Conduit

A RESEARCH AND ALUMNI NEWS MAGAZINE

DEPARTMENT OF COMPUTER SCIENCE
BROWN UNIVERSITY

A Fresh Perspective on Computer Science

Volume 16, Number 2

Conduit

Published twice yearly by the Department of Computer Science at Brown University, Conduit is distributed free to all computer science alumni, faculty, staff, students and industrial partners.

Queries about this publication can be directed to:
conduit@cs.brown.edu or
115 Waterman Street
Campus Box 1910
Providence, RI 02912

PUBLISHER
Industrial Partners Program,
Department of Computer Science

EDITOR-IN-CHIEF
Amy Tarbox

GRAPHIC DESIGN
Amy Kendall

COPY EDITOR
Amy Morton

FEATURED COLUMNISTS
Eugene Charniak
Shriram Krishnamurthi

CONDUIT ON THE WEB
John Bazik
Kathy Kirman

PRINTING & DISTRIBUTION
Brett Turner, Information Services
Brown Graphic Services

BROWN

Cover photo: Amy Tarbox

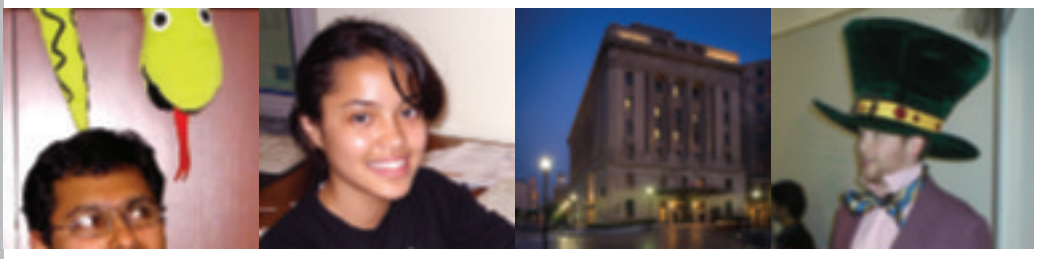

Notes from the Chair

the latest news from 115 Waterman

Greetings to all CS alums, supporters and friends.

Recent months have brought about a number of changes to the Computer Science Department and I am excited to update you.

In May, we awarded an all-time high of 31 Master of Science degrees and 12 Ph.D. degrees. In addition, 41 Bachelor's degrees were conferred in the department's 28th year. The department is thrilled to have had a large and growing number of talented graduates in the last couple of years. Class of 2007— please keep in touch.

In July, I assumed the responsibility of department chair from Eli Upfal, who had held the position for five years. I am very grateful to Eli for his leadership and steadfast efforts on behalf of the department and I wish him an enjoyable and productive time during his well-deserved sabbatical in Italy.

In my first few months as department chair, I have been impressed with the continued stream of research accomplishments of the entire CS faculty. In particular, I am happy to share with you some of the successes of our youngest faculty—Amy Greenwald, John Jannotti, Chad Jenkins, Anna Lysyanskaya, Ben Raphael and Meinolf Sellmann:

- Amy published her first book titled "Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition."
- John received an NSF grant on "Sensing in Three Dimensions with Smart Cameras."
- Chad was selected as one of the recipients of the Presidential Early Career Award for Scientists and Engineers (PECASE) for his research on the development of methods for autonomous robot control and perception.
- Anna has been included in the 2007 Technology Review TR35 for her work in cryptography, identity verification and privacy protection. Each year, the TR35 honors 35 innovators under the age of 35 in science and technology.
- Ben received a grant from the Department of Defense's breast cancer research program to study hormone-regulated gene amplification.
- Meinolf founded the Constraint Programming Society of North America and hosted the International Conference on Principles and Practice of Constraint Programming and the International Conference on Automated Planning and Scheduling.

I am also pleased to share with you that sponsored research expenditures have increased 27% from FY 2006 to FY 2007. This is an amazing accomplishment since obtaining research grants has become extremely challenging in the last several years, with some NSF programs reporting single-digit proposal acceptance rates. The fact that the department has significantly boosted its sponsored research funding in this highly competitive environment is an affirmation of the depth and innovation of our research projects. In addition, supported by Brown Technology Partnerships, our faculty members have established several interesting research and technology transfer collaborations with a variety of companies ranging from established information technology leaders, such as Microsoft and Sun Microsystems, to fast-growing startups, such as IAM Technologies and Vertica Systems.

Feedback we receive from current students and alumni helps drive our course development efforts to ensure the strength and appeal of our concentrations. We have been hard at work fine-tuning the CS curriculum and have introduced a new course (CS19), taught in the fall by Steve Reiss, which covers the material of CS15 and CS16 at an accelerated pace. This

continued on page 3

Industrial Partners Program

The primary goals of the Industrial Partners Program (IPP) are to exceed the expectations of our partner companies in terms of recruiting and outreach; to allow our faculty to engage in challenging and meaningful research collaborations and to provide resources and employment opportunities to our students.

The department wishes to thank our Industrial Partners for their support:

Collaborators
Cisco Systems

Premier Partners
Adobe
Network Appliance
Sun Microsystems

Affiliates
Apple
Facebook
Google
GTECH
ITA Software
Microsoft
Oracle
VMware

Small Business Supporters
IAM Technologies
Linkage Systems
Vertica Systems

Individuals
Jim Baker, Zyasoft
Paul Edelman, Edelman & Associates

For more information about the Industrial Partners Program, contact:

Ugur Cetintemel
Assistant Professor and IPP Co-Director
Telephone: (401) 863-7600
Ugur_Cetintemel@brown.edu

John Hughes
Associate Professor and IPP Co-Director
Telephone: (401) 863-7600
John_Hughes@brown.edu

Amy Tarbox, Program Manager
Telephone: (401) 863-7610
ipp@cs.brown.edu

To learn more about IPP visit:
http://www.brown.edu/industry/ipp

continued from page 2

course is intended for freshmen who have acquired a solid CS background in high school or for students who have completed CS4. We are also developing a new course on creative programming for digital artists in collaboration with Brown's Literary Arts Program and the Digital+Media Department at Rhode Island School of Design.

As I have become more involved in the day-to-day operation of the department, I have increasingly appreciated the efficiency and proficiency of our administrative and technical staff. The faculty and students are very grateful to them for making sure that the department runs smoothly while we are focused on our research and educational endeavors.

We have hired two new staff members, Jane McIlmail and Amy Tarbox. Jane, a long-time Brown staffer and a member of the Class of '81, is our new academic department manager. She was most recently executive assistant to the dean of the faculty. The department is delighted to have Jane's expertise and perspective on board. For more information on Jane, please see page 22. Amy Tarbox comes to us from Millennium Pharmaceuticals in Cambridge, Mass., where she was a senior associate of corporate communications. Amy holds a B.S. in business communications from Bentley College and is the department's manager of the Industrial Partners Program (IPP) and alumni affairs.

I'd like to take this opportunity to also thank Trina Avery for coming out of retirement and acting as the interim academic department manager in the first half of 2007. Her vast knowledge of university procedures and department history was invaluable to us during a transitional time, and we are grateful for her hard work.

The department is very appreciative of a recent generous gift from John V. Guttag (Class of '71), who is the Dugald C. Jackson Professor of Electrical Engineering and Computer Science at MIT, on behalf of the Guttag Family Foundation. This gift will help our initiatives to enhance the participation of women and minorities in computer science. Cisco Systems, Goldman Sachs, Google and Microsoft also supported these efforts by contributing to our Artemis Project, a five-week summer program designed to encourage high school women to pursue college degrees in science and technology (see page 27 for more information).

As we look forward to a bright future for the department, I would like to remember the legacy of our dear friend and colleague Paris Kanellakis. His unbounded energy and outstanding scholarship greatly inspired all those who interacted with him. Personally, I also admired his fierce stand on academic and political freedom. Through a generous gift from his parents, General and Mrs. Kanellakis, the department has funded several bright Ph.D. students who are poised to follow in Paris's steps. Each year we honor and remember Paris with our Kanellakis Lecture, given by distinguished computer scientists.

In closing, I would like to express our gratitude to all of our alumni, friends and industrial partners for their continued generosity, involvement, and support of our students and research. From attending reunions and IPP Symposiums to recruiting students and graduates, we sincerely appreciate your interactions with the department. A gift to the computer science department's fundraising, scholarship and fellowship initiatives directly benefits the faculty members and students who drive our ongoing, innovative work in this discipline. We are fortunate to have such an enthusiastic and involved community—thank you!

Roberto Tamassia
Professor and Chair
Department of Computer Science
Brown University

Optimization at Brown

Pascal Van Hentenryck

Research in combinatorial optimization and constraint programming at Brown's computer science department is thriving.

Pascal Van Hentenryck had two books published in the last two years, *Constraint-Based Local Search* (MIT Press), with his former Ph.D. student Laurent Michel, and *Online Stochastic Optimization* (MIT Press), with his Ph.D. student Russell Bent. As of October first of 2007, the latter made it to rank 13 on the best sellers in technology list compiled by YBP Library Services.

His recent paper with Ph.D. student Luc Mercier, on online anticipatory algorithms for large multistage stochastic integer programs, won an award at IJCAI'07. In this paper, Luc and Pascal explain theoretically why the anticipatory algorithms designed by Pascal and Russell behave so well in practice in applications ranging from dynamic fleet management to dynamic project scheduling. Pascal and Laurent have also been continuing the development of the award-winning Comet system (www.comet-online.org), giving tutorials at the main AI and OR conferences and pushing the boundaries of its parallel and distributed computing abstractions.

Meanwhile, Meinolf Sellmann has been particularly busy this fall. He was the conference co-chair (along with Laurent Michel) for CP 2007 (the premier conference in constraint programming) which was held here in Providence in late September, co-located with ICAPS (the International Conference on Automated Planning and Scheduling). His work won't ease up much, though, as he is the workshop co-chair for AAAI 2008 as well. Earlier this year, Pascal was program chair of CP'AI'OR 2007 with Laurence Wolsey, a conference which is steadily attracting very high-quality work at the intersection of constraint programming, artificial intelligence, and operations research.

On the technical side, Meinolf had a paper at AAAI 2007 on "Propagating Knapsack Constraints in Sublinear Time," along with Irit Katriel, Pascal Van Hentenryck, and Eli Upfal, and a CPAIOR paper, jointly with Luc Mercier and undergraduate Dan Leventhal (now working at IPP partner Microsoft), on "The Linear Programming Polytope of Binary Constraint Problems with Bounded Tree-Width."

Pascal had four papers at AAAI 2007 on topics ranging from sport scheduling to large-scale vehicle routing and the synthesis of search algorithms, as well as two papers at CP'07 on model-based visualizations of local search (with Gregoire Dooms and Laurent) and on the transparent parallelization of constraint programming (with Andrew See and Laurent).

While Pascal and Meinolf form the core of the optimization research group, other faculty with interest in algorithms (including Claire Mathieu and Eli Upfal) are also active participants and have co-authored papers with them on a variety of topics related to combinatorial optimization.

Three post-doctoral students have also joined the group this year. Gregoire Dooms works on online stochastic optimization and the development of the Comet system, while Ivan Dotu and Manuel Cebrian focus primarily on optimization algorithms for tertiary structure prediction.

Watch for further updates here on their progress in increasing the robustness and speed of optimization algorithms, on new anticipatory algorithms for highly stochastic programs, and about exciting applications in dynamic project management, networking, processor validation, and computational biology.

Anticipatory Algorithms? What're Those?

John Hughes, IPP Director

It's an hour before change-of-shift at the hospital, and on one of the surgical units, the nurse manager gets a call. "Nancy? This is Robin. I can't make my shift—I just got a call from my son's school ... his leg got broken during soccer practice. I have to leave right now to meet them at the hospital." And so a well-made plan begins to unravel. A few weeks earlier, all the nurses talked to the nurse manager and gave their constraints for scheduling, such as "I can't work on the 24th—it's my son's wedding," and "I only want to work third shift this month, except for the 11th, when I can't work at all." They also stated some desires—things they'd like if it's consistent with the rest of the schedule—such as, "I'm happy to work a double shift or two, but I'd really like to have a four-day stretch away from work sometime this month." The nurse manager compiled all the requests into a master plan for staffing—a schedule of shifts for all those nurses.

What happens when Robin calls? A panic, mad shuffling, a call to see if one of the per diem workers can come in; calls to others asking if they can arrive early and do a double shift, which then creates a need to reschedule the next morning's shift; and in the end, the disrupted schedule is far worse than the one the nurse manager could have made, if only she'd known that Robin couldn't do this shift.

An anticipatory algorithm can help: instead of optimizing with respect to some constraints and goals, the algorithm can also be told something about the expected classes of disturbances that may arise, and can then formulate a plan with the property that

(a) it's pretty good, and
(b) it remains pretty good under all expected contingencies.

continued on page 26

On Discourse

Eugene Charniak

My research over the last two years has focused on what is called "discourse." The "meaning" aspect of language is divided into "semantics," how the meaning of a sentence is built up out of the meaning of the words, and "discourse," how the meaning of an entire text (we will be discussing news articles) is built up out of the meaning of the sentences. As a computer scientist, I want to understand such processes by building computer programs that can mimic them.

As is almost always the case in studying human abilities, our understanding of what is going on in discourse is tremendously limited. We thus look at the simpler problems hoping to gain some purchase on the more complicated ones. In this spirit, consider the following comparatively clear-cut problem. I take a news article and permute the order of the sentences. I then give you (or my program) the original and the randomly permuted version and ask you (or it) which is which. Obviously, if you can do this, it must be on the basis of discourse knowledge since the individual sentences remain unperturbed.

I don't have actual data, but I would guess that people can get this right, say, 98% of the time. Maybe more. Right now, thanks to joint work with one of my graduate students (Micha Elsner), a former undergraduate (Joseph Ousterwiell) and a current undergraduate (Joseph Browne), my laptop is at 93%. In the rest of this article I describe how how my laptop does it. Then, if you have made it that far, I will tell you why someone might find this useful.

I became interested in this problem because of some papers by Regina Barzilay, Lillian Lee, and Mirella Lapata, who pointed out that one way to attack it is to exploit the fact that in text, entities get mentioned, stick around for a bit, and then drop out of the discourse. They formalized this by imagining a graph showing the entities mentioned along one axis and the sentence numbers along the other. At each point in the graph we put in an X if the entity is mentioned and leave it blank otherwise. (I am simplifying a bit here.) They reasoned that the "clumpy-ness" of when entities are mentioned will show up in the clumpy-ness of this graph. And if you think about it for a bit, you see that if you permute the order of the sentences, the clumpy-ness will be lost. Formalize this some more ("what is the probability that an entity is mentioned in sentence x if it has [or has not] been mentioned in sentence x–1?") and a program can distinguish real from permuted about 85% of the time.

Note, by the way, that we can collect the necessary data just by perusing regular news articles. We have sitting around our file system about 500 million words of the Wall Street Journal, New York Times, LA Times, etc. The idea in this project is that we NEVER hand code any information. If we cannot figure out a way of learning the necessary probabilities directly from this or other available text, we do without. The history of artificial intelligence has taught us that AI systems that depend on extensive hand coding are almost always brittle. As it is, this simple technique lets us throw any news article (or its permutation) and our system will get it right 85% of the time.

It seemed to me that this idea could be pushed a lot further. More generally, to raise our capabilities from the 85% to the 93% we mentioned at the beginning, we just need to add further clues to the mix. Consider pronouns. Which of the following two orderings is more plausible?

"President Bush discussed the troops in Iraq today. He said the troops were fine."

"He said the troops were fine. President Bush discussed the troops in Iraq today."

This one is really obvious. Any program that tries to assign antecedents to pronouns will have a great deal of trouble with he second version. More generally, we expect that a pronoun of a particular gender (he/she/it) and number (he/they) will have an entity in the recent past to which it can refer. The less plausible the antecedents, the less plausible the ordering.

Actually, there is a second clue in our two-sentence example. Notice that the first mention of "troops" is in the noun-phrase "the troops in Iraq" while the second is just "the troops." As you might expect, it is typical to introduce an entity with a full description and then abbreviate it considerably in further mentions. Thus we have a model that looks at the complexity of the entity descriptions: the more complex, the earlier in the text it is probably situated. If we add these two components (pronouns and noun-phrase descriptions) to our model we are now up to about 91%.

We get the last 2% (bringing us to the 93% success rate mentioned at the beginning) by looking at quotations. One really simple clue comes from quotes that start in one sentence and end in the next. If the order of these sentences is reversed, or a few sentences are permuted into positions between them, it is a dead giveaway. A less obvious clue, one I bet you never thought about before, is that news articles almost never (only 2% of the time) start with a full quotation. This is as opposed to a quotation fragment like my use of "discourse" in the first sentence of this essay.

Obviously there is much more we can and, I expect, will do. I wonder if verb tense and other temporal words have certain patterns. A future-tense statement might be more likely if the previous sentence is also in a future tense. The use of temporal modification ("today") is probably more common in first sentences. I think I will ask the students in my graduate statistical natural language processing course to research this for me.

More critically, often entities are implied by previous statements.

"A Boeing 747 crashed today at Chicago's Midway airport. The pilot was intoxicated."

The Boeing 747 implies that there is an associated pilot. The reverse order does not work very well. This ties in with some

continued on page 12

James Bond at the Electronic News Stand

Anna Lysyanskaya

While on a top-secret assignment, James Bond wants to read his favorite local newspaper. This requires a subscription. The standard way to proceed is for the magazine to first ask Bond who he is, have him prove that he is, indeed, who he says he is, and then check that Bond is an authorized user. Let us call this approach "authentication BY identification."

Authentication by identification solves the basic problem of guaranteeing access only to authorized users. However, as far as protecting Bond's privacy (and therefore his secret mission—his reading habits might give away information about it) is concerned, this approach does not do very well: every transaction requires James Bond to reveal his identity.

What if, instead of using his real identity, James Bond just uses some arbitrary username with the magazine? Unfortunately, this is not good enough (in the U.S., anyway). It turns out that in the U.S., 87% of the population is likely to be uniquely identifiable by their zip code, gender and date of birth. He may disclose his zip code in order to obtain the local weather report. He might disclose his date of birth to retrieve his horoscope. His other habits may reveal his gender. If all of these transactions are linked together by the same username, they are sufficient to identify him!

What is needed, instead, is a method for James Bond to convince the magazine that he has a subscription without disclosing who he is, and in such a way that his present session cannot be linked to his previous sessions. Brown's cryptography group is actively involved with this problem—the problem of carrying out authentication WITHOUT identification.

Authentication BY identification can be accomplished, using a standard digital signature scheme, as follows: When James Bond orders a subscription, the magazine can sign his public key (PK), resulting in a certificate (Cert). When logging in, Bond must submit (PK, Cert), and prove that he is the owner of the public key (PK).

To get authentication WITHOUT identification, we need to modify this approach to let Bond convince the magazine that he is a valid subscriber without giving away any information about his identity. A zero-knowledge proof is a theoretical tool that allows a prover to convince a verifier that a statement is true, without revealing any information besides the fact that the statement is true. For example, if a prover knows a satisfying assignment to a Boolean formula Φ, then he can convince the verifier Φ is satisfiable without revealing the input values on which Φ evaluates to TRUE. Applying this to our situation, using a zero-knowledge proof, Bond wants to convince the magazine that he knows a public key, its corresponding secret key, and a certificate on the public key without revealing any of this data.

Even though it has been known for the past twenty years that everything that is provable is provable in zero knowledge, the general protocol is too expensive to use in practice. We need a digital signature scheme designed with our specific application in mind. It must support efficient zero-knowledge proofs of knowledge of a secret key, a public key, and a certificate such that (1) the secret key corresponds to the public key; and (2) the certificate is the magazine's signature on the public key.

Over the last ten years in a series of papers (and jointly with many others, notably Jan Camenisch of IBM Zurich and recently my Ph.D. students Mira Belenkiy and Melissa Chase) we came up with several signature schemes supporting such protocols and various applications for these signature schemes.

Not only can James Bond gain access to the magazine without revealing his public key, he can even obtain a certificate (Cert) on his public key without revealing that public key, or indeed anything about his identity. He can anonymously obtain credentials from various organizations, and, when needed, prove that he possesses them without revealing any other information. (He is operating, as per instructions from his superiors, on a need-to-know basis.) Moreover, credentials may have certain attributes, and James Bond can prove that these attributes satisfy broad classes of relations. The running time for these protocols is comparable to carrying out 10-20 RSA decryptions.

We have also discovered efficient techniques that limit the number of times James Bond can show his credential, or limit the number of times he can show it within a certain time period. Should he exceed this allotted limit, he gets caught. This strikes a balance between privacy and accountability: once a user breaks the rules, he can be identified, but if he behaves properly, his privacy is protected.

Some of this technology has been implemented by IBM, and IBM recently announced an open-source release of this software. The idea is to make this technology the default for how authentication is done online.

Many authentication transactions carried out today require us to disclose more information than is strictly needed just for verification purposes. Modern cryptography provides us with techniques that solve the verification problem without leaking unnecessary personal information. These techniques are fast, secure, and preserve privacy. Let's use them!

Anna Lysyanskaya was recently honored by *Technology Review* as one of the top 35 innovators under 35 (see page 15 for full story). In this article, she describes the research recognized by this award.

Michael Black

In January 2007, Michael made a TV appearance in Discovery Channel's "2057," a special focused on the medical and technological advances we can expect in 50 years.

In April, Michael and his graduate class, Topics in Computer Vision, received commendations from the Henrico County Division of Police in Virginia for their work on analyzing and enhancing surveillance video of a homicide crime scene.

Also in April, Michael attended the IEEE Conference on Neural Engineering in Hawaii where his group reported on the first brain interface for point and click cursor control by paralyzed humans.

In May, two of Michael's Ph.D. students, Stefan Roth and Frank Wood, graduated. Stefan's thesis won the Joukowsky Family Foundation Outstanding Dissertation Award, and he took a faculty position at Technische Universität Darmstadt in Germany. Frank is doing postdoctoral research at the Gatsby Computational Neuroscience Unit of University College London, one of the best groups for machine learning and computational neuroscience in the world.

In May, Michael also gave an invited talk on high-order Markov random fields at the fourth Canadian Conference on Computer and Robot Vision in Montreal.

Michael Black and his graduate class receive commendations from the Henrico County Division of Police in Virginia for their work on analyzing and enhancing surveillance video of a homicide crime scene.

Michael Black (third from left) serves as the keynote speaker at the 6th International Conference on Energy Minimization Methods in Computer Vision and Pattern Recognition, held in Ezhou, China.

June took Michael and his vision group to Minneapolis for the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) where his Ph.D. student, Alexandru Balan, presented a new method for human shape and pose estimation using a graphics-quality body model. Michael and Ph.D. student Leonid Sigal also ran a workshop on the analysis of human motion.

In August, Michael gave a keynote at the 6th International Conference on Energy Minimization Methods in Computer Vision and Pattern Recognition, held in Ezhou, China. This was the first international meeting in Ezhou which, at 300,000 people, is a small city by Chinese standards. Consequently, all the heads of local government showed up to open the meeting, give speeches, and be filmed for the evening news. Ezhou is near Wuhan which is considered one of the five "furnaces" of China—not a good place to go in August!

In September, Michael gave The Gerard Salton Lecture in the department of computer science at Cornell University. Other invited talks included one at the Hong Kong University of Science and Technology and two at the University of British Columbia.

Michael and his collaborators, John Donoghue (Neuroscience) and Arto Nurmikko (Engineering), received a $1.5M award from the Office of Naval Research to study Neurotechnology. Michael also received a gift from Intel Corporation to support his research on optical flow estimation.

Maurice Herlihy

Maurice Herlihy received a research grant from Microsoft to support work on a compiler front-end for the SXM software transactional memory system. He gave talks at Stanford (May 2007), Intel (May 2007), the Technical University of Lisbon (September 2007) and a keynote address at the IEEE International Symposium on Workload Characterization (September 2007).

John Hughes

Spike spent most of the summer working on the new edition of *Computer Graphics: Principles and Practice*. If any former students would like to review/critique chapters, they should contact him. The one break from that was attending SIGGRAPH in San Diego, where he, his former Ph.D. student Lee Markosian, and two co-authors had a paper on line drawing through abstraction.

Sorin Istrail

Last May, at Leiden University in the Netherlands, Sorin gave an Anton van Leeuwenhoek Lecture on BioScience entitled, "The Regulatory Genome and the Computer"; the lecture is named after the "father" of microscopy, Anton van Leeuwenhoek (1632-1723). A paper of the same title, co-authored with Eric Davidson of CalTech, appeared in the October 2007 issue of *Developmental Biology*. Sorin's research in this area recently received a NSF grant, "The cisGRN Browser and Database: cis-Regulatory Information Behind the Network," in the expected amount of $850,000, which will be used to develop algorithms and software tools for building genomic maps of the developmental regulatory circuitry of cells.

Sorin's second main research area is the human genome and the genetic basis of disease, and he has been developing a graduate course with this title for the past two years. Together with his postdoctoral fellow Fumei Lam, he is focusing on the computational foundations of genome-wide disease associations such as ancestral recombination graph reconstruction and maximum discrepancy on the genome for case-control associations. He is also working with Brown Medical School faculty on disease-focused projects: Gideon Koren (sudden cardiac death), Jim Padbury (preterm labor), and Barry Lester, (addiction).

A large portion of Sorin's time is devoted to his responsibilities as Director of the Center for Computational Molecular Biology (CCMB) http://www.brown.edu/Research/CCMB/. Renovations have been taking place in the center's new home on the second floor of the CIT, and at the opening ceremony on September 6, Sorin gave an overview of the center and junior faculty described their research to Provost David Kertzer, Associate Provost Pam O'Neil, and Clyde Briant, vice president for Research. Five of the junior faculty were awarded the 2007 CCMB Scholarship Innovator Award in recognition of their important contributions in research, teaching, and service.

Sorin is also the primary investigator of the NSF grant "Sweatbox Q&A Boot Camp at Brown University: Asking Tough Scientific Questions." In preparation for the second CCMB Symposium, tentatively scheduled for spring 2008, this boot camp teaches students selected by an international competition how to ask tough questions of the symposium distinguished speakers. These speakers are among the world's most prestigious scientists in genomics and biotechnology.

More information about the second CCMB Symposium, "The Genome and the Computational Sciences: The Next Paradigms—When von Neumann met Crick and Watson," as well as video footage and information about the first symposium lectures and sweatbox sessions held in December 2006, is available on the CCMB website.

Chad Jenkins

Robots controlled by Nintendo Wiimote were demoed at events such as AAAI, RoboCup, and RoboBusiness. The Brown Robotics group let kids of all ages play with their Sony AIBO robot dogs and iRobot Create-based platforms via Wiimote. Wiimotes provide an easy to understand interface—this is great for pickup games of robot soccer in a variety of locales.

Philip Klein

Philip received the 2007 Philip J. Bray Award for Excellence in Teaching in the Physical Sciences. Philip is on leave from Brown for academic year 2007-2008 to work at a start-up in the Boston area. He is also a research affiliate at MIT.

Shriram Krishnamurthi

Shriram loves being back in the classroom! He took a break from that to ensure that nothing was, indeed, rotten in the state of Denmark, and got drenched by rain for his troubles. He reports with some alarm that nary a cloud was in the shape of a camel, a weasel, or a whale.

David Laidlaw

Not that David can hope to compete with Michael Black, but he did some traveling this summer. He and his family took an 18-day road trip covering some 2,000 miles and featuring water sports in two countries and four states. A week at Block Island included a first for David—a sand structure that included arches (see photo). Oh, and he also visited Germany for a research-related Dagstuhl seminar on scientific visualization.

In exciting news, this year's NSF/Science "International Science & Engineering Visualization Challenge" winner was an entry by visualization group member Misha Kostandov together with collaborators David Willis, Kenny Breuer, Dan Riskin, Jaime Peraire, David, and Sharon Swartz. The title was "Modeling the Flight of a Bat." More can be seen at the NSF and Science Magazine websites and on page 16.

Work titled "Drawing on Air" by David's student-turned-post-doc Dan Keefe together with Bob Zeleznik was profiled by http://physorg.com. Subsequently, it was slashdot'ed, which was a first for David. Dan and David opined that more e-mail comments came in based on the slashdot piece than from pretty much any other past press coverage.

David is teaching CSCI2370 (nee cs237) "Interdisciplinary Scientific Visualization" again this semester. The ten students are learning about visualization by doing: they are meeting and interviewing scientists around campus, developing and writing "grant" proposals, peer reviewing them, executing the proposed research, and presenting it at a mock conference in December, and all in thirteen weeks.

Anna Lysyanskaya

This semester Anna is teaching CS51, "Models of Computation." According to her, "this is beautiful material; I am especially excited about teaching reductions—they also come up in cryptography to prove security, and are, of course, at the heart of pretty much any CS task: to solve a problem, it helps if you can reduce it to a problem you've already solved. I hope I am up to the task of educating our concentrators about all this."

"Another commitment occupying a lot of my time right now is the fact that Philip, Claire and I are local arrangement chairs for FOCS 2007. FOCS is the main conference in theoretical computer science, and we are really excited that it is coming to Providence this October. It will be preceded by a day of tutorials taking place at Brown."

"As usual, I've been travelling a lot. I just came back from the Emerging Technologies Conference at MIT (I was on a panel discussing security and privacy), and tomorrow I am headed to Smith College (my Alma Mater) to celebrate Smith's launch of a new center for women's math education. And next week I am flying to the west coast for an invited lecture at UC Berkeley. I also have a major conference submission deadline coming up in three days. I should stop writing this and get back to working on my submission," she concludes.

Claire Mathieu

This summer, Claire was on the road for many weeks, trekking in Turkey, attending workshops in Dagstuhl, Germany, and at Koc University in Turkey. She also went hiking in the Pyrenees Mountains.

Highlights for Claire included daily swimming in mountain lakes and learning to play Tarot. This fall, she is away from Brown for a semester-long visit to the Theory Group at Microsoft Research in Redmond, Washington, where she plans to work hard during the week and take long hikes on the weekends.

Barbara Meier

In addition to teaching animation, Barb continues to work slowly on her own animated film about a bird who helps improve the neighborhood birdhouses. "Perhaps it is our own remodeling of our 1904 Victorian farmhouse, number three in a series, that has inspired me. I was also reminded that animation is not the only slow process when our family took a road trip this summer from Providence to southern Ontario, Niagara Falls, Lake Erie, Columbus, Ohio, Baltimore, Maryland, and finally back home. We are so used to instant virtual visits around the world via the internet that we had to relearn the fine art of keeping ourselves and our boys (7 and 11) entertained for hours within the confines of the car. Listening to Harry Potter on CD was a good first step. Staying in a water park hotel along the way made some of the getting there all that much sweeter," Barb explains.

Franco Preparata

Franco spent most of the summer of 2007 in Asia. He first participated in the Expert Panel on Mathematics and Computer Science of the Ministry of Education of the Republic of Singapore. Next, he collaborated with his research group at the National University of Singapore, working on the efficient detection of genomic repeats. Finally, in Lanzhou, Ganssu, China, he chaired an international workshop entitled "Frontiers in Algorithmics," which he promoted and helped establish and where he delivered the opening address. He concluded this active summer with an equally active trip to Lhasa, Tibet, reached after a 30-hour train ride through the permafrost of the Tibetan plateau.

John Savage

John Savage visited Vietnam during the month of August on behalf of the Vietnam Education Foundation, an entity created by Congress to award fellowships to Vietnamese to pursue graduate education in the U.S. During the trip, which involved interviewing candidates for these fellowships, the American interviewers gave talks at Vietnamese academic institutions. The accompanying photo shows John standing before the Post and Telecommunications Institute of Technology in Hanoi. (They apparently were very happy to have him visit.)

John is currently teaching and doing research in nanotechnology. He has started research in a new area for him, coded computation, which is computation done on encoded data. Coded computation is expected to become important as semiconductor devices shrink and error rates increase. In September he received a new NSF grant for research in this area.

Ben Raphael

This year, Ben gave invited talks at the GEM4 Conference on Cancer in Singapore, the Bertinoro Computational Biology Meeting in Italy, the McGill Workshop on Bioinformatics in Barbados, Harvard Medical School, and the University of Virginia. In the spring, he attended the Research in Computational Biology (RECOMB) Conference in Oakland where he and Erik Corona, an undergraduate at UCSD, presented their paper "Identification of Deletion Polymorphisms from Haplotypes," co-authored with Eleazar Eskin at UCLA. Ben also organized the first RECOMB Satellite Workshop on Computational Cancer Biology in San Diego, which brought together computer scientists, mathematicians, and biologists to discuss computational problems arising in cancer biology. Needless to say, it was an incredible year of travel! On the research front, Ben received a grant from the Department of Defense Breast Cancer Research Program with Susan Gerbi and Alex Brodsky in the MCB Department to study hormone-regulated gene amplification. Ben continues to work with Ph.D. student Anna Ritz on cellular signaling networks.

Finally, Ben is teaching his graduate seminar, "Genomes, Networks, and Cancer" this fall, with each student in the class undertaking an original research project in computational biology.

Meinolf Sellmann

Last year was busy for Meinolf: In addition to his more common editorial work—he serves as associate editor for the *Informs Journal on Computing*, editor for the *Constraint Programming Letters*, and worked for the program committees of the CPAIOR'07, AAAI'07, and CP'07 conferences as well as the SymCon'07 and LSCS'07 workshops—he also collaborated with Bernard Gendron and Gilles Pesant from the University of Montreal to edit a special issue for the *Journal on Computers and Operations Research*. The CS department also volunteered him as the interim editor for the spring issue of Conduit. Moreover, he organized sessions on the integration of constraint and mathematical programming at both Informs conferences in Puerto Rico and Seattle this year. Meinolf also supervised the honors thesis of Daniel Leventhal who graduated this summer, and was invited to give a talk at IBM on a novel technique for learning better heuristics in restarted complete search methods.

He is happy to report that the long-term project of his to found the first Constraint Programming Society in North America was finally realized (http://www.cs.brown.edu/people/sello/CPNA). The purpose of the new organization is to foster CP on the North American continent by restoring the original broad and inclusive vision of constraint programming as an area of research that actively incorporates methods from related fields such as algorithm theory, mathematical programming, and artificial intelligence.

By far the most time-consuming task was Meinolf's role as conference chair of CP 2007 (in collaboration with Laurent Michel from the University of Connecticut – http://www.cp2007.org) and as local chair of ICAPS 2007. Both conferences were co-located at Brown this year, and attendees had a very good time and left with a nice image of Brown. Many thanks to the department and to the Brown administration for supporting this effort so very well!

Meinolf is looking forward to leaving all the organizational overhead behind and spending time doing research with his students.

Don Stanford

In the past year Don has been very busy with other activities in addition to teaching CSCI0020 (formerly CS2). He had an opportunity to lecture in the engineering division's PRIME program (Program in Innovation, Management and Entrepreneurship) and will continue doing so in future semesters. Don also completed five years as the president of the Board of Times2 Academy, a Providence public charter school that inspires 645 students in K-12. He is proud to say that a recent graduate of Times2 recently entered Brown as a freshman biomedical engineering concentrator. Don is also serving on the board of the Business Innovation Factory (http://www.businessinnovationfactory.com) and the Science and Technology Advisory Committee that was formed by Governor Carcieri two years ago. He continues to travel to points of interest around the world and to spend quality time at his home in the Virgin Islands.

Debbie, Andy and Frank Tompa (Andy is wearing the "Waterloo pink tie" and his jester's cap made by daughter Elisa)

Andy van Dam

This past spring, Andy van Dam and his team of pen-computing researchers hosted the 2007 Invited Workshop on Pen-Centric Computing on the Brown campus with support from Microsoft and the Intelligence Advanced Research Projects Agency (formerly ODNI's disruptive technology office). The three-day workshop proved to be a real "hands-on," shared learning experience with 32 respected researchers from academe and industry demonstrating their cutting-edge work in progress—flaws and all.

In May, Andy was invited by Dr. Harry Shum, Director of Microsoft Research Asia, to give a lecture on pen-computing at his lab, after which he and his wife Debbie took a nearly two-week, fascinating tour in China with former student David Salesin (Adobe Research and University of Washington) and David's family. David's wife Ondi was a most excellent "cruise director," interpreter, and guide to the sights as well as the wonderful local cuisine.

That trip was followed almost immediately by a visit to the University of Waterloo (in Canada) where Andy received an honorary degree, hosted by former student and former chairman of the top-ranked David R. Cheriton School of Computer Science, Professor Frank Tompa.

Andy continues to be addicted to outdoor activities. In early June, he and some friends, including former student and godson Evan Schrier, did a supported mountain-biking trip on the North Rim of the Grand Canyon, and he and Debbie did a ten-day hiking (alas mostly rained out) and gastronomic (it was great!) tour in France (Chamonix, Annecy) and Switzerland (Montreux, Saas-Fee). In early September, he did his first-ever mini-triathlon, the Hyannis Sprint2, which is a 1/4 mile swim, a 10 mile bike ride, and a 3.5 mile run (well, in his case, a geriatric jog).

Enrollments were slightly up in his fall courses: CS15 (Object-Oriented Programming in Java, with Graphics) with about a hundred students and CS123 (Intro to Graphics) with about 40 students.

Pascal Van Hentenryck

Pascal had a busy summer, lecturing at the INFORMS meeting in Puerto Rico, giving a tutorial and three talks at AAAI in Vancouver, teaching constraint programming at the ECAI summer school in Leuven, and giving an invited talk at MISTA (the multi-disciplinary scheduling conference) in Paris.

His book, *Online Stochastic Combinatorial Optimization*, was reprinted by Prentice-Hall in India. He also graduated two students, Ionut Aron and Yannis Vergados last academic year. He is now recovering. Seeing Tom Brady throw the ball to Randy Moss every Sunday has helped.

On Discourse

continued from page 6

work another undergraduate and I did about eight years ago on finding "parts" of objects. In this work "pilot" would be found as "part" of airplane because it is often found in constructions like "the pilot of the airplane" or "the airplane's pilot."

Moving away from nouns, verbs often suggest entities that are about to be mentioned ("convinced" followed by "the agreement"), and vice versa ("a crash" followed by "died"). More generally, we want to capture inferences that lead from one sentence to the next and thus, in our simple-minded way, sneak up on real "understanding."

OK, now, what use could any of this possibly have? Well, suppose I jumble together the sentences of not just one article, but two. Can you separate them (as well as order the sentences)? This is much harder, but the better we do at the first task, presumably the better we will do at this new one. Now suppose I ask the program to create a single article using two-thirds of the total sentences. This would be interesting to the degree that the final result left out the relatively unimportant stuff. Perhaps when we combine the two articles, the mode of reference should change. That is, in article one we need the sentence "Joe Smith, 46, resident of Seattle ..." but we want to put this into a context where Joe Smith has already been mentioned. We don't know how to do this yet, but we do have a program that can tell us which of several versions is the best, and that is (maybe) halfway there. Obviously any strides made towards programs that can synthesize information are well worth taking.

I like work that has applications because then my students get job offers. I like doing things that others find useful, and I believe in capitalism as a great driving force for research. For me, however, the ideas are most important. I want to understand how language carries meaning. At the moment, I have found that this project is the best path leading in that direction.

Interactive Computation

Springer Verlag Oct 2006

In October 2006, Springer published a book co-edited by Peter Wegner, Dina Goldin and Scott Smolka, titled *Interactive Computation: The New Paradigm*. The book's 18 chapters explored interaction as an emerging paradigm of computer programming and problem solving. Its authors include Robin Milner, Manfred Broy, Moshe Vardi, Peter Denning, and other well-known contributors to interactive research.

Computing has evolved substantially since Turing showed in 1936 that computers could not solve the *Entscheidungsproblem* (decision problem) because the halting problem was unsolvable. Turing's 1936 paper focused on the negative idea that computers were too weak to prove all expressible theorems, supporting Godel's contradiction of Hilbert's hypothesis that all expressible theorems were provable.

The Church-Turing Thesis (CTT), formulated when Turing visited Church in Princeton in 1937-38, asserted that that Turing machines and Lambda calculus both compute the same set of mathematical functions, which corresponds to functions computable with an effective algorithm.

However, this narrow thesis about limited algorithmic computability was extended in the 1960s to the broader "strong CTT" —that a Turing machine can solve all problems expressible as computations for any computer, whether or not they represent a function. We show that the strong CTT is false while the original CTT is true—though narrow in the range of its applicability.

It is surprising that the strong CTT was adopted as a central model of computing when it was so clearly beyond the original CTT. We suggest that this was due to confusion over the notion of "algorithms," which has been widened from its original strictly functional meaning (as specified by Turing and Knuth) to the wider interactive meaning often used today.

While mathematics and algorithms are useful models of computation, they are not complete models, and many applications must go beyond algorithms and mathematics in their specification of computation. The focus on mathematics as a complete description of the world, originally proposed by the Greeks, was central to Newtonian physics until relativity theory and quantum mechanics broadened scientific models to include nonmathematical questions. Computing, like physics, requires a basic model broader than mathematics. In fact there is no complete model for physics or mathematics in spite of numerous attempts during the past 300 years to develop a complete model, and it seems unlikely that we can develop a complete model of computing any time soon. Our model of interactive computing is a useful additional model beyond that of Turing machines, but should not be viewed as a complete model.

In our article we explore philosophical notions of "truth" in both scientific and political disciplines, and suggest that philosophers like Deskartes and Kant specified inappropriate models of truth that have been misused by politicians and scientists. In particular, the rationalist assumption that our mental beliefs are "true" can lead to political as well as scientific mistakes based on inappropriate desires, while empiricist testing of our ideas generally provides a better understanding of the truth or falsity of ideas. Hilbert's rationalist notion that all mathematical proofs can be established by logic was empirically disproved by Godel, Turing, and Church, while a rationalist basis for the Iraq war has been empirically disproved. We believe that much work is needed in both science and politics to develop better principles of truth as a foundation for making the world a better place.

Dina, Scott, and Peter enjoyed working together on this book, to improve their understanding of interactive computing and to invite other researchers to write articles about their work on interaction.

Peter was especially pleased that his co-editors were former Brown Ph.D.s whose substantive research after they had completed their degrees included models of interactive computing.

Peter Wegner & Dina Goldin

This year's faculty-grad retreat, held on May 22nd at Brown's Haffenreffer Museum of Anthropology in Bristol, was planned by students Eric Rachlin and Aparna Das and a-staffer Janet Eager.

2007 marked the second year in which the departmental retreat replaced the annual faculty retreat. Talks were given by Alexandru Balan, Melissa Chase, Amy Greenwald, Spike Hughes, and Sorin Istrail.

Alex discussed his work with Michael Black on people tracking. Using multiple video cameras, they are able to determine the pose and 3-D shape of a person walking around a room.

Melissa discussed an attribute-based encryption scheme, in which each user is identified by a set of attributes, and some function of those attributes is used to determine decryption ability for each ciphertext.

Amy's talk was about no-regret learning. She and her group are studying a general class of online learning algorithms, "regret-matching algorithms," along with a "regret-based framework" for analyzing their performance in ODPs. A regret-based framework is characterized by a set of transformations of the agent's history of actions. Regret is then defined as the difference between the rewards obtained by the agent and the rewards the agent would have obtained by playing according to a transformation of its history of actions.

Spike showcased algorithms for creating sparse computer-generated line drawings, motivated in part by a sparse line drawing by Matisse. Given a 3-D model of an object, Spike explained how to generate a nice-looking line drawing of the object.

Sorin's talk pointed out the interesting similarities between computational gates (logic gates) and biological gates (genes, proteins). He also showed the group (shocking) pictures of himself working in a biology lab with test tubes in his hands.

Unlike last year, the caterer did not burn down the day before the event, which was a definite plus.

The weather was fantastic, so students and faculty discussed the department while enjoying the sun and fabulous view of the bay. After the discussion, people played soccer, frisbee, golf and volleyball.

Anna Lysyanskaya Honored as Top Young Innovator by *Technology Review*

Anna Lysyanskaya has been included in the 2007 Technology Review TR35, an honor given each year to 35 innovators in science and technology under the age of 35 whose inventions and research the magazine finds most exciting. Anna's extraordinary work in cryptography, identity verification and privacy protection led to her inclusion in this elite group of accomplished young innovators.

As evidenced daily in the media, reports of compromised corporate databases and online invasions of privacy are a frequent occurrence. To battle this problem, Anna has developed ground-breaking methods that allow authentication without identification. These ideas are being incorporated into practical software systems. In particular, her work on anonymous credentials in collaboration with Jan Camenisch (IBM Research) is now incorporated into the Idemix software, which will be available through the Eclipse Open Source Foundation's Project Higgins.

"I think it is a big honor to be included on this list," said Anna. "My research builds on a lot of fundamental prior work on digital signature schemes and zero-knowledge proofs, which, luckily for me, was done by people who are currently over 35, so I am the only one to be honored now. The main thrust of my research is methods that would allow a user to access a resource without revealing any information other than that she is, in fact, an authorized user."

"There is still work to be done. I have gotten my Ph.D. students Mira Belenkiy and Melissa Chase involved in this field and they're contributing exciting new results. Anyone who wants to learn more about this topic should take my 'Advanced Topics in Cryptography' course in the spring," concluded Lysyanskaya.

"Loss of confidence in online privacy and security poses a growing threat to internet commerce and information exchange," said Chair Roberto Tamassia. "Anna's creative approach to identity verification helps restore an expectation of privacy to our online lives. The department is thrilled that Anna's talent and innovation have been recognized by such a prestigious magazine and that more people will learn about her outstanding contributions to the field of internet security."

Technology Review, the world's oldest technology magazine (est. 1899), aims to promote the understanding of emerging technologies and to analyze their commercial, social and political impacts.

White House Honors Chad Jenkins

Chad Jenkins was honored at the White House as one of the nation's top young scientists. Jenkins was selected as one of the recipients of the Presidential Early Career Award for Scientists and Engineers (PECASE) for his research on the development of methods for autonomous robot control and perception. His work advances the idea that robot control and computational perception are better learned from human demonstration rather than by explicit computer programming. The PECASE program recognizes outstanding scientists and engineers who, early in their careers, show exceptional potential for leadership at the frontiers of knowledge. This Presidential Award is the highest honor bestowed by the United States Government upon scientists and engineers beginning their independent careers.

Cetintemel, Jannotti, and Taubin Receive Joint NSF Grant for Visual Sensor Networks

The National Science Foundation has awarded funding in the expected amount of $700,000 for research at Brown in the use of cooperative "smart cameras" to sense and act in three dimensions. These camera networks will be capable of tasks such as reconstructing three-dimensional features, producing images from novel viewpoints, and tracking moving objects from camera to camera.

These tasks are more challenging than the simpler data collection tasks performed by existing sensor networks that operate simpler, scalar data. The planned imaging applications require high data rates, remote sensing, and precise calibration. To address these challenges, the project will focus on real-time, in-network feature extraction, processing, and storage to mitigate the need for massive, centralized collection and processing. The team hopes to develop technology at all levels of the sensor network platform. They envision advances from the lowest level to the highest— from power-efficient image processing algorithms that operate in-camera, to scalable data-centric routing protocols that enable cooperation without network congestion, to powerful data abstraction techniques that will allow sensor network applications work above the details of a particular sensor network and operate on a simpler model of the world.

Professors Ugur Cetintemel, John Jannotti, and Gabriel Taubin bring a range of expertise to the three-year project. Taubin, who holds a joint appointment with the Engineering and Computer Science departments, has previously designed smart-camera prototypes of this kind, capable of power-efficient real-time image processing. Jannotti's work in decentralized ad-hoc networks will be used to efficiently route image and feature data among cooperating cameras without central collection. Cetintemel brings experience in the area of data-management, critical to building abstractions that will allow applications to operate conveniently on the large distributed datasets sensed by these networks.

Andries van Dam receives Doctor of Mathematics degree

On June 15, 2007, Andries van Dam received the degree Doctor of Mathematics, honoris causa from the University of Waterloo in Waterloo, Ontario.

David J. Willis and Mykhaylo Kostandov of Brown Receive First Place in the Informational Graphics Category of the NSF 2007 International Science & Engineering Visualization Challenge for their "Modeling the Flight of a Bat" Research

The team of Kenneth S. Breuer, David J. Willis, Mykhaylo Kostandov, Daniel K. Riskin, Jaime Peraire, David H. Laidlaw, and Sharon M. Swartz took first place in the informational graphics category of the challenge where illustrators, photographers, computer programmers, and graphics specialists from around the world were invited to submit visualizations that would intrigue, explain and educate. More than 200 entries were received from 23 countries, representing every continent except the Arctic and Antarctica.

"Breakthroughs in science and engineering are often portrayed in movies and literature as 'ah-ha!' moments. What these artists and communicators have given us are similar experiences, showing us how bats fly or how nicotine becomes physically addictive," said Jeff Nesbit, director of NSF's Office of Legislative and Public Affairs. "We look at their visualizations, and we understand."

A description of the team's visualization follows.—Ed.

Most short-nosed fruit bats (Cynopterus brachyotis) spend their nights flitting about in the jungles of southeast Asia. However, some of the tiny creatures, which weigh less than 50 grams fully grown, lead an altogether different existence: flitting about in wind tunnels under the watchful eyes of aerodynamics researchers. Interested in the tiny mammals' flight dynamics, Brown University engineer Kenneth Breuer used lasers and a sophisticated multicamera motion-tracking system to record how their wings and the air around them distorted as the animals flapped against the wind. Based on the experiments, aeronautical engineer David Willis, who has a joint appointment at Brown and MIT, Brown computer scientist Mykhaylo Kostandov, and their colleagues created a computer model of bat flight—visually conveyed in their poster. "When viewed in slow motion," says Willis, "bat flight is beautiful and complex. The goal of this illustration is to capture that beauty while also adding scientific merit."

Sorin Istrail Receives NSF Grant to Build Database and Next-Generation Browser for the Regulatory Genome

Sorin Istrail, Julie Nguyen Brown Professor of Computational and Mathematical Sciences and Professor of Computer Science, received a NSF grant, titled "The cisGRN Browser and Database: cis-Regulatory Information Behind the Network." Funding in the expected amount of $850,000 will be used by Istrail and his students—Ryan Tarpine and a few others to be added to the team—to develop algorithms and software tools for building genomic maps of the regulatory circuitry of cells. The work will focus on solutions to three major problems in the developmental biology community: building genomics libraries of software tools to reveal the logical principles and gene-network architecture of the "genomic computer," building a lexicon of experimentally validated cis-regulatory modules to improve predictive algorithms for cis-regulatory modules and networks, and designing a next-generation regulatory genomics logic map browser. This work builds on Istrail's extensive collaboration with CalTech's Eric Davidson towards a unified concept of "regulatory information" and the unveiling of the principles of information processing of cell regulation. This collaboration has already been instrumental in: (1) constructing the first sea-urchin genome assembly, (2) identifying a first repertoire of logic functions of the genomic cis-regulatory code, (3) building the first high-resolution transcriptome map of the sea-urchin embryo, and (4) providing the first comprehensive analysis of the principles of the regulatory genome vis-à-vis those of the computer. The goals of this research also build on Istrail's extensive experience with his group at Celera Genomics, developing software libraries for genome assembly, whole-genome alignment, and automatic annotation of genomes.

John Savage Receives NSF Collaborative Research Grant to Study Coded Computation and Storage at the Nanoscale

John Savage (with Andre' DeHon of UPenn) received a collaborative research grant from the National Science Foundation to study coded computation and storage at the nanoscale.

A key challenge before the semiconductor industry is coping with high error rates resulting from the decreasing size of chip features. Transient faults, along with permanent defects and stochastic assembly, make it difficult to implement traditional architectures. Research has been done on routing around defects and coping with large amounts of device variation. Little is known, however, about how to cope efficiently with high rates of transient errors during computation. This research will take a new systematic approach to the tolerance of transient failures. The goal is to help the semiconductor industry to better understand the dimensions of the nanoscale reliability problem. This research has relevance to space-borne applications where error control can serve as an alternative to radiation hardening.

This research employs a sophisticated approach to fault-tolerant computation. First, it exploits differential reliability; it examines the use of a small number of reliable elements to oversee a large number of unreliable elements. Second, it draws on the success of coding theory to explore both special and general methods to encode inputs and outputs of a potentially faulty computation, paying particular attention to a seminal approach taken by Spielman in 1996. By encoding computations, faults at the encoded outputs can then be detected and corrected. Third, it examines the use of small check computations followed by possible rollback, where most of the checking is done using unreliable elements. Allowing a computation to be repeated in time, rather than space, reduces the overhead of fault free computations. The design work is expected to have immediate impact on practice, whereas development of a general theory is expected to have a longer-term impact.

Roberto Tamassia Receives IEEE Computer Society Technical Achievement Award

In May, Roberto Tamassia Received an IEEE Computer Society Technical Achievement Award.

Report on the 37th IPP Symposium

Tom Doeppner

The title and theme of the 37th IPP symposium was "Operating Systems Technology and Brown Alumni." But a major sub-theme was honoring alumni Bryan Cantrill ('96), Adam Leventhal ('01), and Michael Shapiro ('96, Master's '97), winners of the 2006 Wall Street Journal Technology Innovation Award for their work on DTrace at IPP partner Sun Microsystems. The three were former head TAs for our operating systems course, CS167/169, and they suggested we use the symposium to showcase the achievements of other alums of the course. We organized a program featuring six former head TAs (including Bryan, Adam, and Mike) as speakers. I tried to contact all alums who had been TAs for the course, but this was a challenge—both for my memory and for Google. Although I well remember the first time I taught the course in the spring of 1977, my recollections of who TA'd up until the mid-1980s were a bit fuzzy. However, I think I was able to contact most of my former TAs and was able to send email invitations to the vast majority of them. A fair number were able to attend, resulting in not only an interesting and fruitful day, but also a great reunion.

The first speaker was Bryan Cantrill, who gave an entertaining and fascinating talk on "The Inculcation of Systems Thinking." Quoting from his abstract, "Software, even purely algorithmic software, never exists in a vacuum—it becomes a component in a larger system. Such software systems are themselves rarely self-contained: they exist in larger information systems that are in turn components in yet larger economic, geopolitical and social systems." He went on to discuss the larger notion of systems, providing examples and assessments of failures and pointing out the need for high-quality systems education.

Mike Shapiro spoke next, on "Purpose-Built Languages in Systems Design." Quoting from his abstract, "System calls and libraries aren't the only interfaces a system exports. Purpose-built languages have formed a critical component of system design, particularly in UNIX, for many years." He went on to discuss several case studies of purpose-built languages to illustrate design and implementation issues, attributes of successful languages, and open problems in systems that can be addressed with purpose-built languages. Mike showed us many examples of how good language design has made possible many advances in operating system development over the years, as well as entertaining examples of not-so-good design.

Adam Leventhal completed the morning session, speaking on "Run-Time Code Instrumentation: Maintaining the Illusion of Normalcy." This covered some of the innovative work for which Bryan, Mike, and Adam received their Wall Street Journal award. Quoting from Adam's abstract, "Observing software is a bit of a fiction: while other engineering disciplines can evaluate their constructions using physical measurements, it is only through modifying software that it can be 'seen.' Changing code at run-time is a high-wire act as the executing software needs to continue seemingly unperturbed in spite of dramatic modification." Adam showed how they've managed to do this, with impressive results.

Jeff Korn ('94), who went on to get his Ph.D. at Princeton and is now at Google, spoke next. However, before doing so he and his CS169 co-head-TA Scott Raposa ('94) played for us a CD they produced of the new (first

continued on page 20

Hickey '08 Squashes 'Banner' Bug

Brown's goal in implementing the Banner software was to replace eleven systems with a single, integrated student information and class registration system capable of supporting faculty, student and administrative needs. After much fanfare and more than a year of working to implement Banner, a potentially detrimental attack was discovered by Brendan Hickey '08 just days before the registration period began. —Ed.

Reprinted from the
Brown Daily Herald
-Chaz Firestone

Imagine you have just logged into your Banner account. Moments later, you receive an e-mail that reads "check out this cool video!" followed by an innocent-looking hyperlink. You click the link as your Banner window sits in the background, but the site doesn't seem to load, so you shrug it off and continue with Banner, registering for that last class you had been shopping.

But little do you know that you've just become a victim of theft. Your home address, academic transcript and private financial aid information have been sent to a neighboring room or a different state or country. Your classes have been dropped. And as you stroll out the room with naive confidence, your most precious information sails out, too.

Thanks to Brendan Hickey '08, you can rest easy. Hickey's discovery of this potentially detrimental attack and its subsequent patching—just days before registration period began—means students at Brown and the hundreds of other schools using Banner software needn't worry.

Hickey, a computer security enthusiast, discovered the threat in late August while working on a summer project at Brown. After encountering an error message in Banner, Hickey used his intimate knowledge of Web security to identify a loophole that would allow for a bug called CSRF, or a cross-site request forgery.

"The idea is that a student is using Banner, and you send them a link to another site that can execute actions as if it were the student," Hickey said. "It could force someone to drop classes, add classes, print out their transcripts."

Hickey added that if the CSRF made its way onto a professor's machine, it could potentially be used to alter students' grades.

"I'm not familiar with the professors' user interface, but it's certainly possible," he said.

Upon discovering the potential for a CSRF, Hickey went to the Computing and Information Services Help Desk and showed employees there what he had found.

"They asked me to make a test case," Hickey said, referring to a trial attack that would demonstrate the power of the attack without harming anyone. "I put one together in 15 minutes." Within a matter of days, news of Hickey's "Banner bug" had reached University officials and even a senior vice president of SunGard Higher Education, the company that develops Banner.

"They asked me all sorts of questions about it, and then asked me for suggestions on what to do next," Hickey said. "I gave them."

Vice President for Computing and Information Services Michael Pickett, Brown's chief information officer, said Hickey's work is greatly appreciated by University officials and the developers of Banner.

"Brendan made a nice contribution to the strength of the software," Pickett said. "The SunGard folks are real pleased."

The problem was patched up the weekend before registration, Pickett said, and was responsible for some of the downtime in Mocha, the student-run online course catalog.

Pickett said Hickey made the right choice in bringing the security flaw to the attention of CIS officials and was not in violation of any of Brown's codes of conduct.

"It was not hacking. He violated no rules," Pickett said. "In fact, I have a dinner meeting to congratulate him."

Shriram Krishnamurthi, associate professor of computer science and

continued on page 20

continued from page 18

and probably last) CS169 theme song, an amazing piece that entertained us all. [Visit the CS home page to hear it.] After the musical interlude, Jeff spoke on "Infrastructure for Personalized Search," discussing work that he and his group have done at IPP partner Google. He presented Bigtable, a system for storing and managing very large amounts of structured data. Data is organized into tables with rows and columns, but unlike a traditional database system, the row/column space can be sparse. The system is designed to manage several petabytes of data distributed across thousands of machines, with high update and read request rates coming from thousands of simultaneous clients.

Next that afternoon we heard from Jason Lango ('98), now of Iron Port Systems, recently acquired by IPP partner Cisco. Jason, who until recently worked for IPP partner Network Appliances, spoke on "Appliances: High-Performance, Low-Maintenance Systems to Address Real-World Problems." Focusing on the "C10K problem" (handling tens of thousands of clients), he talked about the benefits of using specialized systems running custom operating systems on commodity components rather than using off-the-shelf operating systems. He discussed programming techniques designed to accommodate both large scale concurrency and rapid application development.

The final talk, on "ZFS: Developing a Next-Generation File System" was given by Matt Ahrens ('01) and Eric Schrock ('03), both of Sun. ZFS (Zettabyte File System) handles many problems common to contemporary systems that aren't addressed well by its predecessors and has been adopted not only by Sun, but by Apple as well. They discussed the evolution of ZFS, including how and why decisions were made, underlying design principles, and the challenges faced when developing a file system within the constraints of a complex operating system.

continued from page 19

Hickey's adviser, said he was proud of Hickey's creative thinking and professional attitude in the discovery of the Banner bug.

"He was a total 110 percent pro about the whole thing," Krishnamurthi said. "He could have done nasty or foolish things, but instead he did responsible things."

Krishnamurthi explained how Hickey's interests matched the situation perfectly—Hickey was one of only a handful of people at Brown even capable of discovering the flaw.

"He's studying exactly the security of these sorts of things, so he's in a position to look at these things and understand them immediately," Krishnamurthi said. "There are people who take this world as it is, and there are people who probe every system they encounter. Brendan is most assuredly the latter category."

Krishnamurthi also emphasized the severity of the problem Hickey discovered. While obtaining an address or phone number might not seem particularly threatening, a little information can go a long way and have severe consequences, such as identity theft.

And it doesn't stop there, Krishnamurthi said.

"If someone attacked my machine, they could change the grades of my students," he said. "And this could affect anybody. In principle, the registrar could get hit with this attack, and presumably he has some interesting powers."

Krishnamurthi and Pickett said the turn of events should bode well for Hickey, who plans to apply to graduate school.

"I wouldn't mind writing him a recommendation based on his work," Pickett said.

In the meantime, Hickey has already reaped the benefits of the affair. A self-described "security nut," he couldn't get the details of the Banner bug out of his head, and it visited him in the middle of the night.

"The problems I discovered in Banner are related to some of the problems I've been working on all summer," Hickey said. "I was sitting in bed at 2:00 a.m. and discovered a solution to the problem I'd been working on all summer."

Hickey e-mailed Krishnamurthi at 2:53 a.m. on Aug. 25 with the idea for his senior thesis.

"I guess fortune favors the prepared mind," Hickey said.

Theoretical Computer Scientists Come to Brown

Philip Klein, Anna Lysyanskaya & Claire Mathieu

The Symposium on Foundations of Computer Science (FOCS) is the premier and oldest venue for learning about the latest research results in theoretical computer science. This year was the symposium's 48th edition and took place in Providence, organized by members of the Computer Science Department at Brown. This was the second time FOCS was held in Providence: the previous time was in 1977, exactly thirty years ago, and had been organized by Bob Sedgewick.

Tutorial speakers. From left to right: Dan Boneh, Terence Tao, Dan Spielman.

The conference was preceded by a day of tutorial talks at Brown University, attended by more than two hundred people. Paraphrasing Luca Trevisan's blog (in-theory.blogspot.com): Terence Tao (UCLA) gave an inspiring talk about structure and randomness in combinatorics, discussing the duality between pseudorandomness and efficiency, which is a way to give a unified view of techniques coming from analysis, combinatorics and ergodic theory.

Dan Boneh (Stanford) spoke on pairing-based cryptography, an idea that has grown into a whole rich area, with specialized conferences and, according to Google Scholar, more than 1,200 papers published so far. Dan Spielman (Yale) spoke on spectral graph theory, focusing on results and problems that aren't quite studied enough by theoreticians. He showed some remarkable examples of graph drawings obtained by simply plotting a vertex i to the point $(v(i),w(i))$, where v and w are the second largest and third largest eigenvalues of the laplacian of the adjacency matrix (see photo on page 23).

The conference proper had almost 270 attendees and took place in the Renaissance Providence hotel, a recently opened luxury hotel in the old Masonic temple close to the Capitol building. The program consisted of three days, densely packed with talks, each 20-minutes long. One exception was an hour-long talk by Nancy Lynch, winner of the Knuth Prize. She presented a very accessible description of her famous paper on the impossibility of distributed consensus, and gave an overview of current and future work in distributed computing. As Nicole Immorlica, blogger for the conference (weblog.fortnow.com) wrote: "Everyone in the audience was pretty satisfied with her outlined research agenda involving models of distributed computing on mobile networks, until the air traffic controller example. She suggested that her research could replace traffic lights with virtual traffic lights, which made me tense up slightly. Then she suggested we could even replace human air traffic controllers with virtual ones. While we all understand the benefits (e.g., you can have controllers over the ocean, machines don't get tired, etc.), I think we all had a sort of collective gasp." Scary prospects!

continued on page 22

New Faces:

Jane McIlmail

In July, the Department of Computer Science welcomed Jane McIlmail as its new academic department manager. Jane attended Brown University (Class of '81) and joined the staff in 1985. Her most recent position was as the executive assistant to the dean of faculty. She sits on the Agenda Committee for Departmental Managers, and is an alumna of the President's Staff Advisory Committee. Previously at Brown, Jane was manager of finance and administration for the Scholarly Technology Group, Multimedia Lab & Women Writers Project and served as the administrative assistant and departmental computing coordinator for the Department of Visual Art. When she's not running the CS department, Jane belongs to a weekly Scrabble club, enjoys bird watching and attends an annual 24-hour science fiction film festival/marathon in Boston.

"I am delighted to be part of such a creative, accomplished group of scholars, students and staff. The atmosphere in Computer Science is one of true collegiality. The research being done here is truly impressive, and the degree to which our undergraduate and graduate students are involved in that research, and in their own education, is amazing. While my previous position in the administration provided a fascinating look at how things are run at Brown and a chance to support the academic mission at a high administrative level, it is a pleasure to be back in an academic setting. For me, this is where support of higher education is most rewarding," concluded Jane.

Local organizers. From left to right: Philip Klein, Shay Mozes, Aparna Das, Charalampos (Babis) Papamanthou, Glencora Borradaile, Warren Schudy, Claire Mathieu, Arkady Yerukhimovich, Melissa Chase, Alptekin Küpçü, Ayanna Belton. Missing: Mira Belenkiy, Genie de Gouveia, Crystal Kahn, Anna Lysyankaya.

continued from page 21

We were lucky to have summer weather throughout the October conference, and researchers could be seen sitting on the grass under the sun by the Capitol and going together for a jog. Overall, the conference was a huge success.

FOCS greatly benefited from generous financial help from numerous corporate sponsors: Alcatel-Lucent, Akamai, IBM, Yahoo! Research, Google, and Microsoft Research. Led by the local organizers, Philip Klein, Anna Lysyankaya and Claire Mathieu, the Computer Science Department worked hard to organize the conference, with plenty of help from Ayanna Belton and Genie de Gouveia; finally, essential to the success was the help provided by a big team of Ph.D. student volunteers: Mira Belenkiy, Glencora Borradaile, Melissa Chase, Aparna Das, Crystal Kahn, Alptekin Küpçü, Shay Mozes, Charalampos (Babis) Papamanthou and Warren Schudy from Brown University, and Arkady Yerukhimovich from the University of Maryland.

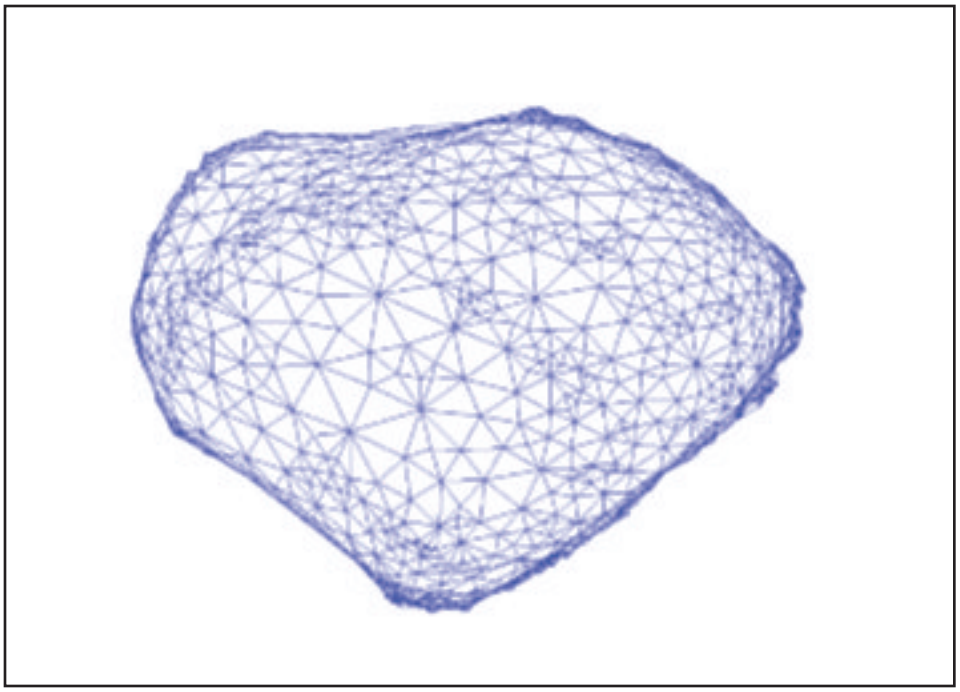

Graph drawing using eigenvectors.

The 2007 Invited Workshop on Pen-Centric Computing

On March 26-28, 2007, a workshop on pen-centric computing hosted by Andy van Dam was held in the CS department; the event was sponsored by Brown University, Microsoft Corporation, and the US Intelligence Community's Disruptive Technology Office. This workshop brought together approximately 30 pen-computing researchers from universities and industrial research labs around the world.

Figure 1: examples show diagram sketching on the left and document editing on the right.

"Pen-centric" computing is a vision of pen computing in which the pen is more than a substitute for the mouse. Pen-centric computing takes advantage of human skill with the pen, particularly the ability to sketch fluidly; while it often focuses on the creation and manipulation of diagrammatic languages, it is appropriate for many tasks in which drawn input is advantageous (see Figure 1).

Today there is increasing demand for software that manipulates complex data—such as images, 3-D geometric models, and video—which may benefit from pen input. Significant improvements in algorithms for handwriting have also created new opportunities. In addition, platforms ranging from electronic whiteboards to tablet PCs to small mobile devices increase the potential application range.

The workshop researchers approached pen-computing from diverse perspectives. Many were engaged primarily in the development of systems that interpret various visual languages, including mathematical expressions (see Figure 2), chemical diagrams (see Figure 3), digital and analog circuits, dynamic systems diagrams, mechanical systems and engineering drawings, as well as concept maps, data structures, probabilistic networks, and graphs.

Others were more engaged in the human side of the interface, investigating new ink application metaphors and low-level interaction techniques.

Within the broad range of perspectives that were presented at the workshop, there was general agreement about the need to identify applications where pen-centric computing could provide a significant advantage over ordinary WIMP GUIs. In addition, participants agreed both that the absence of conventions for low-level interaction techniques is hindering the development of new applications and that interactive correction of the inevitable recognition errors is a common interface design problem.

On the other hand, because of their different perspectives, participants chose different tradeoffs among the many conflicting demands of pen-centric applications. Some of the tradeoffs discussed included:

- Mimicking pen and paper vs. moving beyond pen and paper
- Training vs. "walk up and use" interfaces
- Leaving drawings as rough sketches vs. "beautification"
- Task-specific interfaces vs. reusable components

Recognition techniques engaged many of the participants, who gave details about their approaches to segmentation, recognition, and the determination of 3-D structure from 2-D drawings.

Each presenter's prototype software was not only demonstrated during each talk but was also made available for participant exploration throughout the three-day workshop on tablets obtained with the assistance of Hewlett-Packard. This ability to test-drive other researchers' software, and the ample amount of time provided for informal discussion and investigation of each other's work, was favorably cited by participants as the distinguishing characteristics of the workshop.

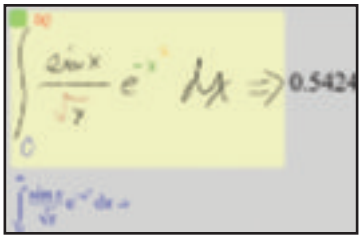

Figure 2

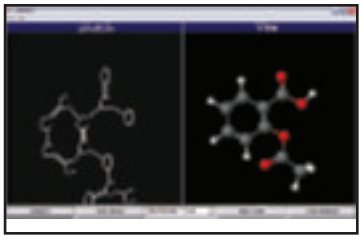

Figure 3

A full report on the workshop is available on the website for the Microsoft Center for Research on Pen-Centric Computing at: http://pen.cs.brown.edu/news.html.

Parenthetically Speaking

The Year of Ignorant Living

Alumni might wonder about the charmed lives faculty lead on sabbatical. To be sure, it's tough to return to civilian life, but not for what might seem to be obvious reasons (I've greatly missed the teaching). Instead, this year has been terrific for me mainly because of what it's meant: a return to a state of ignorance.

For all our talk that research is an activity of constantly confronting ignorance, that's not what we *really* do. Research is more typically a man, a plan, canal panama (women sensibly leave absurd canals out of the picture). We may not know what precise result we're going to get—or even trying to get—but in the big picture we don't flail around very much.

I hadn't *planned* to spend this past year flailing. Now, I regard tenure less as a reward for past activity and more as a recognition of future promise; so the best way to honor it is to do something new, to view the *freedom* to take risks as an *obligation* to do so. Anyway, that's the theory; this runs headlong into (a) having established programs of work in place, (b) not knowing how to achieve ignorance (it's easy to decide to not publish papers or write grants, as we—my wife Kathi and I—did, but harder to decide what to do in its place), and (c) terror.

Proceeding with routine, I spent the summer and early fall working closely with Leo Meyerovich, Greg Cooper, Michael Greenberg, and Alex Bromfield on our new programming language, Flapjax. We finally released it formally in the middle of October 2006 to quite a bit of press coverage. In less than a year the experience of disseminating Flapjax has coughed up several surprises (press coverage for a programming language?—must be slow news days...), some negative in a curious fashion (as a result of which we've come to think of Flapjax not as a *language* but as a *library*), and some pleasantly positive (such as its use at Berkeley). Those would all be subjects for a different article.

We worked overtime on Flapjax last summer in part to have it out before I began my sabbatical travels. Kathi and I had been planning these trips for ages, carefully synchronizing the places we visited to be of mutual interest (since a sabbatical *is* also meant to be a time to recharge personally). Even before we left Providence, however, my carefully laid plans were destroyed by a decision by the Brown administration. In a way, though, it was strangely liberating: if Brown didn't want me to accomplish what I'd set out to do on sabbatical, then I was free to do other things. So I did.

Our first stop was Edinburgh. Kathi was there to visit Keith Stenning, a cognitive scientist she knew from her work on diagrammatic reasoning, while I was there to visit Phil Wadler, one of the designers of Haskell and a pioneer of many programming language concepts. I was, however, also looking forward to talking to the seemingly dozens of other researchers Edinburgh has in programming languages, verification, and other parts of applied logic in which Brown is desperately lacking. When it came to picking an office space, Phil told us that, by coincidence, he and Keith had adjacent offices and the one across the hallway from them was empty; would Kathi and I be willing to share it? It's been a long time since I've had an officemate, but Kathi and I figured we could (just about) survive each others' company, and this way we could reduce our space footprint on their department.

What we *didn't* learn, until our first day in Edinburgh, is that our office neighbors in Edinburgh were Keith, Phil...and nobody else.

Where I'd envisioned a long hallway with logicians in every direction you look, we were in rooms of a small tenement, whose door was locked to the world at large. Nobody was ever going to find us here, nor were we going to find anybody else! (Phil did arrange for me to have another, exclusive, office in the King's Buildings, but distance from home—more than any anti-royalist tendencies—led me to use it only rarely. There I would have been *near* all those logicians, but still in a bit of an odd corner of the world.)

Geography is destiny, they say, and it couldn't be more true here.

Stenning, it transpired, was studying the logical models behind how people reason. His focus, with his collaborator van Lambalgen of Amsterdam, was on the famous Wason experiments in cognitive psychology, which are a kind of card trick that ask the subject to arrive at conclusions and measure how closely they hew to the entailment relation of classical logic; very poorly, it turns out.

This has led some to conclude that logic itself is a poor way to study how people reason. (I hear the hallelujahs from Brown's cognitive scientists.) In contrast, Stenning and van Lambalgen, and others, had revisited the issue with much more detailed studies and found that there were parameterized families of logics that perfectly well explained how the subjects reasoned, and furthermore that environmental characteristics—such as how the prompts were stated—predicted how people set the parameters.

Well! Kathi and I have been spending a lot of effort on the reasoning that goes into access-control security policies; but we've always known that what we're studying is tool support without reference to the underlying cognitive models. It has nagged me for a while now that properly executing this work demands an understanding of these human factors, but I had no idea where to start.And now Stenning had accidentally shown us the world we were looking for. Understanding the consequences of this—and learning how to keep suppressed the repressed memories of my college psychology coursework experiences—has become a new focus since November.

From Edinburgh we went to Oxford and Lausanne for PC meetings, thence to Paris to fly out to India. I've written on my blog at some length about returning home after such a long time. After India came Australia (for a conference, followed by a personal vacation), about which my notes will eventually show up on my blog. For now, even a year later, the memories of that continent are too vivid for words. This was the infamous left-right-left-right period of my life: driving on the right (US), left (UK), right (Continental Europe), left (India), right (US), left (Australia), right (US). This seems to beg for an off-handed wisecrack, but in this department that would be gauche (ha!).

In late January, I attended a Dagstuhl event on Web programming, in which the main thing I learned is confirmation of my opinion that the Semantic Web folks are hopelessly out of touch with reality (perhaps it's a stealth marketing strategy). I was back in Deutschland ten days later at universities in Berlin (see blog!), Tübingen, and Darmstadt, as well as another Dagstuhl, this one on end-user software engineering. Coming as it did after my Damascene conversion to thinking about cognitive models, this was a fantastic opportunity to revel in ignorance and soak up knowledge from the likes of Brad Myers, Mary Shaw, Margaret Burnett, Alan Blackwell, and Stephen Clarke (a UI designer at Microsoft).

In the early spring we visited the programming languages, security, and verification people at Penn, having several enlightening conversations with Insup Lee's group on obligations as a complement to access-control. Readers of faculty notes may recall that we were due to spend the spring at UT Austin; given all this other travel, however, we instead made just two very focused trips to UT. UT recently had the wisdom to hire Brown alum William Cook, who is surely one of the smartest and most tasteful researchers in programming languages; only Will can make even a topic like meta-modeling sound interesting. So a week spent primarily with Will and Don Batory was heavenly.

There were other trips scattered around, but the summer was a good time to consolidate and move forward. Usually I spend much of the school year planning for the summer (and hiring students for that purpose), but this year was obviously exceptional. So it was essentially pure luck that I stumbled upon two outstanding students, Jacob Baskin and Brendan Hickey, who continue in the tradition of Brown undergrads taking me in new directions (not least of all Brendan, thanks to whom I'm talking to vice presidents and lawyers). Combined with two students elsewhere whom I'm co-advising, and my current Ph.D. students—Arjun, who has made strong progress on a very interesting security technique, and Jay, who is feeding me doses of the Coq theorem prover when he's not busy getting married (congrats, Jay!)—it's hard not to realize that sabbatical is over and I'm back.

Now, I regard tenure less as a reward for past activity and more as a recognition of future promise; the best way to honor it is to do something new, to view the *freedom* to take risks as an *obligation* to do so.

continued on page 26

Parenthetically Speaking

continued from page 25

The end of sabbatical doesn't mean I've stopped plumbing the depths of my ignorance. In August, Spike got me excited about graphics for the first time, and I've been programming in Matlab since. Indeed, for the first time in my life I wrote a one-use, throw-away script that actually used trigonometry. This has gotten me interested in research questions related to both images and Matlab. I can only hope that if I lie down for long enough the feeling will pass.

I've also taken the plunge on a few other fronts:

- I've long been skeptical of blogs, which associate a false temporality to thoughts. Largely pushed by Brown alum and Blogger employee Pete Hopkins, I created a blog anyway: Notes from a Sticky Wicket. It will be obvious to readers that I don't "get" the medium, treating it as a repository for essays rather than a dumping ground for thoughts; whether that will change, I don't yet know. I felt obliged to use Blogger, but in retrospect I realize I should have used anything but: that would be the way to test whether Pete was merely trying to drive up Blogger usage or whether he actually cared about what I have to say (my bet, like yours, is not on the latter).

- I finally got my self-published programming languages text into print using lulu.com, which has been impressive. (I actually published the book in three formats: for-pay paper, for-pay PDF *and* free PDF. The beauty of self-publishing is that you can perform any outrageous experiment you want!).

- I dove into understanding Creative Commons licensing—something I've put off for far too long—and found that it offered just the right mix of options for my book. So now, people who've been excerpting parts of it (a.k.a., "remixing") can do so legally.

- I finally learned to use an image-processing application, so I can stop asking Spike and Morgan McGuire how to do what I think they find the equivalent of balancing parentheses (well, for me; I count parens like some sharks count cards).

It's also been a wonderful year personally, from the urban delight that is Edinburgh to the new world being created in real-time in Bangalore, from walking in awe of nature in Australia to biking in Lance's town in Texas, from seeing (from afar) the site of the Burgess Shale to lying on my back on the Scituate Reservoir dam to bask in the Perseids. I've seen, up close and (sometimes a bit too) personal, everything from rattlesnakes to kangaroos, from a platypus to both black and grizzly bears. And as my blog's name suggests, cricket hasn't been too far away, from following a good chunk of the World Cup to fulfilling every fan's dream: watching England play Australia at the Sydney Cricket Ground, even if that verb is a euphemism for the abject surrender of the Three Lions we witnessed that day. Over up!

Algorithms

continued from page 5

Such a plan has a much higher expected value than the "best" plan.

With an anticipatory algorithm helping, Nancy could just say "That's OK, Robin. These things happen. I hope your son is okay, and I'll just switch to plan B." That's an application that can help with nursing performance (fewer overtired nurses is always a good thing) and nursing satisfaction (it is much easier when on-the-job disruptions are less catastrophic).

While Pascal's work has not yet been applied to nursing, he says, "Currently we have applied our work to fleet management, reservation systems, scheduling systems, and project management. Now, we're trying to get data on all these dynamic personnel management problems— nurse scheduling is a example of that. Believe it or not, daycare scheduling is becoming important too." Keep your eyes open for tools using Pascal's work to appear in your daily life!

News & Events

The Artemis Project was created in 1996 with the purpose of opening the field of computer science to young women who otherwise might not have had the opportunity or interest. Since then, the program has grown; the number of coordinators has increased from two to four, the number of students has increased from eight to twenty girls each summer, and a mentoring program was developed to follow the girls through their high school years and beyond.

Artemis students are all between the eighth and ninth grades. This age group has been identified by research sponsored by the National Science Foundation as a time when many young women lose interest in science or succumb to peer pressure in avoiding math and science. Artemis makes it clear that computers and computer science can be fun and that going to college is a worthwhile and exciting goal. Because the Artemis program is entirely free of charge, the program focuses on reaching girls who may be prohibited from attending more expensive summer programs.

Phirum Peang: Weightlifting Champion

Computer science department systems programmer Phirum Peang takes first place in the bench press for the Men's Open 123 Pound Class at the 2007 New England RAW Championship.

In addition to his New England RAW title, Phirum has participated in and taken first place bench press titles in his class at the APA 2007 New England Open, the WNPF New Jersey State Championship, the WNPF New York State Championship, and the RAW 2007 Granite City Iron War. Phirum has also achieved the Rhode Island state bench press record for the federation (230 pounds) and the national bench press record (241 pounds). Most recently, he became the world-record holder in his weight class for bench pressing 242 pounds. Congratulations Phirum!

Halloween 2007

Unplugged

Professor Eugene Charniak's research is in the area of language understanding or technologies which relate to it, such as knowledge representation, reasoning under uncertainty, and learning.

I greatly enjoyed our last issue, particularly my co-columnist's piece. Shriram's comments on the difficulties of being a program chair got me thinking. Although Shriram did not mention it, being asked to be the program committee chair of a good conference indicates a high standing within your community. On the other hand, I have only twice been asked (I accepted both times) and both of these occasions have been in the last ten years. This seems to conflict with another comment Shriram made, that my elevation to university professor indicates that I am a mucky-muck. If I am so important, why have I not yet been asked to be a program chair? There might be an interesting story here—I will let you know after I write it.

Let's start with the basics. Until recently someone considered for academic promotion has been judged on the basis of three kinds of evidence: publications (and the quality of the venues in which they appear), reputation within the local community, and reputation with the professor's world-wide academic community. Lately I, for one, take a look at Google Scholar to gauge impact by using its counts of how often the person's work has been cited. (I will neither confirm nor deny Google scholaring myself.) This is recent because until now citation indices have been terribly biased. They only included journal articles, and in many areas of computer science, journal articles do not play nearly as important a role as do the major conferences (which is why being asked to be program chair is so important).

Ignoring Google, the most important of the three criteria for promotion is outside reputation as indicated (typically) by the letters of reference that are solicited by the university. The publications are, of course, important, but as an indicator of importance they can be overrated. For example, about three years ago a natural-language processing (NLP) researcher submitted a paper to a major NLP conference attempting to quantify the importance of university NLP programs and individual researchers. The quantification was based solely on the number of articles in certain high-prestige conferences. I am pleased to announce that I was ranked sixth in the world. (I figure there are about 2000 people in the pool for which this was selected.) My colleague, Mark Johnson, was ranked fourth! Nevertheless, both Mark and I suggested that the paper be rejected. I have my own personal world-wide ranking, and both Mark and I figure highly in it, but my ranking of other researchers diverge so sharply from what was reported in this paper that I could not give it much credence. I had never heard of the person who ranked first in the world! In a second

continued on page 30

Unplugged

continued from page 29

case I know the person, but I and all of the people I really respect in his sub-area of research think rather poorly of his output. (In this case Google would also get it wrong as his work is highly cited.) The point is this—it is sad to say but even at the best conferences, most papers disappear without a trace—they are impact-free—and even when they do have impact, not all researchers really know what is going on.

As for local reputation, as a means of judging research output it is also quite flawed. Except in the largest departments, any given university has very few people in any academic sub-discipline, often only one. Thus the local people are often not in a position to make key academic judgments and hence the need for outside letters. It is also the case that familiarity breeds contempt, or to cite a few other clichés, no man is a hero to his butler, and no programmer is a hero to his compiler.I mean this as no disrespect to my colleagues, just the opposite, but somehow I am always surprised by the high regard in which they are held at other institutions.

The point is this—it is sad to say but even at the best conferences, most papers disappear without a trace—they are impact-free—and even when they do have impact, not all researchers really know what is going on.

To summarize: I got my promotion because my immediate colleagues thought highly of me, I have sufficient articles in the right places, and most importantly, my close research colleagues also think very well of me. But for a person of my reputation, I have a very poor record of being asked to serve as a program chair. What gives?

The answer lies in the fact that for most of my academic career I have had an uneasy relation to those in my research sub-area(s). I have always been reasonably well known, but my work spanned two academic groups and in neither group was I completely Kosher. I grew up research-wise before artificial intelligence (AI) broke up into sub-specialties. Thus for most of my career I identified myself as an AI researcher rather than as someone in NLP. Nevertheless, as the splits occurred I was clearly an NLP person, rather than vision, machine learning, etc. This was OK but for the fact that at this time NLP was viewed as sort of a "fringe" area by most AI researchers. Thus when looking for a program chair, NLP was not usually the first place one looked.

To make things even worse, within NLP my research was viewed as interesting but quirky. (Needless to say, this is my own analysis. I have not gone around asking people "but what do you REALLY think of me.") Because of my AI background my research had tended to be at the intersection of AI and NLP. Thus I have been concerned with the impact of knowledge representation on NLP, or certain problems of inference as they affect our ability to "understand" text. Meanwhile the NLP community as a whole tended more toward the relation of NLP to linguistics, and to the degree they were concerned with inference, they approached it from a point of view that stressed the use of logic in general and first-order predicate calculus in particular. I have never been completely comfortable within that framework.

I am saying this not to criticize my research community, as I now believe we were all mostly on the wrong track. I say this because what has changed between then and now is what many of us call the statistical revolution. About 15 years ago, I discovered a small group of people who had been applying statistical and probabilistic reasoning to natural-language processing so I dropped my current research and joined that group. At the time this was viewed as being even odder than my previous inclinations, but this time I hit things right. In the early 1990s this kind of research was almost totally absent from the major NLP conferences. Papers that we now regard as groundbreaking where uniformly rejected. Today almost all papers at the Association for Computational Linguistics Conference (the most highly rated conference in my area) are statistics-based. Thus, I am a logical person to ask to chair the conference, and indeed I have. I would not even be surprised if I am asked to chair a major AI conference one of these days, since NLP is gaining more and more respect. Only time will tell.

Yes, I guess this is interesting. At least it kept me interested. And, with any luck, it is not even too self serving.

True Tales from the Tech World

Harry Sparks '76

A Winding Path From Brown

I received my computer science degree so long ago that computer science was a part of the Applied Math Department. I still have in my drawer the first two issues of the (then) famous AM51 tee shirts showing a somewhat frazzled programmer. I've felt that way many times in the intervening years.

On a visit to campus last spring with my daughter who is now looking at schools, I went to look at the computer science building which did not exist when I was

on campus. As I was wandering around, I encountered Professor van Dam who encouraged me to send an update to Conduit. So here it is.

My path has not been very academic as I went from Brown to a programming job with McDonnell-Douglas, (now Boeing) where I did Cobol maintenance for payroll systems that issued 30,000 checks a week. It was a classic information technology dues-paying job. I honestly believe that my education at Brown made me a better programmer than many of my colleagues from other institutions.

From there I veered off into the world of public accounting, joining Coopers and Lybrand (at the time one of the Big 8 accounting firms) as an EDP auditor. My education gave me the background I needed to pick up systems quickly because I understood the underlying concepts better than many of my peers. The IT auditing activity requires knowing a lot about many things, so I rapidly became a generalist. I worked on big IBM main frames, AS/400s, mini-computers, (anybody remember those?) and the occasional HP or DEC system. I saw a lot of different hardware and software and worked at a wide variety of companies from giants like State Farm Insurance and the University of Missouri system to smaller local operations. I knew every restaurant in Bloomington, Illinois, the headquarters town for State Farm.

While working in accounting, I obtained an MBA from the University of Chicago. I am convinced that the MBA, an Ivy League degree and working for a Big 8 firm helped me land my next position, where I ran the business side of IT operations for an engineering/architecture firm. We had DEC equipment. Once I figured out that VMS was keyword and not positional I really began to like it. I am sorry that DEC died. Ken Olson had figured out scalability. I ran the same operating system on my desktop and on my mainframes which is still impossible to do on today's predominant hardware. While I was there, we converted some systems from DEC's codasyl database to the Ingres relational database. Hardly anyone uses that anymore, but I did get the opportunity to pick up some SQL and relational design concepts. I spent a fair bit of time telling one of my older programmers that he had to say we needed to "de-normalize for performance" when he complained about relational models.

That gig lasted about six years; I then went back to the auditing game as director of information technology for Emerson Electric, a Fortune 500 firm. That position was quite interesting as it led me to international travel. I made several trips to Europe and three to Asia. It was clear that labor in Shanghai was cheap. The hotel I stayed in not only had two greeters at the door, but also a person to push the button on the elevator when you got in. Seriously, it was a great opportunity to see

other places and meet some local people. I visited some places that I am sure I would never have seen if I had not had the education and experience to take this job. The Emerson gravy train ended at the end of 2001 with a layoff when IT auditing was outsourced.

I did some freelance consulting for a while and then took a position as a consultant with RSM McGladrey. After a couple of years doing Sarbanes-Oxley compliance for IT and helping companies establish IT controls, I moved to my current position.

It has taken me full circle as I am now installing big systems for big companies. I'm working in what is called the corporate performance, management and business intelligence space. So I am back at least in part to writing programs, mainly in the SAS language at the moment. Having been brought up on third-generation languages, I find myself getting frustrated by the number of components in today's systems. The software I am installing uses java components, but also requires MySQL, two java run-time environments and third-party databases from BEA and Postgres, and a third- party application called Xythos. It is quite messy getting all those components to play well together. Often I have to chase through four or five layers to get to something that might be executable. I'd like to find out if others share this opinion on the current state of programming.

There have been other projects and travels along the way. Professionally, I have been very pleased with my education. From my days at Brown I received enough understanding to be able to keep up with changes and apply underlying concepts to new environments in order to understand the issues involved and make viable recommendations to my clients and employers.

On a personal level, I have been married for 24 years and have two daughters. The oldest is a junior at Trinity College in Hartford and the youngest is a senior at Providence High School. (That is Providence High School in Charlotte, North Carolina, our current residence.) Providence Road is a main street here in Charlotte and it took some getting used to as there are a lot of subdivisions and shopping centers that have Providence in their names. Although I have tried, it does not look like I will be sending any of my children to Brown.

I'd be delighted to hear from old classmates, and any alumni in the Charlotte area. My personal email is hsparks@chicagogsb.edu.

Bill Butler '80

I'm probably the oldest CS graduate as I jumped into Brown and the computer science program in 1977 as a Resumed Undergraduate Education student. I received a Sc. B. in CS in 1980— when I was 41 years old.

I worked as a data processing manager for several years after graduating, followed by a few years of consulting, but I was fortunate enough with my investments that I could retire early.

In 1997 I retired and moved to Durango, Colorado (location of choice). Since then I've had a chance to get in a considerable amount of hiking, and over the years I've been on nine Grand Canyon raft trips.

I have a website at http://www.durangobill.com which is mostly devoted to geology (Grand Canyon and "Peak Oil"), math (especially combinatorics), and a swipe or two at the "Young Earth Creationists." I also run the Yahoo group "Suncor Energy andCanadian Oil Sands Resources" at http://finance.groups.yahoo.com/group/suncor_energy_and_canadian_oil_sands/.

Randy Pausch '82

In September 2007, Carnegie Mellon Professor Randy Pausch delivered what will be his final lecture, entitled "Really Achieving Your Childhood Dreams." During his decade at Carnegie Mellon, he co-founded the Entertainment Technology Center and developed a free software program, called Alice, to get children, especially girls, more interested in computer programming.

This amazing lecture, which was given for his children, drew an enormous amount of media attention and was written about in *The Wall Street Journal,* the *Pittsburgh Post-Gazette* and the *Pittsburgh Tribute-Review*. In addition, Randy was the ABC World News Person of the Week and was talked about on Good Morning America, the CBS Evening News and Oprah.

In the summer of 2006, Randy was diagnosed with pancreatic cancer, the most deadly of cancers, with only a four percent five-year survival rate. He was among the 20 percent of patient eligible for surgery and had Whipple surgery in September 2006 followed by chemotherapy until May 2007. In August 2007, Randy found

out that the cancer had returned. At that time, the doctors estimated that he had three to six months of healthy living left. On October first, he received the good news that the first round of palliative chemotherapy was working and he will have a bit more time to spend with his wife, Jai, and his three children, Dylan (5), Logan (3), and Chloe (1).

To view his inspiring lecture, please visit: http://www.cmu.edu/randyslecture/

Tapan Parikh, '96

Tapan is the winner of this year's TR35 Humanitarian of the Year award for his research at the University of Washington on simple, powerful mobile tools for developing economies.

Parikh has created information systems tailored for small-business people, such as fisherman, in the developing world—systems with the mobile phone, rather than the PC, at their core. His goal is to make it easier for these business owners to manage their own operations in an efficient and transparent way, and to build connections both with established financial institutions and with consumers in the developed world.

Viswanath (Vishy) Ramachandran Ph.D. '98

This is Viswanath (Vishy) Ramachandran (Brown CS Ph.D. 1991-1998). After Brown, I worked at Netscape, Siebel, Friendster and Digeo in Silicon Valley. Around three years ago, I relocated to Mumbai, India (my birthplace) as vice president for engineering at Webaroo (www.webaroo.com). I now live here with my wife and two children.

Regards,
Vishy

Over 100 people gathered for the third annual Computer Science Reunion & Networking Reception at the CIT. It was great to see everyone and we look forward to seeing you all again in 2008!

Commencement 2007

Brown University Department of Computer Science Celebrates Record Number of Master's and Ph.D. Graduates

Brown University's 239th Commencement exercises were held on Sunday, May 27, 2007, and the Department of Computer Science, celebrating its 28th year, awarded an all-time high of 31 Master of Science degrees and 12 Ph.D.s. In addition, 41 Bachelor's degrees were conferred. Graduates gathered after the Commencement ceremony to celebrate with the department. We hope that all of our grads keep us posted on what they are up to. Email conduit@cs.brown.edu with updates.

2007 Ph.D. Graduates included (left to right): Greg Cooper, Glencora Borradaile, Olga Karpenko

Ping!

Where are you and what are you doing?

Let us know what's happening in your life! New job? Received an award? Recently engaged or married? New additions to your family? Use this form to submit your news or email conduit@cs.brown.edu.

First Name _____

Last Name _____ Class Year _____

Address _____

_____

City _____

State _____ Zip _____

Email _____

My news:

Mail to: Conduit, Department of Computer Science, Brown University, Box 1910, Providence, RI 02912

Above: Participants of the workshop on pen-centric computing. The event was hosted by Andy van Dam and sponsored by Brown University, Microsoft Corporation, and the US Intelligence Community's Disruptive Technology Office. Read more on page 23.

Department of Computer Science
Brown University
Box 1910
Providence, RI 02912
USA

BROWN