CS@ILLINOIS RESEARCH BRIEFS

click!

University of Illinois at Urbana-Champaign

what's next in parallel computing ?

College of Engineering

the computer was born
to solve problems
that did not exist before

Bill Gates
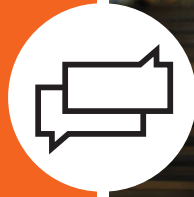
cl**i**ck!

University of Illinois at Urbana-Champaign

CS@ILLINOIS RESEARCH BRIEFS

"
We've been defining the rules of this game for a *long* time.

Rob A. Rutenbar
"

Welcome to CLICK! Research Briefs edition-- the research news magazine for the Department of Computer Science at the University of Illinois at Urbana-Champaign. We use CLICK! Research Briefs to give our friends look at some of the exciting accomplishments of our faculty and students.

I'm particularly excited about this issue, because of the focus on parallelism. This is an area close to the heart of CS, integrating issues across architecture, compilers, operating systems, networks, applications, and theory. It's also an area where Illinois has long been in the lead. Fifty years ago, when a 'serious' mainframe was less powerful than today's cheapest cellphone, we were building Illiac III, a SIMD parallel machine; followed shortly by ILLIAC IV, a 256-node machine with a vector-like architecture. We've been defining the rules of this game for a *long* time.

And the excitement continues today. New centers in Assured Cloud Computing, and Programmable Extreme-Scale Systems. A new campus-wide effort, the Parallel Computing Institute (PCI), which integrates all our broad, diverse efforts in one place. And lets not forget Blue Waters, a collaboration between the National Center for Supercomputing Applications (NCSA) and Cray, which is today standing up a peta-scale machine with just under 400,000 cores. CS@Illinois is a very good place to be doing parallelism right now.

Those of us working in this area (full disclosure: I'm a hardware guy, and a silicon guy) all have our stories for how we got into this business. It's worth sharing mine. As a new grad student working for Prof. Dan Atkins (CS Ph.D., '69) I was assigned to straighten up a hardware lab. In this process, I came across a box filled with tiny, ancient-looking circuit boards. "Hey Dan!" I asked, "What are these?"

"Illiac III," said Dan. "Want some?"

(I should have grabbed more....)

Department Head and Abel Bliss Professor of Engineering

ROBERT
Bocchino:

Alumnus Robert Bocchino (PhD 2010) was awarded the ACM SIGPLAN Outstanding Doctoral Dissertation Award for 2010. The award honors Bocchino's Ph.D. thesis, entitled "An Effect System and Language for Deterministic-by-Default Parallel Programming.î The Outstanding Doctoral Dissertation Award is "presented annually to the author of the outstanding doctoral dissertation in the area of Programming Languages" and includes a prize of $1,000.

Bocchinoís research, advised by Prof. Vikram Adve, describes the design and evaluation of a new, Java-based object-oriented parallel language called Deterministic Parallel Java (DPJ). DPJ is a safe and modular parallel language that helps developers port (parts of) sequential Java applications to run on multicore systems. The parallel language is the first to guarantee deterministic semantics without run-time checks for general-purpose, object-oriented programs. Itís also the first language to use compile-time type checking for parallel operations on arrays of references (ipointersî) to objects, and the first language to use regions and effects for flexible, nested data structures.

Bocchino's current research focuses on using and extending the Plaid programming language, together with related verification techniques, to design and verify abstractions that make it easy for programmers to write correct and efficient parallel code.

TIANYI
WU:

Alumnus Tianyi Wu was named the runner-up for the 2011 ACM SIGKDD Ph.D. Dissertation Award for his thesis "A Framework for Promotion Analysis in Multi-Dimensional Space".

Wu's work involves the use of a novel class of data mining problems, called promotion analysis, for marketing purposes. His thesis discusses how promotion analysis can be used for promoting a given product in a multi-dimensional space by leveraging object ranking information. Wuís key observation is that while most products may not be highly ranked in the global space, where all products are compared by all aspects, there often exist interesting and meaningful local spaces in which the given product becomes prominent. Therefore, Wu's goal is to break down the data space and discover the most interesting local spaces in an effective and efficient way.

The promotion analysis problem is highly practical and useful in a wide spectrum of decision support applications, says Wu. Typical application examples include merit discovery, product positioning and customer targeting, object profiling and summarization, identification of interesting features, and explorative search of objects.

Wu currently works as an Applied Researcher at Microsoft, focused on the Bing technologies. His Ph.D. advisor was computer science professor Jiawei Han.

RAY
OZZIE:

Alumnus Ray Ozzie was inducted into the College of Engineering Hall of Fame, October 14, 2011. Ozzie is a software visionary and entrepreneur who influenced communications and productivity in business, first as creator of LotusNotes and founder of Groove Networks, and most recently as Chief Software Architect of Microsoft.

Raymond E. Ozzie is an independent software entrepreneur and pioneer in social productivity, an area widely known in the field as computer-supported cooperative work. Through late 2010, he was Chief Software Architect of Microsoft, the company's most senior technical strategy and architecture role previously held by Bill Gates.

Ozzie came to Microsoft in 2005 through the acquisition of Groove Networks, a company he founded in 1997 to focus on software and services for small-team dynamic collaboration. Prior to Groove, in 1984 Ozzie founded and led Iris Associates, the creator and developer of Lotus Notes. A decade later, Iris was acquired by Lotus and then by IBM. Under his leadership during that period, Lotus Notes grew to be used for communication and social productivity by hundreds of millions at most major enterprises worldwide. Before creating Notes, he worked on 1-2-3 and Symphony at Lotus, on VisiCalc and TK!Solver at Software Arts, and on operating systems at Data General.

Ozzie studied computer science and engineering at the University of Illinois. While earning his degree, he also worked as a systems developer on the seminal PLATO project. He credits that work with helping him to understand the significance of online community and social interactive systems.

Honored as one of seven Windows Pioneers by Microsoft, Ozzie was named "Person of the Year" in 1995 by PC Magazine, and he has been inducted in to the Computer Museum Industry Hall of Fame. He was selected for an Engineering at Illinois Alumni Award for Distinguished Service in 1997. He received the 2000 Institute for Electrial and Electronics Engineers Computer Society's W. Wallace McDowell Award and the 2005 SDForum Visionary Award. In 2001, he was honored as a World Economic Forum Technology Pioneer.

He is a member of the National Academy of Engineering and was named a fellow of the American Academy of Arts and Sciences last year.

Ozzie has served as a member of the National Research Council's Computer Science and Telecommunications Board, and he was a member of the NRC committee that produced the landmark CRISIS report on the societal impact of cryptography.

Illinois CS Alumnus Ray Ozzie To Receive Honorary Degree

BY JEFF UNGER, UNIVERSITY OF ILLINOIS NEWS BUREAU
Raymond Ozzie, former chief software architect for the Microsoft Corp., has been chosen to receive an honorary doctor of engineering degree at the 2 p.m. campus-wide Commencement ceremony May 13 at the Assembly Hall.

In the early 1990s, his creation, Lotus Notes, was the first networked groupware application for the personal computer, revolutionizing business computing. In 1994, he was named one of seven "Windows Pioneers" by Microsoft because of the impact he and Lotus Notes had on the development of the Windows operating system.

In 1995, Ozzie, was named "Person of the Year" by PC Magazine. Two years later, the U. of I. College of Engineering honored Ozzie with its alumni award for distinguished service. Ozzie earned a bachelor's degree in computer science at Illinois in 1979.

In 2000, Ozzie was awarded the IEEE Computer Society's W. Wallace McDowell Award for his "vision, determination, and programming skill in the development of Lotus Notes, a program that enables groups of people to collaborate over computer networks."

did you know?

Blue Water Stats:

Memory: >1.5 Petabytes

Peak Performance: >11.5 Petaflop

Number of cores: >380,000

Number of NVIDIA GPUs: >3,000

parallel
computing

Encyclopedia of Parallel

University of Illinois computer science professor David Padua served as Editor-in-Chief for a new encyclopedic effort to survey the concepts behind the significant shift towards parallel computing in todayís computer industry. The 4 volume Encyclopedia of Parallel Computing contains more than 300 entries on topics related to the critical demand for continued advances in parallel programming. Contributors to the effort were leading international experts in the field, including Illinois computer science professors William Gropp and Marc Snir, professor emeritus David J. Kuck, and Jack Dongarra, Michael Flynn, David E. Shaw, and Guy L. Steele Jr. among others.

"Parallel computing has already impacted or will soon impact everyone who uses a computing device, from supercomputers to laptops, tablets, and smart phones," said Padua.

Today's supercomputers are massively parallel machines with thousands of processors. The fastest supercomputer today uses 705,024 processing cores, capable of 10.51 quadrillion calculations per second.

Ten years ago the worldís fastest supercomputer used a total of 8,192 processing cores and was only capable of 12.3 trillion calculations per second, almost one thousand times less powerful. This type of accelerated parallelism is critical to science and engineering, enabling discoveries and designs that would not be possible otherwise. For consumer and mobile devices, parallelism is the only viable strategy for continued performance gains, while also allowing chipmakers to optimize for energy efficiency.

With the need for parallelism at an all-time high, the Encyclopedia of Parallel Computing provides researchers and developers with an authoritative reference that pulls together the tools necessary to take advantage of these pioneering concepts.

"This monumental work will be an invaluable resource for practitioners and students in all areas of computer science," said Alex Nicolau of UC Irvine. "In today's world where parallel computing is ubiquitous--from desktops to cell phones and game consoles--this reference work will be indispensable."

Key concepts in the Encyclopedia for professionals, researchers and students of Parallel Computing include:

- Programming models and programming languages

- Debugging and race detection

- Laws of parallel computing

- Theoretical models of computation

- Supercomputer/high-performance computing machines

- Interconnection networks

DeNovo

University of Illinois computer science researchers have won the Best Paper Award at the Parallel Architectures and Compilation Techniques (PACT 2011) conference for their paper, "DeNovo: Rethinking the Memory Hierarchy for Disciplined Parallelism." The research team, led by Illinois computer science professor, Sarita Adve, is working on the DeNovo project that takes a new approach to building multicore hardware. DeNovo exploits emerging software trends in disciplined parallel programming to make hardware simpler, higher performance, and lower energy, all at the same time.

Most multicore programs use a shared memory programming model. Shared-memory programs have many advantages, but they are known to be difficult to program, debug, and maintain. At the same time, shared-memory hardware is complex and inefficient, leading to unnecessary energy consumption and performance bottlenecks. After decades of trying, researchers have found it difficult to even develop satisfactory shared-memory semantics for common languages such as C++ and Java. A recent article, co-authored by the DeNovo lead, calls the research community to rethink how we design both parallel languages and parallel hardware.

At the root of these problems is what the Illinois team refers to as "wild shared memory" behavior. Shared-memory programs tend to exhibit unstructured parallelism with implicit communication and side effects, leading to hard-to-debug data races and ubiquitous non-determinism. The Illinois team believes general-purpose languages must enforce more discipline and eliminate such wild behavior by design if parallel computing is to become tractable for mass programmers. Such a discipline would enforce more structured parallelism and make side effects of a parallel task be more explicit.   Many software researchers today are working on such an approach,

including pioneering work by an Illinois team on the Deterministic Parallel Java (DPJ) language, led by Vikram Adve. The DeNovo team, working closely with the DPJ team, has shown that the same disciplined parallel programming features that simplify software can also enable more performance-, energy-, and complexity-scalable hardware. As their first step, they have developed a cache coherence protocol and consistency model that takes an order of magnitude less time to verify and runs some applications in less than half the time with less than half the network traffic and cache misses than the state-of-the-art. The simplicity and low network and cache traffic means that the performance increases come with significant power and energy benefits. It is rare in computer architecture that a hardware design improves complexity, performance, and power all at once.

According to Sarita Adve, "this paper is a first step towards an ambitious vision. While it presents significant new technical contributions that we hope will eventually be adopted, it also opens up many new questions. We hope that the largest impact of this paper will be in inspiring a broad research agenda anchored in a more disciplined approach to parallel systems. The paper motivates hardware research driven by disciplined programming models and also seeks to inspire architects to extend their influence on the development of such models. We believe this is an opportune time for such a co-designed evolution of the current hardware-software ecosystem."

CS Course Leads to Groundbreaking Health Informatics Text

Bruce Schatz, an affiliate faculty member in Computer Science, has co-authored a groundbreaking book on Health Informatics based on his popular computer science course. The text is the first book combining the solutions of modern computer science with the problems of modern medical science. The book is expected to be a key reference for professionals working in health management, from information to healthcare executive, health information technologist to computer scientist, and physician to patient.

Healthcare Infrastructure: Health Systems for Individuals and Populations describes the new healthcare infrastructure that will gather personal health records from every individual and correlate each longitudinal record across whole populations. The book explains the problems of personal medicine and public health, then the solutions possible with information technology.

"What is needed [in health care] is a vast backbone, a health care infrastructure consisting broadly of health and deeply of medical information, which is recorded through personal sensors, analyzed on supercomputers, communicated through networks, and accessed through computers," says Schatz. His new text is about that infrastructure: who will use it, what problems it solves, where it will be used, why it chooses its designs, and how it works.

writing the books

Hoiem Authors Computer Vision Text

A new text in computer vision co-authored by professor Derek Hoiem introduces the subject of 3D object recognition and scene interpretation in depth. The text, Representations and Techniques for 3D Object Recognition & Scene Interpretation has a primary focus on recent efforts to fuse models of geometry and perspective with statistical machine learning.

The text aims to make the latest developments in 3D scene understanding and object recognition accessible to newcomers to the field. With existing research scattered across journal and conference papers, the subject was left to the purview of experts. Hoiem's text organizes this research and provides an historical and technical background so that newcomers to the field can learn about this emerging area.

"In recent years, the rigid, algebraic view of 3D geometry has given way to a more statistical, probabilistic view. In consequence, we've seen amazing new abilities to reconstruct 3D scenes and recognize 3D objects from a photograph," said Hoiem. "These technologies could have far-ranging impact, from robotics, to vehicle safety, to content creation and photo enhancement."

Har-Peled Authors Text on Geometric Approximation Algorithms

A new text on geometric approximation algorithms authored by University of Illinois computer science professor Sariel Har-Peled is the first to cover the subject in detail. Geometric Approximation Algorithms describes key techniques in geometric approximation algorithms and also surveys some of the more traditional computational geometry techniques, such as sampling and linear programming.

The field of geometric approximation algorithms is a subfield of both computational geometry and approximation algorithms, explains Har-Peled.

"Exact algorithms for dealing with geometric objects are complicated, hard to implement in practice, and slow," says Har-Peled. "Over the last 20 years, a theory of geometric approximation algorithms has emerged. These algorithms tend to be simple, fast, and more robust than their exact counterparts."

What began as a collection of class notes on the subject was expanded by Har-Peled into his full length text.

Illinois-Developed Data Creation Tool a Boon for Genetic Biologists

BY: LIZ AHLBERG, ILLINOIS NEWS BUREAU

With the BeeSpace Navigator, University of Illinois researchers have created both a curation tool for genetic biologists and a new approach to searching for information.

The project was a collaboration between researchers at the Institute for Genomic Biology and the department of computer science. Led by Bruce Schatz, professor and head of medical information science at the U. of I. and affiliate professor in computer science, the team described the software and its applications in the web server issue of the journal Nucleic Acids Research. The research team also included CS faculty members Chengxiang Zhai and Saurabh Sinha.

When biologists need information about a gene or its function, they turn to curators, who keep and organize vast quantities of information from academic papers and scientific studies. A curator will extract as much information as possible from the papers in his or her collection and provide the biologist with a detailed summary of what's known about the gene — its location, function, sequence, regulation and more — by placing this information into an online database such as FlyBase.

"The question was, could you make an automatic version of that, which is accurate enough to be helpful?" Schatz said.

Schatz and his team developed BeeSpace Navigator, a free online software that draws upon databases of scholarly publications. The semantic indexing to support the automatic curation used the Cloud Computing Testbed, a national computing datacenter hosted at U. of I.

While BeeSpace originally was built around literature about the bee genome, it has since been expanded to the entire Medline database and has been used to study a number of insects as well as mice, pigs and fish.

The efficiency of BeeSpace Navigator is in its specific searches. A broad, general search of all known data would yield a chaotic myriad of results – the millions of hits generated by a Google search, for example. But with BeeSpace, users create "spaces," or special collections of literature to search. It also can take a large collection of articles on a topic and automatically partition it into subsets based on which words occur together, a function called clustering.

"The first thing you have to do if you have something that's simulating a curator is to decide what papers it's going to look at," Schatz said. "Then you have to decide what to extract from the text, and then

what youíre going to do with what you've extracted, what service you're going to provide. The system is designed to have easy ways of doing that."

The user-friendly interface allows biologists to build a unique space in a few simple steps, utilizing sub-searches and filters. For example, an entomologist interested in the genetic basis for foraging as a social behavior in bees would start with insect literature, then zero in on genes that are associated in literature with both foraging and social behavior – a specific intersection of topics that typical search engines could not handle.

This type of directed data navigation has several advantages. It is much more directed than a simple search, but able to process much more data than a human curator. It can also be used in fields where there are no human curators, since only the most-studied animals like mice and flies have their own professional curators.

Schatz and his team equipped the navigator to perform several tasks that biologists often perform when trying to interpret gene function. Not only does the program summarize a gene, as a curator would, but it also can perform analysis to extrapolate functions from literature.

For example, a study will show that a gene controls a particular chemical, and another study will show that chemical plays a role in a certain behavior, so the software makes the link that the gene could, in part, control that behavior.

BeeSpace can also perform vocabulary switching, an automatic translation across species or behaviors. For example, if it is known that a specific gene in a honeybee is analogous to another gene in a fruit fly, but the function of that gene has been documented in much more detail in a fruit fly, the navigator can make the connection and show a bee scientist information on the fly gene that may be helpful.

"The main point of the project is automatically finding out what genes do that don't have known function," Schatz said. "If a biologist is trying to figure out what these genes do, they're happy with anything. They want to get as much information as possible."

The project included the work of several CS students, now alumni, including Qiaozhu Mei, Jing Jiang, Xu Ling, Xin He, and David Arcoleo.

The BeeSpace Navigator, now in its fourth version, is available free online. Overview documentation is available as well. The National Science Foundation supported this work as the bioinformatics flagship of the Frontiers in Integrative Biological Research program.

Search Engine Optimazation

Search for the term "computer virus", and itís likely that some of your search results will contain only computer, or only virus. Why do these results show up instead of other more relevant results?

PhD student Yuanhua Lv sought the answer to this question as he began exploring new methods to improve search algorithms. The performance of a search engine is mainly determined by its retrieval model which formally specifies how to score and rank documents optimally for a user's query. Optimization of retrieval models is a fundamentally important research problem because an improved retrieval model would lead to improved performance for all search engines.

Lv first analyzed the deficiencies in current models, and revealed a previously unknown common deficiency in all current retrieval methods, namely that the component of term frequency normalization by document length is not properly lower-bounded. As a result of this deficiency, long documents which do match the query term can often be scored by search engines unfairly as having a lower relevancy than shorter documents that do not contain the query term at all. For example, for a query such as "computer virus", a long document matching both "computer" and "virus" can easily be ranked lower than a short document matching only "computer".

Lv's discovery led him to create a novel methodology for introducing a sufficiently large lower bound for term frequency normalization, which can be used as a plug-and-play patch to multiple current retrieval models to eliminate the problem. Lab results indicate that Lv's methodology incurs almost no additional computational costs while delivering more precise search results, particularly in cases where the queries are verbose.

Until Lv's work, it has proven difficult to improve the current retrieval models. The well-known BM25 retrieval function was proposed in 1994, but efforts to improve its performance since then have been fruitless.

"[Lv's work] makes a breakthrough contribution in improving the general search algorithms used in all search engines," said Illinois computer science professor ChengXiang Zhai, Lv's advisor and the co-author of a paper explaining Lv's discovery and new methods. "The new models proposed can improve over many strong baseline methods that have been the state of the art for more than a decade. These new models can be immediately used to replace the old models in all search engines to improve search accuracy."

The paper describing Lv's work, "Lower-bounding term frequency normalization" received the Best Student Paper Award at the 20th ACM Conference on Information and Knowledge Management (CIKM 2011).

Building a Better Internet

University of Illinois computer science professor Brighten Godfrey was among a select group of academic researchers and Internet visionaries chosen to participate in Verisign's "Building a Better Internet Symposium". Godfrey's project was one of four chosen internationally to receive a $75,000 infrastructure grant that Verisign awarded as part of its 25 Years of .Com commemorations.

The University of Illinois project, a collaboration with Ph.D. students Wenxuan Zhou and Qingxi Li and Professors Matthew Caesar and Brighten Godfrey, developed methods to accelerate the Web and other interactive networked applications, via secure, deployable extensions to the domain name system (DNS) and transport control protocol (TCP). The team created the Accelerated Secure Association Protocol, or ASAP, which establishes a connection between client and server quickly and securely. The protocol enables the server to verify the key security property that the client's source address is not forged, yet avoids the delay of TCP's "handshake" method of verification.

"What I'm really excited about is how do we make the other side of the world feel like it's right at our fingertips," said Godfrey. "The exciting thing is that this work can have broad impact. If ASAP is widely deployed, it would make every connection on the web faster."

Pre-Social Networks

A technology that can tell where users are going to be, how long will be there, and who they will meet.

Sound like a sci-fi movie?

At Professor Klara Nahrstedt's lab, it's a reality. Nahrstedt and computer science graduate student Long Vu's new technology Jyotish draws up maps of people's movements by monitoring the connections their smart phones make to WiFi and Bluetooth networks. Over time, the system is able to determine the patterns in users activities and movements, and can make predictions on where people will be in the future, and what other people might be nearby during the same time frame.

The project began as an effort by Boeing to find better ways to track and predict the movements of work crews in its aircraft manufacturing facilities.

"It is well known that people movement exhibits a high degree of repetition since people visit regular places and make regular contacts for their daily activities," says Vu. "Our work constructs a predictive model by exploiting the regular pattern of people movement found in real joint Wifi/Bluetooth trace."

The model constructed by Jyotish is able to answer three fundamental questions: (1) where the person will stay, (2) how long she will stay at the location, and (3) who she will meet.

In order to construct the predictive model, Jyotish includes an efficient clustering algorithm to exploit regularity of people movement and cluster Wifi access point information in Wifi trace into locations. Then, a Naive Bayesian classifier assigns these locations to records in Bluetooth trace. Next, the Bluetooth trace with assigned locations is used to construct predictive model including location predictor, stay duration predictor, and contact predictor.
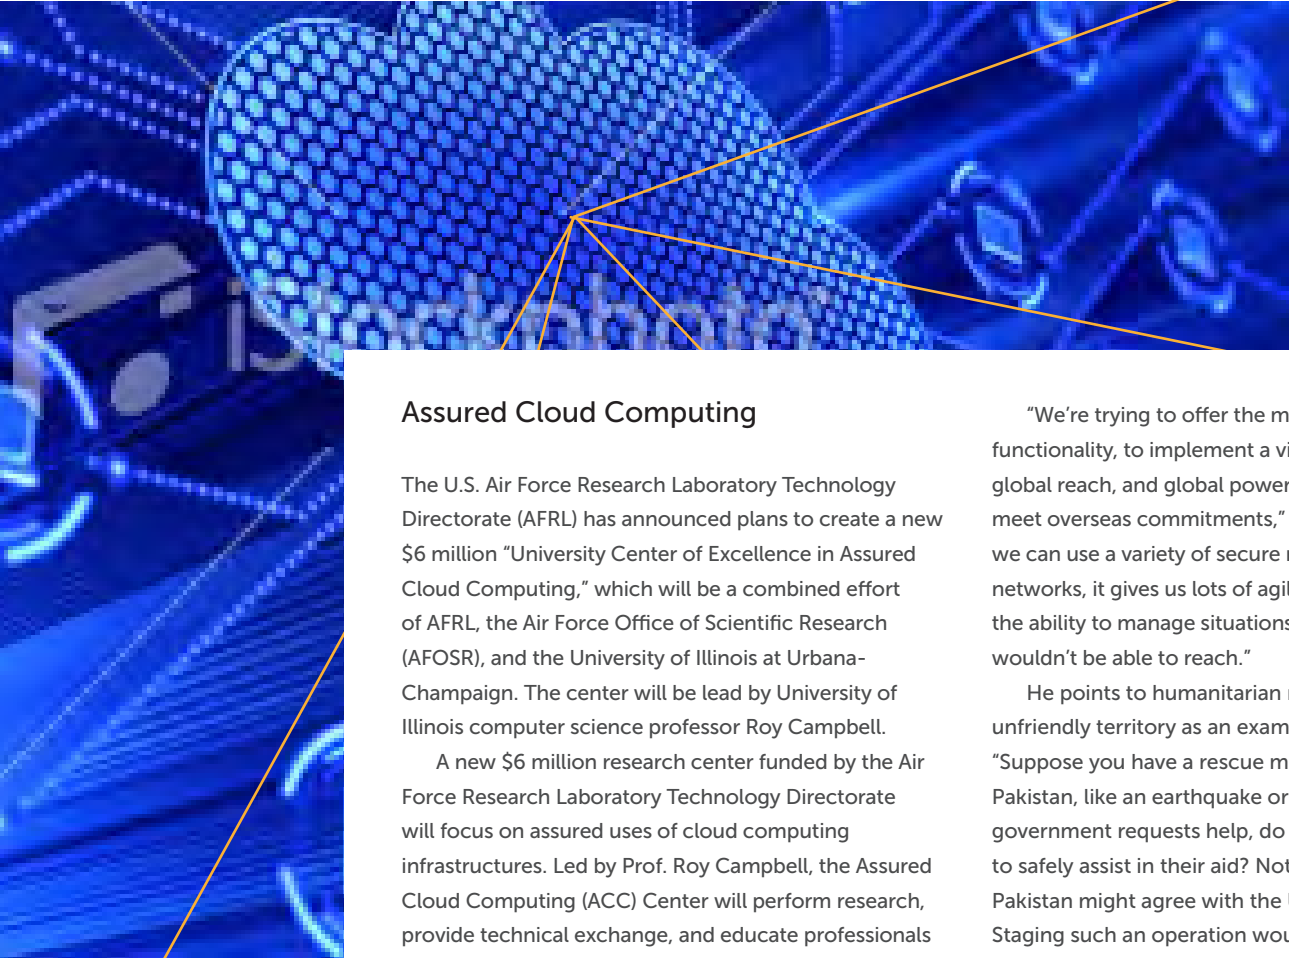
Assured Cloud Computing

The U.S. Air Force Research Laboratory Technology Directorate (AFRL) has announced plans to create a new $6 million "University Center of Excellence in Assured Cloud Computing," which will be a combined effort of AFRL, the Air Force Office of Scientific Research (AFOSR), and the University of Illinois at Urbana-Champaign. The center will be lead by University of Illinois computer science professor Roy Campbell.

A new $6 million research center funded by the Air Force Research Laboratory Technology Directorate will focus on assured uses of cloud computing infrastructures. Led by Prof. Roy Campbell, the Assured Cloud Computing (ACC) Center will perform research, provide technical exchange, and educate professionals and students in the secure cloud computing sciences and technologies that are needed to allow the Air Force to succeed in air, space, and cyberspace missions. ACC's research activities will focus on developing technology for assured, mission-critical cloud computing across "blue" and/or "gray" networks that ensures the confidentiality and integrity of data and communications, job completion in the presence of cyber attacks and failures, and timely completion of jobs to meet mission requirements.

A computational cloud used in military applications may include both blue and gray networks, where "blue" networks are U.S. military networks, which are considered secure, and "gray" networks are those in private hands, or perhaps belonging to other nations, which are considered insecure. In order to reach mission goals, it will sometimes be necessary to coordinate computation across a mixture of these blue and gray resources. Thus, cloud computing in a military context presents special security challenges. Specifically, assured mission-critical cloud computing across blue and/or gray networks requires the realization of "end-to-end" and "cross-layered" security, dependability, and timeliness.

"We're trying to offer the military extended functionality, to implement a vision of global vigilance, global reach, and global power giving us the ability to meet overseas commitments," explains Campbell. "If we can use a variety of secure networks plus insecure networks, it gives us lots of agility and mobility and the ability to manage situations where normally we wouldn't be able to reach."

He points to humanitarian missions in potentially unfriendly territory as an example application. "Suppose you have a rescue mission, say a disaster in Pakistan, like an earthquake or a river flooding. If their government requests help, do we have the capabilities to safely assist in their aid? Not all the people in Pakistan might agree with the U.S. providing assistance. Staging such an operation would be risky without a cloud infrastructure that has secure properties. So how do you assure a successful mission in a possibly hostile environment? How do you benefit from the cloud, its communication, computations, and data in missions to help people in need?"

Ilesanmi Adesida, the Dean of the College of Engineering at Illinois, observed that the planned research has broad implications. "Although the new project will primarily benefit the Air Force over the short term, assured use of the cloud will be a tremendous benefit to humanity. Today, you can't really trust cloud computing because of the security issues that remain to be addressed," he explains. "No one has been able to use cloud computing in a task-oriented way before. The work of the Assured Cloud Computing Center will make it possible to deploy cloud computing in task-based, mission-oriented human activity. For that reason, this work will be groundbreaking."

In addition to Roy Campbell, the Center's research team includes computer science faculty members Gul Agha, Indranil Gupta, and JosÈ Meseguer among other Illinois faculty. The center is part of the Information Trust Institute at Illinois.

NCSA, Cray partner on sustained-petascale Blue Waters supercomputer

BY NCSA COMMUNICATIONS

The National Center for Supercomputing Applications at the University of Illinois has finalized a contract with Cray Inc. to provide the supercomputer for the National Science Foundationís Blue Waters project. Cray replaces IBM, which terminated its contract in August because, IBM said, the technology required more financial and technical support than originally anticipated.

Cray will begin installing hardware in the UIís National Petascale Computing Facility within the next few weeks, with an early science system expected to be available in early 2012. Blue Waters is expected to be fully deployed by the end of 2012.

This new Cray supercomputer will support significant research advances in a broad range of science and engineering domains, meeting the needs of the most compute-intensive, memory-intensive and data-intensive applications. Blue Waters is expected to deliver sustained performance, on average, of more than one petaflops on a set of benchmark codes that represent those applications and domains.

"We are extremely pleased to have forged a strong partnership with Cray. This configuration will be the most balanced, powerful and useable system available when it comes online. By incorporating a future version of the XK6 system, Blue Waters will also provide a bridge to the future of scientific computing," said Thom Dunning, NCSA director.

Physically Asynchronous Logically Synchronous Architecture

In networked cyber-physical systems, real time global computations require consistent views, consistent actions, and synchronized state transitions across network nodes. However, the convergence of sensing, control, communication and coordination in cyber-physical systems (CPS) such as modern airplanes, power grid, train and medical device networks poses enormous challenges.

Computer science researchers at the University of Illinois, led by Professors Lui Sha and Jose Meseguer, are developing formal complexity reduction architecture patterns aimed at addressing these challenges. The team has developed the Physically Asynchronous Logically Synchronous (PALS) architecture pattern that supports real time global computation. The objective of the PALS protocol is to provide the optimal real time logical (virtual) synchronization protocol.

The groups are working in collaboration with industry and research partners Rockwell Collins Inc., Lockheed Martin Corporation, Software Engineering Institute, and the University of Oslo. Researchers and engineers of Rockwell Collins Inc., in cooperation with Illinois computer science researchers, led the study on the implementation of PALS for an Integrated Modular Avionics application. The results showed that PALS reduced verification time from 35 hours to less than 30 seconds as compared with traditional design in a dual redundant flight guidance system prototype. This PALS study, authored by Steven P. Miller, Darren D. Cofer, Lui Sha, Jose Meseguer and Abdullah Al-Nayeem, won the David Lubkowski Award for the Advancement Digital Avionics.

AADL is an open standard architecture analysis and description language supported by the avionics community. Because of this remarkable success in complexity reduction, Illinois researchers, in cooperation with Software Engineering Institute and University of Oslo, are now working on a special AADL Annex called Synchronous AADL to support PALS based designs. In addition, Illinois researchers, led by Dr. Cheolgi Kim, in cooperation with Lockheed Martin, are developing a production quality implementation of PALS library.

New Illinois center to develop smarter infrastructures, smarter publics

From smart utilities like the smart grid and intelligent transportation systems to social networks on sites like Facebook and YouTube, the infrastructures of tomorrow will heavily utilize information technology. While these "smart" infrastructures promise many benefits, they often require new kinds of interaction between people and the machines meant to serve them. Yet the social, cultural, economic and political side of these relationships often receives little attention.

The new Center for People and Infrastructures at the University of Illinois at Urbana-Champaign seeks to address these issues by better understanding social norms, market structures, public policies and human capabilities that shape and are affected by the development of smart infrastructures. The center, part of the Coordinated Science Laboratory, brings together experts in engineering, design, the social sciences and computer science.

The Center will initially focus on research about broadband telecommunications and energy. Researchers will work with fiber networks like UC2B (Urbana-Champaign Big Broadband), which will deliver high-quality broadband connections to several thousand households in Champaign and Urbana, to understand the consequences of fiber infrastructure for education, the economy, health and community participation. In addition, the Center will work to help identify the next generation of broadband Internet applications enabled by fiber infrastructure.

"Infrastructures are about computers, wires and pipes but they are also about human relationships, economics and justice," said Co-Director Karrie Karahalios, associate professor of computer science. "We want to see infrastructures that not only work, but that help humans to flourish."

Machine Learning and Perception

Three University of Illinois engineering faculty members will be investigating machine learning and perception via stochastic computation as part of a new Intel-funded center on embedded computing. Computer science professors Rob Rutenbar and Paris Smaragdis join electrical and computer engineering professor Naresh Shanbhag as part of the Intel Science and Technology Center on Embedded Computing headquartered at Carnegie Mellon University.

A key area of research for the center is to make it easier for these everyday devices to continuously collect, analyze and act on useful data from both sensors and online databases in a way that is timely, scalable and reliable. For example, in cars, this data could be used to customize in-vehicle entertainment options when specific passengers are recognized, and provide them better routing, retail, dining, and entertainment recommendations while on-the-road.

Tapping into the expertise of leading researchers from across the country, the ISTC for embedded computing forms a new collaborative community to drive research to transform experiences in the home, car and retail environment of the future.

The team's research will be informed by new applications of machine listening in embedded computing environments for cars, retail environments, and homes. For example, the Illinois team's efforts may lead to cars that can listen for driver activity such as sleepiness, mechanical issues, or external traffic cues such as accidents. In a retail environment, the Illinois team envisions applications that can listen for customer movement and patterns linked to in-store advertising or product placement, or for emergency situations. Home applications may include listening for accidents and emergencies for the elderly or those needing assistance, or for building systems malfunctions like broken pipes.

"These new ISTCs are expected to open amazing possibilities," said Justin Rattner, Intel Chief Technology Officer. "Imagine, for example, future cars equipped with embedded sensors and microprocessors to constantly collect and analyze traffic and weather data. That information could be shared and analyzed in the cloud so that drivers could be provided with suggestions for quicker and safer routes."

International Coalition Aims to Enable Climate Simulation at Extreme Scale

Policy decisions for mitigating climate change or adapting to it are subjects of great discussion throughout the world. Uninformed decisions will impose a heavy cost on future generations, both financial and human. Therefore, it is essential to reduce the current uncertainties about future climate changes and their impact by running climate simulations at 1,000 times larger scales than today. Exascale supercomputers (those capable of 1018 operations per second) may appear in 2018-2020, featuring a hierarchical design, gathering hundreds of millions of computing cores. The numerical models of the physics, chemistry, and biology affecting the climate system need to be improved to run efficiently on these extreme systems. Without improvement, these codes will not produce simulation results required to respond to the societal and economical challenges of climate change.

The objective of the G8 ECS project is to investigate how to efficiently run climate simulations on future exascale systems and get correct results. This project gathers top researchers in climate and computer science to focus on three main topics:

- how to complete simulations with correct results despite frequent system failures
- how to exploit hierarchical computers with hardware accelerators close to their peak performance
- how to run efficient simulations with 1 billion threads.

This project also aims to educate new generations of climate and computer scientists about techniques for high-performance computing at extreme scale.

This project is a direct result of the collaboration between the University of Illinois at Urbana-Champaign, the National Center for Supercomputing Applications and INRIA, the French National Institute for Research in Computer Science and Control, through their Joint Laboratory for Petascale Computing. In addition to these three institutions, other partners on the project come from Canada's University of Victoria, the German Research School, Japan's Tokyo Tech and University of Tsukuba, Spain's Barcelona Supercomputing Center, the University of Tennessee and the National Center for Atmospheric Research. The project will employ the top supercomputers to experiment with new techniques in the previously described three topics.

This three-year project receives G8 coordinated funding from the Natural Sciences and Engineering Research Council of Canada, French National Research Agency, German Research Foundation, Japan Society for the Promotion of Science (JSPS) and U.S. National Science Foundation. This project, together with five other projects, was funded as part of the G8 Research Councils Initiative on Multilateral Research, Interdisciplinary Program on Application Software towards Exascale Computing for Global Scale Issues. This is the first initiative of its kind to foster broad international collaboration on the research needed to enable effective use of future exascale platforms.

Picture in Picture: Adding 3D Objects to 2D Images

Inserting objects into existing photographs or videos is extremely difficult because the object's appearance depends on the complex lighting and geometry of the environment. Researchers from University of Illinois show how to realistically render 3D objects into an existing photograph, requiring no special equipment and much less effort from the artist than previous approaches.

Researchers at Illinois have shown how get the lighting right, by combining methods from computer vision, human computer interaction, and computer graphics. A user who wanted to insert an object or an animation into a picture would sketch the main layout of the space (or use a computer vision program to estimate it) and the lights using an interface, and use a standard modeling tool to make the object or the animation.

"Detecting photo fraud is often easy, because it is very difficult to ensure that the new object has the right shadows, highlights, and shading; casts shadows in the right places; and is reflected in mirrors," said computer science professor David Forsyth. "In contrast, experimental work conducted at Illinois suggests that this new system (a) produces accurate estimates of shading, tested by comparison with real photographs and (b) fools human observers into believing that an edited image is real."

The technical breakthrough is getting the lighting right by taking advantage of small amounts of annotation to recover a simplistic model of geometry and the position, shape, and intensity of light sources. The team's approach applies a computer vision

program to estimate strong directional sources of light like the sun through windows; applies a computer vision program to estimate how much light the surfaces in the picture reflect, and to correct the user sketch of light; and then uses a standard computer graphics program to generate the picture from detailed physical models of light behavior.

In addition to the overall system, the research team also developed a semiautomatic algorithm for estimating a physical lighting model from a single image. The Illinois method is able to generate a full lighting model that is demonstrated to be physically meaningful through a ground truth evaluation. As part of their work, the team also introduced a state-of-the-art image decomposition algorithm for single image reflectance estimation.

"With a single image and a small amount of annotation, our method creates a physical model of the scene that is suitable for realistically rendering synthetic objects with diffuse, specular, and even glowing materials, while accounting for lighting interactions between the objects and the scene," said computer science PhD student Kevin Karsch, lead author of the approach. "Our approach has applications in the movie and gaming industry, as well as home decorating and user content creation, among others. Imagine being able to slay dragons right in your own living room."

Kevin's work, joint with Varsha Hedau, Derek Hoiem, and David Forsyth, will appear in this year's SIGGRAPH Asia conference. The authors' version of the paper can be found at http://kevinkarsch.com/publications/sa11.html, and they have released a video at http://vimeo.com/28962540. They plan to produce a web-server version of this application in the near future.

National Petascale LEED

The University of Illinois' National Petascale Computing Facility has been certified LEEDÆ Gold in the U.S. Green Building Council's Leadership in Energy and Environmental Design (LEEDÆ) rating system, which is the recognized standard for measuring sustainability in construction.

The building, which opened in summer 2010, is home to supercomputers and other high-performance systems operated by the National Center for Supercomputing Applications and used by scientists and engineers across the country.

The LEED system awards points in a half-dozen categories, including energy and water efficiency, indoor environmental quality, and innovation in design. NPCF's energy-saving features include:

- A highly efficient power distribution system that is based on 480 V power for the computational equipment.

- Focusing on water-cooled computational and storage equipment. Liquid cooling is two times more efficient that air cooling.

- External cooling towers that let Mother Nature chill the water needed for cooling the building and the supercomputers a large part of the year. This is expected to cut the chilled water costs for the facility by about $1 million per year.

- Low-impact landscaping with native prairie plants that thrive without frequent watering or mowing.

- Using best practice construction methods to improve the air quality environment within the facility.

Best Paper for Work on Regenerating Codes

A paper co-authored by University of Illinois computer science professor Brighten Godfrey received the 2010 IEEE Communications Society Technical Committee on Data Storage Best Paper Award.

The award is given by the Signal Processing for Storage Technical Committee of the IEEE Communications Society to a conference or journal paper published in 2010 that presents novel research on data storage.

The paper, "Network Coding for Distributed Storage Systems", was authored by Alex Dimakis, Brighten Godfrey, Yunnan Wu, Martin Wainwright, and Kannan Ramchandran, and appeared in the IEEE Transactions on Information Theory.

Illinois computer science professor Brighten Godfrey

The paper created a new sub-area of coding theory called regenerating codes, which can make reliable storage in data centers more efficient.

Large data stores, which often appear in today's data centers and cloud services, use erasure coding to provide redundancy across multiple disks so if part of the file is lost it can be reconstructed. But the erasure-coded fragments must also be replenished after a failure, to maintain the same level of reliability. Traditionally a lost fragment is replaced by first reconstructing the entire file, and then distilling it into a small additional fragment to replace the lost fragment. The paper introduced a technique called regenerating codes which can produce new fragments much efficiently, without reconstructing the entire file.

The award was presented at the IEEE Global Communications Conference (Globecom) 2011 in Houston, Texas, in December.

Citation:

Alexandros G. Dimakis, P. Brighten Godfrey, Yunnan Wu, Martin J. Wainwright, and Kannan Ramchandran. Network Coding for Distributed Storage Systems. IEEE Transactions on Information Theory, vol. 56., no. 9, September 2010

Yu Receives Best Paper for Data-Driven Graphics Models

University of Illinois computer science professor Yizhou Yu received a Best Paper Award for his work to recreate physical imperfections in scanned models of a human hand. The work, "Controllable Hand Deformation from Sparse Examples with Rich Details" received the Best Paper award at the ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA 2011).

Illinois computer science professor Yizhou Yu

"Recent advances in laser scanning technology have made it possible to faithfully scan a real object with tiny geometric details, such as pores and wrinkles," said Yu in the paper. "However, a faithful digital model should not only capture static details of the real counterpart, but be able to reproduce the deformed versions of such details."

The team used models of a human hand to test their approach. Hand models tend to be among the most challenging, because human hands have such a large degree of freedom of motion, and highly deformable wrinkles.

The team developed data-driven methods to reproduce such deformities both on a large-scale, and in high resolution. The team's framework was capable of synthesizing high-resolution hand mesh animations with rich and varying details from as few as 14 training examples. The team's approach is able to be applied both to keyframe animation systems as well as performance-driven animation.

The papers coauthors included Haoda Huang and Xin Tong (Microsoft Research Asia), Ling Zhao and Yue Qi (Beihang University), and KangKang Yin (National University of Singapore).
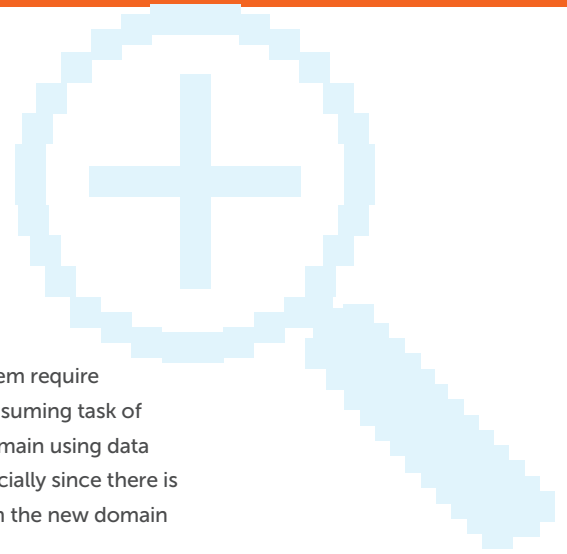
Best Paper Award for Work in Natural Language Processing

University of Illinois computer science graduate student Gourab Kundu received the Best Student Paper Award at the 2011 Conference on Computational Natural Language Learning (CoNLL 2011) for his work on open domain natural language processing. His paper, "Adapting Text Instead of the Model: An Open Domain Approach" introduces a new training approach that allows Natural Language Processing (NLP) researchers to adapt existing systems to new domains without the need for retraining their system on labeled data from the new domain.

Researchers in NLP develop algorithms and automated systems that analyze and interpret natural language text. For example, given the sentence "Gourab wrote two excellent papers and submitted them to 2012 conferences", a Semantic Role Labeling program would analyze the sentence at the level of "who does what to whom, when, where", and determine, for example, that Gourab is the Writer and also the Submitter, that two papers were submitted, that the submission happened in 2012, etc. While this task is trivial to humans, automating this process is very difficult due to the inherent ambiguity and variability of natural language. Modern NLP approaches therefore use machine learning and inference approaches to accomplish this task. Consequently, these systems must be "trained" or fed a lot of sample text, along with the corresponding analysis, in order to learn how to analyze new, previously unobserved text. However, while natural language systems trained on labeled data perform well on the domain in which they were trained, they adapt poorly to other domains due to changes in vocabulary, sentence structure, etc . For example, a system that was trained to parse natural language sentences on the Wall Street Journal may perform very poorly when parsing language from fiction novels.

Current approaches to this problem require researchers to conduct the time-consuming task of training a new model for the new domain using data from that domain. This is costly especially since there is a need to first annotate the data from the new domain with the correct analysis.

Kundu proposes looking at the problem from another angle. Rather than adapting the old model for new text, why not adapt the new text for the old model? The result is ADUT (ADaption Using label-preserving Transformation), an approach that avoids the need for retraining and does not require any knowledge of the new domain. ADUT applies multiple label-preserving transformations to the text it is attempting to analyze in order to make the target text "more similar" (in terms of vocabulary and structure) to the text on which it was initially trained. The ADUT system then applies the existing model on the transformed sentences, and combines the predictions using a constrained optimization inference approach to produce the desired prediction on the target text. In this manner, AUDT is able to work on the fly to analyze new text.

This way, Kundu's ADUT approach can use a single trained model to analyze text from multiple domains, even when the notion of a domain is not well defined (e.g., streaming data from the web). The approach was shown to yield a 13% reduction in error for the task of Semantic Role Labeling when adapting from news wire text to text from fiction.

Kundu conducts his research in computer science professor Dan Roth's research group. Roth was co-author of the paper. One of Professor Roth's former students, Scott Wen-tau Yih, also won the Best Paper Award (non-student) at the same CoNLL conference for his paper, "Learning Discriminative Projections for Text Similarity Measures."

Rather than adapting the old model for new text, why not adapt the new text for the old model?

Gropp Leads New Parallel Computing Institute

With the computing world in the midst of a parallelism revolution, researchers are seeking to take advantage of unprecedented new speed and power potential for applications ranging from astronomy to zoology.

But meeting challenges in new programming technologies, power consumption and performance for scientific applications will require major interdisciplinary engineering efforts. The Parallel Computing Institute (PCI) is designed to enable Illinois researchers from across campus to come together in new, application-focused research centers and achieve their scientific goals using the latest and most efficient parallel computing technologies.

Led by Illinois computer science professor Bill Gropp, PCI will facilitate their work by providing an incubation and growth mechanism for new interdisciplinary centers and initiatives that will benefit from parallel computing; by expanding access to computing resources and infrastructure; by teaching critical skills; by identifying more opportunities for funding; and by helping to establish key external partnerships. The ultimate goal is to advance Illinois research capabilities by enabling high-impact collaborations and providing access to cutting-edge resources and professional instruction and assistance — all built around the revolutionary new possibilities offered by parallelization.

Research thrusts include:

Inpatient MRI

Through the Illinois Massively Parallel Acceleration Toolkit project, researchers have honed in on a way to speed up the acquisition and reconstruction of MRI images, while maintaining accuracy. The project, which draws on Brute Force and Toeplitz technologies, has enabled MRI processing in hours instead of weeks. The research is sponsored in part by the National Institutes of Health.

DNA sequencing

DNA sequencing, an important process in the bioengineering field, is largely used in gene analysis, extraction and adjustment. However, Illinois researchers believe it could also be a powerful tool in identifying disease. This research thrust explores ways to speed up the de novo genome assembly process using high-performance GPUs. In practice, clinicians could take a sample of DNA and run it through a program of coded diseases to find a match.

Productive heterogeneous exascale computing

The "Vancouver Project," funded by the Department of Energy, is designing a next-generation software infrastructure for productive heterogeneous exascale computing. The project uses GPUs to address issues of realized performance, scalability and programming productivity in scalable heterogeneous computing (SHC). With attention to issues of energy-efficiency and resiliency of these systems, the goal is to provide tools and infrastructure for the transition to SHC.

Data layout transformations

As applications require increasing amounts of memory, researchers are exploring new ways to store data. Because data layout is not usually considered during programming, data are frequently saved in a default layout, which is inefficient because the data the programs need at a particular time may not reside nearby in the memory space. This research seeks to change how programs save data, focusing on how it will later be accessed. By matching the code to data layout, data can be accessed more quickly and programs will run faster.

Scalable, implicitly parallel programming

As programming shifts from sequential to parallel modes, researchers want to modify the way coding is accomplished. But instead of requiring programmers to learn an entirely different method of coding, researchers suggest that it is more practical to modify sequential coding techniques. With this approach, programmers can more easily spot problems and donít need to relearn their trade. Illinois research is focused on creating programs that run on parallel systems in order to fully utilize all of a machineís resources. This method is termed ìimplicitly parallelî because the coding presents itself externally as sequential code,

" The Center for Programmable Extreme Scale Computing aims to create the next generation of computers that are 1000x more energy-efficient, yet deliver the same performance and physical footprint of current computers. "

Center for Programmable Extreme Scale Computing

Enabling programmable, high-performance and highly power/energy efficient many-cores is a major step toward solving the true grand challenge of popular parallel programming. It is key to our industryís ability to build future-generation computers, keeping our nation competitive in IT and, therefore, ensuring our future quality of life.

Current petascale architectures are hardly scalable. They consume several megawatts (MW) of power and waste about 20% of their capacity because they have typical MTBFs of a few days. Consequently, as we look toward building exascale machines, power/energy-efficiency and resiliency are key challenges. In addition, exascale machines need to support over one billion threads, and operate with an even worse ratio of flops to memory bytes than todayís already challenging 10/1. This makes concurrency and locality major challenges as well. Finally, providing a programmable environment — one that assists the programmer to express, reveal, and exploit an application's parallelism — is an additional key challenge.

The Center for Programmable Extreme Scale Computing aims to create the next generation of computers that are 1000x more energy-efficient, yet deliver the same performance and physical footprint of current computers. The Center aims to create hardware and software technologies that fulfill this mission that are applicable to the whole range of computing systems: from terascale portable devices, to petascale departmental servers, and exascale data centers.

"We envision the Extreme-Scale many-core circa 2020 to be very high performance, and yet very power- and energy efficient," said Josep Torrellas, director of the Center. "Designing a usable many-core, however, introduces many challenges. The biggest one is programmability."

Although programmability is a notoriously difficult metric to define and measure, we suggest it implies two things: (i) the system should be able to attain high efficiency while relieving the programmer from managing low-level tasks, and (ii) the system should help minimize the chance of (parallel) programming errors.

Research at the Center for Programmable Extreme Scale Computing is organized around three major efforts..

Ubiquitous High Performance Computing: Runnemede
Funded by: DARPA

As part of the $49 million Intel-led effort to develop a new platform for ubiquitous high performance computing, the Centerís researchers are investigating new compute and memory system architectures that aim to dramatically optimize energy efficiency. To this end, researchers will be leading efforts to create new architectures for fast synchronization and communication; innovative new memory hierarchy organizations; aggressive new techniques for power and resiliency management; and in-memory state versioning, among other advances.

In addition, the Center's researchers are participating in the design of programming systems that facilitate coding and enable automatic program optimization. Very high level APIs and compiler strategies that support static and dynamic optimization characterize the programming system of extreme-scale computer systems. The multiple optimization dimensions that must be addressed in extreme-scale computing — including execution speed, energy efficiency, and resiliency — complicate tuning and demand the support of novel and sophisticated compiler strategies.

Thrifty: An Exascale Architecture for Energy Proportional Computing
Funded by: Department of Energy (DOE)

Thrifty aims to provide a platform for highly-efficient, energy-proportional computing — a platform where energy is spent proportionally to the work done, and in the resources that truly are the bottleneck. The Thrifty concept and its software stack simultaneously innovate in power/energy efficiency, resiliency, and performance. Each topic is tackled at multiple layers of the computing stack: circuits, architecture, compilation, and runtime/application.

The Thrifty research involves the development of a detailed architecture simulator of a 1K-core Thrifty system and its evaluation; the implementation and evaluation of a power-aware optimizing compiler based on ROSE; the development of libraries for application characterization and tuning, and the fabrication of a large test chip in 22nm technology by Intel to test novel circuit technologies.

The goal of this project is to attain exascale performance at about 20MW by year 2018-2020 at 8nm technology.

Designing Programmable Many-Core for Extreme Scale Computing
Funded by: National Science Foundation

This project aims to address the challenge of Programmable Extreme-Scale Computing through a multidisciplinary and ambitious effort that cuts across the architecture, compiler, operating system, and correctness/performance tools areas. The Centerís goal is this work to make fundamental advances toward designing a programmable many-core for Extreme-Scale Computing in mobile platforms (netbooks and smart-phones) of the year 2020. Researchers focus on a key concept in this work: that cores and all of the software continuously operate in Chunks (i.e., atomic blocks) of instructions at a time — eliminating the need for in-order, single-instruction-at-a-time commit. To implement this scheme, researchers are working on innovations in architecture, compiler support, operating systems, and correctness/performance tools.

ADSC Working to Analyze Biomedical Data without Risking Individuals' Privacy

Professor Marianne Winslett is leading a new project that will eliminate significant privacy risks that currently impede researchers' ability to analyze biomedical data about individuals. Immense amounts of valuable data now exist that are unusable by the research community due to the lack of an effective method for concealing individuals' identities.

The project is taking place at the Advanced Digital Sciences Center (ADSC) and is funded by A*STAR. The new ADSC work will generate new publication schemes for the results of data analyses, thus making detailed summaries of health data available that can offer unprecedented insight into a vast range of medical conditions and provide useful input for urban planners, public health officials, and researchers.

The $2 million project, "Enabling Mining of Medical Records through Differential Privacy," is led by Winslett. Her co-principal investigators include Prof. Xiaokui Xiao of Nanyang Technological University, Prof. Jiawei Han of the Department of Computer Science at the University of Illinois at Urbana-Champaign, Dr. See Kiong Ng of the Institute for Infocomm Research, and Prof. Nikita Borisov of the Department of Electrical & Computer Engineering at Illinois. The team also includes biomedical researchers Dr. E Shyong Tai from the National University of Singapore and Dr. Edison Liu from the Genome Institute of Singapore.

The widespread availability of biomedical data, ranging from reports of the locations of new cases of dengue fever to individuals' genomic variations, appears to offer researchers a tremendous opportunity. Statistical analysis of such data can help researchers and public health officials better understand a disease and its transmission patterns, gain new insights into the human body, and develop new treatments and services that can improve the quality of life of millions of people.

Unfortunately, privacy concerns make it infeasible to provide researchers with unlimited access to biomedical information. Previous attempts to solve this problem have tried to anonymize data by removing personally identifiable information from medical records, but this does not provide sufficient protection. The main problem is that external knowledge can be used to re-identify individuals whose data appear in supposedly anonymized data sets. Many ideas for mitigating the problem have been proposed, but all of them have made the unrealistic assumption that adversaries had limited prior knowledge.

"In fact, this has been shown to be a fundamental barrier," explains Winslett. "An anonymized database will either reveal private information, given certain external knowledge -- or will be useless for answering some questions."

To the extent that databases of patient information have already been made available, they have made many lifesaving discoveries possible. For example, a University of San Antonio study involving data collected from over 9,000 breast cancer patients showed that amplification of the HER-2 oncogene was a significant predictor of both overall survival and time to relapse in patients with breast cancer. This information subsequently led to the development of Herceptin (trastuzumab), a targeted therapy that is effective for many women with HER-2-positive breast cancer. Likewise, it was medical records research that led to the discovery that supplementing folic acid during pregnancy can prevent neural tube birth defects (NTDs), and population-based surveillance systems later showed that the number of NTDs decreased 31 percent after mandatory fortification of cereal grain food products. No one doubts that additional valuable findings would follow if a way to tackle the privacy limitations can be found, so that far more patient data can be made available to researchers.

To that end, medical studies funded by the National Institutes of Health (NIH) in the U.S. are required to make the data they collect, as well as summaries of analysis results, available to other researchers. Originally, the statistical summaries were freely available to other researchers via NIH's dbGaP database (http://www.ncbi.nlm.nih.gov/gap), while access to the detailed patient records required researchers to undergo a rigorous and fairly arduous approval process with their Institutional Review Boards (IRBs). Privacy concerns subsequently led NIH to restrict dbGaP access, so that today many of the statistical summaries cannot be viewed without IRB approval. The need for IRB approval is a significant hurdle for researchers who want to access the

summary statistics from old studies to help them plan their future work.

To find a practical solution, the ADSC team is using the recently developed concept of "differential privacy." Differential privacy works by adding a small amount of noise to the results of statistical analyses of sensitive data sets. Under differential privacy, the contributions of any one individual's data towards the outcome of an analysis are negligible; analysis results will be essentially identical regardless of whether a particular person's data are included. This should not limit the usefulness of the results, since in a large and well-designed medical study, the history of a single individual should not have a significant impact on overall results. When analysis of a data set begins, its owners decide on a total "privacy budget" for the entire data set. Each published analysis result uses up a little bit of the privacy budget, and once the budget has been exhausted, no more results can be published, as they could open the possibility of at least one individual's data having a non-negligible impact on overall results.

"Differential privacy offers us the tantalizing possibility of being able to do privacy-preserving data analysis that is both useful and secure," says Winslett. "It's such a new concept, but the implications are immense. Whoever comes up with a practical approach to differentially private access to biomedical data -- which is what we aim to develop with this new project -- will set off a free-for-all. It will open up so many new opportunities to revolutionize treatments and reduce health care costs."

The new project is starting by analyzing privacy issues in the statistics released by Singapore's Ministry of Health (MOH). Due to the potential for privacy breaches, MOH currently publishes detailed statistics only for highly dangerous infectious diseases, and only very sketchy information, such as the total numbers of male and female patients in Singapore, for other types of diseases. For instance, one report says that there was exactly one male patient aged 30-39 with relapsed tuberculosis in 2010. The team's goal is to make it possible to publish detailed statistics

for all diseases, but with strong privacy-preservation guarantees.

The project's next step will be to investigate ways to re-enable open access to the summary information in dbGaP by making the summary tables differentially private. The researchers will also target other custodians and users of health-related statistics in Singapore. That work is projected to include applications in pharmacoeconomics and in analysis of hospital records to reveal the effectiveness of different treatments for a disease.

Winslett is quick to point out that several fundamental research challenges remain before differentially private analyses will be practical, but she is optimistic that ADSC has advantages that make it an ideal location for this research. In particular, Singapore is unique in its close cooperation among the government, the medical fraternity, and research institutes. This will give the ADSC researchers exceptionally good access to the parties who have a vested interest in broader dissemination of health data summaries. This concerted effort to bring together medical researchers, computer scientists, and medical records could one day enable Singapore to be a world leader in technologies for analyzing sensitive data.
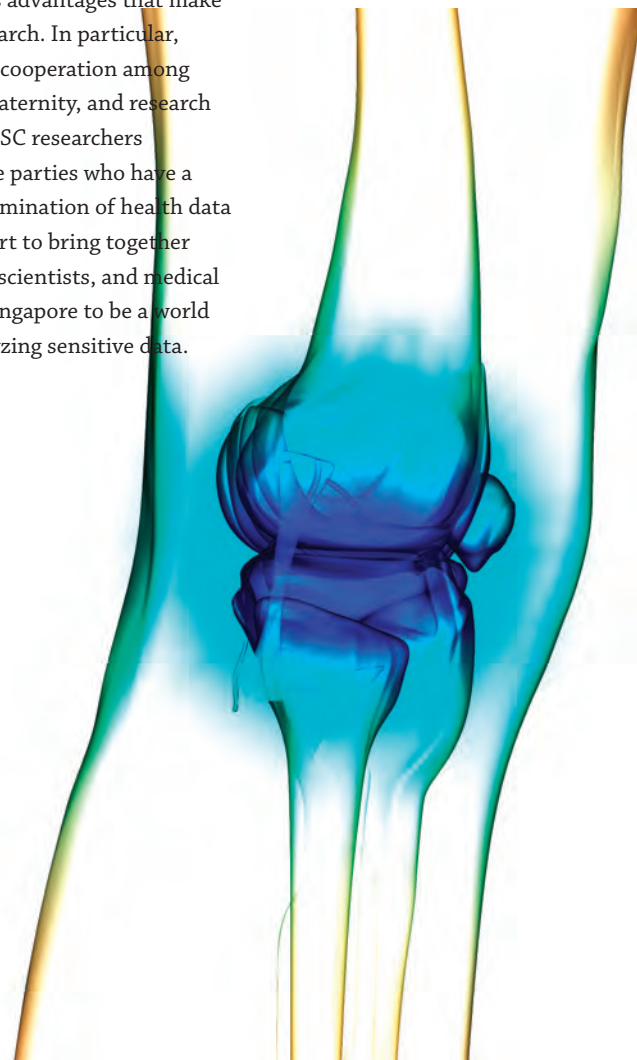
Science of Security "Lablet" Established at the University of Illinois

The University of Illinois at Urbana-Champaign will receive an initial $1 million in grant funds from the U.S. National Security Agency (NSA) to stimulate the creation of a more scientific basis for the design and analysis of trusted systems.

It is widely understood that critical cyber systems must inspire trust and confidence, protect the privacy and integrity of data resources, and perform reliably. To tackle the ongoing challenges of securing tomorrow's systems, the NSA concluded that a collaborative community of researchers from government, industry, and academia is a must.

To that end, the NSA grant will seed an academic "Lablet" focused on the development of a Science of Security (SoS) and a broad, self-sustaining community effort to advance it. A major goal is the creation of a unified body of knowledge that can serve as the basis of a trust engineering discipline, curriculum, and rigorous design methodologies. The results of SoS Lablet research are to be extensively documented and widely distributed through the use of a new, network-based collaboration environment. The intention is for that environment to be the primary resource for learning about ongoing work in security science, and to be a place to participate with others in advancing the state of the art.

The Illinois Lablet, which will be housed in the Information Trust Institute at Illinois, will contribute broadly to the development of security science while leveraging Illinois expertise in resiliency, which in this context means a system's demonstrable ability to maintain security properties even during ongoing cyber attacks. David M. Nicol, the Illinois Lablet's principal investigator, explains, "The complexity of software systems guarantees that there will almost always be errors that can be exploited by attackers. We have a critical need for foundational design principles that anticipate penetrations, contain them, and limit their effects, even if the penetration isn't detected."

The Lablet's work will draw on several fundamental areas of computing research. Some ideas from fault-tolerant computing can be adapted to the context of security. Strategies from control theory will be extended to account for the high variation and uncertainty that may be present in systems when they are under attack. Game theory and decision theory principles will be used to explore the interplay between attack and defense. Formal methods will be applied to develop formal notions of resiliency. End-to-end system analysis will be employed to investigate resiliency of large systems against cyber attack. The Lablet's work will draw upon ideas from other areas of mathematics and engineering as well.

Nicol, the project's principal investigator, is a professor of Electrical and Computer Engineering (ECE) at Illinois and the director of the Information Trust Institute. The Lablet's leadership is shared with co-principal investigators William H. Sanders, who is an ECE professor and director of the Coordinated Science Laboratory at Illinois, and Josè Meseguer, a professor of Computer Science.

> **It is widely understood that critical cyber systems must inspire trust and confidence, protect the privacy and integrity of data resources, and perform reliably.**

Finding Lost Data

Misplace your keys and the advice you'll be given is to "try and remember where you were the last time you had them." Why shouldn't finding your lost digital data work the same way?

Thanks to new technology developed by University of Illinois computer science PhD student Joshua Hailpern, it can. Hailpern's software, YouPivot, helps you retrace your steps back to your lost files, websites, or documents by tapping into how the human memory works.

YouPivot keeps track of your activities on your computer at any given point in time - music you're playing, websites you're browsing, files you're working with. Using this data as a reference, YouPivot allows you to browse, search and "pivot" to see what else was happening on your computer at the time and rediscover what you are looking for. As Hailpern told IEEE Spectrum in a recent article, "What was that website or PDF I was looking at when AC/DC's 'Hells Bells' was playing?"

YouPivot works with even a vague recollection of what activities or data might be associated with the information you are seeking. Unlike current search methods for email and browser history, YouPivot does not require specific keywords or dates to recall information. What's more, pivoting is only one of the many new features Hailpern created to allow users to find their "digital car keys."

YouPivot is currently available in alpha as a free extension to the Google Chrome browser, though Hailpern and his team are working on a beta release by end of 2011, and full release in spring of 2012.

Hailpern's work on YouPivot was recently profiled in an IEEE Spectrum article "A Prototype of Pivot Searching." Learn more about Hailpern's technology by reading the article, or by viewing a YouTube video on the software.

awards

Sarita Adva
IEEE Fellow
ACM Fellow

Robert Bocchino
ACM SIGPLAN Dissertation Award

Matthew Caeser
NSF Career Award

Brighten Godfrey
NSF Career Award

Bill Gropp
IEEE Fellow
SIAM Fellow
National Academy of Engineering

Julia Hockenmaier
NSF Career Award

Derek Hoiem
NSF Career Award

Forrest Iandola & Rohan Sharma
CRA Outstanding Undergraduate Researcher Award, Finalists

Brett Jones & Raj inder Sodhi
Qualcomm $100,000 Innovation Fellowship

David Liu, emeritus
Phil Kaufman Award for Distinguished Contributions to Electronic Design Automation

Dan Roth
ACM Fellow

Josep Torrellas
ACM Fellow

Tianyi Wu
ACM SIGKDD Dissertation Award, Runner-Up

Best Paper Awards

2011 ACM Multimedia Conference (MM 11): Best Student Paper Award, November 2011

"Color-plus-Depth Level-of-Detail in 3D Teleimmersive Video: A psychophysical Approach" W. Wu, A. Arefin, G. Kurillo, P. Agarwal, K. Nahrstedt, R. Bajcsy

2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA 2011): Best Paper Award, August 2011

"Controllable Hand Deformation from Sparse Examples with Rich Details" H. Huang, L. Zhao, K. Yin, Y. Qi, Y. Yu, X. Tong

2011 IEEE Symposium on Reliable Distributed Systems (SRDS 2011): Best Paper Award, October 2011

"Modeling the Fault Tolerance Consequences of Deduplication" E. Rozier, W. H. Sanders, P. Zhou, N. Mandagere, S. M. Uttamchandani, M. L. Yakushev

2011 ACM Conference on Information and Knowledge Management (CIKM 2011): Best Student Paper Award, October 2011

"Lower Bounding Term Frequency Normalization" Y. Lv, C. Zhai

2011 International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2011), August 2011

"Evolutionary Clustering and Analysis of Bibliographic Networks" M. Gupta, Y. Sun, C. Aggarwal, J. Han

2011 Parallel Architectures and Compilation Techniques (PACT 2011): Best Paper Award, October 2011

"DeNovo: Rethinking the Memory Hierarchy for Disciplined Parallelism" B. Choi, R. Komuravelli, H. Sung, R. Smolisnski, N. Honarmand, S. Adve, V. Adve, N. Carter, C-T. Chou

2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2011): Best Student Paper Award, September 2011

"Simplicial Dijkstra and A* Algorithms for Optimal Feedback Planning" D. Yershov, S. M. LaValle

2011 IEEE Computer Vision and Pattern Recognition (CVPR 2011): Best Student Paper, June 2011

"Recognition Using Visual Phrases" A. Farhadi, M. A. Sadeghi

2011 Conference on Computational Natural Language Learning (CoNLL 2011): Best Student Paper, June 2011

"Adapting Text Instead of the Model: An Open Domain Approach" G. Kundu, D. Roth

2011 IEEE Transactions on Computer-Aided Design: Donald O. Pederson Best Paper Award, April 2011

"Statistical Blockade: Very Fast Statistical Simulation and Modeling of Rare Circuit Events and Its Application to Memory Design" R. Rutenbar, A. Singhee

2011 IEEE International Conference on Data Engineering (ICDE): Best Student Paper Award, April 2011

"How Schema Independent Are Schema Free Query Interfaces?" A. Termehchy, M. Winslett, Y. Chodpathumwan

2011 IEEE Pervasive Computing and Communications Conference (PERCOM 2011): Mark Weiser Best Paper Award, March 2011

"Jyotish: A Novel Framework for Constructing Predictive Model of People Movement from Joint Wifi/Bluetooth Trace" L. Vu, Q. Do, K. Nahrstedt

David Liu

Professor Emeritus David Liu Receives Phil Kaufman Award

Professor emeritus Dr. C. L. David Liu was the winner of the 2011 Phil Kaufman Award for Distinguished Contributions to Electronic Design Automation (EDA). The Phil Kaufman Award honors individuals who have made an impact on the field of EDA and pays tribute to Phil Kaufman, the late industry pioneer who turned innovative technologies into commercial businesses that have benefited electronic designers.

Early in his career, Dr. Liu led the transformation from ad hoc EDA to algorithmic EDA. He was an early advocate for more rigorous design automation, arguing that powerful, formal algorithmic techniques were essential to the effective solution of complex design automation problems. His technical contributions are at the foundation of a multitude of current EDA tools within several disciplines, including behavioral synthesis, logic synthesis and physical design.

His technical impact includes the first floorplanning algorithms and scheduling algorithms for hard real-time tasks. His research on floorplanning received DAC's Best Paper Award in 1986 and has been widely regarded as seminal. Dr. Liu's work on Rate Monotonic Scheduling (RMS) is a cornerstone of modern scheduling theory, applicable in the design of real-time operating systems. As of today, his 1973 paper on the subject has over 7,000 citations.

"We should not be surprised when one of our community's leading technical contributors has remarkable leadership and business skills as well," observes Rob A Rutenbar, Abel Bliss professor and head the Department of Computer Science at the University of Illinois at Urbana-Champaign, a colleague of Dr. Liu's for more than 25 years. "But neither should we fail to notice when a colleague like David Liu manifests in such an impressive sphere of activities, visible worldwide."

Over the past 12 years, Liu's contribution to

Taiwan's semiconductor industry has been broad and significant. He serves as chairman of the board of TrendForce, a market intelligence provider in the DRAM, LED, LCD and Solar Energy technical segments. He is a member of the board of Powerchip Semiconductor Corp., United Microelectronics Corp., MediaTek and Macronix International Co., Ltd. Additionally, he is a member of the board of Anpec Electronics Corporation, Andes Corporation, and Cadence Methodology Service Company.

For the last six years, he has hosted a weekly radio show on technology, natural science, social science and literature in Taiwan. He has published three essay collections based on the presentations in the show. One of them is a 2011 winner of a book award in the category of popular science.

Kolla and Lazebnik Join CS@ILLINOIS

Two new faculty will be welcomed at Illinois beginning in 2012! Alexandra Kolla (Microsoft Research) and Svetlana Lazebnik (University of North Carolina - Chapel Hill) have joined the faculty beginning in January 2012.

Tackling some of the hardest problems in their field, these accomplished young researchers and educators are exploring innovative new avenues in complexity theory and computer vision.

Alexandra Kolla research lies at the intersection of theoretical computer science and discrete mathematics, with a particular interest in the mathematical tools used to solve problems in graph theory, convex optimization, approximation algorithms, and complexity theory. Inspired by spectral graph theory and semidefinite programming (sDP), Kolla aims to understand how these tools can be used to solve outstanding questions in the field of algorithms and complexity theory.

Kolla adds terrific depth to the College's expertise in complexity theory. She is working on some of the most difficult and most high-impact problems, trying to answer fundamental questions about what we can and cannot compute efficiently.

Svetlana Lazebnik's research interests in computer vision have led her to explore topics ranging from modeling and organizing large-scale Internet image collections to developing effective image representations for recognition and comprehensive 2D and 3D scene descriptions. her work has implications for assistive technologies, security and surveillance, navigational systems, field biology, and more.

Lazebnik joins the world-leading computer vision group at Illinois. Her innovative work to integrate statistical and geometric techniques to describe images, and her use of information theoretic models to model famous landmarks and entire cities from very large-scale image collections, will open up exciting new collaborations at Illinois.

Gropp Co-Founds ACM High Performance Computing Group

University of Illinois computer science professor Bill Gropp was one of the driving forces behind the launch of a new ACM Special Interest Group on High Performance Computing. Gropp and several others founded the group to address the needs of students, faculty, and practitioners of high performance computing. SIGHPC is the first such group within a major professional society to address this emerging and important topic in computing.

SIGHPC will help address this challenge for professionals and students alike by:

- disseminating research and experiences by those using computing resources to tackle our society's toughest problems;

- promoting the mentoring and education of the next generation of HPC professionals; and

- serving as a source of information about the field to the larger computing and scientific communities.

Snir Named HPC Person to Watch in 2012

Professor Marc Snir has been named one of HPCWire's People to Watch in 2012.

Each year, HPCwire announces its annual "People to Watch" list, comprised of an elite group of community leaders selected from academia, government, business, and industry who we predict will be impacting the world in 2012 and beyond.

Illinois computer science professor Marc Snir was named an HPC Person to Watch for 2012

"The recent flurry of news last fall surrounding IBM's departure from the 10 petaflop NCSA petaflop "Blue Waters" Project and Snir's role as a co-PI for the project will share the focus of a lot of attention this year as the supercomputer nears its delivery date, slated for fall 2012. Combined with his role as a one of the original co-founders and as a co-chair for the relatively new Graph 500, all but guarantees that Snir will be a person of interest to follow this year," writes HPCWire in announcing Snir's selection.

Snir, the Michael Faiman and Saburo Muroga Professor in computer science, has been a leader in shaping high performance computing (HPC) architectures and parallel programming, including contributions to IBM's SP and Blue Gene systems and to MPI, the standard communications library used in HPC.

Snir to Lead Mathematics and Computer Science Division of Argonne

Professor Marc Snir has been chosen to direct the Mathematics and Computer Science Division (MCS) at Argonne National Laboratory (ANL). Snir, the Michael Faiman and Saburo Muroga Professor in computer science, has been a leader in shaping high performance computing (HPC) architectures and parallel programming, including contributions to IBM's SP and Blue Gene systems and to MPI, the standard communications library used in HPC. At MCS, he will be directing over 200 researchers and staff members, who are working on projects ranging from algorithm development and software design in key areas like optimization, to exploration of new technologies such as distributed computing and bioinformatics, to numerical simulations in challenging areas like climate modeling. Snir will continue to hold his appointment as professor of computer science. He will divide his time between MCS and the University of Illinois and will continue to be associated with the Blue Waters project.

"Argonne is one of the most prestigious national labs, conducting significant research across many key disciplines. This is a rare and outstanding opportunity, not only for Marc, but also for the Department of Computer Science and the College of Engineering. The potential for new collaborations between our two world-class institutions is exciting," said Rob A. Rutenbar, the Abel Bliss Professor of Engineering and computer science department head.

A distinguished researcher and scholar, Snir chaired the Department of Computer Science at the University of Illinois, Urbana-Champaign from 2001 to 2007. While at Illinois, he has also co-directed the Intel and Microsoft Universal Parallel Computing Research Center, was the first director of the Illinois Informatics Institute, is the Associate Director for Extreme Scale Computing at NCSA, and is co-PI of the Blue Waters petascale computing project. In addition, Snir co-chaired the National Research Councilís Committee to Study the Future of Supercomputing, and he is a co-author of its influential 2004 report, "Getting Up to Speed: The Future of Supercomputing."

Han named Bliss Professor

Professor Jiawei Han was among six College of Engineering faculty members to be recognized with Bliss Professorships.

The generous Bliss bequest, established by Helen Eva Bliss in memory of Abel Bliss Jr., is used to advance scholarly activities in the College of Engineering. Holders of college professorships are nominated by the dean upon recommendation of the College Advisory Committee on Endowed Appointments and approval of the Provost. Faculty members with named professorships are bestowed to recognize leaders who are among the most talented and accomplished on our faculty.

Han's groundbreaking and highly influential research

has made him one of the top computer scientists in the world. With a focus on knowledge discovery and data mining, data warehousing, and database systems, he is recognized as a pioneer in the field. Han was the first to introduce a pattern-growth methodology for mining frequent, sequential, and structured patterns, as well as the first to develop a set of important algorithms for mining such patterns. These methodologies have been extremely influential in subsequent research and are widely used. Google Scholar lists Han as the second most cited author in the field of data mining with his FP growth algorithm cited more than 3,700 times. This algorithm or its variations have been introduced in most data mining textbooks and has been used at various companies such as Google and Microsoft.

Torrellas Named Program Chair for ISCA 2012

Professor Josep Torrellas has been named the Program Chair for the 39th Annual International Symposium on Computer Architecture (ISCA) in 2012. ISCA is the premier forum for new ideas and experimental results in computer architecture, covering topics ranging from parallel and multi-core systems, interconnection networks, dependable architectures, power and energy efficient architectures, and more.

As Program Chair, Torrellas will be charged with the technical contents of the conference: the papers presented, keynote speeches, and panels. Torrellas will select the program committee members and external reviewers that will review the papers, prepare the final program, and choose the keynote speakers and the panels in the conference, among other responsibilities.

"I am honored at this appointment because ISCA is the premier conference in computer architecture. It allows me to contribute with my efforts to the continued excellence of our research field," said Torrellas. "I am

particularly excited at the fact that our field is now bursting with new concepts that range from data centers to hand-held devices, and from novel memory and interconnect technologies to DNA-based and brain computing. I hope the final program captures the most exciting aspects of this research."

Torrellas and his students have been a major contributor to the ISCA conference for many years. At ISCA 2011, he and his ex-PhD students authored 6 of the 40 papers presented. The papers were authored by Jose Renau from Univ. of California Santa Cruz, Milos Prvulovic from Georgia Tech, Yan Solihin from North Carolina State Univ, Michael Huang from Univ of Rochester, and by Josep Torrellas.

Other Illinois researchers and alumni will be making contributions to ISCA 2012 as part of the organizing and program committee, including Sarita Adve, Wonsun Ahn, Pradip Bose, Luis Ceze, Tom Conte, Xin Fu, Scott Mahlke, Jose Martinez, Trevor Mudge, Jose Renau, Yan Solihin, and James Tuck.

Roth Named ACM Fellow for 2011

Professor Dan Roth was named a 2011 ACM Fellow for his "contributions to machine learning and natural language processing." The ACM Fellows Program celebrates the exceptional contributions of the leading members in the computing field. These individuals have helped to enlighten researchers, developers, practitioners and end-users of information technology throughout the world. These new ACM Fellows join a distinguished list of colleagues to whom ACM and its members look for guidance and leadership in computing and information technology.

Dr. Roth's research interests focus on the computational foundations of intelligent behavior. He has made seminal contributions to several aspects of this problem — from theoretical questions in Learning and Reasoning, to new models and algorithms that have contributed significantly to the use of machine learning and inference techniques in natural language processing. Advances made by Roth have changed how computer scientists develop algorithms and programs for natural language understanding, and how they think about computational modeling of learning and reasoning.

In his research, Roth has pursued several interrelated lines of work that span multiple aspects of this problem — from fundamental questions in learning and inference and how they interact, to the study of a range of natural language processing (NLP) problems, including multiple disambiguation problems, shallow parsing, semantic role labeling, co-reference, question answering and textual entailment, to large scale Natural Language Processing and Information Extraction system development - resulting in a number of software packages for NLP tools.

He has published broadly in machine learning, natural language processing, knowledge representation and reasoning and learning theory, and has developed advanced machine learning based tools for natural language applications that are being used widely by the research community, including an award winning Semantic Parser.

Dr. Roth is the Director of the DHS Institute of Discrete Science Center for Multimodal Information Access & Synthesis Research (MIAS). He is a fellow of the Association for the Advancement of Artificial Intelligence, the premier AI professional society and was an NSF CAREER Award winner. Roth received his Ph.D. in Computer Science from Harvard University in 1995 and is a Willett Faculty Scholar of the University of Illinois College of Engineering.

Adve Named IEEE Fellow for 2012

Professor Sarita Adve has been named an IEEE Fellow for the class of 2012. Adve was selected for her "contributions to shared memory semantics and parallel computing." Adve's research in computer architecture and systems, parallel computing, and power and reliability-aware systems focuses on a full-systems view and is notable for its multidisciplinary collaborations.

Adve's broadest impact has been in hardware and software memory consistency models. She received the 2008 SIGARCH Maurice Wilkes award for this work, specifically "for formalization of memory consistency models, especially data-race free models, and their influence on both hardware and high-level languages."

The memory consistency model lies at the heart of the semantics of any threaded software or hardware. Arguably, it has been one of the most challenging and contentious areas in concurrent hardware and software specification for many years. There is finally now a convergence in both the hardware and software communities. Adve has been a common thread and a leader in the multiple community-scale efforts that have driven this convergence.

The memory consistency model is a hardware/software interface, affecting programmability and performance. Unfortunately, designing a model satisfying all desirable properties has proven difficult. Sequential consistency is simplest to program, but most systems do not provide it for performance reasons. Instead, we had divergent models (often ambiguously specified) for different hardware.

Adve's early work departed from the prevalent hardware-centric approaches to use a combined hardware/software view more appropriate for an interface. She observed that for well-synchronized programs (which she formalized as data-race-free), both sequential consistency and high performance can be provided. She developed a comprehensive framework to specify memory models as providing "sequential consistency for data-race-free programs." Adve's data-race-free model forms the foundation of the memory models for Java and C++.

Adve has been a leader in power- and reliability-aware architectures. Her group was among the first to recognize that significant power reductions required breaking traditional system boundaries in favor of a collaborative, cross-layer system-wide power management framework. The GRACE project she led was the first to demonstrate a prototype system where the hardware, operating system, network, and applications all adapted collaboratively to minimize energy while still meeting real-time quality of service requirements.

Most recently, Adve's work has challenged the research community to rethink how both parallel languages and parallel hardware are designed. Adve recently won honors for her DeNovo system, which exploits emerging software trends in disciplined parallel programming to make hardware simpler, higher performance, and lower energy, all at the same time. As their first step, Adve and her research team have developed a cache coherence protocol and consistency model that takes an order of magnitude less time to verify and runs some applications in less than half the time with less than half the network traffic and cache misses than the state-of-the-art. The simplicity and low network and cache traffic means that the performance increases come with significant power and energy benefits, a rarity in hardware design.

Professor Adve received the Maurice Wilkes Award in 2008, an IBM Faculty Award in 2005, was named a UIUC University Scholar in 2004, and received an Alfred P. Sloan Research Fellowship in 1998, an IBM University Partnership award in 1997 and 1998, and a NSF CAREER award in 1995.

She served on the National Science Foundation's CISE directorate's advisory committee from 2003 to 2005 and on the expert group to revise the Java memory model from 2001 to 2005. She co-led the Intel/Microsoft funded Universal Parallel Computing Research Center (UPCRC) at Illinois as its director of research in its founding year (2008-09). She currently serves on the board of directors for ACM SIGARCH and the Computing Research Association (CRA).

Department of Computer Science
University of Illinois at Urbana-Champaign
Thomas M. Siebel Center for Computer Science
201 N Goodwin Ave
Urbana, IL 61820

All gifts make a difference in upholding our past legacy, preserving our present excellence, and reaching out to future generations. If you would like to make a contribution, please visit **http://cs.illinois.edu/alumni/giving** or scan this QR code

click!

cs.illinois.edu