

The ICSI GAZETTE

volume seven | issue one | september 2008

icsi celebrates its 20th anniversary



INTERNATIONAL
COMPUTER SCIENCE
INSTITUTE

The International Computer Science Institute was inaugurated twenty years ago this September. ICSI was created to allow for international cooperation and advanced research in computer science to benefit from each other. Over these twenty years, ICSI has made many significant contributions to the field of computer science, and has hosted over 500 outstanding researchers with its international programs.

ICSI began in 1985 with a conversation between Professor Norbert Szyperski, director of the German Federal Laboratory for Computer Science (GMD) and Ron Kay, a senior manager at IBM's San Jose Research Laboratory. They set out to create an international research center for computer science which could serve as a "scientist's playground." With the help of UC Berkeley Professors Domenico Ferrari and Richard Karp, ICSI soon found itself a home three blocks from the Berkeley campus. The German visitor program was quickly augmented by similar visitor programs with both Switzerland and Italy, and later programs from Finland and Spain as well.

A testament to ICSI's continued long-term success is how many of its original group leaders, researchers, and trustees have stayed on board since its inception. Many of ICSI's founders are still actively involved here. Three of the four original research groups are still guided by their original leaders; moreover, these group leaders have trained new generations of researchers that are paving new ground at ICSI.

AI

The AI group was founded by ICSI's first director, Jerry Feldman, who is still an active member of the group. As the work done in the group, formerly called Applications, changes domains, the group is thematically tied together by Feldman's original work on neural computing. This work focused around neural connectionist models, particularly in regards to language modeling, spawned the L_0 project in 1990 which developed into the still active Neural Theory of Language (NTL) project.

Feldman's work on connectionist models of language in NTL was supported by graduate student Srinu Narayanan, who would go on to enjoy a postdoc appointment at ICSI and a research position at partner institution SRI before he returned as a senior researcher and eventual AI group leader.

While Narayanan continues to employ neural modeling of language acquisition and use in the NTL project, he also takes techniques central to neural connectionist work and applies them to new domains. He is developing interactive intelligent question-answering systems for use in rural India. Also, working with doctoral student Joe Makin, the connectionist techniques have expanded beyond neural modeling to modeling biological processes such as the blood clotting mechanism.

in this issue:

page 2
As I See It
by Nelson Morgan, Director

page 3
News Briefs

page 7
Featured Alumni

page 9
Visiting Scholars

page 10
Publications

Continued on Page 4



as i see it by Nelson Morgan, Director

The past is never dead. It's not even past.
-- William Faulkner

This issue is a singular one, as it marks the 20th anniversary of ICSI's official inauguration. ICSI was incorporated in June 1986, but by September 2008 staffing had reached critical mass and we formally launched our enterprise with an event. This October we will celebrate this anniversary with another event, which I'll describe below; but first, a bit of reminiscing.

ICSI was the brainchild of Norbert Szyperski, Managing Director of the German National Center for Computer Science (GMD) in Germany, in collaboration with Ron Kay, formerly from IBM in the U.S. Prof. (Dr.) Szyperski felt that he had benefited greatly from the year he spent in the U.S. on an Eisenhower Fellowship (an early visit), and also saw that German computer science was poorly represented in the American scene. The two colleagues saw an opportunity to start something of great significance for the advancement of computer science and international collaboration. And so, after an exploratory phase leading to the decision to affiliate with UC Berkeley, they founded ICSI (with critical assistance from a number of American and German colleagues).

Twenty years ago, ICSI was significantly smaller; we started on a single floor at 1947 Center Street, but within a few years we expanded to a second floor. At first we relied exclusively on our international program, but soon our funds were supplemented by domestic peer-reviewed grants. We had roughly 40 people then, including administrative staff, and now we have about 130. Half of the Institute was organized around a single coherent theme (massive parallelism), while now the focus has morphed over the years based on the interests of senior ICSI researchers.

Despite all of these changes, ICSI has not lost its idiosyncratic flair. While our research and the corresponding sources of support have diversified greatly, we have just as much of an international environment – we have at least as many visitors from at least as many countries as we did 20 years ago. A few of the same people are still here (Jerry Feldman, Dick Karp, and myself), but after 20 years there are many new faces. ICSI staffers currently include a number of people who were here as graduate students in earlier

years (Sally Floyd, Vern Paxson, Srinu Narayanan, Krste Asanovic, Andreas Stolcke). ICSI's mission, furthering computer science research through international collaboration (and furthering international collaboration through computer science research) has remained the same. And while we have a diverse range of topics, the basic program of theory and application of algorithms to machine perception, cognition, and communication (networking) has remained more or less the same.

Perhaps most importantly, ICSI really has much of the same “feel” as it did 20 years ago. It is part of campus, and yet is separate from it. It is (usually) an island of calm. We're still having “teas” twice a week, and we still have a highly effective administrative staff. Space is still comparatively ample (although look out for those plans to put in lofts to squeeze in more people!) We still have some of the most accomplished researchers in each of their respective fields, great students and postdocs, a cosmopolitan environment, and we still have strong bonds to a number of far-flung labs and researchers. And over these years, there have been many notable achievements, including: the world's first single chip vector microprocessor; seminal work on object-oriented software that had a significant impact on Java; speech processing algorithms incorporated in millions of cell phones; and one of the first and most complete solutions to provide performance guarantees for real-time Internet traffic; Tornado codes, which allow any number of heterogeneous clients to acquire bulk data with optimal efficiency at times of their choosing, even in the face of high channel loss rates; and the eXtensible Open Router Platform (XORP), an innovative and robust open-source routing platform that (we hope) will break open the closed nature of networking hardware. And individual ICSI researchers have won extremely prestigious awards throughout this time, as described in the lead article in this issue.

It is personally extremely gratifying that, not only have we survived 20 years of significant peaks and valleys in each of our support mechanisms, but that we have done so without significantly sacrificing our research lifestyle. ICSI research is still overwhelmingly investigator-driven, despite the many pressures to become more of a contract research house.

This year, we are well along our path to further expansion of our international visitor program, one of the keys to our continued success. In April, we worked

with our Brazilian colleagues to begin to set up a new partnership, one that we hope will be formalized by the end of the year. And a new extension of our German visitor program was announced as well.

We will celebrate this 20th birthday on October 17. The day will include talks from ICSI's founders, feedback from representatives of new and continuing sponsor countries, and keynote lectures by a few of the accomplished ICSI researchers, past and present (and sometimes both!). The first of these will be Professor Dick Karp. Dick was here at the start, and except for a few years when he apparently was lost in Seattle, has been here since. A second talk will

be from Professor Vern Paxson, who leads ICSI's activity on internet security, and who was also part of Domenico Ferrari's networking research group during ICSI's early days. Ben Gomes, an early Google employee, who received his PhD working with Jerry Feldman here at ICSI, will give a third. The final talk will be given by Krste Asanovic, who received his PhD working at ICSI, became an MIT professor, and ultimately returned to a dual role as an ICSI staff member and a UCB professor.

And what about the next 20 years? I'm putting my bet on ICSI becoming an even better place to be and work. Any takers?

news briefs

ICSI's **KRSTE ASANOVIC** is working to set up the Universal Parallel Computing Research Center (or ParLab) at UCB, which was featured in the March 19th New York Times article [Industry Giants Try to Break Computing's Dead End](#).

Professor **RICHARD M. KARP**, Head of the Algorithms Group, received the 2008 Kyoto Prize in the Advanced Technology Category for "fundamental contributions to the development of the theory of computational complexity". More information on page 8.

Two new research groups were formed in the summer of 2008. Professor **KRSTE ASANOVIC**, a former ICSI graduate student researcher and frequent visitor, is heading the new Architecture Group. Professor **TREVOR DARRELL**, formerly of MIT, is heading the new Vision Group.

Professor **KRSTE ASANOVIC**, head of the Architecture Group, is quoted in an article called "Taking advantage of multicore PCs" on [InfoWorld.com](#) about parallel processing. Asanovic is an expert on this new technology, and advises programmers to be cautious in their

choice of programming language for parallel programming. Because the technology is so new, there isn't yet an industry standard programming platform, so many of the existing platforms may become obsolete.



ICSI alum **CHRIS BREGLER** and Emeritus Trustee **JITENDRA MALIK** have won the 2008 Longuet-Higgins prize for "Fundamental Contributions in Computer Vision That have Stood the Test of Time".

ICSI announced that it has transferred its eXtensible Open Router Platform (**XORP**) technology to [XORP, Inc.](#), a startup founded by the leaders of the [XORP.org](#) project. The same day, XORP, Inc. debuted the XORP 1.5 Release.

ICSI had some excitement this summer **UPGRADING THE MACHINE ROOM** by installing a new air-conditioning unit. The installation involved drilling through the ceiling, and lifting the a/c unit in by crane.

Dr. **EERO SILVENNOINEN** of Tekes joined ICSI's Board of Trustees in June 2008. Dr. Silvennoinen is Director of Software and Telecommunications Technologies at Tekes, and has been at Tekes since 1988. He replaced Mr. Jouko Salo as the Finnish representative to ICSI's Board.



Congratulations to Eran Halperin, his wife Leticia, and their son Yanai on the birth of Helio Halperin Ortiz on September 27th at 9:17 a.m.

icsi's 20th anniversary

Continued from Page 1

The AI group's focus expanded significantly when prominent linguists Charles J. Fillmore and Paul Kay retired from their positions in the UCB Linguistics department and took up residence at ICSI. Fillmore's work with former graduate student and current PI Collin Baker has produced FrameNet, a rich online database of lexical semantics. Kay has worked extensively with fellow ICSI linguist and former graduate student Terry Regier on linguistic patterns derived from vision, including color naming and lateralization of the Whorfian hypothesis.

What started with Jerry Feldman's investigation into "brain-like computing" has evolved with the improvement of understanding of real neural activity. The AI group has moved past trying to simulate neural actions, towards recreating the results of actual neural activity. This work has been readily embraced by three generations of researchers at ICSI and continues to expand; much of the NTL group's work is explained in Feldman's recent book [From Molecule to Metaphor](#).

ALGORITHMS

Just as the AI group retains Feldman, the Algorithms group is still led by one of its founders, Turing Award winner Richard Karp. Considering Karp's joint appointments in the Computer Science, Mathematics, Bioengineering, and Operations Research Departments on campus, it shouldn't be surprising that the Algorithms group has addressed problems from a wide array of disciplines.

Karp has been involved with ICSI since the mid-1980's when he helped to write Berkeley's winning proposal to be the host site for ICSI, and he held a research appointment when ICSI was formed in 1988. By that time, he had already graduated two doctoral students who would play major roles at ICSI, Michael Luby and Ron Shamir. Luby would go on to lead the group (then called Theory of Computing) for a portion of the 90's while Karp worked at the University of Washington. Luby pioneered Digital Fountain, possibly the group's most significant achievement; it is a revolutionary way of encoding large files for transmission over the Internet that has the miraculous

property that the receiver only needs to receive some of the packets to decode and receive the entire file. Ron Shamir has been a long-time collaborator with Karp, and their work illustrates ICSI's ongoing partnership with Tel-Aviv University. ICSI and Tel-Aviv have developed an informal researcher exchange, with Shamir's doctoral students Roded Sharan and Gad Kimmel recently enjoying appointments at ICSI, and with PI Eran Halperin beginning to split his time between ICSI and Tel-Aviv. A major outcome of this collaboration is highly fruitful work on human genetic variation - a field hailed as Science Magazine's Breakthrough of the Year in 2007.

In his time at ICSI, Karp also advised the Ph.D. of Sally Floyd, who would return to ICSI to be a PI in the Networking group. Karp and the Networking group collaborated extensively upon his return with the formation of ACIRI, explained below, during the late 90's and early 2000's. This collaboration underscores the interdisciplinary nature of Karp's work; he explains that "if you peel back the layers and look at the core of a problem, it's almost always a simple, concise algorithm."

It is this willingness to innovatively explore new areas that has led the Algorithms group to confront

problems as assorted as congestion control in computer networks, peer-to-peer networking, file encoding and transmission, genetic variation within disease populations, randomized algorithms, game theory, and computational finance. Karp's tradition of innovation and collaboration recently

made him the recipient of the prestigious Kyoto prize, described later in this newsletter.

NETWORKING

ICSI's work in large distributed computer systems began with the Tenet group under the supervision of Domenico Ferrari. Ferrari, then Chair of the Computer Science department on campus, was another member of UCB's faculty who "immediately loved

ICSI has made significant contributions to the field of Computer Science, and has hosted over 500 outstanding researchers with its international programs.

the idea” of ICSI’s creation when Ron Kay initially discussed it with him in the mid-1980’s. His hearty support played a major role in ICSI’s establishment. Accordingly, Ferrari served as the first deputy director of ICSI, and founded the Tenet group, heading it well into the 90’s. Under his direction, the group designed and built two real-time protocol suites. His “very deep belief in the internationality of science” is evidenced by Ferrari’s successful efforts to establish one of ICSI’s first visitor programs, in particular with Italy.

In 1995 Ferrari retired and relocated to Italy to head an institute at the Università Cattolica. In 1999, Scott Shenker secured generous funding from AT&T that revitalized the group with the creation of ACIRI, the AT&T Center for Internet Research at ICSI. This enabled researchers to work on Internet issues as they saw fit, free from the usual time- and topic- based constraints so often encountered in both academia and industry. The promise of unfettered research allowed a “dream team” of Internet researchers to gather at ICSI, bringing together Shenker, Vern Paxson, Sally Floyd, Mark Handley, and Richard Karp (returning to ICSI from Washington). Among many other research developments, the Networking group gave birth to the eXtensible Open Router Platform XORP, which has recently spun out into a new startup company. After AT&T’s funding was discontinued, the group was rechristened as ICIR, the ICSI Center for Internet Research.

The free-thinking innovation of ICSI’s Networking group is apparent by the attention they’ve received. The SIGCOMM Award - the highest honor in the field of networking research - has gone to Shenker in 2002, Ferrari in 2006, and Floyd in 2007. Paxson received SIGCOMM’s Test of Time award in 2006, and the ACM’s Grace Murray Hopper award in 2007. All are ACM Fellows. The Networking group has contributed to the fields of Internet architecture, traffic, congestion, measurement, and security. The group’s efforts in Internet security, guided by Paxson, were the subject of the [last issue of the ICSI newsletter](#).

SPEECH

ICSI’s Speech group is another group that is still being led by its founder, current ICSI director Prof.

Nelson Morgan. Initially designated “Realization of Massively Parallel Systems,” the group’s early endeavors spawned a number of notable hardware innovations. One example is the development of the Ring Array Processor (RAP) in 1989, a high-speed parallel computing processor. The RAP arose out of the group’s desire to do large-scale computation with neural nets; no hardware existed that was well-suited to the job, so they invented one. That same problem-solving resourcefulness led to the development of a vector microprocessor system comprised of the microprocessor T0, designed by then-graduate student Krste Asanovic and others, and SPERT-II, designed by James Beck. All while creating innovative hardware, the group was using their new inventions to tackle the multi-faceted problems facing computer speech recognition.

The late 1990’s saw the rate of improvement of commercial machines rapidly increase. It reached a point where after taking the requisite years to develop and implement new special-purpose hardware, the resultant hardware may only have been marginally better than mainstream developments. This obviated the group’s need to build faster machines to perform their work, and permitted them to refocus primarily on the use of neural nets in speech processing. Accordingly, their name changed from the Realization group to the Speech group.

This shift coincided with the graduation of Asanovic, who would return from MIT a few years later to start the Architecture group (see below). Other early members of the Speech group have also played very long-term roles in ICSI’s development. Chuck Wooters (ICSI’s first graduate student) returned from his work in industry to enjoy a seven-year stint as a staff researcher working on speaker diarization. David Johnson, a former software engineer for the group’s processors, is ICSI’s current Systems Manager.

The focus of the group over the last decade has been speech recognition, speaker recognition, and speaker diarization: what is being said, who is saying it, and when is each person speaking. Senior researchers in the ICSI Speech Group, including Morgan and new member Dilek Hakkani-Tur, hold a number of speech-recognition patents, one of which may be implemented on your current cell phone.

Continued on Page 6

icsi's 20th anniversary

Continued from Page 5

NEW GROUPS

Krste Asanovic recently returned to ICSI in 2005 to revive hardware development by forming the Architecture group. His research focus is energy-efficient high-performance computing, building upon his early work with the Realization group and later projects at MIT. Asanovic is also closely involved with UCB's Par Lab, which is investigating ways to simplify the challenging task of designing parallel software for the new wave of multicore architecture. He is also exploring the use of silicon photonics to provide sufficient memory bandwidth to feed the higher demands of future manycore systems.

Another research area that has seen a recent revival at ICSI is computer vision. Work in this area began in the 1980's as part of the AI group with Steve Omohundro. Omohundro, a staff scientist at ICSI for its first six years, synthesized work done from a number of different disciplines. Utilizing machine-learning techniques, he co-designed computer vision methods to augment speech recognition. In 2008, a decade after ICSI wrapped up its initial efforts in computer vision, Prof. Trevor Darrell joined ICSI and the UC EECS faculty, starting a Vision group. Specific research areas include object and activity recognition, with application to human-robot interaction, visual surveillance, and multimodal human computer interfaces.

... AND BEYOND!

ICSI continues to benefit from the wide array of past and present visitors, not only those who have remained to guide its day-to-day activities, but

also those who are present in other capacities. ICSI founder Ron Kay is still a frequent visitor and familiar face at ICSI talks and conferences. Hervé Bourlard, ICSI's very first postdoctoral researcher and current Board member, is another example of ICSI reaping long-term benefits from its training of emerging scientists. As Director of IDIAP and IM2, Bourlard oversees the Swiss visitor program and matches ICSI with many promising graduate students and postdocs. Former graduate student and postdoc Dan Jurafsky is still in the Bay Area (on the Stanford faculty) and is a frequent collaborator with the AI and Speech

groups. Terry Regier, mentioned above with AI, is currently taking a one-year sabbatical at ICSI.

In its time, ICSI has opened its doors – both physical and electronic – to over two thousand researchers, international visitors, students, and staff members.

The more interconnected that science becomes, the more ICSI has to offer because of its unique environment promoting collaboration in innovative, international, open computer science research.

The rate at which ICSI retains those researchers, and the rate at which they are recognized for their work, speaks volumes about the quality of ICSI's research environment. Having made remarkable achievements in international macro-collaboration, ICSI is also looking to foster interdisciplinary micro-collaboration within its groups. Speaker recognition can benefit from intelligent computer vision; speech recognition can be augmented by a computer that understands the context of a discussion and can intelligently understand distorted speech; networking can benefit from the development of robust algorithms. The more interconnected that science becomes, the more ICSI has to offer because of its unique environment promoting collaboration in innovative, international, open computer science research.

featured alum: eddie kohler

Professor Eddie Kohler is ICSI's featured alum for this issue. Kohler received his Ph.D. in Computer Science from MIT in 2000, and then joined ICSI as a postdoc in the Networking group. Much of his early work focused on DCCP, the Datagram Congestion Control Protocol. Kohler's work on the congestion control mechanisms led him



ICSI Alumnus Eddie Kohler

to be the project's main architect during his three-year stay at ICSI, which was initially engineered by Sally Floyd and Mark Handley. While the TCP protocol is standard for situations where all data must be received, DCCP addresses the problem with unreliable transfer in time-sensitive settings like streaming media or Internet telephony. DCCP was first implemented in a Linux kernel release in 2005.

In 2004, Kohler accepted a position as Assistant Professor at UCLA's Computer Science Department. While there, he was instrumental in the development of Asbestos. Asbestos is a prototype operating system designed to limit and contain the effects of exploitable software flaws, such as private data security. Asbestos is designed to allow one process to act on behalf of multiple users without compromising the private data of any individual user. The way that Asbestos labels and processes information is also embodied in the related Asbestos web server, which correspondingly isolates individual users' data among Web application traffic.

Kohler's longest-running project, though, is the Click modular router. The design of Click comprised Kohler's thesis work at MIT, and the implementation of Click has been a focus of his work ever since. Click

is used in the framework for XORP, the eXtensible Open Router Platform which was conceived of at ICSI. XORP has the potential to do to computer-networking equipment what Linux has done to operating systems.

Kohler has recently returned to the Bay Area for another new utilization of Click. While Google and Earthlink have received press in recent years by promising to bring free wireless Internet to San Francisco, these plans have not yet seen fruition. This is where Meraki, the start-up where Kohler recently began working, comes in. They are a grassroots wireless networking company that is currently distributing repeater nodes to broaden their free wireless network currently active in the city's Mission District. Their work has been acknowledged by the [New York Times](#), the [Wall Street Journal](#), and many [other major news organizations](#).

Kohler himself has been the recipient of some very favorable attention as well. Receiving a [PECASE](#) (Presidential Early Career Award for Scientists and Engineers) from NSF in 2007 led him to be described as "wonderfully creative and talented" by UCLA Engineering Dean Vijay K. Dhir. Kohler's work on Asbestos was recognized with Technology Review's [TR35 Young Innovator](#) award for 2006, and he also received a 2007 [Sloan Research Fellowship](#).

Kohler's work overseeing the implementation of DCCP often brings him back to ICSI to work with Sally Floyd, 2007 winner of the prestigious SIGCOMM award. Floyd, a long-time collaborator, says that "working with Eddie has been lovely because he brings insight and energy to the collaboration." Networking group leader Scott Shenker describes him as "one of those rare individuals who follows his own internal compass, setting trends rather than following them. As

a result, Eddie's research doesn't just add to our body of knowledge, it also significantly broadens our perspective. He is a gift to our field."

Kohler reflects fondly on his three-

year stay at ICSI and enjoys his return visits because "it provides a home for people who are great researchers to come and do work that doesn't necessarily fit into academia."

"Eddie's research doesn't just add to our body of knowledge, it also significantly broadens our perspective. He is a gift to our field."

richard karp wins kyoto prize



ICSI is delighted to announce that Professor Richard M. Karp has been awarded the prestigious 2008 Kyoto Prize for Advanced Technology, in the field of Information Sciences. The prize,

given by the Inamori Foundation in Japan, is sometimes referred to as the “Japanese Nobel”, and is one of the highest honors a scientist can receive.

The Inamori Foundation only issues an award for information technology every four years. [In recognition of Karp’s achievements, the Inamori Foundation explains](#) that “Dr. Richard M. Karp has had a profound influence on the guiding principles for analysis and design of algorithms which are being used for a broad range of applications today by establishing the theory of NP-completeness,” and that “the magnitude of his fundamental contributions to the development of the theory of computational complexity is immeasurable, and holds potential for even further development beyond the framework of information science.”

Karp has been an integral part of the ICSI community since our inauguration. Not only was he instrumental in seeing that Berkeley would become ICSI’s home, but he founded the Algorithms group, and still serves as both the group leader and a long-time member of ICSI’s

Board of Trustees. But beyond the specific accomplishments, Karp embodies the spirit of ICSI at its finest. He is a brilliant and innovative researcher at the leading edge of his field; moreover, he is a gracious leader and team-builder.

ICSI’s Director, Professor Nelson Morgan, says “we are thrilled to have this valued member of our community receive this richly deserved award. His previous work was a key part of the development of the field of computer science, and he continues to generate important results, currently in bioinformatics and computational biology. We’re honored to have him in our community.”

Professor Shankar Sastry commented “The Kyoto Prize recognizes not only outstanding achievements, but also the stellar personal characteristics that have shaped those achievements. I can imagine no more appropriate recipient for this honor than Dick Karp.” Sastry is Chairman of ICSI’s Board and UC Berkeley’s Dean of Engineering.

On July 14, a ceremony was held at UC Berkeley’s Hearst Memorial Mining Building in congratulation of Karp’s achievements. The ceremony was emceed by Dean Sastry, and included comments by UC Berkeley Chancellor Robert J. Birgeneau.

Karp’s other notable achievements include a Turing Award, the U.S. National Medal of Science, and eight honorary doctorates.



Chancellor Robert J. Birgeneau

visiting scholars

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

FROM CHINA

Bin Dai (Networking)
Bo Xu (Networking)

FROM EUROPEAN UNION (AMIDA)

Joan Isaac Biel (Speech)
Thomas Kleinbauer (AI-FrameNet)
Rosemary Orr (Speech)
Korbinian Riedhammer (Speech)
David VanLeeuwen (Speech)
Oriol Vinyals (Speech)

FROM FINLAND

Yoshia Hirase (Industrial)
Jyri Kivinen (Algorithms/Campus Affiliation)
Tommi Lampikoski (Campus Affiliation)
Annukka Näyhä (Campus Affiliation)
Boris Nechaev (Networking)
Teemu Roos (Campus Affiliation)
Ville-Pekka Seppä (Campus Affiliation)

FROM GERMANY

Jan Baumbach (Algorithms)
Gerald Friedland (Speech)
Tobias Friedrich (Algorithms)
Martin Gairing (Algorithms)
Martin Hilpert (AI-FrameNet)
Birte Lönneker-Rodman (AI-FrameNet)
Andreas Maletti (AI)
Christian Müller (Speech)
Andreas Raabe (Architecture)
Ulrich Rueckert (Algorithms)
Felix Salfner (Industrial)
Thomas Sauerwald (Algorithms)
Guido Schryen (Networking)

FROM SPAIN

Oscar Ferrandez (AI)
Eduardo Lopez (Speech)
Carlos Subirats (AI-FrameNet)
Enrique Torres (Architecture)

FROM SWITZERLAND (IM2)

Nikhil Garg (Speech)
Bao-Lan Huynh (Speech)
David Imseng (Speech)
Adish Singla (Speech)

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

AI

Alberto Amengual

ALGORITHMS

Lucia Conde
Ron Shamir
Roded Sharan

ARCHITECTURE

Chris Batten
Rose Liu
Heidi Pan

NETWORKING INTERNS

Gregor Maier

FRAME.NET

Martina Johnson
Alessandro Lenci
Kyoko Ohara
Hiroaki Sato

NETWORKING

Leo Juan
Chih-Hung Lin
Po-Ching Lin
Jamon Liu

SPEECH

Umit Guz
Petr Motlicek



Bao-Lan Huynh



Bogdan Posaniuc



Oscar Ferrandez

publications listing

- M. ALLMAN, C. KREIBICH, V. PAXSON, R. SOMMER, AND N. WEAVER.** *Principles for Developing Comprehensive Network Visibility. Proceedings of USENIX Workshop on Hot Topics in Security, San Jose, CA, 29 July 2008.*
- L. BARRETT, J. FELDMAN, AND L. MACDERMED.** *A (Somewhat) New Solution to the Binding Problem. Neural Computation, Vol. 20, Issue 9, pp. 2361–2378, Sept. 2008 (previously published as ICSI Technical Report in 2006).*
- L. BARRETT AND S. NARAYANAN.** *Learning All Optimal Policies with Multiple Criteria. Proceedings of International Conference on Machine Learning (ICML), pp. 41–47, Helsinki, Finland, 5–9 July 2008*
- C. BATTEN, A. JOSHI, J. ORCUTT, A. KHILO, B. MOSS, C. HOLZWARTH, M. POPOVIC, H. LI, H. SMITH, J. HOYT, F. KÄRTNER, R. RAM, V. STOJANOVIC, AND K. ASANOVIC.** *Building Manycore Processor-to-DRAM Networks with Monolithic Silicon Photonics. Proceedings of Annual Symposium on High-performance Interconnects (Hot Interconnects 2008), Stanford, California, August 2008.*
- K. BOAKYE, G. FRIEDLAND, AND O. VINYALS.** *Two's a Crowd: Improving Speaker Diarization by Automatically Identifying and Excluding Overlapped Speech Authors. To appear in Proceedings of Interspeech 2008, Brisbane, Australia, September 2008.*
- K. BOAKYE, G. FRIEDLAND, B. TRUEBA-HORNERO, AND O. VINYALS.** *Overlapped Speech Detection for Improved Speaker Diarization in Multiparty Meetings. Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Las Vegas, Nevada, pp. 4353–4356, April 2008.*
- B.-G. CHUN, E. KOHLER, AND S. RATNASAMY.** *NetComplex: A Complexity Metric for Network System Designs. Proceedings of USENIX Symposium on Networked Systems Design and Implementation (NSDI 2008), San Francisco, CA, pp. 393–406, April 2008.*
- B.-G. CHUN, P. MANIATIS, AND S. SHENKER.** *Diverse Replication for Single-Machine Byzantine-Fault Tolerance. To appear in proceedings of 2008 USENIX Annual Technical Conference, Boston, MA, pp. 287–292, June 2008.*
- E. W. DE LUCA AND B. LONNEKER-RODMAN.** *Integrating Metaphor Information into RDF/OWL EuroWordNet. Proceedings of International Conference on Language Resources and Evaluation (LREC-2008), Marrakech, Morocco, 28–30 May 2008.*
- H. DREGER, A. FELDMANN, V. PAXSON, AND R. SOMMER.** *Predicting the Resource Consumption of Network Intrusion Detection Systems. To appear in proceedings of International Symposium on Recent Advances in Intrusion Detection, Boston, MA, September 2008*
- A. FARIA AND N. MORGAN.** *Corrected Tandem Features for Acoustic Model Training. Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Las Vegas, Nevada, pp. 4734–4740, April 2008.*
- A. FARIA AND N. MORGAN.** *When a Mismatch Can Be Good: Large Vocabulary Speech Recognition Trained with Idealized Tandem Features. Proceedings of the ACM Symposium on Applied Computing, Fortaleza, Brazil, pp. 1574–1577, March 2008.*
- B. FAVRE, R. GRISHMAN, D. HILLARD, H. JI, D. HAKKANI-TUR, AND M. OSTENDORF.** *Punctuating Speech For Information Extraction. Proceedings of IEEE ICASSP, Las Vegas, NV, April 2008.*
- L. FERRER, M. GRACIARENA, E. SHRIEBERG, AND A. ZYMNIS.** *System Combination Using Auxiliary Information for Speaker Verification. Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Las Vegas, Nevada, pp. 4853–4856, April 2008.*
- C. J. FILLMORE.** *The Merging of "Frames". Book Chapter in Frames, Corpora, and Knowledge Representation. R. Rossini Favretti, Ed., Bononia University Press, pp. 1–12, 2008.*
- S. FLOYD AND M. ALLMAN.** *RFC 5290: Comments on the Usefulness of Simple Best-effort Traffic. Request for Comments 5290, Informational, July 2008.*
- G. FRIEDLAND, E. HENSLEY, R. JAIN, AND J. SCHUMACHER.** *Appscio: A Software Environment for Semantic Multimedia Analysis. Proceedings of IEEE Conference on Semantic Computing, Santa Clara, CA, pp. 456–459, August 2008.*
- G. FRIEDLAND, W. HUERST, AND L. KNIPPING.** *Educational Multimedia. Special Section in IEEE Multimedia, v. 15, n. 3, pp. 54–74, July–September 2008.*
- G. FRIEDLAND, W. HUERST, AND L. KNIPPING.** *Multimedia Education in Computer Science -- A Little Bit of Everything is Not Enough. IEEE Multimedia Magazine, v. 15, n. 2, pp. 78–82, April–June 2008*
- G. FRIEDLAND, W. HUERST, L. KNIPPING, AND M. MUHLHAUSER.** *Can We Escape the Trough of Disillusionment?—A Perspective on E-Learning Technology Research from the ACM Workshop on Educational Multimedia and Multimedia Education. To appear in ACM E-learn Journal, 2008.*
- D. GILLICK, D. HAKKANI-TUR, AND M. LEVIT.** *Unsupervised Learning of Edit Parameters for Matching Name Variants. Proceedings of International Conference on Spoken Language Processing (Interspeech), Brisbane, Australia, 22–26 September 2008.*
- M. HAMPTON AND K. ASANOVIC.** *Compiling for Vector-Thread Architectures. Proceedings of International Symposium on Code Generation and Optimization (CGO-2008), pp. 205–215, Boston, Massachusetts, April 2008.*
- M. HILPERT.** *The English comparative - language structure and language use. English Language and Linguistics, 2008.*
- M. HILPERT AND S. GRIES.** *The Identification of Stages in Diachronic Data: Variability-Based Neighbor Clustering. Corpora, issue 3, volume 1, 2008.*
- H. HUNG, G. FRIEDLAND, D. GATICA-PEREZ, AND Y. HUANG.** *Estimating the Dominant Person in Multi-party Conversations Using Speaker Diarization Strategies. Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Las Vegas, Nevada, pp. 2197–2200, April 2008.*
- L. JUAN, C. KREIBICH, C.-H. LIN, AND V. PAXSON.** *A Tool for Offline and Live Testing of Evasion Resilience in Network Intrusion Detection Systems. Proceedings of GI International Conference on Detection of Intrusions and Malware and Vulnerability Assessment, Paris, France, pp. 267–278, July 2008*

- K. KAMANGAR, D. HAKKANI-TÜR, G. TÜR, AND M. LEVIT.** *An Iterative Unsupervised Learning Method for Information Distillation. Proceedings of IEEE ICASSP, Las Vegas, NV, April 2008.*
- M. KNOX.** *Automatic Laughter Segmentation. Master's report, 22 May 2008.*
- M. KNOX, N. MIRGHAFORI, AND N. MORGAN.** *Getting the Last Laugh: Automatic Laughter Segmentation in Meetings. To appear in Proceedings of Interspeech 2008, Brisbane, Australia, September 2008.*
- R. KRASHINSKY, C. BATTEN, AND K. ASANOVIC.** *Implementing the Scale Vector-thread Processor. ACM Transactions on Design Automation of Electronic Systems (TODAES), Vol. 13, Issue 3, pp. 41:1-41:24, July 2008.*
- C. KREIBICH, B. ENRIGHT, C. KANICH, K. LEVCHENKO, V. PAXSON, S. SAVAGE, AND G. VOELKER.** *On the Spam Campaign Trail. Proceedings of USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET), San Francisco, CA, April 2008.*
- J. W. LEE, M. C. NG, AND K. ASANOVIC.** *Globally-synchronized Frames for Guaranteed Quality-of-service in On-chip Networks. Proceedings of International Symposium on Computer Architecture, Beijing, China, June 2008.*
- H. LEI AND D. V. LEEUWEN.** *ICSI System Description for SRE2008 Submission. Speaker Recognition Evaluation 2008, National Institute of Standards and Technology.*
- H. LEI AND N. MIRGHAFORI.** *Comparisons of Recent Speaker Recognition Approaches based on Word Conditioning. Proceedings of Odyssey 2008, Stellenbosch, South Africa, 2008.*
- B. LOENNEKER-RODMAN.** *The Hamburg Metaphor Database Project: Issues in Resource Creation. Language Resources and Evaluation, 2008.*
- G. MAIER, H. DREGER, A. FELDMANN, V. PAXSON, F. SCHNEIDER, AND R. SOMMER.** *Enriching Network Security Analysis with Time Travel. Proceedings of ACM Special Interest Group in Data Communications (SIGCOMM), Seattle, WA, pp. 183-194, August 2008.*
- A. MALETTI.** *Minimizing Deterministic Weighted Tree Automata. Proceedings of 2nd International Conference on Language and Automata: Theory and Applications, 2008.*
- A. MALETTI.** *Myhill Nerode Theorem for Recognizable Tree Series --- Revisited. Proceedings of 8th International Conference on Latin American Theoretical Informatics, 2008.*
- C. MÜLLER, J. BIEL, E. KIM, AND D. ROSARIO.** *Speech-Overlapped Acoustic Event Detection for Automotive Applications. To appear in Proceedings of Interspeech 2008, Brisbane, Australia, September 2008.*
- M. PETRUCK.** *Framing Motion in Hebrew and English. Book Chapter in Frames, Corpora, and Knowledge Representation. R. Rossini Favretti, Ed., Bononia University Press, pp. 43-52, 2008.*
- K. RIEDHAMMER, D. GILLUCK, B. FAVRE, AND D. HAKKANI-TÜR.** *Packing the Meeting Summarization Knapsack. Proceedings of Interspeech 2008, Brisbane, Australia, September 2008.*
- S. SANKARARAMAN, G. KIMMEL, E. HALPERIN, AND M. I. JORDAN.** *On the Inference of Ancestries in Admixed Populations. Genome Research, Vol. 18, pp. 668-675, 2008.*
- A. SINGLA AND D. HAKKANI-TÜR.** *Cross-Lingual Sentence Extraction for Information Distillation. To appear in Proceedings of Interspeech 2008, Brisbane, Australia, September 2008.*
- C. F. SKIBOLA, P. M. BRACCI, E. HALPERIN, A. NIETERS, A. HUBBARD, R. A. PAYNTER, D. R. SKIBOLA, A. AGANA, N. BECKER, P. TESSLER, M. S. FORREST, S. SANKARARAMAN, L. CONDE, E. A. HOLLY, AND M. T. SMITH.** *Polymorphisms in the Estrogen Receptor 1 and Vitamin C and Matrix Metalloproteinase Gene Families Are Associated with Susceptibility to Lymphoma. PLoS ONE, Vol. 3, Issue 7, 30 June 2008.*
- A. STOLCKE, L. FERRER, AND S. KAJAREKAR.** *Nonparametric Feature Normalization for SVM-Based Speaker Verification. Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Las Vegas, Nevada, pp. 1577-1580, April 2008.*
- S. STOYANCHEV, G. TÜR, AND D. HAKKANI-TÜR.** *Name-aware Speech Recognition for Interactive Question Answering. Proceedings of IEEE ICASSP, Las Vegas, NV, April 2008.*
- T. STRAYER, M. ALLMAN, G. ARMITAGE, S. BELLOVIN, S. JIN, AND A. W. MOORE.** *IMRG Workshop on Application Classification and Identification Report. Editorial contribution to ACM Computer Communication Review, Vol. 38, Issue 3, pp. 87-90, July 2008.*
- Z. TAN, K. ASANOVIC, AND D. A. PATTERSON.** *An FPGA Host-multithreaded Functional Model for SPARC v8. Proceedings of International Symposium on Computer Architecture, Workshop on Architectural Prototyping, Beijing, China, June 2008.*
- D. VERGYRI, A. MANDAL, W. WANG, A. STOLCKE, J. ZHENG, M. GRACIARENA, D. RYBACH, C. GOLLAN, R. SCHLATER, K. KIRCHOFF, A. FARIA, AND N. MORGAN.** *Development of the SRI/Nightingale Arabic ASR system. Proceedings of Interspeech 2008, Brisbane, Australia, September 2008.*
- O. VINYALS AND G. FRIEDLAND.** *Modulation Spectrogram Features for Speaker Diarization. To appear in proceedings of Interspeech 2008, Brisbane, Australia, September 2008.*
- O. VINYALS AND G. FRIEDLAND.** *Towards Semantic Analysis of Conversations: A System for the Live Identification of Speakers in Meetings. Proceedings of IEEE International Conference on Semantic Computing, Santa Clara, CA, pp. 426-431, August 2008.*
- M. VUTUKURU, H. BALAKRISHNAN, AND V. PAXSON.** *Efficient and Robust TCP Stream Normalization. Proceedings of IEEE Symposium on Security and Privacy, Oakland, CA, pp. 96-110, May 2008.*
- H. WANG, D. GELBART, W. HEMMERT, AND H.G. HIRSCH.** *The Value of Auditory Offset Adaptation and Appropriate Acoustic Modeling. To appear in proceedings of Interspeech 2008, Brisbane, Australia, September 2008.*
- F. YANG, E. SHRIBERG, AND G. TÜR.** *Exploiting Dialog Act Tagging and Prosodic Information for Action Item Identification. Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Las Vegas, Nevada, pp. 4941-4944, April 2008.*
- S. ZHAO AND N. MORGAN.** *Multi-Stream Spectro-Temporal Features for Robust Speech Recognition. To appear in Proceedings of Interspeech 2008, Brisbane, Australia, September 2008.*



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.

BOARD OF TRUSTEES

Javier Aracil
Hervé Bourlard
Beth Burnside
Adele Goldberg
Greg Heinzinger
Clifford Higgerson
Richard Karp
Nelson Morgan
David Nagel
Prabhakar Raghavan
Stuart Russell
Shankar Sastry
Scott Shenker
Eero Silvennoinen
David Tennenhouse
Wolfgang Wahlster

The ICSI Gazette is published twice a year by the International Computer Science Institute (ICSI). All material in The ICSI Gazette is copyrighted by ICSI and should not be used without explicit prior written permission.

EDITOR-IN-CHIEF

Nelson Morgan

SENIOR EDITOR

Maria Quintana

DESIGN

envision ID

LAYOUT

Jacob Wolkenhauer

COPYWRITER

Jacob Wolkenhauer

ADDITIONAL COPYWRITING

Leah Hitchcock

PHOTOGRAPHY

Diane Starr

INTERNATIONAL
COMPUTER
SCIENCE
INSTITUTE

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