistant Professor of the School of Electrical and Computer Engineering and Member of the Graduate Field of Computer Science

sam@csl.cornell.edu

http://www.csl.cornell.edu/~sam

Sally McKee received her bachelor's degree in computer science from Yale University (1985), master's from Princeton University (1990), and doctorate from the University of Virginia (1995). Before graduate school, she worked for Digital Equipment Corporation's Personal Computer Division and Microsoft Corporation. She has also held internships at Digital Equipment Corporation's Systems Research Center and the former AT&T Bell Labs (now Lucent Technologies Bell Labs).

McKee was a post-doctoral research associate in the University of Virginia computer science department from May of 1995 through May of 1996.

Her research is primarily in the area of computer architecture, particularly in memory-system design and analysis. She worked at Intel's Oregon Microcomputer Research Lab in 1996 through 1998, during which time she also taught at the Oregon Graduate Institute of Science and Technology and Reed College. In 1998, she moved to the University of Utah, where she performed research on the Impulse Adaptable Memory System Project and her own research (partly in collaboration with colleagues at Lawrence Livermore National Laboratories [LLNL]) for four years. She was a participating guest at LLNL during much of 2001–2002. In 2002, she moved to the School of Electrical and Computer Engineering at Cornell University, where she works with four Ph.D. students and one research associate on infrastructure for introspective (ultimately autonomic) computer systems, intelligent memory systems, and other aspects of high-end computing.

PUBLICATIONS

"Formal Hardware Specification Languages for Protocol Compliance Verification". *ACM Transactions on Design Automation of Electronic Systems* 9(1) (January 2004). [accepted for publication in 2003] (With A. Bunker and G. Gopalakrishnan).

"A Cost Model for Integrated Restructuring Optimizations". *Journal of Instruction Level Parallelism* 5 (2003). (With B. Chandramouli, W. Hsieh, and J. Carter).

"Restructuring Computations for Temporal Data Cache Locality". *International Journal of Parallel Programming, Kluwer Academic Press,* 31(4): 306–338 (August 2003). [accepted for publication in 2003] (With V. Pingali, W. Hsieh, and J. Carter).

LECTURES

"Perspectives on the Memory Wall Problem". Invited Lecture, Salishan Department of Energy High Performance Computing Meeting, Glen Eden, Oregon (April 2003).

"Young Researcher's Panel, Student Days". Invited Lecture, Fifteenth ACM/IEEE Conference on Supercomputing, Baltimore, Maryland (November 2002).

"A Parallel Vector Access Memory System". Invited Lecture, Department of Computer Science and Engineering, the Pennsylvania State University, State College, Pennsylvania (October 2002).

"An MPEG–4 Performance Study". Proceedings of the Third IEEE International Symposium on Performance Analysis of Systems and Software, Austin, Texas (March 2003).

"A Parallel Vector Memory Subsystem for the Impulse Adaptable Memory Controller". Computer Science Colloquium Series, Cornell University (February 2003).

PROFESSIONAL ACTIVITIES

Member, Steering Committee, International Conference on Parallel Architectures and Compilation Techniques (2002–present).

Member, Advisory Committee, IEEE Technical Committee on Computer Architecture (2001–present).

Co-chair, Program Committee, IEEE International PACT Conference (September 2002) (with E. Altman, IBM Research).

Officer, ACM Micro Special Interest Group on Computer Microarchitecture (2002–present).

Guest Editor, *Journal of Instruction Level Parallelism* Special Issue on the Best of PACT 2002 (January 2003). (With E. Altman).

AWARDS AND HONORS

Junior Faculty Fellow, Frontiers in Education Conference, Boulder, Colorado, November 2003.

Cornell University, President's Council of Cornell Women Affinito-Stewart Junior Faculty Award, June 2003.

AAAS/NSF Women's International Science Cooperation Travel Grant, with Universitat Politecnica de Catalunya, 2003.

Cornell University, Cornell Information Technology Innovation in Teaching Grant, January 2000.

Greg Morrisett

Associate Professor

jgm@cs.cornell.edu

http://www.eecs.harvard.edu/~greg/

Greg Morrisett obtained his Ph.D. in computer science from Carnegie Mellon University in 1995. In 2004, he was 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 Typed Assembly Language and Certifying Compilation. These are important mechanisms that can be used to automatically verify a wide 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.

PUBLICATIONS

"Compiling for Runtime Code Generation". *Journal of Functional Programming* 13(3): 677–708 (May 2003). (With F. Smith, D. Grossman, L. Hornoff, and T. Jim).

"Achieving Type Safety for Low-level Code". In *Proceedings of the Eighth Asian Computing Science Conference* 1–2, Mumbai, India (December 2003).

LECTURES

Typed Assembly Language (Background), Dagstuhl Conference on Language-based Security, October 2004.

Research Challenges in Language Support for Trustworthy Computing, Information Security Research Council, Washington D.C. (January 2004).

Towards Type-safety for Low-level Code, Danish Technical University, Copenhagen (January 2004).

Achieving Type-safety for Low-level Code, Asian Computing Conference, Mumbai, India (December 2003).

Beyond Regions in Cyclone, New Jersey Programming Language Seminar, Princeton, New Jersey (September 2003).

Type-safe Memory Management in Cyclone, Rome Labs, New York (August 2003).

Next-generation Low-level Languages, ONR MURI Principal Investigators Meeting, Ithaca, New York (July 2003).

Next-generation Low-level Languages, Yale University, New Haven, Connecticut (June 2003).

Tutorial on Language-based Security, Programming Language Design and Implementation, San Diego, California (June 2003).

Next-generation Low-level Languages, AFOSR Principal Investigators Meeting, Rome, New York (June 2003).

Cyclone Memory Management, Microsoft Research Ltd, Cambridge, England (May 2003).

Language-based Security: A Brief Overview, Cambridge University, Cambridge, England (May 2003).

Cyclone Memory Management, University of Edinburgh, Edinburgh, Scotland (April 2003).

Next-generation Low-level Languages, Queen Mary College, London, England (March 2003).

Next-generation Low-level Languages, University of Birmingham, England (February 2003).

PROFESSIONAL ACTIVITIES

Editor, *Journal of Functional Programming*.

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

Program Chair, *ACM Symposium on Principles of Programming Languages*.

Program Committee, *ACM Workshop on Semantics, Program Analysis, and Computing Environments for Memory Management*.

Program Committee, Second EATCS Workshop on Foundations of Global Computing.

Steering Committee, ACM Conference on Generative Programming and Component Engineering.

Member, Trustworthy Computing Academic Advisory Board, Microsoft.

AWARDS AND HONORS

Allen Newell Medal of Research Excellence, 2001.

Ralph Watts Excellence in Teaching Award, 2000.

Presidential Early Career Award for Scientists and Engineers, October 2000.

National Science Foundation Career Award, 1999.

Alfred P. Sloan Fellow, 1998.

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

PUBLICATIONS

"Enforcing Robust Declassification". *Proceedings of the Seventeenth IEEE Computer Security Foundations Workshop (Pacific Grove, California)* (June 2004). (With A. Sabelfeld and S. Zdancewic).

"A Model for Delimited Release". *Proceedings of the International Symposium on Software Security* (October 2003). (With A. Sabelfeld).

"Observational Determinism for Concurrent Program Security". *Proceedings of the Sixteenth IEEE Computer Security Foundations Workshop (Pacific Grove, California)* 29–43 (July 2003). (With S. Zdancewic).

LECTURES

Building Distributed Systems Secure by Construction. Carnegie Mellon University (January 2004).

Secure Program Partitioning and Replication. University of Arizona (October 2003).

Introduction to Language-based Information Flow. Stevens Institute of Technology (November 2003).

Using Information Flow Policies to Construct Secure Systems. Griffiss Institute University–Industry Conference (New Paltz, New York) (November 2003).

Using Information Flow Policies to Construct Secure Systems. Dagstuhl Seminar on Language-based Security (Dagstuhl, Germany) (October 2003).

PROFESSIONAL ACTIVITIES

Program committees:

IEEE Symposium on Security and Privacy (2004).

Computer Security Foundations Workshop (2004).

Network and Distributed System Security Symposium (2004).

AWARDS AND HONORS

Alfred P. Sloan Research Fellowship, 2002.

Abraham T. C. Wong '72 Excellence in Teaching Award, 2002.

CAREER Award, 2001.

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 Gödel, at the Institute for Advanced Study from 1957 to 1958; visiting assistant professor with Alfred Tarski, Berkeley from 1958 to 1959; 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 to 1987, and was director of the Mathematical Sciences Institute from 1987 to 1996. He also served as director of the Center for Foundations of Intelligent Systems from 1996 to 2001.

Nerode's research areas include mathematical logic, recursive functions, computer science, mathematics of AI, and control engineering.

PUBLICATIONS

Control in Hybird Systems. (2004). (With W. Kohn, V. Brayman, and P. Cholewski).

"Topological Semantics for Intuitionistic Modal Logics I: Completeness". In *Gödel Translations, and Bisimulations* (2004). (with J. Davoren, V. Coulthard, R. Goré, and T. Moor).

"Topological Semantics for Intuitionistic Modal Logics II: Applications to Approximate Model-checking in Classical Modal and Tense Logics". *[to appear in 2004]*. (with J. Davoren, V. Coulthard, R. Goré, and T. Moor).

LECTURES

AMS Special Sessions Lecture for CS Colloquium for New York City Area (January, 2004).

Invited Address, Knowledge Fusion Research Workshop, St. Michaels, Maryland (November 11–13, 2003).

UNIVERSITY ACTIVITIES

Chair, University Benefits Committee.

Member, CS Department Curriculum Committee.

Member, Graduate Admissions Committee.

PROFESSIONAL ACTIVITIES

Editor, *Mathematics and AI; Pure and Applied Algebra; International Hybrid Systems Journal; Documenta Mathematica*.

Board Member, several math institutes in New Zealand and in Russia.

Co-organizer, First IEEE Multi-agent Security Workshop, Philadelphia, Pennsylvania (August 30–31, 2004).

Keshav Pingali

India Professor of Computer Science

Director of Undergraduate Studies

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.) at Kanpur in 1978, and a Sc.D. degree 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. In 2003, he became the first India Chair of Computer Science at Cornell University. 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 HP's PA–RISC compiler. In his current research, he is investigating language-based fault-tolerance, and highly adaptive software systems for large-scale computational science 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.

PUBLICATIONS

"Computational Science Simulations Based on Web Services". *International Conference on Computational Science* 299–308 (2003). (With P. Chew, N. Chrisochoides, S. Gopalsamy, G. Heber, T. Ingraffea, E. Luke, J. Neto, A. Shih, B. Soni, P. Stodghill, D. Thompson, S. Vavasis, and P. Wawrzynek).

"Collective Operations in an Application-level Fault Tolerant MPI System". *International Conference on Supercomputing* 234–243 (2003). (With G. Bronevetsky, D. Marques, and P. Stodghill).

"Automated Application-level Checkpointing of MPI Programs". *Principles and Practice of Parallel Programming* 84–94 (2003). (With G. Bronevetsky, D. Marques, and P. Stodghill).

"A Comparison of Empirical and Model-driven Optimization". *ACM Symposium on Programming Language Design and Implementation (PLDI 2003)*. (With K. Yotov, X. Li, G. Ren, M. Cibulskis, G. DeJong, M. Garzaran, D. Padua, P. Stodghill, and P. Wu).

"Algorithms for Computing the Static Single Assignment Form". *Journal of the ACM* 50(3): 375–425 (2003). (With G. Bilardi).

"A System for Automating Application-level Checkpointing of MPI Programs". *Sixteenth International Workshop on Languages and Compilers for Parallel Computers (LCPC' 2003)* 357–373 (2003). (With G. Bronevetsky, D. Marques, and P. Stodghill).

LECTURES

"A Comparison of Empirical and Model-driven Optimization"

Bertinoro, Italy (April 2003).

ACM Symposium on Programming Language Design and Implementation (June 2003).

University of Padua, Italy (April 2003).

IBM Watson research Center (November 2003).

PROFESSIONAL ACTIVITIES

Associate Editor, *IEEE Transactions on Parallel and Distributed Systems.*

Member, Editorial Board, *Discrete Mathematics and Theoretical Computer Science.*

Member, Program Committee: International Conference on Supercomputing (ICS) (2003); International Conference on Parallel Programming (2004); Code Generation Optimization (2004); International Conference on High-performance Computing (2005).

UNIVERSITY ACTIVITIES

Member, Strategic Planning Committee, College of Engineering (2003–2004).

Director of Undergraduate Studies, Department of Computer Science (2003–).

Member, Faculty Recruiting Committee, Department of Electrical and Computer Engineering (2003–2004).

Member, Advisory Committee, Cornell Theory Center.

AWARDS AND HONORS

Distinguished lecturer, Department of Computer Science, Texas A&M University (2003).

Robbert van Renesse

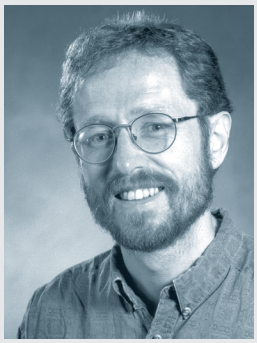

Senior Research Associate

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 Professor 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 fault-tolerant and secure distributed applications. Jointly with Ken Birman and Werner Vogels he is running the Quicksilver project, which seeks to develop reliable technologies for large-scale enterprises. Most recently, van Renesse developed the Willow protocol, which is a novel P2P protocol that supports DHT routing, multicast, publish–subscribe, as well as aggregation. Jointly with Fred Schneider, van Renesse is investigating how to build strongly consistent replicated data repositories under high sustained read and write loads, and possibly scaling to very large data sets. Finally, with Bob Constable and Mark Bickford, he is interested in automatically synthesizing code for network protocols.

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

PUBLICATIONS

"Adding High Availability and Autonomic Behavior to Web Services". *Proceedings of the ICSE,* Edinburgh, Scotland (May 23–28, 2004). (With K. Birman and W. Vogels).

"Willow: DHT, Aggregation, and Publish/Subscribe in One Protocol". *Proceedings of the International Workshop on Peer-to-peer Systems (IPTPS 2004)* San Diego, California (February 2004). (With A. Bozdog).

"P6P: A Peer-to-peer Approach to Internet Infrastructure". In *Proceedings of the IPTPS,* San Diego, California (February 2004). (With L. Zhou).

Heterogeneity-aware Peer-to-peer Multicast. *Proceedings of the Seventeenth International Symposium on Distributed Computing (DISC 2003)* Sorrento, Italy (October 2003). (With K. Birman, A. Bozdog, D. Dumitriu, M. Singh, and W. Vogels).

"SelectCast—A Scalable and Self-repairing Multicast Overlay Routing Facility". *Proceedings of the First ACM Workshop on Survivable and Self-regenerative Systems* Fairfax, Virginia (October 31, 2003). (With A. Bozdog and D. Dumitriu).

PROFESSIONAL ACTIVITIES

Program co-chair, IPTPS (2004).

DISC (2004).

Workshop on Dependable Distributed Data Management (part of Symposium on Reliable Distributed Systems [SRDS 2004]) (October 2004).

Program co-chair, International Workshop on Large-scale Group Communication (part of SRDS 2003) (October 2003).

Workshop on Adaptive Distributed Systems (part of DISC 2003) (October 2003).

Mirek Riedewald

Research Associate

mirek@cs.cornell.edu

http://www.cs.cornell.edu/~mirek

Mirek Riedewald obtained his M.S. in computer science from the University of the Saarland, Saarbrücken, Germany in 1998, and his Ph.D. in computer science from the University of California at Santa Barbara in 2002. He has worked as a research assistant at University of California at Santa Barbara for the ADEPT digital library project (including online analytical processing and data warehousing) and the University of California at Santa Barbara Digital Campus project. In 1997 he was a research assistant at the Max Planck Institute of Informatics in Saarbrücken, Germany, where he worked on a software re-engineering project for a computational biochemistry application. Currently, Riedewald leads a data-management project in collaboration with Cornell's physics department. He is also investigating novel data-mining approaches for ornithology and high-energy physics jointly with Johannes Gehrke and Rich Caruana.

Riedewald joined CS in 2002. His main area of research is database and information systems, especially data-stream processing.

PUBLICATIONS

"Semantic Approximation of Data Stream Joins". In *IEEE TKDE* [to appear]. (With A. Das and J. Gehrke).

"Approximation Techniques for Spatial Data". In *Proceedings of the ACM SIGMOD International Conference* (2004). (With A. Das and J. Gehrke).

"Research Issues in Mining and Monitoring of Intelligence Data". In *Data Mining: Next Generation Challenges and Future Directions* (H. Kargupta, ed.) AAAI Press (2004). (With A. Demers and J. Gehrke).

"The Architecture of the Cornell Knowledge Broker". In *Proceedings of the Symposium on Intelligence and Security Informatics* (2004). (With A. Demers and J. Gehrke).

"Approximate Join Processing over Data Streams". In *Proceedings of the ACM SIGMOD International Conference* 40–51 (2003). (With A. Das and J. Gehrke).

"Accessing Scientific Data: Simpler is Better". In *Proceedings of the International Symposium on Spatial and Temporal Databases* 214–232 (2003). (With D. Agrawal, A. El Abbadi, and F. Korn).

"Exploiting the Multi-append-only-trend Property of Historical Data in Data Warehouses". In *Proceedings of the International Symposium on Spatial and Temporal Databases,* pages 179–198 (2003). (With H.-G. Li, D. Agrawal, and A. El Abbadi).

PROFESSIONAL ACTIVITIES

Member, Program Committee:

ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (2004).

ACM SIGMOD International Conference on Management of Data (2004).

Symposium on Intelligence and Security Informatics (2004).

International Conference on Machine Learning (2003).

International Conference of Asian Digital Libraries (2003).

NSF National Institute of Justice Symposium on Intelligence and Security Informatics (2003).

Reviewer:

ACM Transactions on Database Systems (TODS).

ACM Transactions on Information Systems (TOIS).

Journal of Very Large Data Bases.

IEEE Transactions on Knowledge and Data Engineering.

Information Processing Letters.

Encyclopedia of Information Science and Technology, vols. I–III.

Idea Group Publishing.

AWARDS AND HONORS

President's Work Study Award, University of California at Santa Barbara, 2001–2002.

Scholarship for Young Researchers to attend the International Conference on Database Theory, European Union, 2001.

Two-year scholarship and teaching assistantship on entry to University of California at Santa Barbara, 1998.

Scholarship, "Studienstiftung des Deutschen Volkes" (German National Merit Foundation), 1992–1998.

Second and fourth places in German national mathematics competitions, 1988–1992.

Mats Rooth

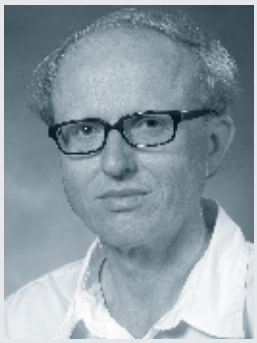

Professor

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 that 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 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. 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.

PUBLICATIONS

"Topic Marking on Quantifiers". in *The Partee Effect* (G. Carlson and J. Pelletier, eds.) M.I.T. Press. *[expected to appear in Fall 2004]*

"Comments on Krifka's Paper", in *Context Dependence in the Analysis of Linguistic Meaning* (H. Kamp and B. Partee, eds.) Elsevier. *[expected to appear Fall 2004]*

"Empty Domain Effects for Presuppositional and Non-presuppositional Determiners" in *Context Dependence in the Analysis of Linguistic Meaning* (H. Kamp and B. Partee, eds.) Elsevier. (with D. Abusch).

LECTURES

Covert Focus Variables, Italian Society for Analytic Philosophy, Milan Meeting (MM '04) (June, 2004).

PROFESSIONAL ACTIVITIES

Member, Editorial Board, *Natural Language Semantics*.

UNIVERSITY ACTIVITIES

Director, Computational Linguistics Laboratory.

Member, Founders Committee. CIS Council.

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 are in programming languages and compilation techniques for program understanding, verification, transformation and optimization. He is particularly interested in the area of program analysis, with emphasis on the analysis of pointer-based data structures. His current projects include: analysis algorithms for checking quantitative properties of recursive structures; practical pointer and shape analyses for detecting memory errors in low-level programming languages such as C; and an analysis, transformation, and run-time system that provides compiler and language support for region-based memory management in Java programs.

In his past work, Rugina has developed pointer and symbolic analyses for recursive and multithreaded programs that heavily manipulate pointers. These techniques have been used to enable the static detection of data races in pointer-based multithreaded programs, and the automatic parallelization of complex divide-and-conquer programs.

PUBLICATIONS

"Region Analysis and Transformation for Java Programs". In Proceedings of the ACM SIGPLAN 2004 *International Symposium on Memory Management,* Vancouver, British Columbia (October 2004). (With S. Cherem).

"Quantitative Shape Analysis". In *Proceedings of the International Static Analysis Symposium* (August 2004).

"Static Analysis of Accessed Regions in Recursive Data Structures". In *Proceedings of the 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): 70–116. (With M. Rinard).

"Region Inference for Imperative Languages". *Technical Report TR2003–1914, Computing and Information Science, Cornell University* (October 2003). (With S. Cherem).

PROFESSIONAL ACTIVITIES

Member, Program Committee, ACM SIGPLAN 2005 Conference on Programming Language Design and Implementation.

Member, Program Committee, ACM SIGPLAN 2003 Symposium on Principles and Practice of Parallel Programming.

Fred B. Schneider

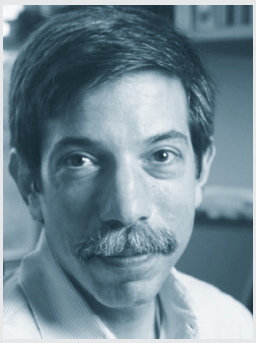

Professor

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.
- the design of 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 project, with attention now focused on implementing secure and scalable publish/subscribe protocols.

PUBLICATIONS

"Least Privilege and More". *IEEE Security and Privacy* 1(3):55–59 (September/October 2003) 253–258. Also appears in *Computer Systems: Theory, Technology, and Applications* (A. Herbert and K. Jones, eds.), Springer-Verlag, New York.

"Lifting Reference Monitors from the Kernel". *Formal Aspects of Security,* London, United Kingdom (December 2002) (A. Abdullah, P. Ryan, and S. Schneider, eds.). *Lecture Notes in Computer Science* 2629: 1–2, Springer-Verlag, New York (2003).

"The Next Digital Divide". [editorial] *IEEE Security and Privacy* 2(1): 5 (January/February 2004).

"Peer-to-peer Authentication with a Distributed Single Sign-on Service". *Proceedings of the IPTPS,* San Diego, California, ACM (February 2004).

LECTURES

Trustworthy Services from Untrustworthy Components. Invited Speaker, Future Directions in Distributed Computing II—Survivability: Obstacles and Solutions. Bertinaro (Forli), Italy (June 2004).

Scalable Replication with Chains. IAI@Rome Lecture Series. AFRL Rome Laboratories, Rome, New York (June 2004).

Trustworthy Services from Untrustworthy Components. The City University, London, England (June 2004).

Paying for Security. Invited Speaker, IBM Architecture of On-demand Business Summit. Yorktown Heights, New York (May 2004).

Trustworthy Services from Untrustworthy Components: Overview. Invited Speaker, NSF–RPI Workshop on Pervasive Computing and Networks. Rensselaer Polytechnic Institute, Troy, New York (April 2004).

AFRL/Cornell Information Assurance Institute. Rome Laboratories, Rome, New York. (April 2004).

Scalable Replication with Chains. Department of Computer Science, University of Trömso, Trömso, Norway (March 2004).

CorSSO: A Road to Single Sign-on. Department of Computer Science, University of Trömso, Trömso, Norway (March 2004).

Cyber-terrorism: Yesterday, Today, and Tomorrow. Invited Lecture, Cornell ALS 481 (Global Conflict and Terrorism). Cornell University, Ithaca, New York (February 2004).

Implementing Trustworthy Distributed Services. Graduate Student Brown-bag Lecture Series. Department of Computer Science, Cornell University, Ithaca, New York (March 2004).

Trustworthy Search. Keynote Address, Search Leaders' Summit. New York City, New York (March 2004).

Language-based Security for Malicious Mobile Code. ONR Principal Investigators Meeting. Office of Naval Research, Arlington, Virginia (February 2004).

The Mechanics of Paper Submission, Review, and Publication. Graduate Student Brown-bag Lecture Series. Department of Computer Science, Cornell University, Ithaca, New York (February 2004).

Bolt-on Fault-tolerance: Fact, Fantasy, and Future. Invited Lecture, Information Science and Technology Study on Bolt-on Security. Alexandria, Virginia (February 2004).

What's the Fuss About E-voting. Kendal of Ithaca, Ithaca, New York (January 2004).

Advice on ECE Hiring: Strategy and Directions. ECE Seminar, Cornell University, Ithaca, New York (January 2004).

Implementing Trustworthy Distributed Services. Keynote Address, International Association of Science and Technology for Development International Conference on Communication, Network, and Information Security. Uniondale, New York (December 2003).

Two Lectures: Computer Security Principles; Quick Summary of Cryptography. Computer Security Webinar. eCornell. Ithaca, New York (November 2003).

Systems That Heal Themselves: Fault-tolerance and Attack-tolerance. Air Force Science Advisory Board Review. Rome Laboratories, Rome, New York (November 2003).

Building Trustworthy Systems from Untrustworthy Components. Distinguished Lecture Series. Department of Computer Science, University of Wisconsin at Madison (November 2003).

Language-based Security for Malicious Mobile Code. ONR Briefing. Washington, D.C. (October 2003).

New York State's Information Security Policy Gap Analysis. New York State Taskforce on Cybersecurity and Critical Infrastructure Protection. Albany, New York (October 2003).

Excursions to Two New Frontiers in Computer Security Research. Advanced School on Mobile Computing, Scuola Normale Superiore. Pisa, Italy (September 2003). Four Lectures: (a) Computer Security Principles; (b) Quick Summary of Cryptography; (c) Trustworthy Services from Untrustworthy Components; (d) The Way of Language-based Security.

Infosec Research Hard Problems. Infosec Research Council, Washington, D.C. (October 2003).

Building Trustworthy Systems from Untrustworthy Components. Fast Search and Transfer, Oslo, Norway (August 2003).

APSS: Asynchronous Proactive Secret Sharing. National Institute of Science and Technology, Gaithersburg, Maryland (August 2003).

Trustworthy Computing Academic Advisory Panel. Microsoft Faculty Summit. Redmond, Washington (July 2003).

Least Privilege and More. MURI Mobile Code Principal Investigators Meeting. Ithaca, New York (July 2003).

PROFESSIONAL ACTIVITIES

Director, AFRL/Cornell Information Assurance Institute.

Chief Scientist, Griffiss Institute.

Editorial:

Editor, *Distributed Computing*.

Editor, *Information Processing Letters*.

Editor, *High Integrity Systems*.

Editor, *ACM Computing Surveys*.

Associate Editor-in-chief, *IEEE Security and Privacy*.

Editor, *IEEE Transactions on Dependable and Secure Computing*.

Co-managing Editor, *Springer-Verlag Texts and Monographs in Computer Science*.

Program Committee:

CRA workshop on Grand Challenges in Information Security.

Industrial Advisory:

Member, Technical Advisory Board, CIGITAL.

Member, Technical Advisory Council, Cloakware Corporation.

Member, Technical Advisory Board, Fast Search and Transfer.

Member, Advisory Board, Intel Microprocessor Research Lab.

Member, Advisory Board, IBM Autonomic Computing.

Co-chair, Trustworthy Computing Academic Advisory Board, Microsoft Corp.

Member, Technical Advisory Board, Packet General Networks.

Other Advisory Committees:

Member, ACM Advisory Committee on Security and Privacy.

Member, Computer Science and Telecommunications Board, National Research Council.

Member, Advisory Committee, NSF CISE.

Member, ST Search Committee, AFRL Rome.

Member, IFIP Working Group 2.3 (Programming Methodology).

UNIVERSITY ACTIVITIES

Member, Faculty Advisory Board in Information Technology.

Member, Founders Committee, Faculty of Computing and Information Science.

Member, IT Workforce Planning Team, Cornell University.

Member, Advisory Committee, Engineering Research Centers, College of Engineering.

Member, Teaching Awards Committee, College of Engineering.

AWARDS AND HONORS

Doctor of Science (*honoris causa*), University of Newcastle-upon-Tyne, U.K. (May 2003).

Daniel M. Lazar Excellence in Teaching Award, 2000.

Professor-at-large, University of Trömso, Trömso, Norway, 1996–2006.

Fellow, Association for Computing Machinery, 1994.

Fellow, American Association for Advancement of Science, 1992.

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.

This year, collaboration began on the Cornell Library Collaborative Learning Computer Laboratory (CLCLCL), which will host introductory programming labs and Schwartz's newly-approved computer-game design course.

PUBLICATIONS

"First Programming Course in Engineering: Balancing Tradition and Application". *Computers in Education Journal* 3 (2004). (With D. Fan).

Introduction to Maple, Second Edition. Prentice Hall (2003).

Introduction to Unix, Second Edition. Prentice Hall. *[expected 2004]*

UNIVERSITY ACTIVITIES

Member, Engineering Admissions Committee.

Member, Computing Policy Committee.

Member, Lab Futures Committee.

Member, Cornell Library Collaborative Learning Computer Laboratory Development Team.

Director, The Game Design Initiative at Cornell University.

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 previously worked 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 a 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.

PUBLICATIONS

"Towards Efficient Sampling: Exploiting Random Walk Strategies". *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI 2004),* San Jose, California (With W. Wei and J. Erenrich).

"Tracking Evolving Communities in Large Linked Networks". *Proceedings of the National Academy of Sciences.* (With J. Hopcroft, B. Kulis, and O. Khan).

"Equilibria for Design Tradeoffs in Autonomous Trading Agents". *Third International Conference on Autonomous Agents and Multi-agent Systems.* New York, New York, 2003. (With J. Vetsikas).

"From Spin Glasses to Hard Satisfiability Formulas". *Proceedings of the Seventh International SAT Conference,* Vancouver, British Columbia, 2004. (With H. Jia and C. Moore).

"Recent Progress in Propositional Reasoning and Search". *Discrete Applied Mathematics,* 2004. (With H. Kautz).

"Natural Communities in Large Linked Networks". *Proceedings of the Ninth International Conference on KDD,* Washington, D.C., 2003. (With J. Hopcroft, B. Kulis, and O. Khan).

"Ten Challenges Redux: Propositional Reasoning and Search". *Proceedings of the Ninth International Conference on Constraint Programming (CP 2003).* (With H. Kautz).

"Backdoors to Typical Case Complexity". *Proceedings of the IJCAI,* Acapulco, Mexico, 2003. (With R. Williams and C. Gomes).

LECTURES

"Phase Transitions and Algorithm Design", Annual Meeting, American Association for the Advancement of Science, Seattle, Washington, 2004.

"Phase Transitions, Complexity, and Algorithm Design", Santa Fe Institute, Santa Fe, New Mexico, 2004.

"Recent Advances in Fast Large-scale Reasoning Method"

City University of New York, 2004.

Department of Computer Science, Carnegie–Mellon University, 2003.

"Connections Between Statistical Physics and Computational Complexity", Department of Physics, Michigan State University, 2003.

PROFESSIONAL ACTIVITIES

Program committees (2004): International Conference on Automated Planning and Scheduling, International Conference on Principles of Knowledge Representation and Reasoning, International Conference on Theory and Applications of Satisfiability Testing, American Association for Artificial Intelligence.

AWARDS AND HONORS

Fellow, American Association for the Advancement of Science (2003).

Fellow, American Association for Artificial Intelligence (2001).

Alfred P. Sloan Research Fellow (1999–2000).

CAREER award (1998–2002).

Stephen '57 and Marilyn Miles Excellence in Teaching Award, College of Engineering, Cornell University (2002).

Cornell Outstanding Educator Award (selected most influential Cornell professor by a Merrill Presidential Scholar) (2001).

Executive Council, American Association for Artificial Intelligence (1999–2002).

Best Paper Award, Thirteenth National Conference on Artificial Intelligence (AAAI '96), 1996.

Best Paper Award, Tenth National Conference on Artificial Intelligence (AAAI '92), 1992.

Best Paper Award, First KR, 1989.

Best Paper Award, Seventh Biennial Conference of the Canadian Society for the Computer Studies of Intelligence (CSCSI '88), 1988.

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. degree 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 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 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 on emotional and social experiences, and new techniques for designing systems, including the use of self-experiment in design and new forms of evaluation for open-ended systems.

PUBLICATIONS

"The Agents of McDonaldization". In *Agent Culture* (S. Payr, ed.), Lawrence Erlbaum (2004).

"Culturally Embedded Computing". *Pervasive Computing* 3(1) (2004). (With J. Kaye, K. Boehner, J. Fairbank, G. Gay, Y. Medynskiy, and S. Wyche).

LECTURES

"Experience as Interpretation". CHI 2004 Workshop on Cross-dressing and Boundary Crossing: Exploring Experience Methods Across the Disciplines. Vienna, Austria, (April 2004). (With K. Boehner, G. Gay, J. Kaye, M. Mateas, B. Gaver, and K. Höök).

"Designing Subjects: Artificial Intelligence, Human–Computer Interaction, and the Politics of Identity & Design." Rensselaer Polytechnic Institute Department of Art and Science and Technology Studies, Integrated Electronic Arts Series (December 2003).

"Making It by Making It Strange: The Politics and Design of Domestic Technologies". 4S Conference (November 2003). Also presented at the Conference for the Society for Literature and Science (November 2003). (With G. Bell and M. Blythe).

PROFESSIONAL ACTIVITIES

Member, Editorial Board, *Interaction Studies*.

Co-chair, CHI 2004 Workshop on Reflective HCI.

UNIVERSITY ACTIVITIES

Co-organizer, Configurations Conference, Department of Science and Technology Studies.

AWARDS AND HONORS

NSF Career Award, 2002–2007.

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

Fulbright Award, 1998–1999.

Cornell Faculty Innovation in Teaching Grant (for COM S/INFO 130, Web Design and Programming) (April 2004).

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 two 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. Quark also provides 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.

PUBLICATIONS

"XRANK: Ranked Keyword Search over XML Documents". *Proceedings of the ACM SIGMOD International Conference.* San Diego, California (June 2003). (With L. Guo, F. Shao, and C. Botev).

"Index Structures for Querying the Deep Web". *WebDB,* San Diego, California (June 2003). (With J. Qiu, F. Shao, and M. Zatsman).

"TeXQuery: A Full-text Search Extension to XQuery". *World Wide Web Conference,* New York, New York (May 2004). (With S. Amer-Yahia and C. Botev).

"P–Tree: A P2P Index for Resource Discovery Applications". *World Wide Web Conference* (poster), New York, New York (May 2004). (With A. Crainiceanu, P. Linga, and J. Gehrke).

"An Indexing Framework for Peer-to-peer Systems". *World Wide Web Conference* (poster), New York, New York, May 2004. (With A. Crainiceanu, P. Linga, A. Machanavajjhala, and J. Gehrke).

"Querying Peer-to-peer Networks using P–trees". *WebDB,* Paris, France (June 2004). (With A. Crainiceanu, P. Linga, and J. Gehrke).

LECTURES

TeXQuery: A Full-text Search Extension to Xquery, IBM Almaden Research Center, March 2004.

Database Research for the Current Millennium, ICDE, March 2004.

Information Retrieval and Databases: Synergies and Syntheses, NSF Information and Data Management Workshop, September 2003.

XRANK: Ranked Keyword Search over XML Documents

University of Washington, July 2003.

Microsoft Corporation, July 2003.

BEA Systems, June 2003.

Oracle Corporation, June 2003.

PROFESSIONAL ACTIVITIES

Vice Chair, Program Committee, XML and Semistructured Data Track, ICDE 2005.

Invited Expert, World Wide Web Consortium XQuery Full-text Task Force.

Member, Program Committee, International Workshop on XQuery Implementation, Experience, and Perspectives, Paris, France (2004).

Member, Program Committee, WebDB, Paris, France (2004).

Member, Program Committee, ICDE, Boston, Massachusetts (2004).

Member, Program Committee, ACM Symposium on Applied Computing, Nicosia, Cyprus (2004).

Member, Program Committee, International VLDB Conference, Berlin, Germany (2003).

Member, Program Committee, ACM SIGMOD Conference, San Diego, California (2003).

Member, Program Committee, International Conference on Data Base Programming Languages, Potsdam, Germany (2003).

Member, Program Committee, WebDB, San Diego, California, (2003).

AWARDS AND HONORS

CAREER award.

IBM Faculty Partnership Award.

David B. Shmoys

Professor of Operations Research and Industrial Engineering and Computer Science and Member of 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, as well as in stochastic models for clustering and inventory problems.

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 (three times). He is currently on the editorial board of *SIAM Journal on Computing, SIAM Journal on Discrete Mathematics,* and *Mathematical Programming*.

PUBLICATIONS

"Primal–dual Algorithms for Deterministic Inventory Problems". In *Proceedings of the Thirty-sixth Annual Symposium on Theory of Computing* 353–362 (June 2004). (With R. Levi and R. Roundy).

"Lagrangian Relaxation for the *k*-median Problem: New Insights and Continuity Properties". In *Proceedings of the Eleventh Annual European Symposium on Algorithms,* pages 31–42 (September 2003). (With A. Archer and R. Rajagopalan).

"Improved Approximation Algorithms for the Uncapacitated Facility Location Problem". *SIAM Journal on Computing* 33: 1–25 (2003). (With F. Chudak).

"An Improved Approximation Algorithm for the Partial Latin Square Extension Problem". *Operations Research Letters* 32: 479–484 (2004). (With C. Gomes and R. Rommel).

"LP–based Approximation Algorithms for Capacitated Facility Location". In *Proceedings of the Tenth Mathematical Programming Society (MPS) Conference on Integer Programming and Combinatorial Optimization* 206–218 (June 2004). (With R. Levi and C. Swamy).

"The Design and Analysis of Approximation Algorithms: Facility Location as a Case Study". In *Trends in Optimization. American Mathematical Society (AMS) Proceedings of Symposia in Applied Mathematics* (S. Hosten, J. Lee, and R. Thomas, eds.) (June 2004).

"Facility Location with Service-installation Costs". In *Proceedings of the Fifteenth Annual ACM–SIAM Symposium on Discrete Algorithms* 1081–1090 (January 2004). (With C. Swamy and R. Levi).

LECTURES

"The Design and Analysis of Approximation Algorithms: Facility Location as a Case Study." Short Course on Trends in Optimization, AMS National Meeting, Phoenix (January, 2004).

"Stochastic Optimization Is (Almost) as Easy as Deterministic Optimization". Bertinoro Workshop on Combinatorial Optimization, Bertinoro, Italy (May 2004).

"Facility Location with Service-installation Costs". Fifteenth Annual ACM–SIAM Symposium on Discrete Algorithms, New Orleans, Louisiana (January 2004).

"Primal–dual Algorithms for Deterministic Inventory Problems", Thirty-sixth Annual Symposium on Theory of Computing, Chicago, Illinois (June 2004).

"LP–based Approximation Algorithms for Capacitated Facility Location," Tenth MPS Conference on Integer Programming and Combinatorial Optimization, New York, New York (June 2004).

"Primal–dual Algorithms for Deterministic Inventory Problems", Theory Seminar at IBM Yorktown Heights (February 2004).

PROFESSIONAL ACTIVITIES

Member, Editorial Board, *SIAM Journal on Computing*; *SIAM Journal on Discrete Mathematics*.

Guest Editor, *Journal of Algorithms, Special Issue for Selected Papers from the Eleventh Annual Symposium on Discrete Algorithms* (appeared February 2004).

Associate Editor, *Mathematics of Operations Research* (until December 31, 2003).

Associate Editor, *Mathematical Programming*.

Steering Committee Chair, Mathematical Programming Society conferences on Integer Programming and Combinatorial Optimization.

UNIVERSITY ACTIVITIES

Ad Hoc Committee of the Dean of the Faculty on Course and Exam Scheduling.

AWARDS AND HONORS

College of Engineering Sonny Yau Excellence in Teaching Award.

Emin Gün Sirer

Assistant Professor

egs@cornell.edu

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

Emin Gün Sirer received his Ph.D. from the University of Washington in 2002. His research is focused on self-organization in peer-to-peer and mobile ad hoc networks, and spans operating systems, networking, and distributed systems.

The recently emerging peer-to-peer paradigm enables novel distributed services. Sirer's group investigates peer-to-peer systems that combine high-performance with scalability and fault-tolerance. In the Beehive project, the group developed a new distributed hash table with strong (O(1)) performance guarantees for Zipf-like query distributions. The Cooperative Domain Name Service (CoDoNs) project is developing an alternative to the Domain Name Service, the system used to translate Internet names to addresses, based on Beehive that provides low latency, resilience against denial of service attacks, and security against active attackers. The Herbivore project has applied the peer-to-peer paradigm to anonymous communication, and provides strong anonymity guarantees to its participants, even against attackers with unlimited wiretapping power. The Meridian project is developing a lightweight yet accurate peer-to-peer framework for solving location-based problems, such as finding the closest node, determining a geographically diverse set, and discovering a good overlay route, that can serve as a building block for large-scale distributed systems.

Another emerging domain where self-organization plays a large role is mobile ad hoc and sensor networks. The MagnetOS project investigates operating system support for this new domain. Specifically, Sirer's group is designing and building a new operating system that improves the longevity and reliability of applications on ad hoc networks through energy-aware, adaptive object migration. The group previously investigated hybrid routing protocols and proposed the SHARP protocol, which automatically finds the optimal mix of proactive route dissemination and reactive route discovery for node-specific performance metrics. The Zoom project investigates cheap, software-based location-discovery mechanisms that can determine the physical location of wireless nodes without expensive and power-consuming GPS receivers.

Sirer's past work focused on operating-system architecture. The SPIN kernel proposed language-based 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.

PUBLICATIONS

"The Design and Implementation of CoDoNs, A Robust, Scalable and High-performance Domain Name Service". To appear in the *Proceedings of ACM Special Interest Group on Data Communications,* Portland, Oregon (August 2004). (With V. Ramasubramanian).

"Staged Simulation: A General Technique for Improving Simulation Scale and Performance". *ACM Transactions on Modeling and Computer Simulation* (April 2004). (With K. Walsh).

"Peer-to-peer Authentication With a Distributed Single Sign-on Service". In *Proceedings of the International Workshop on Peer-to-peer Systems,* San Diego, California (February 2004). (With W. Josephson and F. Schneider).

"Beehive: O(1) Lookup Performance for Power-law Query Distributions in Peer-to-peer Overlays". In *Proceedings of Networked System Design and Implementation,* San Francisco, California (March 2004). (With V. Ramasubramanian).

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

"Staged Simulation for Improving the Scale and Performance of Wireless Network Simulations". In *Proceedings of the Winter Simulation Conference,* New Orleans, Louisiana (December 2003). (With K. Walsh).

"KARMA: A Secure Economic Framework for P2P Resource Sharing". In *Proceedings of the Workshop on the Economics of Peer-to-peer Systems,* Berkeley, California (June 2003). (With V. Vishnumurthy and S. Chandrakumar).

"Comprehensive Synchronization Elimination for Java". *Science of Computer Programming* 47(2–3): 91–120 (May–June 2003). (With J. Aldrich, C. Chambers, and S. Eggers).

"Herbivore: A Scalable and Efficient Protocol for Anonymous Communication". *Technical Report TR2003–1890, Computing and Information Science, Cornell University* (February 2003). (With S. Goel, M. Robson, and M. Polte).

LECTURES

Peer-to-peer Authentication With a Distributed Single Sign-on Service." International Workshop on Peer-to-peer Systems, San Diego, California (February 2004).

CoDoNS: A High-performance, Fault-resilient Replacement for the Legacy Domain Name Service. University of Pennsylvania, Philadelphia, Pennsylvania (March 2004).

Microsoft Research, Palo Alto, California (March 2004).

Running a Self-organizing Domain Name Service 24/7. Planetlab Workshop, HP Labs, Palo Alto, California (March 2004).

AWARDS AND HONORS

Ralph S. Watts '72 Excellence in Teaching Award, 2003.

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 University in 1997. Since 1997, he has been a post-doctoral research associate and research associate in CS.

With deployment of high-bandwidth networks, computational science is entering a new era of distributed and collaborative computing. Stodghill'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.

PUBLICATIONS

"Implementation and Evaluation of a Scalable Application-level Checkpoint-recovery Scheme for MPI Programs". *Supercomputing 2004* (November 2004). (with M. Schulz, G. Bronevetsky, R. Fernandes, D. Marques, and K. Pingali).

"O'SOAP—A Web Services Framework for DDDAS Applications". *Workshop on Dynamic Data-driven Application Systems, International Conference on Computational Science* 2004 (June 2004). (With K. Pingali).

"Computational Science Simulations Based on Web Services". *Workshop on Dynamic Data-driven Application Systems, International Conference on Computational Science 2003* (June 2003). (With P. Chew, N. Chrisochoides, S. Gopalsamy, G. Heber, T. Ingraffea, E. Luke, J. Neto, K. Pingali, A. Shih, B. Soni, D. Thompson, S. Vavasis, and P. Wawrzynek).

"C3: A System for Automating Application-level Checkpointing of MPI Programs". *Sixteenth LCPC* (October 2003). (With G. Bronevetsky, D. Marques, and K. Pingali).

"Collective Operations in an Application-level Fault Tolerant MPI System". *ICS,* San Francisco, California (June 23–26, 2003). (With G. Bronevetsky, D. Marques, and K. Pingali).

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

"A Comparison of Empirical and Model-driven Optimization". *ACM PLDI Symposium.* (July 2003). (With K. Yotov, X. Li, G. Ren, M. Cibulskis, G. DeJong, M. Garzaran, D. Padua, K. Pingali, and P. Wu).

LECTURES

Web Services for Distributed Simulation: The Adaptive Software Project". Argonne National Laboratory (December 3, 2003).

PROFESSIONAL ACTIVITIES

Co-chair, Grid Checkpointing and Recovery Working Group, Global Grid Forum.

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.

PUBLICATIONS

"Maximizing the Spread of Influence through a Social Network". In *Proceedings of the Ninth ACM SIGKDD International Conference* (August 2003). (With D. Kempe and J. Kleinberg).

"Group Strategyproof Mechanisms via Primal–Dual Algorithms". In *Proceedings of the FOCS (2004)*. (With M. Pal).

"Bounding Braess's Paradox". In *Proceedings of the Fifteenth Annual ACM–SIAM Symposium on Discrete Algorithms* (January 2004). (With H. Lin and T. Roughgarden).

"Approximate Classification via Earthmover Metrics". In *Proceedings of the Fifteenth Annual ACM–SIAM Symposium on Discrete Algorithms* (January 2004). (With A. Archer, J. Fakcharoenphol, C. Harrelson, R. Krauthgammer, and K. Talvar).

"Bounding the Inefficiency of Equilibria in Nonatomic Congestion Games". In *Games and Economic Behavior* 47(2): 389–403 (May 2004). (With T. Roughgarden).

"Network Games". In *Proceedings of the Thirty-sixth ACM Symposium on the Theory of Computing* (June 2004).

LECTURES

"Network Design Games". Dagstuhl seminar on Algorithmic Game Theory and the Internet (July 2003).

"Price of Anarchy". Annual Meeting of the Packard Fellows (September 2003).

"Network Games and Approximation Algorithms". Invited Address at the Annual meeting of the American Mathematical Society (January 2004).

"Price of Anarchy in Network Games". IMA Workshop on Control and Pricing in Communication and Power Networks, Minneapolis, Minnesota (March 2004).

Euler Institute for Discrete Mathematics and its applications minicourse on "Approximation Algorithms and Games on Networks". Eindhoven, Netherlands (May 2004).

"Games in Networks". Invited talk at the ACM Symposium on Theory of Computing, Chicago (June 2004).

PROFESSIONAL ACTIVITIES

Editor-in-chief, *SIAM Journal on Computing,* since January 2004

Editor, *Journal of the ACM*; *Combinatorica*; *SIAM Journal on Discrete Mathematics*; *Mathematics of Operations Research*.

Member, Program Committee, ACM Symposium on Theory of Computing, 2004.

AWARDS AND HONORS

Best paper award, Ninth ACM SIGKDD International Conference, 2003.

Fellow, American Academy of Arts and Sciences, 2001.

Fellow, ACM, 1998.

F. I. Li '78 and D. Li '75 Excellence in Teaching Award, 2003.

Cornell Association of Computer Science Faculty of the Year, 2001.

John Simon Guggenheim Memorial Foundation Fellowship, 1999–2000.

Miller Visiting Professorship, University of California at Berkeley, 1999–2000.

Alfred P. Sloan Research Fellowship, 1991–1993.

Presidential Young Investigator (NSF), 1991–1996.

David and Lucille Packard Foundation Fellowship in Science and Engineering, 1990–1995.

Fulkerson Prize for ``A Strongly Polynomial Minimum-cost Circulation Algorithm" awarded jointly by the Mathematical Programming Society and the American Mathematical Society for a paper in *Discrete Mathematics*, 1988.

Grunwald Prize, J. Bolyai Mathematical Society (Hungary) [awarded to young researchers in mathematics], 1986.

Farkas Prize, J. Bolyai Mathematical Society (Hungary) [awarded to young researchers in applied mathematics], 1985.

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 M.I.T. 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 source and binary code. The objective is a new generation of tools that provide precise and complete information about the structure of complex software. 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 re-engineering of legacy code, security-assurance and safety-assurance inspection, and software anti-tamper protection.

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 co-founder and chairman of GrammaTech, Inc.

PUBLICATIONS

"Design and Implementation of a Fine-grained Software Inspection Tool". In *IEEE Transactions on Software Engineering* 29(8): 721–733(August 2003). (With P. Anderson and T. Reps).

"Tool Support for Fine-grained Software Inspection". *IEEE Software* 20(4): 42–50 (July–August 2003). (With P. Anderson, T. Reps, and M. Zarins).

AWARDS AND HONORS

Listed by Thomson ISI as one of the 238 most cited computer-science authors of 1981–1999.

Charles Van Loan

Joseph C. Ford Professor of Engineering 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.

PUBLICATIONS

Matrix Computations (with G. Golub).

Computational Frameworks for the Fast Fourier Transform.

Introduction to Scientific Computation—A Matrix Vector Approach Using MATLAB.

"Nineteen Dubious Ways to Compile the Exponential of a Matrix, Twenty-five Years Later". *SIAM Review* 45: 3–19 (2003). (With C. Moter).

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

LECTURES

Compression of 3D tensors using the SVD–SIAM Conference on Linear Algebra, Williamsburg, Virginia (July 2003).

Block Matrix Computations. I.I.T. at Kanpur, India (February 2004).

PROFESSIONAL ACTIVITIES

Program Chair, Householder Conference (2005).

Organizer, ARCC Tensor Decomposition Workshop, Palo Alto, California (July 2004).

UNIVERSITY ACTIVITIES

Council on Mental Health.

AWARDS AND HONORS

McCormick Advising Award, College of Engineering, 2003.

Paul Advising Award, College of Arts and Sciences, 1997.

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 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, Research Institute for Advanced Computer Science, 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. Jónsson, now at deCode Genetics), he has developed the first resultant-based algorithm for solving polynomial equations with provable accuracy guarantees.

PUBLICATIONS

"Time Continuity in Cohesive Finite-element Modeling". *International Journal for Numerical Methods in Engineering* 58(5): 679–701 (October 2003). (With K. Papoulia and C.-H. Sam).

"A Mesh Warping Algorithm Based on Weighted Laplacian Smoothing". *Proceedings of the Twelfth International Meshing Roundtable,* Santa Fe, New Mexico (September 2003). (With S. Shontz).

PROFESSIONAL ACTIVITIES

NSF Panelist.

Member, NRC Board for Assessment, National Institute of Standards and Technology.

Member, Editorial Board, *Mathematical Programming.*

Vice-chair, SIAM Activity Group on Linear Algebra.

Co-organizer, Foundations of Computational Mathematics 2005 Workshop on Numerical Linear Algebra.

UNIVERSITY ACTIVITIES

Member, Faculty Senate.

Chair, Ph.D. Admissions Committee.

AWARDS AND HONORS

Ip-Lee Teaching Award, 1999.

Guggenheim Fellowship, 1996–1997.

Presidential Young Investigator Award, 1990–1995.

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. He received his Ph.D. from the Vrije Universiteit in Amsterdam in 2003.

Vogels's research interest is in communication technologies for scalable distributed systems, with a focus on the interactions among 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.

PUBLICATIONS

"Technology Challenges for the Global Real-time Enterprise". *Journal of Knowledge Management* 8(4)(2004).

"Adding High Availability and Autonomic Behavior to Web Services". *Proceedings of the Twenty-sixth Annual International Conference on Software Engineering (ICSE 2004),* Edinburgh, Scotland (May 23–28, 2004). (With K. Birman and R. van Renesse).

"Web Services are not Distributed Objects". *IEEE Internet Computing* 7(6) (November/December 2003).

"Tracking Service Availability in Long Running Business Activities". *Proceedings of the First International Conference on Service Oriented Computing (ICSOC 2003)* Trento, Italy (December 2003).

"HPC.NET—Are CLI-based Virtual Machines Suitable for High-Performance Computing?" *Proceedings of the Fifteenth Supercomputing Conference (SC 2003),* Phoenix, Arizona (November 2003).

"Overcoming Communications Challenges in Software for Monitoring and Controlling Power Systems". *Proceedings of the IEEE, "Energy Infrastructure Defense Systems"* (October 2003). (With K. Birman, J. Chen, K. Hopkinson, R. Thomas, J. Thorp, and R. van Renesse).

"Benchmarking CLI-based Virtual Machines". *IEEE Proceedings—Software* 150(6) (October 2003).

"Scalability and Robustness in the Global Real-time Enterprise". *Proceedings of the 2003 Workshop on High-performance Transaction Processing (HPTS 2003),* Asilomar, California (October 2003).

"Heterogeneity-aware Peer-to-peer Multicast". *Proceedings of the Seventeenth International Symposium on Distributed Computing (DISC 2003)* Sorrento, Italy (October 2003). (With R. van Renesse, K. Birman, A. Bozdog, D. Dumitriu, and M. Singh).

PROFESSIONAL ACTIVITIES

Program Committee Co-Chair, International Symposium on Distributed Objects and Applications.

Program Committee Member, ACM/Usenix/IFIP Middleware (2004).

Program Committee Member, IEEE International Conference on Web Services.

Program Committee member, Usenix '04 - Usenix Annual Technical Conference.

Program Committee member, WWW 2004.

AWARDS AND HONORS

Doctor of Philosophy at the Division of Mathematica and Computer Science, Faculty of Sciences, Vrije University, Amsterdam, The Netherlands. Topic: Scalable Cluster Computing Technologies for Mission-Critical Enterprise Computing.

Recipient of a 2003 DURIP grant to develop a highly scalable cluster for experimentation with mission-critical distributed systems.

David P. Williamson

Professor

CIS, joint with the School of Operations Research and Industrial Engineering

dpw@orie.cornell.edu* and *dpw@cs.cornell.edu

http://www.orie.cornell.edu/~dpw/

David Williamson received his Ph.D. in computer science from M.I.T. in 1993. He spent several years as a research staff member at the IBM T. J. Watson Research Center, and then as a senior manager at the IBM Almaden Research Center. He joined Cornell University in January 2004 with a joint appointment in CIS and the School of Operations Research and Industrial Engineering.

His primary research interest has been in the design and analysis of polynomial-time algorithms for the approximate solution of hard problems in discrete optimization, especially problems arising in network design, scheduling, facility location, and routing. He focuses on the use of techniques from the area of mathematical programming for designing such algorithms, including such techniques as the primal–dual method and semidefinite programming.

PUBLICATIONS

"Approximation Algorithms for MAX 3–CUT and Other Problems Via Complex Semidefinite Programming". *Journal of Computer and System Sciences* 68: 442–470 (2004). (With M. Goemans).

"A Faster, Better Approximation Algorithm for the Minimum Latency Problem". Preliminary version appeared in *Symposium on Discrete Algorithms* (*SODA 2003)*. (With A. Archer and A. Levin).

"Searching the Workplace Web". Appeared in *World Wide Web Conference 2003*. (With R. Fagin, R. Kumar, K. McCurley, J. Novak, D. Sivakumar, and J. Tomlin).

PROFESSIONAL ACTIVITIES

Area Editor, *Discrete Optimization; Mathematics of Operations Research.*

Associate Editor, *SIAM Journal on Computing; SIAM Journal of Discrete Mathematics.*

Member, Program Committee, 2004 IEEE Symposium on the Foundations of Computer Science.

AWARDS AND HONORS

Mathematical Programming Society Fulkerson Prize, 2000.

SIAM Optimization Group Prize, 1999.

SIAM DiPrima Prize, 1996.

Mathematical Programming Society A. W. Tucker Prize, 1994.

Golan Yona

Assistant Professor

golan@cs.cornell.edu

http://ww.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 Biozon (*http://biozon.cornell.edu)*.

PUBLICATIONS

"Distributional Scaling: An Algorithm for Structure-preserving Embedding of Metric and Nonmetric Spaces". *Journal of Machine Learning Research* (2004). (With M. Qusit).

"Automatic Prediction of Protein Domains from Sequence Information Using a Hybrid Learning System". *Bioinformatics* (2004). (With N. Nagarajan).

"Prediction of Protein–protein Interactions and the Interaction Site from Sequence Information—An Extensive Study of the Co-evolution Model". *Technical Report TR2004–1919, Computing and Information Science, Cornell University* (2004). (With J. Davis).

"The URMS–RMS Hybrid Algorithm for Fast and Sensitive Local Protein Structure Alignment". *Technical Report TR2004–1922, Computing and Information Science, Cornell University* (2004). (With K. Kedem).

LECTURES

"BIOZON: A Unified Knowledge Resource on DNA Sequences, Proteins, Complexes, and Cellular Pathways".
Stanford University (February 2004).
University of California at Berkeley (February 2004).
University of Toronto (February 2004).
Cornell University (January 2004).

Institute for Pure and Applied Mathematics, University of California at Los Angeles (March 2004).

PROFESSIONAL ACTIVITIES

Member, Program Committee, Eighth Annual International RECOMB Conference, San Diego, California (April 2004).

Member, 'Faculty of 1000' (*http://www.facultyof1000.com*).

Reviewer: *Bioinformatics; Journal of Computational Biology; Journal of Molecular Biology; Proteins: Structure, Function, and Bioinformatics; Protein Science; NSF; The Binational Science Foundation; Israel Science Foundation; Conference on Neural Information Processing Systems.*

UNIVERSITY ACTIVITIES

Member, Planning Committee, Physical Sciences/Life Sciences Interface.

Member, Planning Committee, Computational and Statistical Genomics.

AWARDS AND HONORS

NSF CAREER award, 2002.

Burroughs–Wellcome Fellow from the Program in Mathematics and Molecular Biology, 2000.

Ramin Zabih

Associate Professor

rdz@cornell.edu

http://www.cs.cornell.edu/~rdz/index.htm

Ramin Zabih received undergraduate degrees in computer science and in mathematics from M.I.T., and a Ph.D. in computer science from Stanford 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 Weill Medical College in New York City.

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.

PUBLICATIONS

"A Multiprocessor Scheduling Implementation of the Simultaneous Multiple Volume (SMV) Navigator Method". *Magnetic Resonance in Medicine* [in press]. (With V. Kolmogorov, T. Nguyen, A. Nuval, P. Spincemaille, M. Prince, and Y. Wang).

"Spatially Coherent Clustering Using Graph Cuts". *Proceedings of the IEEE Computer Vision and Pattern Recognition Conference* (June 2004). (With V. Kolmogorov).

"What Energy Functions Can Be Minimized via Graph Cuts?". *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26(2): 147–159 (2004). (With V. Kolmogorov).

"Simultaneous Multiple Volume (SMV) Acquisition Algorithm for Real-time Navigator Gating". *Journal of Magnetic Resonance Imaging* 21(9): 969–975 (November 2003). (With V. Kolmogorov, R. Watts, M. Prince, and Y. Wang).

"Automatic Segmentation of Contrast-enhanced Image Sequences". *Journal of X-ray Science and Technology* 11(4): 241–251 (2003). (With J. Kim).

"A Segmentation Algorithm for Contrast-enhanced Images". *ICCV* 502–509 (October 2003). (With J. Kim).

"Visual Correspondence Using Energy Minimization and Mutual Information". *ICCV* 1033–1040 (October 2003). (With J. Kim and V. Kolmogorov).

"Generalized Multicamera Scene Reconstruction". *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition* 501–516 (July 2003). (With V. Kolmogorov and S. Gortler).

LECTURES

"Fast Energy Minimization for Computer Vision via Graph Cuts".

Microsoft Research (April 2004).

University of Washington (April 2004).

PROFESSIONAL ACTIVITIES

Program Committee, CVPR (June 2004); International Conference on Pattern Recognition (June 2004); European Conference on Computer Vision (ECCV) (May 2004); Asian Conference on Computer Vision (2004); ICCV (October 2003); CVPR (June 2003); International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition (July 2003).

Reviewer for:

Journal of Lisp and Symbolic Computation; AAAI/IJCAI Conferences; International Joint Conference on Artificial Intelligence; Journal of Logic Programming; Journal of the ACM; IEEE International Conference on Robotics and Automation; Image and Vision Computing; Artificial Intelligence; ACM Conference on Multimedia; Springer-Verlag Lecture Notes in Computer Science; IEEE Transactions on Pattern Analysis and Machine Intelligence; IEEE Transactions on Medical Imaging; International Journal of Computer Vision; National Science Foundation *(CAREER award and regular proposals); ACM SIGGRAPH Proceedings.*

UNIVERSITY ACTIVITIES

Member, Computer Security Committee, Department of Radiology.

Field Memberships: Computer Science, Cognitive Studies, Electrical and Computer Engineering.

Minor Member of five Ph.D. Committees, in the Fields of Electrical Engineering, Mechanical and Aerospace Engineering, and Geological Sciences.

AWARDS AND HONORS

Best Paper Award, European Conference in Computer Vision (ECCV), 2002.

Abraham Wong Teaching Award, College of Engineering, 1995.

All photos by Cornell University Photography (Robert Barker, Frank DiMeo, Charles Harrington, Nicola Kountoupes) unless otherwise credited.

Printed on recycled paper.

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

Communication and Marketing Services
9/04 1M EL 04103

Explanations of Acronyms Used in This Report

AAAI	American Association for Artificial Intelligence
AAAS	American Association for the Advancement of Science
ACL	Association for Computational Linguistics
ACM	Association for Computing Machinery
ACS	American Chemical Society
ACSU	Association of Computer Science Undergraduates (Cornell)
ACT	Algorithms and Computational Theory
ADEPT	a Sequel (SQL) tool
AF	U.S. Air Force
AFOSR	U.S. Air Force Office of Scientific Research
AFRL	U.S. Air Force Research Laboratory
AI	artificial intelligence
ALOHA	a product of Aloha Networks
AMIA	American Medical Informatics Association
AMS	American Mathematical Society
ARDA	Advanced Research and Development Activity
ART	Artificial Intelligence
ASE	Advances in Science and Engineering
ASP	Adaptive Software Project
ATC	Air Training Corps
BGP	Border Gateway Protocol
BOOM	Bits On Our Minds
BSCB	Biological Statistics and Computational Biology (Cornell)
BTR	Bridging the Rift Center
CAD	computer-aided design
CALS	New York State College of Agriculture and Life Sciences (Cornell)
CAREER	NSF Faculty Early Career and Development Award
CASE	Council for Advancement and Support of Education
CAVE	computer-aided virtual environment
CBSU	Computational Biology Service Unit (part of CTC at Cornell)
CBT	core-based tree
CCB	Chemistry and Chemical Biology (Cornell)
CDMA/ NC–PRMA	a protocol for wireless multimedia communication
CHI	Computer–Human Interaction
CIPIAF	critical infrastructure protection and information assurance
CIS	Computing and Information Science
CISE	NSF Directorate for Computer and Information Science and Engineering
CLI	Command Line Interface
CLR	common language runtime
CMP	chip multiprocessing
CNS	Computer and Network Systems
COCA	Cornell Online Certification Authority
CoDoNs	Cooperative Domain Name Service
COM S	[Computer Science course acronym] (Cornell)
COMSOC	Communications Society
CorSSO	Cornell Single Sign-On
CP	constraint programming
CRCNS	Collaborative Research in Computational Neuroscience
CS	Department of Computer Science (Cornell)
CS&E	Computational Science and Engineering
CSE	Computer Science Education
CSL	Computer Systems Laboratory (Cornell)
CSP	communicating sequential processors **or** Cognitive Studies Program (Cornell)
CTC	Cornell Theory Center
CURIE	Cornell University Research in Engineering
CvaR	Conditional value at Risk
CVPR	Conference on Computer Vision and Pattern Recognition
DA&G	Digital Arts and Graphics Program (Cornell)
DARPA	Defense Advanced Research Projects Agency
DDDAS	dynamic data-driven application systems
DDoS	Distributed Denial of Service
DISC	International Symposium on Distributed Computing
DLKC	International Symposium on Digital Libraries and Knowledge Communities in Networked Information Society
dmc	innovative approaches to the integration of data, models, communications, analysis and/or control systems, including dynamic, data-driven applications for use in prediction, risk-assessment and decision-making
DMCA	Digital Millennium Copyright Act
DNA	deoxyribonucleic acid
DRAM	dynamic random access memory
DSM	distributed shared memory
DTN	Disruption Tolerant Networking
DURIP	Defense University Research Instrumentation Program
EATCS	European Association for Theoretical Computer Science
ECE	School of Electrical and Computer Engineering (Cornell)
ECCV	European Conference on Computer Vision
ECDL	European Conference on Digital Libraries
ELCAP	International Early Lung Cancer Action Project
EMNLP	Empirical Methods in Natural Language Processing
EPA	Environmental Protection Agency
EPR	electronic paramagnetic resonance
ETH	Eidgenößische Technische Zürich [Swiss Federal Institute of Technology in Zürich]
EVS	Ethics and Social Values
FEDORA	Flexible and Extensible Digital Object and Repository Architecture
FIRST	For Inspiration and Recognition of Science and Technology
FOCS	Foundations of Computer Science
4S	Society for the Social Studies of Science
FPGA	field-programmable gate array
GDIAC	Game Design Initiative at Cornell
GE	General Electric
GECCO	Genetic and Evolutionary Computation Conference
GNS	Gene Network Sciences
GRAPH	Graphics and Interactive Techniques
HCI	human–computer interaction
HLT	Human Language Technology
HMM	hidden marker model
HOL	higher-order logic
HPCA	high performance computer architecture
IAI	Information Assurance Institute (Cornell)
ICCV	International Conference on Computer Vision
ICDE	International Conference on Data Engineering
ICDM	International Conference on Data Mining
ICML	International Conference on Machine Learning
ICS	International Conference on Supercomputing
ICSE	International Conference on Software Engineering
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IF	Information Directorate
IGERT	Integrative Graduate Education and Research Traineeship
IISI	Intelligent Information Systems Institute (Cornell)
IJCAI	International Joint Conference on Artificial Intelligence
IJCAR	International Joint Conference on Automated Reasoning
IMA	Institute for Mathematics and its Applications
int	integration of computing, networking, human–computer interfaces, and information management to support reliable, complex, distributed systems
IP	Internet protocol
IPSN	Information Processing in Sensor Networks
IPTPS	International Workshop on Peer-to-peer Systems
IR	Information Retrieval
ISI	Symposium on Intelligence and Security Informatics
ITR	Information Technology Research for National Priorities
Jif	Java Information Flow
KDD	Knowledge Discovery and Data Mining
K–MODDL	Kinematic Models for Design Digital Library
KR	International Conference on Principles of Knowledge Representation and Reasoning
LCPC	International Workshop on Languages and Compilers for Parallel Computers
LICS	IEEE Conference on Logic in Computer Science
LLNL	Lawrence Livermore National Laboratories
LOOPP	Learning, Observing, and Outputting Protein Patterns
LPE	logical programming environment
M&AE	Mechanical and Aerospace Engineering (Cornell)
MAC	medium access control
MagnetOS	(a distributed operating system for ad hoc and sensor networks)
MATLAB	[software from Mathworks]
MAX 3-CUT	(a mathematical problem)
MBG	Molecular Biology and Genetics (Cornell)
M.Eng.	Master of Engineering degree (Cornell)
METRICS	Measurement and Modeling of Computer Systems
ML	machine learning
MOD	Management of Data
MOIL	MOlecules at ILlinois [a program used to view and analyze molecular structures and dynamics]
MPI	message passing interface
MPS	Mathematical Programming Society
MRI	magnetic resonance imaging
MSPA–MCS	Mathematical Sciences Priority Area in Mathematical and Computer Science
MURI	Multidisciplinary University Research Institute
NAACL	North American Association for Computational Linguistics
NASA	National Aeronautic and Space Administration
NAT	network address translation
NCBC	National Center for Biomedical Computing
NeTS	Network Technology and Systems
NHS	National and Homeland Security
NIJ	National Institute of Justice
NIH	National Institutes of Health
NIPS	Neural Information Processing Systems
NLP	natural language processing
NOSS	Networking of Sensor Systems
NP	[part of a problem in mathematics: the P vs NP Problem]
NR	Networking Research
NRRC	Northeast Regional Research Council
NSDL	National Science Digital Library
NSF	National Science Foundation
NUTSS	(an Internet architecture system for determining protocols)
NY-ELCAP	New York Early Lung Cancer Action Project
O(1)	(a function that converges to zero)
ONR	Office of Naval Research
OR	operations research
OR&IE	School of Operations Research and Industrial Engineering (Cornell)
ORIS	Object Repository Integrated Security
PACT	International Conference on Parallel Architectures and Compilation Techniques
PARC	Palo Alto Research Center
PA–RISC	(a computer architecture used primarily by Hewlett Packard)
PCG	Program of Computer Graphics (Cornell)
PECASE	Presidential Early Career Award for Scientists and Engineers
PERCS	Productive, Easy-to-use, Reliable Computing Systems
PLDI	Programming Language Design and Implementation
PR–DP	penalized regression dynamic programming
PRF	Petroleum Research Fund
PRL	proof/program refinement logic
PVS	(specification and verification system software)
RECOMB	International Conference on Research in Computational Molecular Biology
RMS	root mean square
RPI	Rensselaer Polytechnic Institute, Troy, New York
SAT	propositional satisfiability **or** International Conference on Theory and Applications of Satisfiability Testing
SCorES	Synthesis of Correct Embedded Systems
SEI	Software Engineering Institute
SIAM	Society for Industrial and Applied Mathematics
SIG	Special Interest Group
sim	innovation in computational modeling or simulation in research or education
SOAP	Simple Object Access Protocol
SMETE	Science, Math, Engineering, and Technology Education
SMV	Simultaneous Multiple Volume
SNAP	sensor network asynchronous processor
soc	interactions and complex interdependencies of information systems and social systems
SPIN	(linear temporal logic model checking software)
SPKI	Simple Public Key Infrastructure
SRI	Stanford Research Institute
SQL	Structured Query Language
SRDS	Symposium on Reliable Distributed Systems
SVD	Singular Value Decomposition
SVM	support vector machine
SY	System Design and Implementation
TAL	Typed Assembly Language
TKDE	Transactions on Knowledge and Data Engineering
TPDS	Transactions on Parallel and Distributed Systems
URMS	unit-vector root mean square
USDA	United States Department of Agriculture
VLDB	very large databases
VLSI	very large scale integration
VPC3	Value Predictor for Compression Algorithm
VTC	IEEE Vehicular Technology Conference
WebDB	International Workshop on the Web and Databases
WNL	Wireless Networks Laboratory (Cornell)
WoLLIC	Workshop on Logic, Language, Information, and Computation
XML	extensible markup language

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