the newsletter of the International Computer Science Institute

The ICSI GAZETTE

volume three | issue one | september 2004

# featured alum: Hervé Bourlard

Professor Herve Bourlard, Professor at EPFL and the director of the Swiss research institute IDIAP, was one of the first visitors to ICSI – and for many years, appeared at ICSI so often that we began to view his time away as his vacations. While his many European responsibilities no longer permit long visits here, his collaboration with us has continued. This is evident both in speech research (as in the DARPA EARS project, featured in an earlier issue of the Gazette) and in visitor programs. He has been influential in ICSI's inclusion in several projects sponsored by European and Swiss funders, most recently the Swiss "Interactive Multimodal Information Management" (IM2), which he heads, and the European "Augmented Multi-party Interaction" (AMI), which he co-manages with fellow ICSI alum Steve Renals of the University of Edinburgh.

Bourlard's first visit to ICSI was almost an accident. He was working for Philips in Belgium when Geoffrey Hinton of the University of Toronto asked him to come to Toronto for a two-year visit. Because Philips wouldn't approve such a long absence, Hinton arranged for Bourlard to apply for a one-year visit to ICSI instead. Jerry Feldman, who was the director of ICSI at the time, asked Nelson Morgan to review Bourlard's application. Morgan says that his first thought was "this guy is either crazy or brilliant. After he arrived, I realized that he was both – and I wanted to work with him for both reasons." Morgan headed the Realization Group at that time, which was focused on computer architectures for connectionist algorithms. Morgan wanted to have an applications focus, and Bourlard had ideas for using neural networks in novel ways that were applicable to speech recognition. Morgan had previously worked on speech processing, so Bourlard's "daring" ideas for speech were a perfect fit. Bourlard arrived at ICSI in 1988, and he and Morgan began to develop the theory and practice of hybrid hidden Markov model/neural network systems. The success of their work led to the eventual creation of the Speech group, as well as to a lifelong friendship.

Photo Courtesy of IDIAP

*Hervé Bourlard*

Bourlard's enthusiasm for his work was a good match to the spirit of ICSI. At the time he and Morgan began work on their neural network speech algorithms, workstations were not fast enough to test the theory. "I still remember, arriving at ICSI full of new ideas and theories, explaining why those ideas were potentially so important but, at the same time, why there was no chance to have them implemented on desktop computers. The reaction at ICSI (mainly coming from Nelson Morgan, I have to say) was immediate: 'You've convinced us that these are good ideas, so limited

*Continued on Page 8*

## in this issue:

Howel Shaw

# as i see it by Nelson Morgan, Director

I was recently reading about US Federal Budget projections in a UC Berkeley publication, which drew upon an AAAS study for its information. The projections showed significant reductions in funding for basic research, pretty much across the board. Even in those agencies that were expected to have more funding (notably the Department of Defense), the components that were associated with more basic research (the so-called "6.1" funds) were likely to be reduced in favor of more immediately applicable development projects. It will come as no surprise to anyone that a researcher at an institution like ICSI would be opposed to this turn of events; such a view is, of course, self-serving, as we derive significant revenue from Federal grants. However, it is more than the stability of one institution that is at stake. An intellectual environment with a long-term perspective has been the primary breeding ground for much of the technological progress that we have been enjoying over the last few decades. Retreating from a national commitment to these things is dangerous, both in terms of our security and our economic well-being. Industrial sponsors can help, but they too have increasingly turned towards shorter-term goals.

"I have great confidence in the ability of ICSI researchers to adapt, to find new sources of support for the work, and above all to excel in their research. Nowhere is this ability to adapt more apparent than in cybersecurity, the featured project for this issue."

Of course, policy makers have tough budgetary choices to make, and I don't envy them in their role. Nor am I qualified to seriously assess the economic and political fallout from choices that would favor basic research. But I think I can say this: if the federal government significantly cuts back on its historic post-WW II role of support for basic research, and for support of training of graduate students at research universities, the US risks becoming second-rate in innovation.

I know that these thoughts are considerably less upbeat than my usual columns, but they reflect my growing concern. In a way they are the flip side of my continued enthusiasm for ICSI, which, like the best university research departments, remains a bastion of basic research. I have great confidence in the ability of ICSI researchers to adapt, to find new sources of support for the work, and above all to excel in their research. Nowhere is this ability to adapt more apparent than in cybersecurity, the featured project for this issue. Most of the national focus in the post-9-11 US has been on physical security, which certainly deserves significant attention. However, the growing dependence of so many parts of our lives on computers in general and on the Internet in particular makes network security a critical concern as well. This area certainly can and does attract support given the immediate need; however, Vern Paxson, along with colleagues like Nick Weaver, are in it for the long haul. They have created a new Worm Center at ICSI, which has just been funded by the National Science Foundation as part of their CyberTrust program. Vern is also the current head of the Internet Research Task Force (IRTF), reflecting his commitment to research.

Our featured alumnus this issue is Hervé Bourlard, the Director of the Swiss lab IDIAP. In the previous issue of the Gazette, I briefly described his early time at ICSI, and the seminal role he played in the start of ICSI's speech research. As an ICSI Board Member and a lead researcher on Swiss and EU projects that ICSI is on (IM2 and AMI respectively), he stays current with ICSI events. He is also very involved with the DARPA EARS project described in the last issue. For instance, he recently took an active part in an EARS research retreat here in Northern California in which we both knocked around wild research ideas and worked collaboratively to complete the development of more mature ideas.

I would like to take this occasion to welcome a new key member of our administration: Associate Director Marcia Bush. Marcia comes from a background in both research (for instance, as a faculty member at Brown) and in technology licensing (at Xerox), and will be heading our efforts in business development. Given our focus on open technology and basic research, "business development" at ICSI requires sensitivity to our style of work so that we do not lose our primary focus. Marcia's background makes her a great match to this requirement, and I greatly look forward to working with her in the coming months and years.

Finally: once again, ICSI speech researcher Qifeng Zhu has added another underage scientist to our community – Ryan Ming Zhu was born on July 24th.

# news briefs

## Individual Accomplishments

Congratulations to **Sally Floyd** of ICIR, recipient of the **2005 IEEE Internet Award.** This award is presented annually for exceptional contributions to the advancement of Internet technology for network architecture, mobility and/or end-use applications. Floyd was chosen for "contributions to the Internet architecture, particularly in the areas of congestion control, traffic modeling, and active queue management."


Sally Floyd

Congratulations to **Mark Allman** of ICIR, who received a Public Service Medal from NASA on September 1, 2004. This award is given to non-government employees for significant contributions to the mission of NASA. Allman's contributions are in the area of space networking protocols.

**Jerome Feldman,** of ICSI's AI Group and BCIS, was named Program Director of Cognitive Science at UC Berkeley on July 1, 2004.

ICSI alum **Krste Asanovic**, a member of the Computer Science faculty at MIT, was selected to become an External Fellow of ICSI in June, 2004. Professor Asanovic said that he is delighted to accept the invitation and looks forward to future collaborations with ICSI researchers.

ICSI Trustee **Shankar Sastry** has been elected to the **American Academy of Arts and Sciences** (AAAS). He is one of seven UC Berkeley professors selected this year.

President Bush announced in May, 2004 that ICSI Faculty Associate **Ion Stoica** is one of 57 recipients of the **2002 Presidential Early Career Award for Scientists and Engineers** (PECASE).

A new book, "**The Success of Open Source**" by ICSI Faculty Associate **Steve Weber**, was published in April, 2004.

On July 1, 2004, **Steve Weber** was appointed Director of the Institute for International Studies at UC Berkeley.

**John Moody**, a senior researcher in the Algorithms Group, was interviewed by Erin Burnett on the **Bloomberg Financial News** channel on Monday, March 1, 2004. Moody's academic research centers on machine learning algorithms and computational finance. He is also the founder of J E Moody & Company, LLC, an alternative investment management firm.


John Moody

**In** the live television interview he discussed how his firm's approach to quantitative asset management differs from traditional approaches due to his academic background. Moody uses machine learning and statistics to more accurately predict market trends and avoid "data snooping".

**In** conjunction with the television interview, Moody was quoted extensively in an article entitled "From Harvard to Hedge Funds" in the **April 2004 issue of Bloomberg Markets.**

Former post-doc **Nicholas Weaver** has been hired as a full time senior researcher with the Networking Group. A major focus of his work will be the new Center for Internet Epidemiology and Defenses, which is described on page four, where he will continue his research on Internet worms and play a leadership role for the Center.

## Events

On September 1, 2004, **Marcia Bush** joined ICSI as an Associate Director. Marcia will be working on business development, including technology licensing and working with industry partners. She worked at Xerox for almost fifteen years, initially in a technical position managing a speech and image processing group, and later in marketing and business development. Under Marcia's guidance, ICSI hopes to make future expansion into new areas of research possible.


Marcia Bush

ICSI's FrameNet Project hosted a **TRANSCOOP workshop** with the German SALSA Project March 22-26th, 2004.

ICSI was the host of the 2004 **EARS Novel Approaches retreat.** This year's retreat was at the Bodega Bay Lodge and Spa from

*Continued on Page 6*

# ICSI and UCSD Open Center for

**A worst-case worm attack could cost the United States tens of billions of dollars. How do we combat epidemic-style Internet attacks? Researchers at the new Center for Internet Epidemiology and Defenses will be working hard to answer that question over the next five years.**

The Center for Internet Epidimiology and Defenses (CIED) opens this month at ICSI and UC San Diego. The Center, led by Vern Paxson of the Networking Group and Stefan Savage of UC San Diego, is sponsored by a multi-million dollar grant from the National Science Foundation's Cyber Trust program.

Diane Starr

*Principal Investigator Vern Paxson*

Paxson and Savage point out in their description of the Center that, "The combination of widespread software homogeneity and the Internet's unrestricted communication model creates an ideal climate for infectious pathogens, and, worse, each new generation of outbreaks demonstrates increasing speed, virulence, and sophistication. Indeed, a recent Computing Research Association panel framed eliminating epidemic-style attacks from the Internet within 10 years as a *Grand Challenge* problem."

The Center aims to combat this problem. It will work towards understanding how the Internet's open communications and software vulnerabilities permit worms and viruses to propagate, devising a global-scale early warning system to automatically detect epidemics in their early stages, developing forensics capabilities for analyzing wide-ranging infections, and developing techniques and devices that can suppress outbreaks before they reach pandemic proportions.

Specific technologies to be studied at the Center include the use of "honeyfarms" and "network telescopes" to automatically detect and analyze network intrusions, the use of the XORP router software platform as a means for deploying defense mechanisms within enterprise networks, and the continued analysis of worm behavior, including empirical studies of worm attacks as well as theoretical worm research.

Honeyfarms consist of vulnerable servers ("honeypots") used to detect worms and study their behavior. Analysis of honeyfarm behavior will be used to then program worm suppression devices. A top priority for the first year of operation is the construction of a "network telescope" of unprecedented size. A network telescope is the front-end that feeds the honeypots. It also provides broad visibility into Internet-scale events.

Diane Starr

*Co-PI Nicholas Weaver*

The goals of the Center go beyond researching methods to prevent Internet attacks. CIED will also encourage the growth of the emerging field of "Internet epidemiology" through outreach and education, encouraging students to pursue degrees and careers in cybersecurity. As technology continues to develop at a rapid pace, cybersecurity research must keep pace with new developments, and attracting new talent to the field is vital to the continued success of the Internet.

Paxson has led the cybersecurity research effort at ICSI for several years. The new funding for the Center will allow ICSI to increase its network security efforts by expanding the focus of existing network security research and investigating new directions of research in this exciting field.

Former ICSI postdoc Nicholas Weaver will play an integral role at the new Center. Weaver is an expert on Internet Worms who has been working with Paxson for several years, and will now become a staff researcher at ICSI, spending roughly half of his time working for CIED.

The NSF grant was awarded at a critical point in time for ICSI's cybersecurity effort: DARPA (Defense Advanced Research Projects Agency), one of the top

# Internet Epidemiology and Defenses

**Prior to the opening of the Center for Internet Epidemiology and Defense, ICSI's cybersecurity effort was already well under way. Vern Paxson has led the cybersecurity effort, which includes the study of Internet epidemics and detecting Internet attacks in general.**

g

U.S. government funders of cybersecurity research, has decided to classify its entire worm-research program. ICSI's commitment to an open research environment, including international collaboration, prevents researchers from pursuing grants for classified U.S. government work. Therefore, DARPA's position not only affects the amount of funding available to ICSI researchers doing security research, but also puts ICSI in tighter competition for available grants.

## PREVIOUS CYBERSECURITY RESEARCH AT ICSI

Cybersecurity is not a new field of research at ICSI. Paxson, along with Weaver and numerous graduate student interns, has been actively involved in the study of both Internet epidemics and detecting Internet attacks in general for years.

### WORM RESEARCH

Network telescopes, honeyfarms, and scan detection are all in use in ICSI worm research. Empirical studies have been performed on real-life worms, such as the Witty worm, and theoretical studies are being used to estimate potential damages of a pandemic worm attack. Using "scale-down" techniques, which involve a reduced-scale model of key aspects of the Internet, researchers endeavor to simulate and study worm behavior. Paxson and Weaver estimate that in the worst case, well-engineered worms could cause tens of billions of dollars in damages, and fully propogate in under two seconds.

### INTRUSION DETECTION RESEARCH

Paxson, who holds a joint appointment at the Lawrence Berkeley National Laboratory, is the designer of the "Bro" network intrusion detection system. Intrusion detection research differs from worm research and prevention in that it looks at all unusual network activity, including some that is very "low profile", and aims to discriminate harmless activity from possibly dangerous intrusions. Bro, which is in 24/7 operational use at LBNL, UC Berkeley and ICSI, tracks network activity and in its LBNL deployment is able to block potentially threatening hosts from making any type of contact with the network, thus preventing possible attacks.

Leah Hitchcock

*Cybersecurity Interns*
*Left to Right: Sarang Dharmapurikar, Chema Gonzalez, Abhishek Kumar, Christian Kreibich, Holger Dreger. Not pictured: Vinod Yegneswaran, Stefan Kornexl*

### STAFF

Staff working for CIED at ICSI will include Principal Investigator Vern Paxson, co-Principal Investigator Nicholas Weaver, XORP Principal Investigator Atanu Ghosh, and Staff Researcher Mark Allman.

Diane Starr

*Staff Researcher Mark Allman*

In addition to the permanent staff members, the students and interns who have worked on cybersecurity at ICSI this year include Sarang Dharmapurikar, Holger Dreger, Chema Gonzalez, Stefan Kornexl, Christian Kreibich, Abhishek Kumar, and Vinod Yegneswaran.

# news briefs

*Continued from Page 3*

August 4-7, 2004. In addition to researchers from the ICSI Speech Group, University of Washington, SRI, Columbia University, and IDIAP were represented at the retreat.

## Project News

ICSI's **XORP** Project (eXtensible Open Router Platform) was featured in the **March 2004 issue of Network Magazine**.

**XORP** was featured in a CNET News.com article on April 19, 2004.

In April, ICSI began participation in data collection for the Linguistic Data Consortium's **Mixer Project**. The Mixer project is collecting telephone conversations for use in speech recognition research.

**Scott Shenker** was awarded a subcontract from Yale University for a cybersecurity project entitled **An Economic Approach to Security**.

**Manny Rayner** has received a subcontract from NASA Ames Research Center to work on a project called **Reusable robust speech recognition for spoken dialogue applications**.

**Mark Allman** and **Vern Paxson** of ICSI's Networking group were awarded a grant from Cisco Systems to study **Reactive Network Measurement**. Reactive Measurement, or REM, is a measurement technique that attempts to diagnose the root cause(s) of network problems by using observations or measurements to decide which (if any) additional diagnostic measurements are necessary to uncover the root cause of an anomaly. Allman and Paxson intend to develop a generic framework for conducting such measurement studies.

On September 1st, ICSI welcomed the first two visiting researchers for the European Union Integrated Project called "**Augmented Mulitparty Interaction**" (AMI). **Xavi Anguera** and **Marc Ferras** of Spain will be working with Speech Group researchers on this project, which aims to improve computer enhanced multi-modal interaction. AMI is managed by two ICSI alumni, Herve Bourlard of IDIAP and Steve Renals of University of Edinburgh.

**BFOIT, the Berkeley Foundation for Opportunities in Information Technology**, has just concluded its most successful summer program. Executive Director **Orpheus Crutchfield** said that "both qualitatively and quantitatively, this summer was well above the rest." This year, BFOIT expanded its summer program from serving 25 high school students to 75 middle school and high school students.
**BFOIT** works year-round with young people, providing education, motivation and support to female and underrepresented minority students, encouraging them to pursue careers in computer science and engineering. The Summer Institute, held at Soda Hall on the UC Berkeley campus, is the finale of the year-round program. During the week-long Summer Institute, students take programming classes, become a tight-knit cohort and support network, and attend motivational talks by engineers and computer scientists.

Photo courtesy of Orpheus Crutchfield

*BFOIT students at work*

The **eXtensible Open Router Platform (XORP)** has released its open source router software. XORP differs from commercially available routers in several aspects, which are designed to make it more user-friendly, adaptable, and stable.
**XORP** is designed for extensibility, so that it can simultaneously satisfy several user groups: network researchers needing a platform for experimentation, network operators needing a low-cost stable routing

platform on commodity hardware, network equipment vendors with special purpose hardware, and network application writers looking for an open platform to support their applications. XORP is ideal for experimentation, and has the ability to 'understand' scripting in practically any language, which allows a user to come up with new routing applications using existing code. In addition, the XORP architecture is compartmentalized so that should something go wrong in one area, the other areas are not compromised. This

Howell Shaw

*The XORP development team. Back row: Eddie Kohler, Orion Hodson, Javeir Macias, Fred Bauer Front row: Pavlin Radoslavov, Atanu Ghosh*

offers more security than alternative router platforms. **Creators of XORP** hope that it will one day become an attractive alternative to commercial stacks for network equipment vendors. XORP is licensed under a BSD-style license, which allows the user to use it for any non-commercial purpose. The XORP development team believes that many users will contribute changes and improvements to XORP much like Linux users. Additionally, they envision a new class of software that doesn't currently exist: the router application. They believe that XORP's extensible architecture is the means to create third party software for mainstream commercial router platforms. The possibility for novel network functionality that this would enable is a long-term goal that XORP developers hope to explore in the future.

**Vern Paxson** was awarded a two-year grant from the United States Department of Energy for **Detecting and Blocking Network Attacks at Ultra-High Speeds**.

**Vern Paxson**, along with **Stefan Savage** of UC San Diego, were awarded a grant from the National Science Foundation's Cyber Trust program to open a new **Center for Internet Epidemiology and Defenses**. The complete story is on page four.

## ICSI Community News

**Kaia Alessandra Penilla** was born on March 11, 2004 to proud parents Mary and Martin Penilla.

**Rebecca Siyao Chang** was born on March 9, 2004. Her parents are ICSI alum Shawn Chang and his wife Phoebe.

**Anna Nierhoff**, daughter of Till Nierhoff and Andrea Frömbgen and sister of Lisa Nierhoff, was born on Monday, July 12, 2004 at 4:53 p.m. She weighed nine pounds two ounces and was 22.25 inches long at birth.

**Ryan Ming Zhu**, second son of Qifeng Zhu and his wife Xuemei Feng, was born on July 24, 2004 at 7:02 a.m. Ryan weighed nine pounds nine ounces at birth and was 21 1/2 inches long. His Chinese name is Zhu Ming Run.

**Siina Rantala** was born on August 16, 2004. She is the daughter of visiting researcher Juuso Rantala and his wife Eveliina, and the younger sister of Isla.

Leah Hitchcock

*Foosball Champions Pedro Ruiz and Abhishek Kumar*

Congratulations to **Pedro Ruiz** and **Abhishek Kumar**, champions of the first **ICSI Foosball Tournament**. Eight teams participated in the heated tournament which concluded on Friday, July 16. The tournament was enjoyed by participants and spectators, and we look forward to the next one. Anyone wishing to compete is invited to practice on Tuesdays and Thursdays at 3:30 p.m. on the sixth floor.

# featured alum: Herve Bourlard

*Continued from Page 1*

computer performance is not a valid excuse. Let's build a faster machine.' And they actually did it!" The machine he refers to, the Ring Array Processor, or RAP, was completed at ICSI in 1989, and was fast enough to successfully develop working versions of the new methods. Because of this, Bourlard says that he "realized that ICSI was really an exceptional place where everything was possible."

Bourlard and Morgan look back on their first joint research projects fondly. Morgan recalls "being so excited that we couldn't sleep at night, we would stay up thinking about the ideas, then come to work and talk about them all day". He says that they worked together "not without ego, because there was plenty of that, but without self interest. If we were writing a paper together, we didn't worry about who would be first on the author list – sometimes we would just take turns. It was often hard to even pick apart which idea was mine and which was his."

"To be among the very best, it is a must to collaborate with the best institutions, and to benefit from international networks, which allows the institute to hire rising stars and attract the best students...Successful international collaborations are always based on long-term relationships, mutual respect, and mutual benefit."

Bourlard still believes that ICSI provides an ideal place for research: "Its international nature, its multi-disciplinary approach, the exceptional quality of its researchers, and the quality and inquisitive nature of its scientific environment...But all this would still not be unique to ICSI if it were not combined with a reasonable size institution, resulting in real opportunities for cross-area collaborations and an exceptionally dynamic and inspiring environment. With its affiliation with UC Berkeley, ICSI provides all the advantages of a university environment, while preserving the dynamism of a smaller institution."

Over the last sixteen years, Bourlard's career has progressed from a visiting industrial researcher at ICSI to a professor at EPFL and Director of IDIAP, a research institute in Switzerland where researchers are working on topics in speech and speaker recognition, computer vision and machine learning. He has become a figure of some renown in European research circles. He does manage to continue his own research in speech and related areas. His current interests include statistical pattern recognition, neural networks and applied mathematics with applications to spoken language processing, biometric authentication, and multimodal processing. As mentioned above he manages or co-manages the Swiss research network IM2 and the European Integrated Project AMI.

Bourlard says, "IDIAP in Switzerland is now building upon ICSI's model and, interestingly enough, is also resulting in a unique scientific spirit and exceptional dynamism." Bourlard has been able to recreate aspects of the research environment at ICSI that he feels set ICSI apart. As a result, IDIAP has a similar philosophy to ICSI, and the two institutes have formed a mutually beneficial relationship. According to Bourlard, this kind of connection is extremely important for the success of research efforts. "Research is becoming more and more complex, and also more and more competitive. International collaboration is thus necessary to ensure continued leadership and a "critical mass" of effort. As with industries, research institutions cannot [and] should not survive if they are not among the very best. And to be among the very best, it is a must to collaborate with the best institutions, and to benefit from international networks, which allows the institute to hire rising stars and to attract the best students."

Over the years, Bourlard has contributed to the growth of ICSI's international visitor program. After the first ICSI-Swiss visitor program eventually ended, Bourlard helped to reactivate it through funding from IM2. The IM2-sponsored Swiss program is still going strong today. The AMI project began in January and the first two visitors for the program arrived at ICSI on September 1. Bourlard says, "Successful international collaborations are always based on long-term relationships, mutual respect, and mutual benefit (win-win situations), and this is something that ICSI can provide through its intrinsic international nature." The long-term collaboration between ICSI and IDIAP is certainly a "win-win situation" and one that hopefully will continue for years to come.

# visiting scholars

*Since its inception, ICSI has had a strong international program consisting primarily of ties with specific countries. Current formal agreements continue with Finland, Germany, Spain, and Switzerland.*

### From Finland
Pekka Himanen (BCIS)
Konsta Koppinen (Speech)
Tuomo Pirinen (Speech)
Juuso Rantala (Haas)
Pasi Sarolahti (Networking)
Pertti Tormala

### From Germany
Holger Dreger (Networking)
Jens Gramm (Algorithms)
Stefan Kornexl (Networking)
Till Nierhoff (Algorithms)
Robin Sommer (Networking)
Till Tantau (Algorithms)
Oliver Wendt (Algorithms)

### From Spain
Alberto Amengual (Artificial Intelligence)
Xavier Anguera (Speech)
Marc Ferras (Speech)
Chema Gonzalez (Networking)
Carmen Pelaez (Speech)
Pedro Ruiz (Networking)
Francisco Valverde (Artificial Intelligence)

### From Switzerland
Sebastien Coquoz (Speech)
Vincenzo Pallotta (Artificial Intelligence)

*In addition, we often have visitors associated with specific research and projects.*

### From France
Brigitte Bigi (Speech)

### From Israel
Ron Shamir (Algorithms)
Roded Sharan (Algorithms)

### From the United Kingdom
Christian Kreibich (Networking)
Bruce Simpson (Networking)
Phil Woodland (Speech)

*ICSI is also involved in collaborative research with colleagues at other American institutions.*

### From Georgia Tech
Abhishek Kumar (Networking)

### From MIT
Michael Walfish (Networking)

### From Stanford University
Katerina Argyraki (Networking)
Beverly Yang (Networking)

### From U of Chicago
Terry Regier (Artificial Intelligence)

### From UC Davis
Matt Franklin (Algorithms)

### From U of Colorado
Laura Michaelis (Artificial Intelligence)

### From U of Michigan
Morley Mao (Networking)

### From U of Washington
Sarang Dharmapurikar (Networking)

### From U of Wisconsin
Vinod Yegneswaran (Networking)

*Clockwise from top right: Katerina Argyraki, Alberto Amengual, Christian Kreibich, Holger Dreger. Photos on this page by Diane Starr*

# publications listing

J. Alber, J. Gramm, J. Guo and R. Niedermeier. Computing the Similarity of Two Sequences with Nested Arc Annotations. Theoretical Computer Science, 312(2-3):377-358, 2004.

A. Archer, J. Feigenbaum, A. Krishnamurthy, R. Sami and S. Shenker. Approximation and Collusion in Multicast Cost Sharing. Games and Economic Behavior, 47, pp. 3671, 2004.

H. Balakrishnan, K. Lakshminaryanan, S. Ratnasamy, S. Shenker, I. Stoica and M. Walfish. A Layered Naming Architecture for the Internet. In Proceedings of SIGCOMMM 2004.

A. Ben-Dor, T. Hartman, R.M. Karp, B. Schwikowski, R. Sharan and Z. Yakhini. Towards Optimally Multiplexed Applications of Universal Arrays. Journal of Computational Biology, 2004.

E. Blanton and M. Allman. Using TCP DSACKs and SCTP Duplicate TSNs to Detect Spurious Retransmissions , RFC 3708, February 2004.

K. Boakye and B. Peskin. Text-Constrained Speaker Recognition on a Text-Independent Task. Odyssey 2004 - The Speaker and Language Recognition Workshop, Toledo, Spain, June 2004.

J. Bryant. Scalable Construction-Based Parsing and Semantic Analysis. In Proceedings of ScaNaLU-2004, Boston, MA, 2004.

J. Chuzhoy, S. Guha, E. Halperin, S. Khanna, G. Kortsarz and S. Naor. Tight Lower Bounds for the Asymetric k-Center Problem. In Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC04), 2004.

J. Elson, R.M. Karp, C.H. Papadimitriou and S. Shenker. Global Synchronization in Sensornets. In Proceedings of LATIN, 609-624, 2004.

E. Eskin, E. Halperin and R. Sharan. Optimally Phasing Long Genomic Regions using Local Haplotype Predictions. In Proceedings of the Second RECOMB Satellite Workshop on Computational Methods for SNPs and Haplotypes, pp. 13-16, 2004.

J. Feldman and S. Narayanan. Embodied Meaning in a Neural Theory of Language. Brain and Language 89, pp 385-392, Elsevier Press Incorporated, 2004.

S. Floyd. Limited Slow-Start for TCP with Large Congestion Windows. RFC 3742, Experimental, March 2004.

S. Floyd, T. Henderson and A. Gurtov. The NewReno Modification to TCP's Fast Recovery Algorithm. FRC 3782, Proposed Standard, April 2004.

S. Floyd and J. Kempf, Editors. IAB Concerns Regarding Congestion Control for Voice Traffic in the Internet. RFC 3714, Informational, March 2004.

H. Franco, F. Beaufays, N. Morgan and H. Bourlard. Speech Recognition Technology. Chapter in Handbook of Brain Theory and Neural Networks, 2nd edition, M. Arbig ed. MIT Press, 2002.

M. Galley, K. McKeown, J. Hirschberg and E. Shriberg. Identifying Agreement and Disagreement in Conversational Speech: Use of Bayesian Networks to Model Pragmatic Dependencies. Proceedings of 42nd Meeting of the ACL, July 21-26, Barcelona, 2004.

I. Gat-Viks, R. Shamir, R.M. Karp and R. Sharan. Reconstructing Chain Functions in Genetic Networks. In Proceedings of the Ninth Pacific Symposium on Biocomputing (PSB'04), pp. 498-509, 2004.

B. Godfrey, R.M. Karp, K. Lakshminarayanan, S. Surana and I. Stoica. Load Balancing in Dynamic Structured P2P Systems. Proceedings of INFOCOMM 2004.

J. Gramm. A polynomial-time algorithm for the matching of crossing contact-map patterns. To appear in Proceedings of the 4th International Workshop on Algorithms in Bioinformatics (WABI 2004), Bergen, Norway, September 2004.

J. Gramm, J. Guo, F. H¸ffner and R. Niedermeier. Automated generation of search tree algorithms for hard graph-modification problems. Algorithmica, 39(4):321-347, 2004.

J. Gramm, T. Nierhoff, R. Sharan and T. Tantau. On the Complexity of Haplotyping via Perfect Phylogeny. In Proceedings of the Second RECOMB Satellite Workshop on Computational Methods for SNPs and Haplotypes, pp. 35-46, 2004.

J. Gramm, T. Nierhoff and T. Tantau. Perfect Path Phylogeny Haplotyping with Missing Data is Fixed-Parameter Tractable. To appear in Proceedings of the First International Workshop on Parameterized and Exact Computation (IWPEC 2004), Bergen, Norway, September 2004.

E. Halperin and E. Eskin. Haplotype Reconstruction from Genotype Data Using Imperfect Pylogeny. Bioinformatics 20:1842-1849, 2004.

E. Halperin and R.M. Karp. Perfect Phylogeny and Haplotype Assignment. In Proceedings of RECOMB, 2004.

E. Halperin and R.M. Karp. The Minimum-Entropy Set Cover Problem. In Proceedings of ICALP 2004.

T. Hartman and R. Sharan. A 1.5-Approximation Algorithm for Sorting by Transpositions and Transversals. In Proceedings of WABI, September 2004.

M. Hebert and N. Mirghafori. Desperately Seeking Impostors: Data-Mining for Competitive Impostor Testing in a Text-Dependent Speaker Verification System. Proc. IEEE ICASSP, Montreal, May 2004.

H. Hermansky and N. Morgan. Show what you know: musings on the reporting of negative results in speech recognition research. Journal of Negative Results in Speech and Audio Sciences 2004 Issue.

D. Hillard, M. Ostendorf, A. Stolcke, Y. Liu and E. Shriberg. Improving Automatic Sentence Boundary Detection with Confusion Networks. In Proceedings of HLT-NAACL Conference, April-May 2004, Boston.

A. Janin, J. Ang, S. Bhagat, R. Dhillon, J. Edwards, J. Macias-Guarasa, N. Morgan, B. Peskin, E. Shriberg, A. Stolcke, C. Wooters and B. Wrede. The ICSI Meeting Project: Resources and Research. NIST ICASSP 2004 Meeting Recognition Workshop, Montreal, May 2004.

J. Jung, V. Paxson, A. Berger and H. Balakrishnan. Fast Portscan Detection Using Sequention Hypothesis Testing. Proceedings of IEEE Symposium on Security and Privacy, May 2004.

B. Karp, S. Ratnasamy, S. Rhea and S. Shenker. Spurring the Adoption of DHTs with OpenHash. In Third International Workshop on Peer-to-Peer Systems (IPTPS), 2004.

R.M. Karp, M. Luby and A. Shokrollahi. Finite-Length Analysis of LT-codes. In Proceedings of International Symposium on Information Theory (ISIT) 2004.

P. Kay. Pragmatic Aspects of Grammatical Constructions. In L. Horn and G. Ward, eds., Hanbook of Pragmatics, Oxford and London: Blackwell, 2004.

B.P. Kelley, B. Yuan, F. Lewitter, R. Sharan, B.R. Stockwell and T. Ideker. PathBLAST: A tool for alignment of protein interaction networks. Nucleic Acids Research 32, pp. W83-W88, 2004.

P. Levis, N. Patel, D. Culler and S. Shenker. Trickle: A Self-Regulation Algorithm for Code Propogation and Maintenance in Wireless Sensor Networks. NSDI 2004.

*Y. Liu, A. Stolcke, E. Shriberg and M. Harper. Comparing and Combining Generative and Posterior Probability Models: Some Advances in Sentence Boundary Detection in Speech. To appear in Proceedings of Conference on Empirical Methods in Natural Language Processing, Barcelona, July 2004.*

*N. Mirghafori and M. Hebert. Parameterization of the Score Threshold for a Text-Dependent Adaptive Speaker Verification System. Proceedings of IEEE ICASSP, Montreal, May 2004.*

*E. Mok, J. Bryant and J.A. Feldman. Scaling Understanding up to Mental Spaces. In Proceedings of ScaNaLU-2004, Boston, MA, 2004.*

*N. Morgan, H. Bourlard and H. Hermansky. Speech Recognition and the Auditory Perspective. Chapter in Speech Processing in the Auditory System, S. Greenberg and W. Ainsworth, eds, Springer, 2004.*

*N. Morgan, B. Y. Chen, Q. Zhu, and A. Stolcke. TRAPping Conversational Speech: Extending TRAP/Tandem approaches to conversational telephone speech recognition. Proc. IEEE ICASSP, Montreal, May 2004.*

*I. Pe'er, T. Pupko, R. Shamir and R. Sharan. Incomplete Directed Perfect Phylogeny. SIAM Journal on Computing 33(3), pp. 590-607, 2004.*

*T. Pirinen and J. Yli-Hietanen. Time delay based failure-robust direction of arrival estimation. Proceedings of IEEE SAM 2004, Sitges, Barcelona, Spain, July 2004.*

*T. Pirinen, J. Yli-Hietanen, P. Pertilä and A. Visa. Detection and compensation of sensor malfunction in time delay based direction of arrival estimation. In Proceedings of IEEE ISCAS, Vancouver, May 2004.*

*P. Radoslavov, C. Papadopoulos, R. Govindan and D. Estrin. A Comparison of Application-Level and Router-Assisted Hierarchical Schemes for Reliable Multicast. IEEE/ACM Transactions on Networking (TON), Volume 12, Issue 3, pp. 469-482, June 2004.*

*T. Regier and P. Kay. Color naming and sunlight. Psychological Science 15, 288-289, 2004.*

*E. Segal and R. Sharan. A Discriminative Model for Identifying Spatial Cis-Regulatory Modules. To appear in Proceedings of RECOMB 2004, pp. 141-149, 2004.*

*R. Shamir and R. Sharan. Optimal Fully Dynamic Algorithms for Cograph Recognition. Discrete Applied Mathematics 136, pp. 329-340, 2004.*

*R. Shamir, R. Sharan and D. Tsur. Cluster Graph Modification Problems. Discrete Applied Mathematics, 2004.*

*R. Sharan, A. Ben-Hur, G.G. Loots and I. Ovcharenko. CREME: Cis-Regulatory Module Explorer for the Human Genome. In Nucleic Acids Research 32, pp. W253-W256, 2004.*

*R. Sharan, A. Ben-Dor and Z. Yakhini. Multiplexing Schemes for Generic SNP Genotyping Assays. In Proceedings of PSB'04, pp. 140-151, 2004.*

*R. Sharan, T. Ideker, B.P. Kelley, R. Shamir and R.M. Karp. Identification of Protein Complexes by Comparative Analysis of Yeast and Bacterial Protein Interaction Data. In Proceedings of RECOMB 2004, pp. 282-289, 2004.*

*L. Shastri. From Transient Patterns to Persistent Structures: A model of episodic memory formation via cortic-hippocampal interactions. Behavioral and Brain Science, in revision, 2004.*

*E. Shriberg, R. Dhillon, S. Bhagat, J. Ang, and H. Carvey. The ICSI Meeting Recorder Dialog Act (MRDA) Corpus. Proc. HLT-NAACL SIGDIAL Workshop, Boston, April-May 2004.*

*A. Stolcke, C. Wooters, N. Mirghafori, T. Pirinen, I. Bulyko, D. Gelbart, M. Graciarena, S. Otterson, B. Peskin and M. Ostendorf. Progress in Meeting Recognition: The ICSI-SRI-UW Spring 2004 Evaluation System. In NIST ICASSP 2004 Meeting Recognition Workshop, Montreal, May 2004.*

*L. Subramanian, V. Roth, I. Stoica, S. Shenker and R. Katz. Listen and Whisper: Security Mechanisms for BGP. NSDI 2004.*

*A. Tanay, R. Sharan, M. Kupiec and R. Shamir. Revealing Modularity and Organization in the Yeast Molecular Network by Integrated Analysis of Highly Heterogeneous Genome-Wide Data. In Proceedings of the National Academy of Sciences USA 101, pp. 2981-2986, 2004.*

*T. Tantau. Comparing Verboseness for Finite Automata and Turing Machines. Theory of Computing Systems, Vol. 31, No. 1, pp. 95-109, 2004.*

*T. Tantau. A Logspace Approximation Scheme for the Shortest Path Problem for Graphs with Bounded Independence Number. In Proceedings of STACS 2004, Lecture Notes on Computer Science, Springer-Verlag, 2004.*

*M. Walfish, H. Balakrishnan and S. Shenker. Untangling the Web from DNS. NSDI 2004.*

*N. Weaver and D. Ellis. Reflections on Witty: Analyzing the Attacker. ;login: pp. 34-37, June 2004.*

*N. Weaver, D. Ellis, S. Staniford and V. Paxson. Worms vs. Perimeters: The Case for Hard-LANs. Hot Interconnects 12, August 2004.*

*N. Weaver and V. Paxson. A Worst-Case Worm. In Proceedings of Third Annual Workshop on Economics and Information Security (WEIS04), May 2004.*

*N. Weaver, S. Staniford and V. Paxson. Very Fast Containment of Scanning Worms. In Proceedings of USENIX Security Symposium, 2004.*

*S. Weber. The Success of Open Source. Cambridge, MA, Harvard University Press, 2004.*

*S. Wernicke, J. Alber, J. Gramm, J. Guo and R. Niedermeier. Avoiding Forbidden Submatrices by Row Deletions. In Proceedings of SOFSEM 04, LNCS, Vol. 2832, pp. 349-360, Springer, 2004.*

*C. Wooters, N. Mirghafori, A. Stolcke, T. Pirinen, I Bulyko, D. Gelbart, M. Graciarena, S. Otterson, B. Peskin and M. Ostendorf. The 2004 ICSI-SRI-UW Meeting Recognition System. In Proceedings of the Joine AMI/Pascal/IM2/M4 Workshop on Meeting Recognition, 2004.*

*E.P. Xing and R.M. Karp. MotifPrototyper: A Bayesian Profile Model for Motif Families. In Proceedings of the National Academy of Sciences, 2004.*

*E.P. Xing, R. Sharan and M.I. Jordan. Bayesian Haplotype Inference via the Dirichlet Process. In Proceedings of the Second RECOMB Satellite Workshop on Computational Methods for SNP and Haplotypes, pp. 99-112, 2004. Also in Proceedings of ICML 2004.*


INTERNATIONAL
COMPUTER SCIENCE
INSTITUTE

**THE INTERNATIONAL COMPUTER SCIENCE INSTITUTE (ICSI)** is the only independent, non-profit US lab solely conducting open, non-proprietary, pre-competitive research in computer science. Affiliated with the University of California campus in Berkeley, ICSI's mission is to further research in computer science through international collaboration, and further international collaboration through research in computer science. ICSI provides a haven for computer scientists to conduct long-term research without commercial limitations and with few faculty pressures. ICSI has significant efforts in four major research areas: Internet research, including Internet architecture, related theoretical questions, and network services and applications; theoretical computer science, including applications to bioinformatics; artificial intelligence, particularly for applications to natural language understanding; and natural speech processing.




**The ICSI Gazette** is published twice a year by the International Computer Science Institute (ICSI).




### Editor-in-Chief
Nelson Morgan

### Senior Editor
Maria Quintana

### Design
envision ID

### Layout
Leah Hitchcock

### Copywriter
Leah Hitchcock

### Photography
Diane Starr
Leah Hitchcock



INTERNATIONAL
COMPUTER
SCIENCE
INSTITUTE

1947 Center Street Suite 600
Berkeley, CA 94704 USA
v 510 666 2900
www.icsi.berkeley.edu