UMIACS
UNIVERSITY OF MARYLAND INSTITUTE FOR ADVANCED COMPUTER STUDIES

InterConnections

AUGUST 2008

Contents

InterConnections is published by the University of Maryland Institute for Advanced Computer Studies (UMIACS).

Visit UMIACS on the web at: *www.umiacs.umd.edu*.

Director: V.S. Subrahmanian

Contact information
UMIACS
2119 A.V. Williams Bldg.
University of Maryland
College Park, MD 20742-3251
Phone: 301.405.6722

Written by Karen Jegalian
Edited by Jennifer Newlin

Mihai Pop: The Next Generation of Bioinformatics

On the way to completing a Ph.D. in computer science at Johns Hopkins University in 2000, Mihai Pop became interested in bioinformatics. "I wanted to do research on something that might have an impact on human life," he said. After five years at The Institute for Genome Research (TIGR), Pop joined UMIACS in 2005.

Since the early days of graduate school, Pop has worked on what he calls infrastructure work—writing genome assembly software and working to automate the process of assembling DNA sequence data. That problem appeared to have been more or less solved, allowing more and more species' genomes to be fully sequenced over the last decade. Recently, though, new sequencing technologies have presented fresh challenges. New techniques rapidly generate large amounts of sequence data, which is a boon for allowing the analysis and comparison of DNA, but the new, faster sequencing techniques produce shorter segments of data. A single rapidly churning machine produces an overwhelming amount of data—a terabyte an hour. "Most of our computers can't handle one terabyte of data per hour," notes Pop. "Just transferring data from a sequencing machine to a disk causes computing headaches now."

The new technologies are pushing high-performance computing to the forefront of genome researchers' attention again. One challenge with analyzing the new data is distinguishing real DNA sequence differences from sequencing errors. But computational techniques for assembling sequences are not only faster but often better than relying on humans to analyze and finish sequence assembly. "We often find that the computational method is better," said Pop, "eliminating human error."

Pop collaborates with Steven Salzberg, director of the UMIACS Center for Bioinformatics and Computational Biology (CBCB), on these problems, funded by the Naval Medical Research Center and startup funds from the University.

From Single Species to Populations

Another area of innovation in genome science is an increasing ability to study the DNA of mixed populations of organisms rather than of individual species. Pop is collaborating with UMIACS/Cell Biology and Molecular Genetics associate professor Najib El-Sayed to develop bioinformatics tools for analyzing mixed populations.

Metagenomics is the name for the emerging field which allows researchers to gauge similarities and differences

continued on page 2.

Director's Message

UMIACS Director V.S. Subrahmanian

The first half of 2008 has been an extraordinarily successful period for UMIACS, with an unprecedented set of success stories.

UMIACS Professor Uzi Vishkin garnered world headlines with his invention of a new desktop supercomputer. Vishkin's success comes after over 25 years of dedicated research following the PRAM paradigm that had initially been dismissed by researchers as impractical. An article on Vishin's invention is contained in this issue.

The second major breakthrough in 2008 was the establishment within UMIACS of the first WIMAX Forum lab in the western hemisphere. Called the Maxwell Lab, and headed by UMIACS Professor Ashok Agrawala, the Maxwell Lab will provide an infrastructure to test out applications that use the next generation of wireless technology. The MAXWELL Lab will be featured in the next issue of this newsletter.

The third major breakthrough was the creation of the world's first social network environment for the study of terrorism. Developed within the Lab for Computational Cultural Dynamics, the SOMA Terror Organization Portal (STOP) has registered users from several defense organizations and has tools to learn and forecast behaviors of terror groups. STOP will also be featured in the next issue of this newsletter.

In addition, UMIACS faculty and students received numerous awards during this period, delivered numerous invited lectures around the world, and received considerable coverage of their work.

Pop, continued from page 1.

among populations of cells. Such comparisons call for intensive analysis. "New statistical methods will be needed," Pop states.

"As sequencing becomes cheaper, we can look at whole populations of bacteria instead of cultured bacteria," explains Pop. Only about 1 to 5 percent of bacteria can be cultured, and without culturing bacteria into dense populations, it is impossible to do conventional genome sequencing. In studies at TIGR, Pop and collaborators began bacterial population studies by looking at micro-organisms from the human gut.

"There are ten times as many bacteria cells in the human body as human cells. They play an important role in who we are. They produce a lot of the nutrients we need. We couldn't live without them," Pop said.

The field of microbial community study is at its beginning. Studies have looked at populations of bacteria in the ocean and in highly acidic environments, and new studies are attempting to gain insights on a range of diseases, including obesity, autism, and Crohn's disease. "With new sequencing techniques, we can afford to do very deep sampling of many people," Pop said.

Pop is collaborating with a team at the University of Maryland School of Medicine that recently received funding from the Gates Foundation to study bacterial populations in the human gut and their relationship to diarrhea. The researchers will be comparing bacterial populations in the intestines of children who have diarrhea with those who do not. Although diarrhea is one of the biggest causes of childhood mortality, the range of bacteria that causes diarrhea is not well understood. Finding the causative agents could help in the development of better treatments. The newly launched study will examine bacteria samples from children in Bangladesh, Mali and Kenya. The researchers will sample about 100 DNA sequences from bacteria taken from approximately 1200 children. The team plans to see whether there are systemic differences in the bacteria of those with diarrhea.

"The big challenge in the field right now is that the statistics you can develop for small sample sets don't apply to these large sets of

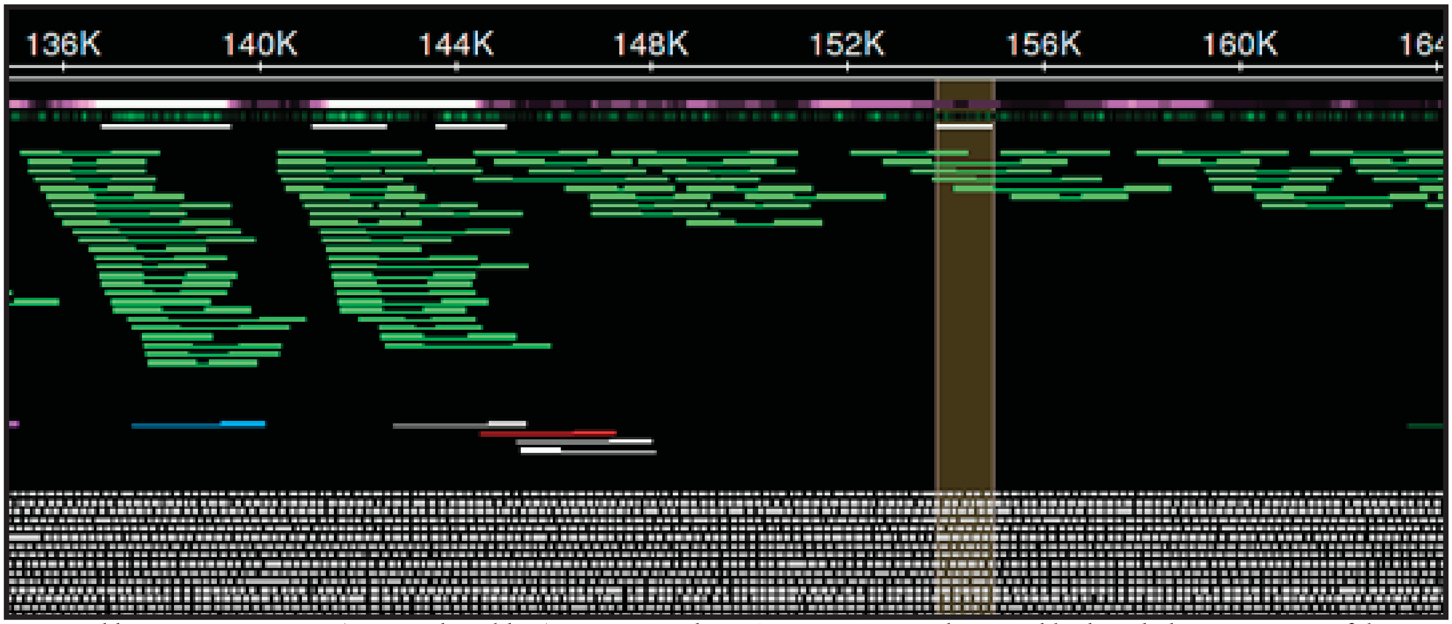

To assemble genome sequence (top, numbered line), conventional DNA sequencing machines yield relatively large segments of data that need to be assembled (green lines), whereas new techniques generate data much more rapidly but yield much shorter segments (white lines), creating a challenging assembly problem.

samples," Pop said. He is working to develop better methods of analysis. Another challenge in the diarrhea project will be the sheer amount of data. The team will have thousands of sequences to align and compare.

In his earlier work on metagenomics of human gut bacteria, (see Science, June 2, 2006), Pop collaborated with Steve Gill, formerly at TIGR and now at SUNY Buffalo. "Mihai Pop is an outstanding computational biologist and collaborator," Gill said. "Mihai is unique in that his focus is not only on computational tools, but on the larger biological questions that are only approachable because of his expertise."

"With metagenomics being relatively new, Mihai has already established himself as a leader in the field and will continue to have a very significant influence on development of the tools urgently needed for analysis of metagenomic data," Gill adds.

In an entirely different sort of metagenomics project, Pop is working with Jocelyne diRuggiero, professor of Cellular Biology and Molecular Genetics, University of Maryland. They are examining the bacteria in the Atacama Desert in Chile, one of the driest places on Earth. The micro-organisms there are able to withstand intense ultraviolet light radiation and survive despite lack of water and nutrients. As metagenomics projects go, this one should be simplified by the relatively low number of species that can survive under such conditions.

"He's been very good about thinking about new ideas and new approaches to this problem," diRuggiero said of Pop. "He's brought great ideas on how to analyze the data." Regarding having CBCB at the University, diRuggiero adds, "Now I can think about doing a project like this without having a hard time finding collaborators. It's helpful to be able to meet face to face. The personal interaction brings the scientific discussion to another level and I think it really works well."

Sorting Through Biological Data

In a new project, Pop and one of his students are building a database of genes that confer antibiotic resistance. The researchers hope to help standardize how such genes are described. Their goal is to be able to predict what kind of genetic changes can lead to antibiotic resistance. For example, cell membranes have various efflux proteins with normal biological roles, such as ridding cells of waste products. The researchers' analysis could identify how normally innocuous efflux proteins become antibiotic resistant with the addition of one or two particular types of mutations that allow the proteins to pump out antibiotics. Pop and his team want to be able to say which specific sequence changes can confer antibiotic resistance. "This is definitely important to human health, and surprisingly little is known," said Pop.

As many biologists and computer scientists note, biology is increasingly becoming an information science. Molecular biology experiments increasingly rely on complex statistical techniques for interpretation. Bioinformatics helps sort through the data. "There will be more and more high-throughput techniques being developed, and someone will need to analyze them," said Pop. His goal is to develop a broad array of tools for analyzing large amounts of biological data and to be actively involved with biologists to make discoveries.

Increasingly, computer scientists can suggest refinements to biological experiments, Pop said, to develop better approaches to solving problems. "Biology and computer science should provide feedback to each other and not be separate from each other," Pop adds.

Of their work together in the past, Gill said, "Mihai was not only responsible for the majority of the data analysis, but he also had significant input on the overall concepts and biological significance of the study."

Uzi Viskin: A Desktop Supercomputer

Most people still think of supercomputers as refrigerator-sized machines whirring away in bare, climate-controlled rooms. Uzi Vishkin is working to make desktop supercomputers as common and unremarkable as personal computers. His vision is to make the next generation of personal computers machines that use parallel processing and compute a hundred times faster than current desktop machines. Faster computers can enhance practically any computing task, from drug discovery to virtual reality. The supercomputers he envisions will be run by single, small chips, essentially succeeding Pentium chips. Last year, Vishkin's research group built a prototype based on a circuit board about the size of a license plate mounted with 64 parallel processors. In the future, Vishkin foresees being able to arrange a thousand processors on a chip the size of a fingernail.

More than simply succeeding in packing together a great deal of computing power on a single circuit board, Vishkin and his team carefully laid out the processors to have a parallel organization. This allows them to work together efficiently and run software that is practical and straightforward to write. For his work, Vishkin was named an "Innovator of the Year" in 2007 by The Daily Record, an award sponsored by the Maryland Department of Business and Economic Development.

"Today's multi-core processors support coarse grain parallelism. Professor Vishkin has defined a new parallel architecture that supports extremely fine-grained threading," said Geoff Lowney, an Intel Fellow. "Professor Vishkin is proposing a compelling alternative design for future microprocessors."

Vishkin uses an analogy to explain his approach to power computing. A single person may be able to clean a house in 5 hours. In theory, the same amount of labor could be done in 3 minutes by 100 workers, but only if the workers were meticulously organized. "You see here the promise of an opportunity but also a significant intellectual challenge in getting this to work," said Vishkin.

"Mathematics prepared us for sequential computing but not for parallel computing," he said. Until 2003, sequential computing chips got progressively faster and faster. The rate at which computers performed operations sped up because of advances in chip fabrication technology and miniaturization. In 1946, the first computer operated at 5000 "ticks," or operations, per second; by 2003, computer chips were able to tick 4 billion times per second. In 2003, the ever-growing speed of conventional computer chips hit a wall due to overheating. Theoretical limitations, i.e., the speed of light, ensure that this is not just a transient problem. Vishkin and others had long anticipated that computers would reach these limits, and he began thinking of ways to bypass the limitations in 1979, while working on his Ph.D. in computer science at the Technion - Israel Institute of Technology.

Throughout his career, Vishkin has worked on improving computer productivity by distributing computational load among multiple processors. In 1979, he began his first stage of this work by contributions to the development of a theory of parallel algorithms—PRAM, for Parallel Random Access Machine. Around 2004, to continue to make more powerful computers, manufacturers started to build desktop computers with parallel processors. The speed of individual calculations is no faster in parallel-processing machines, but more processors on a chip allow simultaneous calculations to occur on parallel tracks.

Building Hardware With Software in Mind

Vishkin, who earned BSc and MSc degrees in mathematics before obtaining a Ph.D. in computer science, combines the perspective of a computer engineer who designs computer chips with that of a computer scientist who writes programs. The problem with the current practice of building increasingly parallel processors, Vishkin explains, is that writing programs for them becomes progressively more complicated. "Between a computer and an application, you have programmers. The single most expensive component in the information technology enterprise is the programmer's time," he said. "The difference between my beliefs and most of the community is that I think it's a waste of time to build machines before we can figure out how to use them."

Now, the technology market faces a challenging time. Manufacturers are selling multicore computers built with parallel processors but are unable to explain how to effectively program the machines so that they actually complete a given task more quickly. Application software makers face the prospect of writing code for fundamentally different hardware but have little incentive to invest in a programming solution when it is unclear which platforms will win over the market. Vishkin argues that computer science students are being taught to program in the same way that they have been taught for years. "We produce young dinosaurs. There is a fog on where we are going. It's a serious predicament," Vishkin said. In fact, at a Microsoft workshop on many-core computing Vishkin decried the lack of clear direction from the computer industry and warned, "There is an appearance of cluelessness."

Vishkin warns that industry is betting on a direction for the future of computing without knowing how that path will take shape. He argues that dramatic changes in hardware should be complemented with a clear vision for future software development. "What we present is a school of thought that is coherent and addresses all the issues," Vishkin said.

"The single-chip supercomputer prototype built by Professor Uzi Vishkin's group uses rich algorithmic theory to address the practical problem of building an easy-to-program multicore computer," comments Charles E. Leiserson, a professor of computer science and engineering at the Massachusetts Institute of Technology.

In the 1980s, Vishkin and his colleagues worked heavily on the PRAM theory, offering a comprehensive answer to how to produce parallel algorithms. The key strength of PRAM is that it suggests a way of orchestrating processors together, and the theory is embedded with the methods for programming such machines.

The theory, however, has had its ups and downs. In the early 1990s, the PRAM approach was considered impractical because of the slow speed of data movement between separate, parallel chips. During this time, Vishkin continued his work on building a chip prototype. Eventually, advances in on-chip interconnection made PRAM feasible after all, and data can move quickly among the processors.

"PRAM was just ahead of its time," he said. "The extra push needed is much smaller than you would guess." However, PRAM methods can't be applied to the types of multicore processors that are currently being built.

"You can think of the theory as the specifications for the machine being built," said Vishkin. When one builds a computer, what is built should complement how it will be used, Vishkin emphasizes. "The way most of the industry operates is to build first and figure out how to program later. At UMIACS, we have been following a different route," he said.

Vishkin gave the computer science community its first chance to test his prototype "explicit-multi-threaded" or XMT computer at the Association for Computing Machinery (ACM) International Conference on Supercomputing. Leiserson commented that the prototype addressed the problem of building an easy-to-program multicore computer. "Vishkin's chip unites the theory of yesterday with the reality of today," he said.

"This system represents a significant improvement in generality and flexibility for parallel computer systems because of its unique abilities," added Burton Smith, a Technical Fellow at Microsoft. "It will be able to exploit a wider spectrum of parallel algorithms than today's microprocessors can, and this in turn will help bring general purpose parallel computing closer to reality."

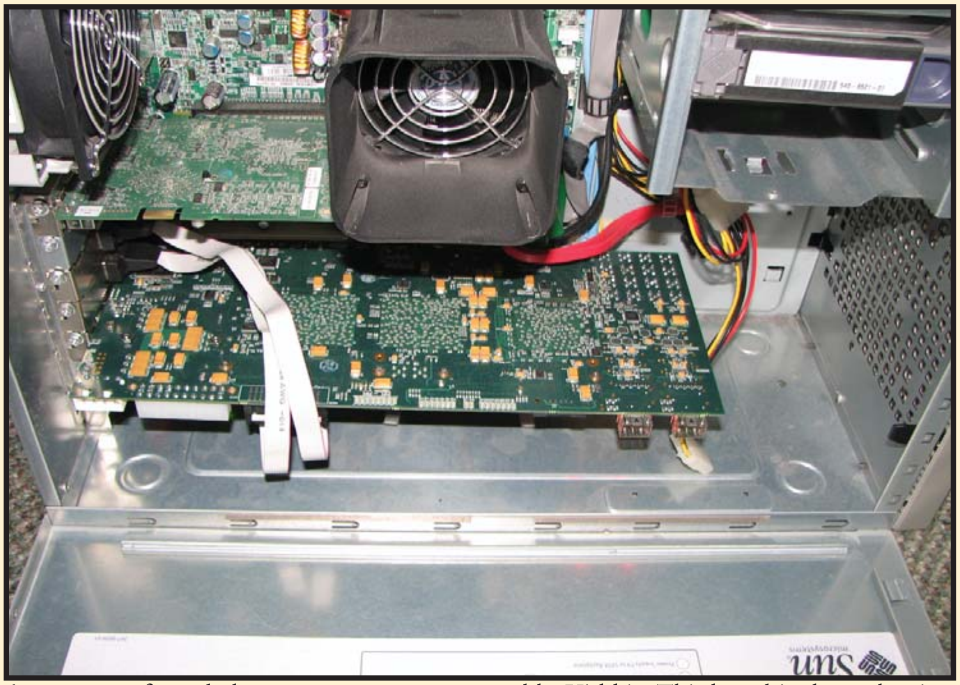

A prototype for a desktop supercomputer created by Vishkin. This board is about the size of a license plate. Manufactured chips would be smaller than a fingernail.

Avoiding a Dead-End

The strategy of placing multiple processors on a single chip has spread widely in the semiconductor industry since 2004. But engineers and computer scientists acknowledge that the software industry lags in the ability to write parallel programs for the chips that are being manufactured and marketed.

Vishkin has long advocated PRAM as a theory that can avoid the dead-end the industry appears to be facing. He notes that it's important for industry to settle on a viable microprocessor architecture before investing more resources into potentially useless devices. Architecture instability, he warns, is bad for business. A coherent solution will only come from examining the viable candidates and picking a winner to invest in. The parallel computing framework that becomes standard ought to be easy to program and maintain good performance with any amount of parallelism.

Vishkin has tested his prototype with students to evaluate and verify its ease to program. He contrasts his approach with the "build-first-figure-out-how-to-program-later" architectures being produced by industry. Any successful approach must answer what will be taught to students in an algorithms class, Vishkin said. Otherwise, the approach will be a dead-end. The alternative he has worked on for nearly three decades could become the new standard that hurdles over this barrier.

Samir Khuller: Finding Cheap Gas and Other Problems

About seven years ago, when gas was one-third the price it is now, and there weren't numerous web sites tallying gas prices, Samir Khuller was walking to class, thinking about applications for the kinds of computer science problems he was teaching his students. He started to think about how one could calculate the most cost-effective strategy for filling a car's gas tank on a trip. The problem was an academic exercise at the time, but as gas prices have soared and the Internet gives access to prices at different stations, the methods Khuller and his students have developed have come to seem eminently applicable.

With websites publishing gas prices, the problem can be fed with real-life data. Khuller said, "Given that I know what the prices are, how can I optimize my cost? If you have a fleet of trucks, the savings would add up."

For his class, Khuller initially outlined a solution involving assumptions such as a fixed driving route and a requirement to buy gas in units of a half gallon. "But I wasn't satisfied with the assumptions," he said. So Khuller suggested that his graduate students Azarakhsh Malekian and Julian Mestre delve more deeply into the problem. Mestre was already working on two-dimensional matching problems such as how to best match applicants to jobs, or how to analyze voting results to allow voters to rank candidates instead of choosing just one.

Khuller speaks frequently about his students. "He takes both teaching and mentoring very seriously," said Leana Golubchik, now at the University of Southern California. She points out that Khuller mentors high school students as well as undergraduates and graduate students. Golubchik has long collaborated with Khuller. "I like to work with people who are nice and people who are very interested in the problem," she said, "and I think of Samir always as really being a true scholar and very interested in the science."

An Zhu, a former student who now works at Google, adds, "Samir is the best professor that I've ever met." She took an entry-level algorithm class with him and then did an undergraduate honors project with him. "Samir is very good at explaining complicated material, and he would also challenge his students to do better. Samir really opened my eyes and showed me that research could be a lot of fun."

Algorithm Design and a Diversity of Problems

"The main problem I work on is algorithm design, where the solution space is vast," Khuller said. He worked for several years on network design, or how to connect many locations in a cost-effective way while ensuring connectivity among all the points, or including enough redundancy to withstand different kinds of disruptions. "These problems are extremely hard because the number of choices is on the order of trillions," Khuller explains. A classic analogous problem is the traveling salesman problem, where, for example, someone wants to visit all 50 state capitals in the United States and wants to know the shortest route. Khuller's work focuses on developing efficient ways to calculate possible solutions and to come up with "approximation algorithms" that may not necessarily give the cheapest solution but do yield one that comes with some guarantees, such as being within 25 percent of the optimum solution.

"For the last 15 years, the theme in my work has been developing a better understanding of approximation algorithms," Khuller said. Heuristics are approaches that can yield a quick answer although one that's not entirely accurate. The question Khuller has focused on is evaluating heuristics and finding better ones. Over time, researchers have found that a small set of methods can provide solutions for the majority of problems. For about one-third of problems, however, one needs to invent specific tools, Khuller said. Each of those remaining problems requires a different approach.

The voting problem his student Mestre was working on, a problem known as Condorcet's paradox, is another classic. "You have to define a pairwise comparison between every pair of candidates," Khuller explains. The difficulty is that allowing voters to rank candidates could, in some situations, yield no clear winner. In the past, asking voters to rank candidates would have yielded an

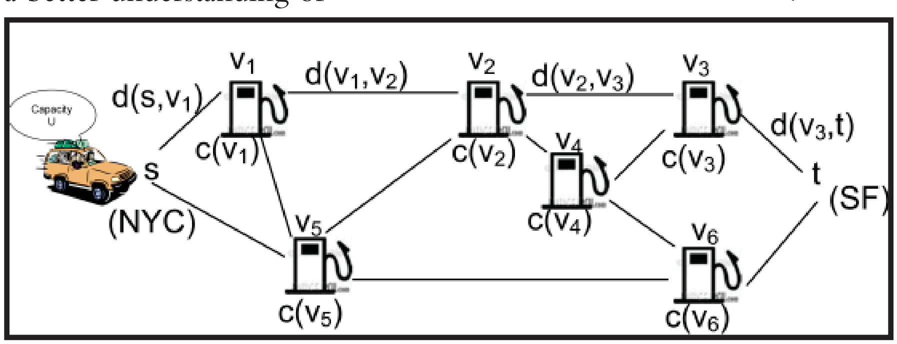

The gas station problem Khuller and his students solved was how to find the minimum cost solution to go from point S to point T.

impractically complex outcome to sort through, but given the speed of modern computing, Khuller and Mestre determined that it could be an opportune time to revisit the question. Giving voters the opportunity to rank their choices or perhaps give different proportions of a single vote to different candidates would give citizens a more nuanced way to choose candidates. The option would also allow voters to choose small-party candidates without unwittingly paving the way for a major-party candidate that has a better chance of winning.

Solutions to Future Problems

Khuller has a history of working on problems out of intellectual interest that eventually turn into economically interesting questions. The solution that he and his students developed for the gas station problem is somewhat intuitive: If the next available refill stop is more expensive, then fill up now. If the next stop is less expensive, fill up only enough to reach the next stop with a near-empty tank. "In the end, the algorithm is just 20 to 30 lines of code, but it took us months to get there," said Khuller. Earlier solutions were more complicated though no more effective than the simpler solution.

Another example of a problem that has turned out to be practical is how a company such as Netflix should organize to most efficiently allow users to watch videos on-line. Khuller started studying the question a decade ago with Golubchik. Their work addressed how a company should decide how many copies of each video to have on hand and how to distribute the copies across their servers. Because servers have limited capacities, a company can always add hardware to satisfy consumer demand, or it can save money by designing more efficient ways to organize its data. Khuller and Golubchik developed an algorithm that works well if all movies are approximately the same length, he said. This summer, a student has worked further on the problem.

More recently, Khuller and Golubchik have revisited the problem of distributing video on demand, this time from the point of view of users. "Let's say we're watching the same movie," said Khuller. "A server will stream data to you. If the data is already being streamed to me, I might be able to forward the data to you instead of you having to go to the server." Such sharing would drastically reduce the load on a server. As with other problems he addresses, Khuller is trying to define the boundaries of how such a scheme could work. For example, data could arrive at one user from multiple other users out of order. An efficient scheme could use buffer space to store data temporarily and place the information in the right order. Mathematical analysis can reveal how much buffer space would be required.

Khuller's research is largely funded by the National Science Foundation and is in the public domain. How public companies use such results isn't generally revealed, he notes.

Invited to speak at Google, Khuller talked about the gas problem. Khuller is happy to share what he and his students have discovered and then stand back while companies build on the work and make it accessible and useful. Google is also funding Khuller's further work on the problem. Said Zhu, "Gas price is one kind of information that almost every family in the U.S. is concerned with. We want to provide this useful information to our users."

News in Brief

Rance Cleaveland was quoted in the Baltimore Sun on the use of static analysis methods for testing medical devices.

Rita Colwell co-authored an article in *Science* on January 25 titled "Mobilizing Science-Based Enterprises for Energy, Water, and Medicines in Nigeria," which addressed the potential for a sustainable approach to supplying these basic services to Nigeria's poor by encouraging private companies to become involved.

Through the work of Allison Druin and the HCI team, Maryland Governor Martin O'Malley recognized May 29-30, 2008 as HCIL's 25th Anniversary and Annual Symposium Days.

Hanan Samet, Houman Alborzi, and Jagan Sanakaranarayanan won the best paper award at SIGMOD 2008 for their paper entitled "Scalable Network Distance Browsing in Spatial Databases".

The Lab for Computational Cultural Dynamics' SOMA Terror Organization Portal (STOP) and social network site for terrorism related analysis and prediction was featured in Computerworld Magazine, UPI News, and several major news media. STOP provides methods for reasoning about terror groups and forecasting what they might do in the future.

UMIACS is participating in the Cloud Computing Initiative. "We're aiming to train tomorrow's programmers to write software that can support a tidal wave of global Web growth and trillions of secure transactions every day." said Samuel J. Palmisano, chairman, president and chief executive officer, IBM. The initiative will involve Carnegie Mellon, Massachusetts Institute of Technology, Stanford University, the University of California, Berkeley, the University of Maryland and the University of Washington.

Neil Spring: Mapping the Internet

When the Internet was run by a small number of academic and government institutions, all the routers and connections among them were known. Now, the Internet has grown into a huge, international maze of wires and routers supported by more than 10,000 organizations including private companies that compete with each other. Large businesses with an internet presence have their own routers, as do internet service providers (ISPs) and telecommunications companies. Together with academic institutions and governments, these disparate entities have cobbled together a network that continues to grow and work reliably. To users, the Internet is vast and seemingly amorphous, but not even the experts know exactly how the whole thing is wired or how many components it has. Neil Spring is working to chart this unknown territory and build a map of the Internet.

"We rely on this network every day. Its operation requires cooperation both between competing companies and across different countries," Spring said. "While you may trust each of them individually to run a really great network, it's unclear why you should trust all of them collectively to run a really great network. The tension is that they all must want the network to operate." Unlike other infrastructure networks, users of the Internet can use internet protocols (IP) to probe how the Internet is structured.

The companies that have largely fed the expansion of the Internet keep information about their networks confidential, so "reverse engineering" through glimpses afforded to users offers an entry point to studying the network. Spring's raw data for mapping comes from the routes that desktop computers show listing the IP address between themselves and various network destinations. This information, a list of letters and numbers where each line represents an interface on a router, can be taken from practically any machine. Each IP address identifies a network handoff taking place at a router, one of the switching devices that directs the flow of data through the network. The list of routers between two destinations traces the path that information followed. Piecing together millions of these routes, as Spring has done, yields a map.

Spring, who joined UMIACS in 2004, began to work on mapping the Internet as a graduate student at the University of Washington. "Neil knows what he's doing," said Walter Willinger, a researcher at AT&T who has collaborated with Spring. "Neil's work has been on the forefront in showing there is hope that if you do it right, you can reverse engineer the Internet to a degree that a few years ago a lot of people would not have believed."

"Neil and his colleagues managed to recover the network topology and routing policy for large-scale ISPs for the first time," said Yan Chen, a computer scientist at Northwestern University. "His work has been widely cited."

Techniques and Tricks

Spring, his student Rob Sherwood, and their collaborators have collected data using PlanetLab (www.planet-lab.org) as a platform. PlanetLab is a partnership of hundreds of universities and research labs that provides access to machines around the world for network research. The maps built by Spring's team chart how PlanetLab machines are connected to each other.

"What we're trying to do is get the most accurate map we can. That's difficult because there are all sorts of devices in the network built by lots of different kinds of companies," said Spring.

Another complication is that individual routers have many interfaces with the network and are identified by many IP addresses. Each interface has a unique IP address. "One of the challenges is to figure out which IP addresses represent the same router," said Spring. "If you get this wrong, you get a bad map." To help group IP addresses into the routers they represent, Spring turned to UMIACS director V. S. Subrahmanian and his student Austin Parker to incorporate artificial intelligence (AI) techniques into the mapping process.

"To get the most accurate map we can, we try to measure twice," said Spring. His team looks at interfaces from opposite directions and then attempts to line up the data into a single path. With enough data and enough tests, the research-

```
sr-rockettrace www.planet-lab.org
1 router-126.cs.umd.edu (128.8.126.1) [27/27] {CollegePark, MD} 1.375 ms 0.240 ms 0.313 ms
2 128.8.239.69 [/27] 0.675 ms 0.463 ms 0.942 ms
3 Gi3-1.ptx-core-r1.net.umd.edu (129.2.0.113) [27/27] {CollegePark, MD} 0.527 ms 0.554 ms 0.655 ms
4 Gi3-1.ptx-fw-r1.net.umd.edu (129.2.0.86) [27/27] {CollegePark, MD}0.725 ms 0.688 ms 0.748 ms
5 Gi2-8.css-max-r1.net.umd.edu (129.2.0.234) [27/27] {CollegePark, MD} 168.990 ms 1.010 ms 1.204 ms
6 clpk-umd-i2.maxgigapop.net (206.196.177.125) [10866/10886] {College Park, MD} 0.788 ms 0.854 ms 0.842 ms
7 xe-7-2-0-0.lvl3-t640.maxgigapop.net (206.196.178.90) [10866/10886] {} 2.698 ms 1.813 ms 1.903 ms
8 i2-lvl3.maxgigapop.net (206.196.178.46) [0/10886] 1.955 ms 1.946 ms 1.999 ms
9 so-0-0-0.0.rtr.newy.net.internet2.edu (64.57.28.10) [0/0] 7.071 ms 7.387 ms 7.584 ms
10 local.internet2.magpi.net (216.27.100.53) [10466/10466] {Philadelphia, PA} 9.139 ms 9.233 ms 9.078 ms
11 remote1.princeton.magpi.net (216.27.98.114) [10466/10466] {Princeton, NJ} 11.155 ms 11.686 ms 11.038 ms
12 gigagate1.Princeton.EDU (128.112.12.21) [0/88] 10.634 ms 10.820 ms 11.009 ms
13 csgate.Princeton.EDU (128.112.12.58) [0/88] 10.837 ms 10.892 ms 10.842 ms
14 chloe.CS.Princeton.EDU (128.112.139.52) [0/88] 10.587 ms code10 10.741 ms 11.478 ms
```

Traceroutes list the network hosts visited between two destinations by a packet of information.

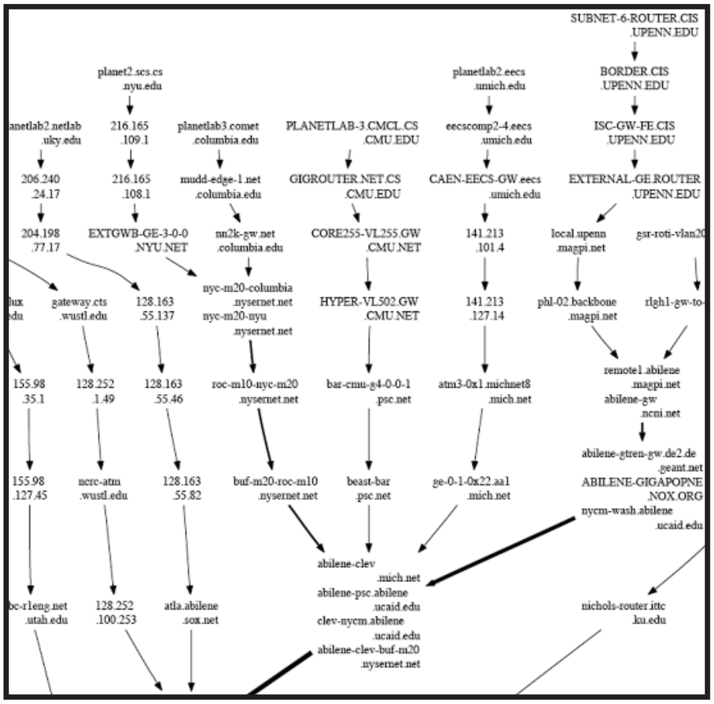

A portion of a reverese-engineered internet route.

ers can gauge the accuracy of the paths they've described. For example, they can send a message directly to a router. "Sometimes it replies with a different IP address," said Spring, which gives the researchers more information. Another strategy his team has used is to "fingerprint" routers. Routers count the packets of information they send, and every packet must be identified uniquely, using tags that are called fragmentation identifiers. "There are tests that can back us up if we're unsure and help the AI, giving hints," Spring said, "or just tell us if we're right."

The work remains challenging because the volume of data can be overwhelming. "The big problem we had was scale," Spring said. "We have millions of paths. We have hundreds of thousands of measured edges." The AI techniques the team tried can process about five paths at a time, and because all the variables can't be accounted for at once, solutions in different batches may be inconsistent. "We need to build a divide-and-conquer technique to plan the execution so we're least likely to end up with conflicts in the end," said Spring. His team has taken to mapping groups of paths that overlap in both directions to assemble more reliable maps. "We're trying to map this puzzle a little bit at a time," Spring said. And then the group tries not to do anything to mess up assembly of the pieces into a larger puzzle.

"In the last five years, a field has emerged called network science," explains Willinger. "It has a very appealing objective—to understand large-scale, complex networks that arise in all kinds of contexts—biology, social networks, airlinesystems. The goal of network scientists is not to study nodes but to study connectivity and the structures of networks." He points out that to really understand a network, one needs to understand the specific field in which it's based—biology if one is studying biological network and computer science if one is studying the Internet. "One very important contribution of Neil's work has been showing all the difficulties associated with inferring an ISP's infrastructure," Willinger comments. Neil and his team "know exactly where they run into problems."

Why Build a Map?

If the Internet has managed to operate for years without anyone really knowing how information gets from place to place, what is to be gained from mapping the Internet? One big benefit is "traffic engineering," said Spring. A map can reveal shorter or less congested routes and ways to avoid dead ends. Having a map of the Internet, or even small parts of it, can reveal the decisions that are made to transmit information and whether parts of the network are unreliable. Maps can reveal inefficiencies and weaknesses and suggest ways to make the network faster and more reliable.

"You can imagine building a network that would make much better decisions and run more efficiently," said Spring. "You want to be able to design a network that's as robust as possible, and you don't want to over- or underestimate the redundancy in the network." A map can help quantify how much room for improvement there is in network routing and policies.

At a very practical level, Spring said, "You can identify problems in the interfaces between organizations. One thing I've been able to study is how cooperative pairs of providers are." Network providers may not consider efficiency or performance when passing information to each other. Sometimes one company does work another could or should have done. Spring recalls an example where providers quarreled over whether they were truly of comparable size and had an equitable relationship. For a while during the disagreement, customers of one company were not able to communicate with customers of the other.

Because companies usually know nothing about the network structure in another organization, choices made by one organization can inadvertently lead to problems in the wider network. "Sometimes this lack of knowledge can cause a network to make a local optimization that causes a problem upstream," notes Spring. "My maps have been used to evaluate protocols for coordinating better," allowing network researchers to better manage their networks and demonstrate how their techniques could help their competitors.

Maps also expose the kinds of relationships in the Internet, which are of two main types: peer links, where network players are of comparable size and cooperate, and provider-customer links. Networking research can reveal the kinds of incentives that could help promote cooperation.

Maps reveal whether the network is being used well and can quantify how much room there is for improvement. They can show how often information has to take a long, circuitous path because of the lack of a direct route; they can reveal bottlenecks and the rea-

continued on page 11.

Duraiswami, continued from backpage

has been primarily addressed by psychologists. Duraiswami aims not only to explain how sound is processed according to physical principles but to build mathematical models that can recreate sounds in virtual reality.

Each Person's Sound-Processing Formula

Duraiswami's work focuses heavily on head-related transfer functions—the mathematical formulas that describe how scattering by one's body selectively amplifies or attenuates sounds at different frequencies, depending on the location and frequency content of the sound. Abbreviated HRTF, a head-related transfer function calculates what the ear perceives versus what would be there if one's head weren't there. Each person's HRTF is unique, because people's heads and ears are as unique as their fingerprints and scatter sound waves in idiosyncratic ways. Depending on the location and frequency of a sound, our heads can make a sound seem sometimes louder and sometimes softer. If one understands an individual's HRTF, one can hope to write virtual reality programs that convey the true feeling of being in a space listening to a sound. Because ear folds and HRTFs are unique to each person, sounds played on headphones would have to be customized to give different people the same perception of a sound experience.

To that end, Duraiswami had developed ways to measure individual's HRTFs. In the past, this was a long, involved process, where a person was presented sounds of different frequencies from different angles. When Duraiswami joined UMIACS in 1998, he and his colleagues initially decided to use photographs of people's heads to compute their HRTFs. Duraiswami, who studied mechanical engineering in school, had worked on numerical simulations before he arrived at UMIACS. His work has been funded by the National Science Foundation, the Department of Veterans Affairs, DARPA and industry.

In his initial work on calculating HRTFs, Duraiswami collaborated with Richard O. Duda, an electrical engineer at University of California, Davis until he retired. "The project that we worked on together was an interdisciplinary one, where his talents really shine," Duda said of Duraiswami. "What struck me most about Ramani was his versatility, optimism and enthusiasm. Ramani has an unusual ability to see research opportunities, and to find effective ways to apply both traditional mathematical analysis and modern computational techniques to important problems. In addition, he is just fun to work with."

After exploring the use of photographs to calculate HRTFs, Duraiswami came up with a second method for obtaining the functions. He and his team developed a way to directly measure HRTFs in a few seconds. The method relies on "the principle of reciprocity," he said. If there is a sound source one place and a receiver at another, swapping the locations of source and receiver should lead to the same perception of sound. Using this principle, one can place tiny speakers in a person's ears, then use a network of receivers around the person's head to map the person's HRTF.

Duraiswami and his team use speakers that emit noise with a broad range of frequencies, and immediately, the surrounding receivers record the sounds they detect. "You have a speaker in the ear and detectors outside," explains Duraiswami.

In the future, Duraiswami said, commercial headphones will be tuned to individuals' HRTFs, giving people customized sound experiences that make listeners feel like they are present in, say, a concert hall. When a listener moves his or her head, the sound could move as well, maintaining the illusion that the sound is outside the head and not right at the ears.

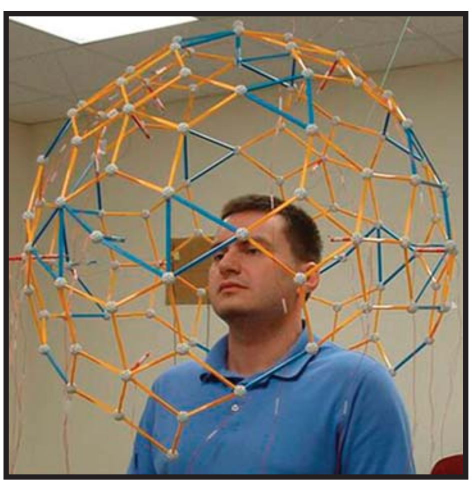

To measure HRTF, an array of receivers detect the sound emitted by speakers placed in the ears..

Besides creating a better way to listen to music, this kind of technology could be a great boon in presenting information to the blind. For example, the Department of Veterans Affairs would like to be able to record the sounds of street scenes in a realistic way that gives blind people a safe way to learn to navigate city streets while in therapy. In fact, refining the recording and play of sounds would enable new ways to present all sorts of information to blind people, for example conveying the spatial information in a map or other graphical forms of data.

Tiny speakers are used to measure HRTF.

"They have done really fabulous work on virtual hearing and developing techniques that recreate realistic perceptions of sounds," comments Shihab Shamma, an electrical engineer at the university who works on how sound is represented in the brain. "There's a great need for that."

Capturing Sound

If playing back sound in a way that captures how we hear it is half the problem, the other half—which Duraiswami also works on—is recording the sound in the first place. "You need to capture sound at many microphones to capture spatial information," he explains.

Duraiswami's group uses a spherical microphone array, placed on a tripod, to capture soundscapes. "We want to capture sound but then also let you as a user be able to move in that scene," he said. He and his team can interpolate sounds between the microphones to give a smooth sense of how something like a street scene would sound as you move through

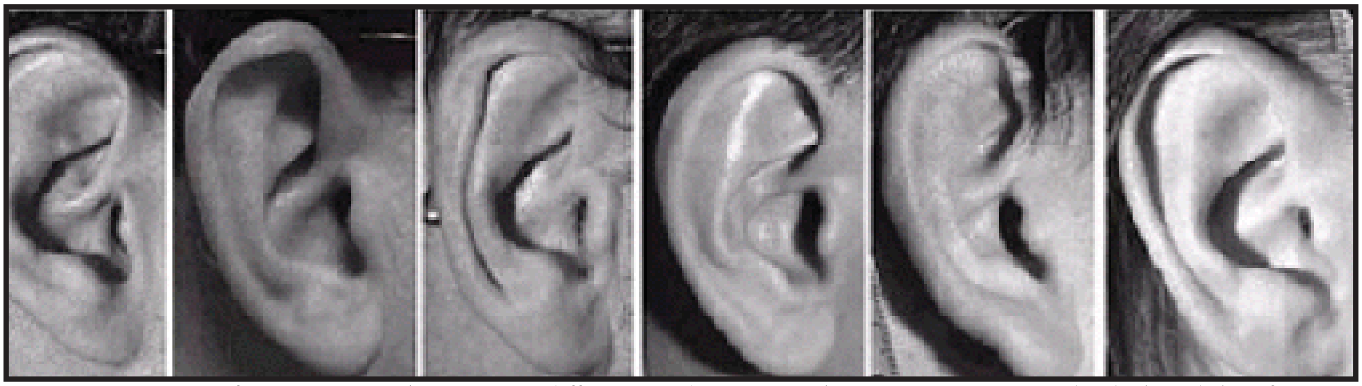

Ears are as unique as fingerprints, resulting in very differnt sound scattering characteristics among individuals and therefore differnt HRTFs.

it. His group is developing the theory and methods to improve the analysis of sounds from such arrays.

Duraiswami's group also uses a hemispherical microphone array. When such an array is placed on a solid surface, sounds reflect from the surface, capturing the effect of recording from a sphere with half the microphones. Duraiswami suggests that such hemispherical arrays could be used in video conferencing.

In a further innovation, Duraiswami's student Adam O'Donovan is working to combine information from microphone arrays and from cameras to produce what O'Donovan and Duraiswami call an "audio camera" that can convey where sounds are coming from in a video image. The audio camera was chosen as the University's Invention of the Year in Information Sciences for 2007.

Using his scientific computing background to work on the complex calculations involved, Duraiswami is also working on what he calls "vision-guided beamforming," or developing ways to pick out a sound coming from a particular direction even in the presence of loud background noise. He is using graphics processors and innovative algorithms to speed up the calculations involved to allow the processing of sound in real time.

Sound and Architecture

One application of sound analysis is helping architects design concert halls, auditoriums and classrooms. In such places, ideally the sound that reaches each seat should have both a component that comes directly from the stage or podium and a few reflections that give the sound warmth and ambiance. To perceive music in a pleasing way, the human brain expects reflections, where the reflections arrive as rhythmic peaks that gradually attenuate. Some people actually feel sick without such background reflections, but the reflections should arrive in a well-paced way that enhances the pleasure of listening. Otherwise, the sound can be jumbled and unclear.

"People who design concert halls have subjective words to describe the quality that they want," said Duraiswami. He and O'Donovan decided to image a concert hall using their audio camera. They generated a panoramic image of the Clarice Smith Performing Arts Center on the university campus by placing a loudspeaker on stage and playing a 10 millisecond "chirp sound" containing all audible frequencies. At individual seats, Duraiswami's team placed their microphone array to record the arriving sound waves—the initial sound and all its reflections. This technique provides a new way of visualizing and analyzing acoustics; Duraiswami's group is considering licensing the technology to architecture firms.

Ramani Duraiswami

Spring, continued from page 9

sons for them. For example, maps can show whether bottlenecks are between providers or at edges where individuals connect, or somewhere else entirely. "Links that are bottlenecks tend to stay bottlenecks," Spring said, and the reasons for them are varied.

Maps can also serve as snapshots over time to track how the network continues to grow and evolve. It would be useful to be able to predict how the network will look in another decade or two—how much bigger it will get, how much more connected. "Maps have been used to validate conjectures about how the network grows and how we could grow a bigger one," Spring said. For example, maps allow one to study node connectivity: Some nodes are highly-connected and others are less-connected, and it's not clear why and whether the degree of connectivity at any given node stays the same. The rules for growth and change differ depending on the specific kind of router, as well.

"The Internet is not a static structure," said Willinger. "It's a dynamic structure shaped by various forces, including economic forces, government forces, forces that vary over time. Maps can predict how changes in those forces will affect the future structure of the Internet."

Adds Chen, "Without this deep understanding, it is very hard, if not impossible, to design new protocols, architectures, or services that will work better for the Internet."

Ramani Duraiswami: Processing Sound

Imagine having a conversation virtually and having an experience that recreates how it would feel to be in the same room as someone even while that person is hundreds of miles away. Capturing the visual experience is relatively straightforward: Images could be embedded in the left and right frames of eyeglasses to allow a user to see a three-dimensional image. But how would one recreate the experience of sound? The brain fuses the inputs from one's ears as well as the inputs from one's eyes, but how one's ears process sound is a layered and complex phenomenon that is far less clearly understood than vision. For example, a listener senses the direction and distance of sounds and also intuits the background—distinguishing whether a speaker is in a carpeted room, say, or in an echoing gymnasium. Ramani Duraiswami studies how humans process audio information and how to use that understanding to recreate audio experiences.

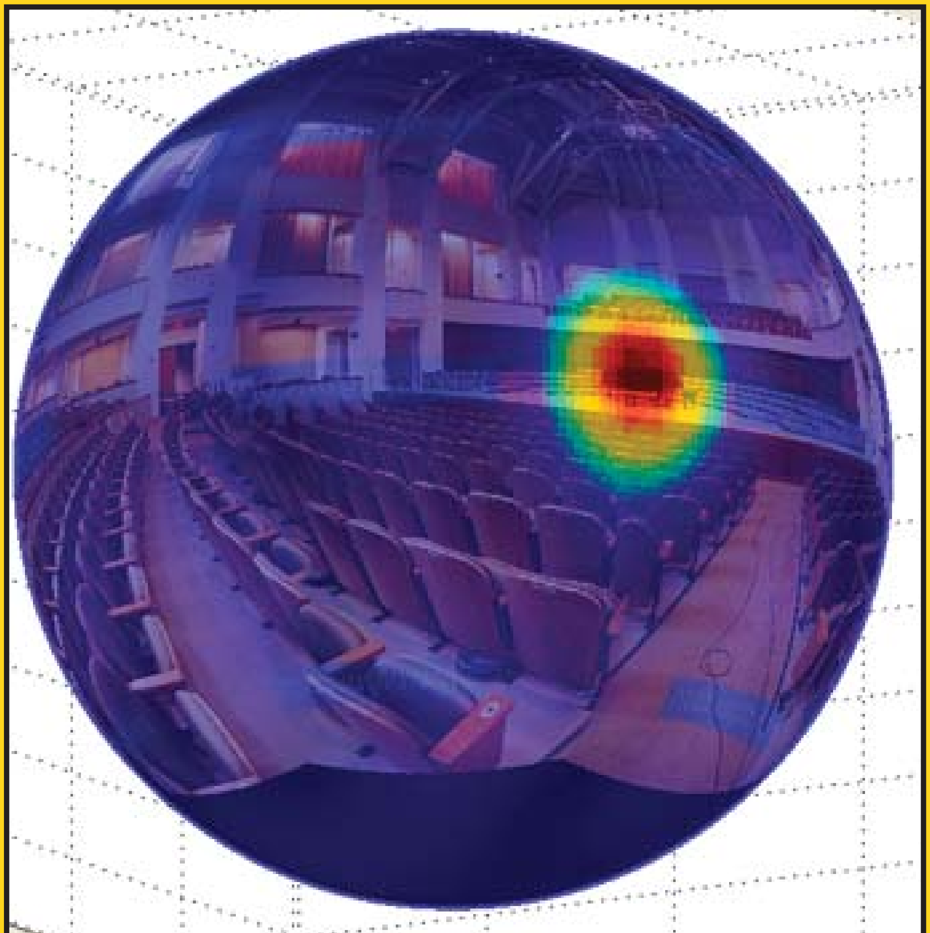

An audio camera shows how sounds reverberate in a concert hall. Courtesy of Adam O'Donovan.

First of all, Duraiswami points out, because sound travels relatively slowly, the ears perceive sound at slightly different times. These differences can give us a sense of the direction of a sound. But analyzing sound is more complex than that because sound bounces off objects around us and, most importantly, our own bodies. The body's geography scatters sound waves.

"It's like the Heisenberg uncertainty principle," Duraiswami said. "When you perceive sound with your ears, you also change it. Your ear is not just a receiver but a part of your body. Sound bounces off your body, head and ears, and that changes the sound." Duraiswami started working on the problem of how we perceive sound more than eight years ago. He brought the perspective of a computational scientist to a question that

continued on page 10.

UNIVERSITY OF MARYLAND
Institute for Advanced Computer Studies
2119 A.V. Williams Bldg.
University of Maryland
College Park, MD 20742-3311

Nonprofit Org.
U.S. Postage
PAID
ICM