

THE LINK

Carnegie Mellon

News from the School of Computer Science
Issue 5.0 / Spring 2010

Better Living through Robotics

PITTSBURGH'S QUALITY OF LIFE TECHNOLOGY
CENTER APPLIES BIG-LEAGUE COMPUTING TO
HUMAN-SIZE PROBLEMS

ALSO INSIDE:

TECHNOLOGY TARGETS CLIMATE CHANGE
SOCIAL DATA MINING OFFERS PROMISE AND PERIL
LENORE BLUM ON PROJECT OLYMPUS

PLUS: TAKE OUR READER SURVEY!



The Link provides a mosaic of the School of Computer Science: presenting issues, analyzing problems, offering occasional answers, giving exposure to faculty, students, researchers, staff and interdisciplinary partners. The Link strives to encourage better understanding of, and involvement in, the computer science community.

Editor-in-Chief

Randal E. Bryant

Managing Editor/Writer

Jason Togyer

Contributing Writers

Jennifer Bails, Tina Carr, Mark Dorgan, Meghan Holohan, Tom Imerito, Byron Spice

Photography

Ken Andreyo, John Barna, Glenn Brookes, Chris Harrison, Iberdrola Renewables, Tim Kaulen, Wade H. Massie

Graphic Design

Melissa Stoebe
Communications Design & Photography Group

Office of the Dean

Gates Center 5113
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

Randal E. Bryant, dean

Tina Carr (HNZ'02), director of alumni relations

Philip L. Lehman (CS'78, '84), associate dean for strategic initiatives

Byron Spice, director of public relations

Jason Togyer (HS'96), writer/editor

Carnegie Mellon University

Phone: 412-268-8721

Email: TheLink@cs.cmu.edu

Web: link.cs.cmu.edu

Facebook: tinyurl.com/linkscs

Twitter: twitter.com/SCSatCMU

Carnegie Mellon University does not discriminate and Carnegie Mellon University is required not to discriminate in admission, employment, or administration of its programs or activities on the basis of race, color, national origin, sex or handicap in violation of Title VI of the Civil Rights Act of 1964, Title IX of the Educational Amendments of 1972 and Section 504 of the Rehabilitation Act of 1973 or other federal, state, or local laws or executive orders.

In addition, Carnegie Mellon University does not discriminate in admission, employment or administration of its programs on the basis of religion, creed, ancestry, belief, age, veteran status, sexual orientation or gender identity. Carnegie Mellon does not discriminate in violation of federal, state, or local laws or executive orders. However, in the judgment of the Carnegie Mellon Human Relations Commission, the Presidential Executive Order directing the Department of Defense to follow a policy of, "Don't ask, don't tell, don't pursue," excludes openly gay, lesbian and bisexual students from receiving ROTC scholarships or serving in the military. Nevertheless, all ROTC classes at Carnegie Mellon University are available to all students.

Inquiries concerning application of these statements should be directed to the Provost, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, telephone 412-268-6684 or the Vice President for Campus Affairs, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, telephone 412-268-2057.

Carnegie Mellon University publishes an annual campus security report describing the university's security, alcohol and drug, and sexual assault policies and containing statistics about the number and type of crimes committed on the campus during the preceding three years. You can obtain a copy by contacting the Carnegie Mellon Police Department at 412-268-2323. The security report is available through the World Wide Web at www.cmu.edu/police/.

Obtain general information about Carnegie Mellon University by calling 412-268-2000.

Produced for the School of Computer Science by the Communications Design & Photography Group, April 2010, 10-437.

©2010 Carnegie Mellon University, all rights reserved. No part of this publication may be reproduced in any form without written permission from the Office of the Dean, School of Computer Science.

CALENDAR OF EVENTS

All events to be held at the Carnegie Mellon University campus in Pittsburgh, unless otherwise noted.

May 3

Final exams begin

May 5

Meeting of the Minds Undergraduate Research Symposium
University Center

May 10–14

Hands-On Workshop on Computational Biophysics
*Sponsored by CMU and the National Resource for Biomedical Supercomputing
Pittsburgh Supercomputing Center*

May 15–16

Commencement weekend

May 17

Summer classes begin

June 28

Summer classes begin for session two

July 3

SCS and ECE Alumni Picnic
*Sam's Chowder House
Half Moon Bay, Calif.*

July 5

Independence Day observed:
University closed

July 8

SCS and ECE Alumni Reception
San Diego, Calif.

July 10

Annual Carnegie Mellon Kennywood Picnic
Arranged by Staff Council

July 11

SCS and ECE Alumni Reception
Los Angeles, Calif.

August 5

Last day for all summer classes

August 23

Fall semester begins

September 6

Labor Day: University closed

FEEDBACK LOOP

This shouldn't come as any surprise, but I read a lot. (It's an occupational hazard for a writer.)

Much of my reading is done online at blogs and news Web sites, but I also buy all manner of print magazines, from *The Economist* to *Hemmings Classic Car*. There's something relaxing about sitting on my porch and flipping through the pages of a printed magazine that a Kindle or a laptop hasn't yet replicated.

Still, the most annoying thing about magazines might be those little "response cards" that urge you to extend your subscription. They fall on the floor, make it hard to turn the pages and generally aggravate me.

I promise that *The Link* doesn't intend to start filling your magazine with those little cards, but we have inserted one into this issue. We need your opinions of the magazine—your likes and dislikes—so that we can judge the effectiveness of *The Link* and improve it where necessary.

If you don't want to fill out the postcard, that's OK. We've put the survey online for you at www.surveymonkey.com/s/QMJJDZ9D. Log on using the password "weanhall." (Please don't take the online survey if you're going to use the card—just do one or the other.)

Of course, your thoughts, comments, criticisms or suggestions are always welcome at TheLink@cs.cmu.edu, or via postal mail at *The Link Magazine*, School of Computer Science, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213. After all, I like to read!

Jason Togyer (HS'96), managing editor

Issue 5.0 / Spring 2010

2 / From the Dean

3 / On Campus

There's an art to creating algorithms, and now algorithms are creating art. Data mining offers new opportunities for social benefits, but can we balance those with our need for privacy? Global software development delivers productivity gains, if communications breakdowns are avoided. And access to Yahoo!'s M45 supercomputer paid big dividends—now we're looking ahead to new frontiers in data-intensive supercomputing.

10 / In the Loop

Lenore Blum, distinguished career professor of computer science, talks about her newest initiative—Project Olympus, which helps innovators turn their research into commercially marketable ideas.

17 / Research Notebook

Eric Paulos, director of the Living Environments Lab in the Human-Computer Interaction Institute, describes the role of "Citizen Scientists" in data collection.

20 / Alumni Director's Message

You can stay connected to Carnegie Mellon in your own backyard. Tina Carr explains.

21 / Alumni Snapshots

Catching up with Eduardo Frias (CS'94) and Ben Ahroni (CS'04).

23 / SCS in the News



7 / Robots for Life

If robots can navigate Three Mile Island or the surface of Mars, why can't they find their way around a suburban bedroom? Shouldn't technology be able to address "quality of life"? The Quality of Life Technology Center aims to answer those questions.

By Tom Imerito



12 / Technology for a Greener Earth

When it comes to conserving energy and reducing production of greenhouse gases, computer science has a powerful role to play. (No pun intended.) At Carnegie Mellon, researchers are looking for practical ways that computer science can reduce our carbon emissions and energy use.

By Jason Togyer

On the Cover:

The Personal Mobility & Manipulation Appliance, or PerMMA, isn't just a "wheelchair with arms." It's more like a semi-autonomous robotic manipulator that also has a seat for a person.

Currently being tested at the Human Engineering Research Laboratories—a joint venture between the University of Pittsburgh and the VA Pittsburgh Healthcare System—erMMA is designed to read its environment and detect what its user is trying to do and where its user wants to go, allowing it to predict needs and provide assistance with a minimum of direct input.

PerMMA is just one of the technologies being developed as part of the Quality of Life Technology Center, a joint venture between Carnegie Mellon and the University of Pittsburgh to apply advanced computer science and robotics to products and techniques that improve the lives of people with disabilities or illnesses.

Shown testing PerMMA's capabilities are Benjamin Salatin, a graduate student in Pitt's School of Health and Rehabilitation Sciences, and Jijie Xu, an SHRS post-doctoral researcher working at HERL.

Our story begins on page 7.

(Glenn Brookes photo)



P.S.: I'm pleased to announce that this year's U.S. News and World Report ranking of

computer science graduate programs has put Carnegie Mellon into first place in a four-way tie with the Massachusetts Institute of Technology, Stanford and the University of California at Berkeley.

We also scored high in specialty areas such as programming languages (No. 1), artificial intelligence (No. 2), systems (No. 3) and theory (No. 5). We're very pleased to be recognized as leaders in all facets of computer science.

People need computers. Computers need people. Those are the themes of many activities of the School of Computer Science, including ones we report on in this issue of *The Link*.

With our longstanding tradition of research and education into human-computer interaction, we've devised many ways to make technology better serve the needs of people.

Our Quality of Life Technology Center—subject of the cover story—challenges the traditional thinking that robots should be as autonomous as possible. Instead, it proposes a model in which robots work with people with disabilities or infirmities to help them live more independently. But rather than taking control or replacing the humans they're supposed to help, these robots are being designed to sense what people are trying to do, and then bridge the gap between their capabilities and their intentions.

Software development requires human and computer interaction of a different sort. It requires extensive human effort, with hundreds of people spread across many locations and time zones. Keeping these projects on track requires understanding people—how they communicate, how they can negotiate effectively and how to avoid inefficiencies. This issue includes an update on the research of Jim Herbsleb and Marcelo Cataldo into ways to improve communication and negotiation on large, international software development projects.

And computers keep tabs on humans in many ways. Our lives are increasingly tracked and recorded by electronic devices, through our cell phones, our credit card purchases and our online activities. As data collection becomes more pervasive, we must also be sure to find the right balance between the common good and individual privacy. The new capabilities provided by machine learning and data mining raise many technical, legal and policy issues. Tom Mitchell, head of our machine-learning department, had a thought-provoking article on this topic in *Science* last December, and you can read about that in this issue as well.

Wherever humans and computers interact in the real world, there are challenges to overcome. But we view these challenges as opportunities for future initiatives, and we look forward to tackling them head on.

Randal E. Bryant
Dean and University Professor
School of Computer Science

Algorithms as Art

Installations in the Gates and Hillman Centers draw inspiration from technology

By Meghan Holohan

“Do you want me to stand in the art?” asks Scott Draves (CS’97).

He positions himself in front of a large screen in a tiny, dark room at the Michael Berger Gallery on Pittsburgh’s South Side. Blobs that resemble miniscule jellyfish pulsate and swirl across his face, highlighting his white shirt. Reds, yellows, greens and blues dance across the screen, tattooing Draves’ skin and clothes.

“It’s like Ray Bradbury’s *The Illustrated Man*,” says Bill Scherlis, professor and director of SCS’s Institute for Software Research.

After a few moments, Draves steps out of the art and exits the room. In the main gallery, several HD monitors show different installations by Draves. A knot of computer science professors crowds around Draves, asking him which equations and algorithms he used. One of the algorithms is embedded in an open-source program called “Electric Sheep” that anyone can download. The program produces screensaver images on users’ computers that pop up whenever their computers go to sleep (hence the name). Users then vote on their favorite designs.

Images that receive the most votes “mate” (in Draves’ words) with other images to create new art. The art evolves, changing based on users’ interactions.

One of Draves’ pieces employing Electric Sheep is on display in the Gates Center for Computer Science. Called “Generation 243,” it’s located near the fifth-floor exit to the Pausch Bridge. Displayed on a large-screen TV set against a lime-green wall, “Generation 243” feels like a passage down a rabbit hole. Patterns snake across the screen, and geometric shapes and ribbons of yellow, green and red bob and weave.

“Watching the video is like walking through a garden—walking one path, you see one sight,” Draves says. “You could walk the same path twice but there are parts you haven’t seen. There are main thoroughfares that appear most commonly

The “Electric Sheep” algorithm developed by Scott Draves (CS’97) is at the heart of “Generation 243,” one of several pieces of computer-generated art on display at the Gates Center for Computer Science.



WADE H. MASSIE PHOTO

and then there are hidden corners that only show up every couple of weeks.”

He calls “Generation 243” a distillation process: “I am just one final filter and the input to it is from all over the world—but then the starting point is that genetic code that is a piece of mathematics that I put out there on the Internet.”

“Generation 243” isn’t the only art installation in Gates. Just down the hallway on the fifth floor, between two cones hung from opposite walls, low voices murmur faintly in a hallway saying things like: Laptops. Microsoft. Dinosaurs. NASA. Walk too quickly and the voices vanish. But linger between the cones (they look a little bit like Princess Leia’s hair) and thoughts might emerge.

Carlos Guestrin and Osman Khan created this piece, called “We Synthesist.” One cone has an algorithm that trolls the Web for blogs about science, as the other searches blogs about computers. Each cone whispers words mentioned most often on the respective blogs. Guestrin, an assistant professor of machine learning, and Khan, a former visiting assistant professor in the School of Art, created the piece in the hopes that the whispered words would spark discussions about the nature of computer science.

“Computer science is neither about laptops or dinosaurs, but it gets people thinking,” Guestrin says. “When you ask people on the street about computer science they say laptops or Microsoft. By creating the public experience, you can ask

‘what does [computer science] mean to me?’”

Last spring, Guestrin and Khan taught a class called “New Media Installation: Art that Learns,” which used interactive tools to create new works of art. Computer science and art students adapted algorithms, such as those to detect fraud or suggest purchases on online shopping Web sites, to create art that’s thinking and learning.

Several student-created installations decorate the new building. On the seventh floor rests a clear box. Known as “The Curator,” this piece of art decides if a drawing is “original” or not. It retains drawings that are new and exciting as “art,” and shreds those it finds unoriginal. This art snob employs a credit card fraud algorithm to detect new patterns—indicating originality.

Another piece uses the algorithm that retail Web sites employ to predict a user’s preferences. The piece has three paths—red, green and blue—with each path representing the Koran, Talmud or New Testament. Passages from each of these holy texts float across the screen and the algorithm tries to predict which line belongs to which book. It turns it out that the books are so similar that the algorithm cannot accurately make those predictions.

Guestrin says that algorithms are all around us, embedded in all sorts of computer applications—why shouldn’t they also be used to create art? “They’re a new medium,” he says. “Now we need to ask: ‘What else can we do with them?’”

Mining our Lives

Data mining offers ways to improve society—but does it also dilute our privacy?

> By Meghan Holohan

Before heading to the doctor for a diagnosis, many people Google their symptoms. Searching terms such as “cold,” “cough” and “swollen glands” are a good indication that the person doing the search is suffering from the flu—or at least flu-like symptoms.

One woman with flu symptoms is no big deal, but if hundreds or thousands of people in Pittsburgh are searching for these terms, there’s a trend emerging. Based on these searches, Google can actually provide estimates of flu activity throughout the world. And it turns out that Google’s Flu Trends application is almost as effective at predicting an outbreak as models used by the Centers for Disease Control and Prevention.

Google Flu Trends is one of the more altruistic examples of data mining. Data mining has many other, more profit-driven uses, such as targeted marketing. Businesses ranging from national credit card companies to the chain grocery store in your neighborhood employ data mining. Every time a customer uses a rewards card at the grocery store, for instance, the company learns more about that user’s favorite items, and can send ads targeted to those specific preferences. Amazon.com mines past purchases to determine what books and music shoppers might enjoy, for instance.

During daily life everyone leaves electronic breadcrumbs, making it easy for any organization to gather data about our social lives. Tom Mitchell feels these digital imprints could be used in ways to benefit society. In his paper “Mining our Reality” in the Dec. 18 issue of *Science* magazine, he explores some of these uses.

“There is this amazing trend going on and it is just accelerating,” Mitchell says in an interview. “There is a lot of acquisition of data for commercial use, but there is not that much discussion about how to use it as a society.”

Mitchell suggests several ways of using data mining to promote public health and safety. An increasing number of cell phones have a GPS chip in them. Mitchell suggests using this GPS data to

help provide real-time traffic reports without sacrificing a user’s privacy. When phone companies see a cluster of cell phone users not moving in one area, they could assume all of those cell phones belong to people stuck in traffic.

“And there is no privacy cost because that data would be anonymized,” he says. “The cell phone company doesn’t have to say ‘Tom Mitchell is on the Fort Pitt Bridge,’ just that there are five times as many phones (as usual) on the Fort Pitt Bridge.”

But Lee Tien, a senior staff attorney at the Electronic Frontier Foundation, worries that using data mining applications actually will lead to violations of the users’ rights to privacy.

“There’s nothing particularly anonymous about your cell phone,” he says. “It’s almost impossible to get the data without having the data show which cell phone it is. It’s true they don’t need to know (whose) cell phone is crossing the bridge at that time to do the application, but it evades the question of whether (companies) are learning about you, and what they might do with the data once they have it.”

Traffic monitoring isn’t the only area where Mitchell suggests data mining could be used for the public good. “Mining our Reality” suggests another application that could help stem the spread of infectious diseases, such as the H1N1 flu virus. Imagine, he says, that a cell-phone user goes to the hospital, where doctors diagnosis him with H1N1. That GPS chip in his phone allows his phone company to know where he had been in the last three days and to what other phones—that is, what other people—he was in close proximity.

Each phone owner could receive a call informing them they were in contact with someone who had H1N1, Mitchell says, and that if they start feeling flu symptoms they should consult a physician.

But Tien says that’s exactly the sort of application privacy advocates are worried about. “Who I have been near in the last 24 to 72 hours is actually pretty sensitive info,” he says. “Just like the records of everyone I made phone calls to and sent emails to over the last two or three days would really give someone an incredibly good picture of who I am and what my social life is. This is really sensitive information.”

There are methods of restricting access to data so that hostile or unauthorized people can’t get to it, Mitchell points out. In his H1N1 example, he says, people don’t need to be told who has the illness—just that they were exposed.

Tien argues that current methods of securing data are inefficient. “The fact is, when we look at the entities that have our data, giant data brokers have information about hundreds of millions of people—most Americans,” he says. “They have our credit reports, our phone calls, and currently we are unable to control the spreading of this kind of data.”

Mitchell is hardly denying that privacy issues have to be discussed before new social data mining applications are developed. In fact, it’s a discussion he and others are trying to spark. Deciding whether or not to use data mining techniques to help society is subtle, he says, and doesn’t lend itself to simple “yes” or “no” answers. “When we have these discussions we need to understand that (privacy) is an issue,” he says. “Different individuals have different ideas about privacy—these are social and political discussions that society needs to have.”

Small World, Big Problems

Global software development can deliver productivity gains, as long as communications breakdowns are avoided

> By Jennifer Bails

A team of software engineers in Pittsburgh is almost finished building a data access interface for a new program that will be used to operate HVAC and lighting systems. Their job is critical because all of the other components of the software—being developed by groups working for the same industrial conglomerate in several different countries—depend on this interface.

Two days before a deadline, the Pittsburgh team realizes it must change the design specifications to satisfy a key software requirement. The developers send an urgent email about the proposed updates to a central team overseeing their project. But they get no response, and with pressure mounting, they

go ahead and submit the changes anyway.

In turn, the other remote teams in India, Germany and Ireland experience significant delays in their work as they scramble to amend their code to account for the unexpected modifications. They also express frustration with the Pittsburgh group, claiming they didn't understand why the changes were necessary.

The coordination breakdown began with a small decision, but the consequences were big. And problems like these—how to smoothly orchestrate technical work over geographic distances—are occurring more and more as software development goes global, says Marcelo Cataldo, a post-doctoral associate at SCS's Institute for Software Research.

"The technical and organizational complexity of software systems has increased by an order of magnitude," Cataldo says. "Twenty or 30 years ago, software projects used to be done by two or three people. Systems today can easily involve hundreds of people in different buildings, different cities or different countries. So all of a sudden your work depends on a lot more people and what you are doing might affect the person sitting next to you—or someone on the other side of the world."

Originally from Argentina, Cataldo earned his Ph.D. at Carnegie Mellon in 2007 under Jim Herbsleb and Kathleen Carley, professors of software research. Now Cataldo and Herbsleb—director of the Software Industry Center at ISR—are working together to find ways to improve globally distributed software development—specifically, how to help hundreds or even thousands of developers collaborate successfully despite being several time zones apart.

The stakes are high for commercial software projects, where production snags can spell enormous financial losses. The above example comes from a case study published in a 2006 paper by Cataldo and Herbsleb titled "Managing Complexity in Software Development: On the Limits of Modularity." Co-authors include Matthew Bass, an instructor in the master of software engineering program at Carnegie Mellon, and Len Bass, a senior member of the technical staff at the Software Engineering Institute.

The researchers participated in the three-year Global Studio Project launched by Siemens Corporate Research in 2005 as a test bed to gain a better grasp of the issues associated with global software development. The experiment simulated a real-world geographically distributed project by calling on teams of software engineering master's students at Carnegie Mellon and several other universities



RACHEL ARNOLD SAGER ILLUSTRATION

worldwide to develop software. "The idea was to try out various ways of managing globally distributed teams in an environment where an actual product was not at risk and see what practices worked best," Herbsleb says.

Herbsleb and his colleagues decided to take a cue from open source software development, which often seems immune to many of the collaboration problems that plague commercial efforts. They gave the students access to a host of communications tools—including a wiki, email, instant messaging and a discussion forum—often used in successful open source projects.

But their analysis showed that although these tools and usage guidelines were available, coordination breakdowns—like the one caused by the last-minute interface changes—still occurred. "People tended not to use a lot of these technologies, or there was great variance in how they used them and whether they trusted them," Cataldo says.

For instance, in another case study, unforeseen changes to the development schedule occurred that more closely tied together the work of two remote teams. Yet those developers weren't accustomed to sharing information, and as a result, many changes to the code had to be reverted because they broke the build.

Cataldo observed similar failures time and again in an actual corporate setting while working in research at industrial giant Robert Bosch GmbH before returning to Carnegie Mellon last year. "One of the classic criticisms when you do experiments with students is ... that this kind of stuff doesn't happen in the real world," he says. "But it happens all the time."

One surprising finding to emerge from the Global Studio Project is that having a central managerial team to coordinate a project—a traditional practice in software engineering—can actually create unwanted bottlenecks. Instead, it might be better to put more decision-making power directly in the hands of developers.

"We've learned from this work that transparency is hugely important," Herbsleb says. "Traditionally, we've kind of stressed control and actively managing things. But what we are finding out is that, much in the spirit of open source, it often works best to give people the resources and information they need to manage themselves and coordinate with each other."

For that approach to succeed, developers must first identify and then manage all of the dependencies—the parts of a project that impact each other—among the different components of software systems. "A lot of people assume that troubles are mainly caused by very complicated dependencies between pieces of software," Herbsleb says. "But even very simple dependencies can cause huge problems."

Forecasting these dependencies at the outset of a software project is no easy task, especially since they can change significantly over time. Cataldo and Herbsleb are now working to design tools that measure dependencies, and then use those data in a predictive fashion to tell developers on an hour-by-hour basis about changes being made elsewhere that might impact their work.

"Having this kind of dynamic computation of dependencies could be a big win for global software development," Herbsleb says. "Requirements change. Design ideas change. So inevitably the interfaces in a project change. In the face of that kind of dynamism, we need additional ways to coordinate the work."

Eye on the Clouds

> By Jennifer Bails

From voting records to Web traffic, datasets are getting bigger and bigger, meaning scientists need bigger and bigger computers to make sense of them. For more than a year, School of Computer Science researchers had exclusive access to Yahoo!'s 4,000-processor M45 supercomputing cluster for cloud-computing research—a rare opportunity in academia to do Internet-scale analysis.

Not surprisingly, the results were ... well, big.

"M45 has inspired a lot of research at Carnegie Mellon and has become an important way for us to do science," says Garth Gibson, professor of computer science and electrical and computer engineering. "People are now wondering how they are going to make progress when this resource goes away."

M45—named for the Pleiades star cluster—was reportedly one of the world's 50 most powerful supercomputers when it debuted in 2007. It runs an open-source distributed file system named HDFS and a parallel programming environment called Hadoop, allowing users to extract meaningful information from colossal amounts of data.

For example, Machine Learning Department head Tom Mitchell built a system called "Read the Web" that uses M45 to extract bits of knowledge with high (85 to 90 percent) accuracy from hundreds of millions of Web pages—processing more data than previously thought possible by two orders of magnitude. "Information extraction at this quality and scale has not been reported before," Mitchell says. "All of the results generated by our system would have been impossible without M45."

Likewise, SCS professor Christos Faloutsos harnessed the computing power of the Yahoo! cluster to develop PEGASUS, a peta-scale graph mining system that can work on enormous graphs with billions of edges and nodes.



Other scientists have been studying the effectiveness of M45 itself and developing ways to improve its performance. Gibson and his students are working on a project called DiskReduce that lowers storage overhead of Hadoop's file system by as much as a factor of three. It works by compressing redundant data where possible and deleting excess copies after the information has been encoded. "Right now you get one-third of the storage capacity you pay for," Gibson says. "We are trying to make the data go further."

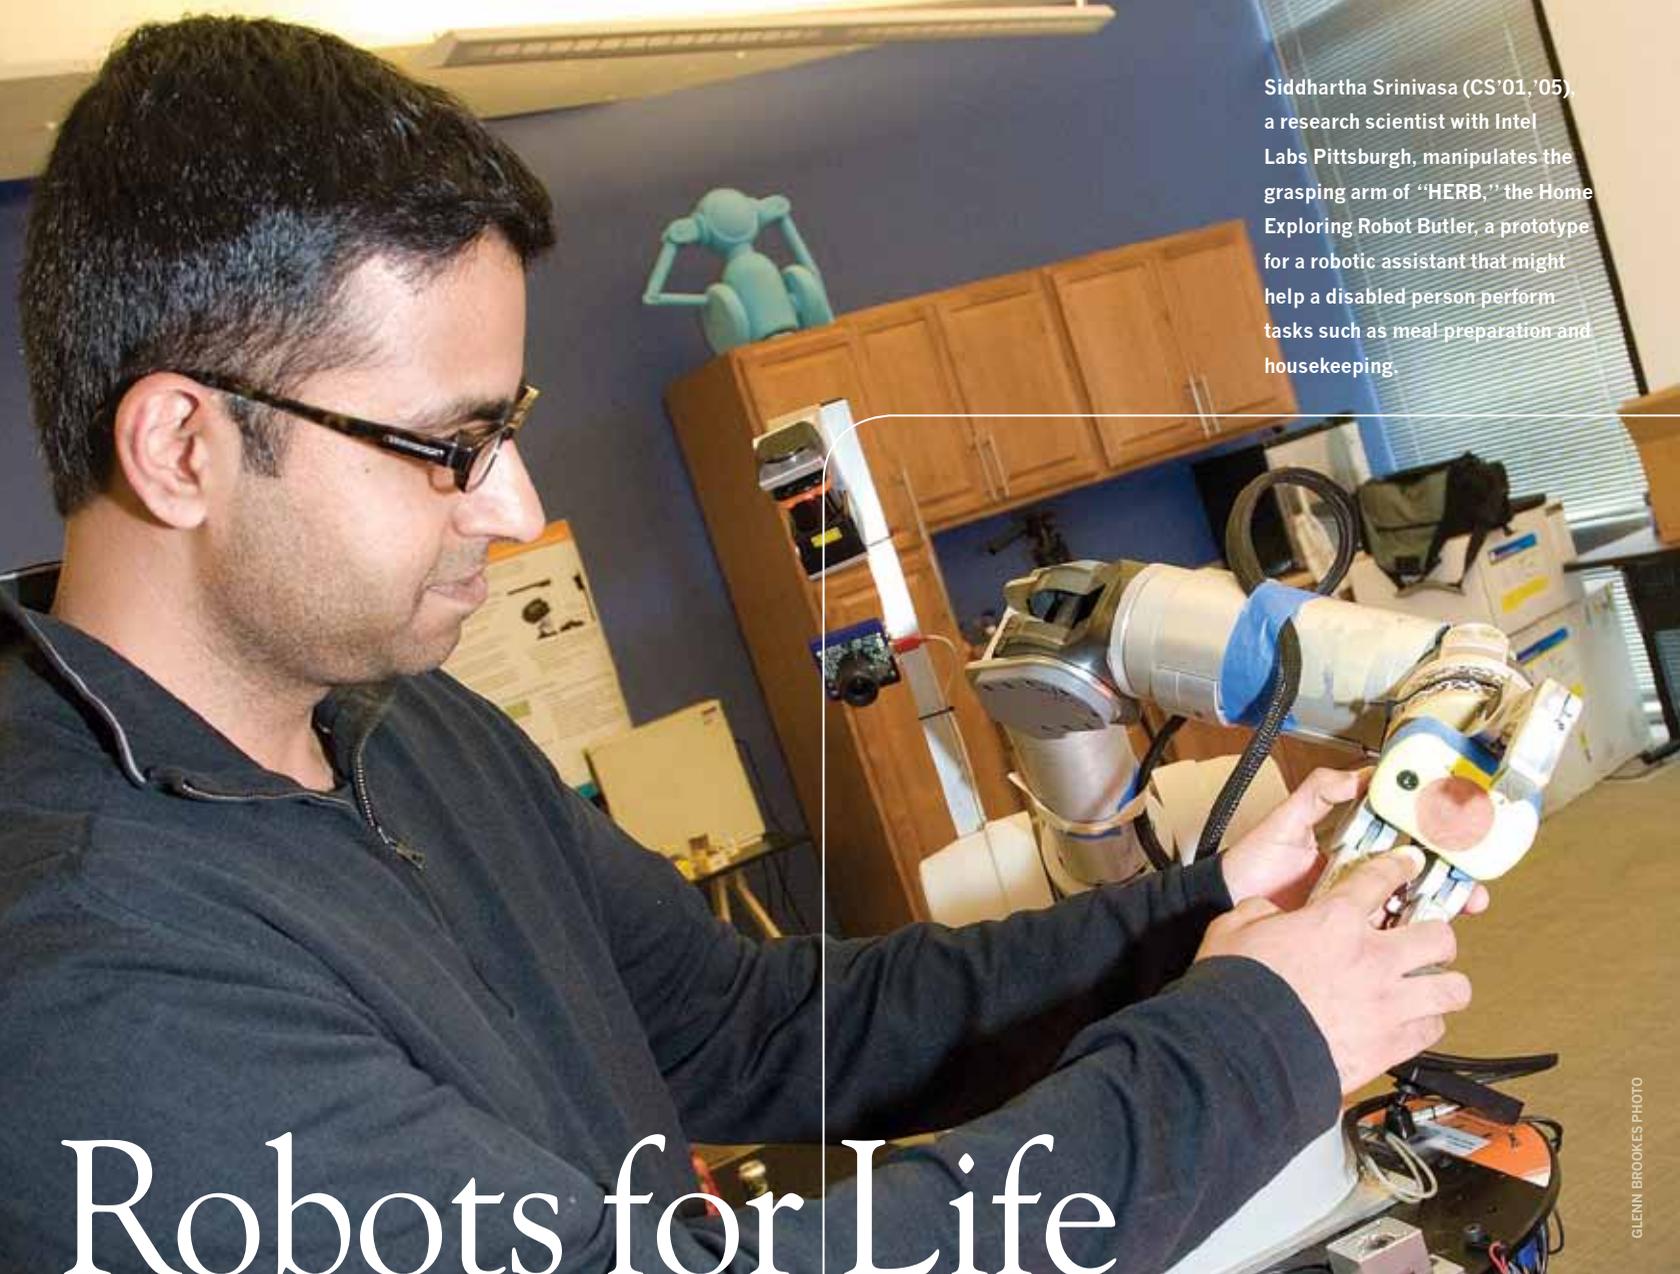
So far, access to the M45 cluster has resulted in the publication of more than two dozen academic papers from SCS in fields as diverse as natural language processing, astrophysics and machine translation. More advances are expected since Yahoo! expanded access to the supercomputer last spring to include University of California at Berkeley, Cornell University and the University of Massachusetts at Amherst.

Still, after less than three years in operation, M45 is starting to show its age—or at least other systems are starting to catch up. "It has moved from being a prime, new computer to getting a little old," Gibson says.

But research into cloud computing hasn't slowed down. In fact, a 64-node cluster funded by the National Science Foundation and a 78-node cluster sponsored by Intel have been set up on the Pittsburgh campus in the Parallel Data Laboratory, a collaborative effort of SCS and CMU's engineering school, Carnegie Institute of Technology. Both clusters are considerably smaller than M45, however, so researchers will continue to use the Yahoo! supercomputer as long as possible.

"We believe in data-intensive, scalable computing," Gibson says. "Two years ago we thought it would be important, and now we're doing it. And lots of important science in the future will be derived from analysis of large amounts of data in much the same way."

RACHEL ARNOLD SAGER ILLUSTRATION



Siddhartha Srinivasa (CS'01,'05), a research scientist with Intel Labs Pittsburgh, manipulates the grasping arm of "HERB," the Home Exploring Robot Butler, a prototype for a robotic assistant that might help a disabled person perform tasks such as meal preparation and housekeeping.

GLENN BROOKES PHOTO

Robots for Life

> By Tom Imerito

"Quality of life." It's one of those stubbornly vague terms that is hard to define. For some people, having a good "quality of life" might mean getting a job that allows them to spend lots of time with family. For others, good "quality of life" might mean a certain possession, like a luxurious home.

But for someone with limited mobility—a veteran with a missing limb, a car-crash survivor, an older person with severe arthritis—having a high "quality of life" is a lot simpler. It's about getting dressed in the morning, navigating the kitchen and going to work or school.

Yet hundreds of thousands of Americans every day face the prospect of becoming dependent on people or machines due to age-related illnesses or conditions. According to the U.S. Centers for Disease Control and Prevention, about 1.5 million people reside in skilled nursing facilities, and about 88 percent of them are age 65 or older. Another 76 million Americans born during the so-called "Baby Boom" of 1946 to 1964 are starting to reach retirement age.

Conditions related to aging aren't the only thing that can make simple, everyday tasks more difficult or force someone

Pittsburgh's Quality of Life Technology Center is solving the human-machine equation and spawning a new industry in the process

into an assisted-living facility. About 250,000 Americans are living with spinal cord injuries, according to the National Institutes of Health, while 2.2 million veterans have a service-related injury, reports the U.S. Bureau of Labor Statistics.

It's a strange dichotomy that technologists like Jim "Oz" Osborn (E'81,'86), longtime researcher and scientist in the Robotics Institute, have watched with frustration. At the same time that computers were making the world smaller and robots were exploring other planets, millions of other people were watching their lives shrink to just one small room.

If computers can run the anti-lock braking systems on our cars and put streaming video on our cell phones, why can't >>>

they make sure that an elderly person is receiving his medication?

If robots can navigate Three Mile Island or the surface of Mars, why can't they find their way around a suburban bedroom?

Shouldn't technology be able to address "quality of life"?

The Quality of Life Technology Center (QoLT) aims to answer those questions. Founded in 2006 as a joint project of Carnegie Mellon and the University of Pittsburgh, it's merged CMU's leadership in developing and deploying robotics and sensor technology with Pitt's excellence in health care and rehabilitation sciences. The QoLT is one of five Engineering Research Centers in biotechnology and health care that's recognized and funded by the National Science Foundation. The center receives \$4 million annually from the NSF and was recently awarded a three-year, \$1.5 million NSF Innovation Award to get "quality of life technology" into the hands of people who need it as quickly as possible.

Highly interdisciplinary, the QoLT draws on the talents of roboticists, computer scientists,

engineers, industrial designers, psychologists, social scientists and other researchers at both universities. Its commercialization counterpart, the Quality of Life Technology Foundry, has already spun off three companies based on ideas created at CMU and Pitt, says Rick McCullough, Carnegie Mellon vice president for research and former dean of the Mellon College of Science. With the new funding, he says, "we envision even greater success."

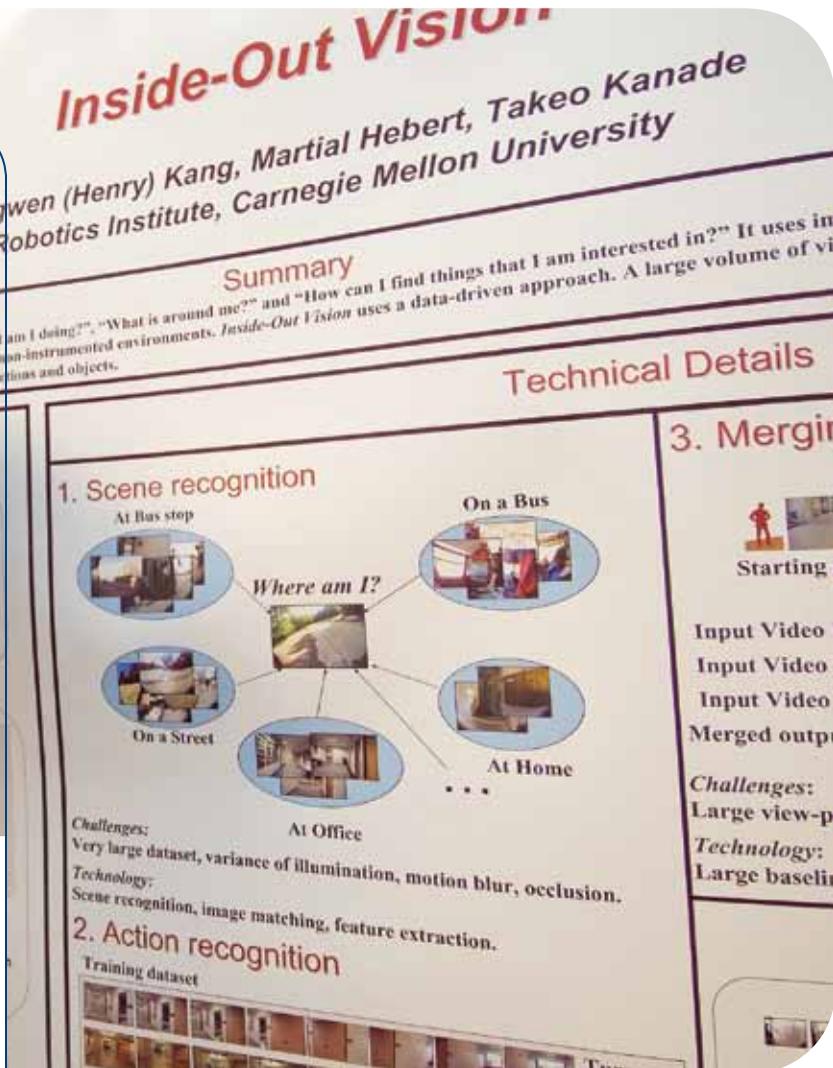
"Everybody wants to live independently, enjoy their life, feel secure," says Osborn, who serves as executive director of the Quality of Life Technology Center. "Robots have taken us into outer space, under the sea and into places too dangerous to go. QoLT gives us the opportunity to merge these technologies into rehabilitation and geriatrics—and at the same time spawn a new industry."

In various stages of development at QoLT are more than 20 assistive devices and systems ranging from things as simple as a computer application called "Lean and Zoom" that automatically zooms a computer screen in response to its viewer's body movements; to a full-scale assistive home that allows persons of limited abilities to live independently.

And thanks to its new funding, the QoLT Foundry is expanding its efforts to accelerate the new-product commercialization process by looking at broader consumer possibilities for assistive technology. "In the past, the typical assistive device has been produced in very limited quantities and paid for by the government," says Curt Stone, director of the QoLT Foundry. "We're trying to develop a business model where large consumer demand results in higher production and lower pricing, so consumers can pay for these products on their own."

For instance, a vest designed to allow people with hearing impairments to enjoy music by means of vibration has morphed into a consumer product called "VibeAttire." It allows the wearer to "feel" the music from an iPod, MP3 player or any other handheld audio device.

Integrating "man" and "machine" requires finesse, says Takeo Kanade, director of the QoLT Center and Carnegie Mellon's U.A. and Helen Whitaker University Professor of Computer Science and Robotics. "Until now, the key goal of robots has been to reduce human involvement," he says.



"First-Person" or "Inside-Out" vision devices that can see exactly what a person is looking at can help a robotic assistant to predict what tasks that person is trying to perform. Wearing a prototype eyepiece is Michael S. Devyver, a member of the research staff at the Robotics Institute.

“However, I don’t think people want to live with robots that do everything for them.”

For Kanade, the “magic equation” of a quality of life technology system is solved by taking what a person wants to do, and subtracting what he or she can do. “The difference between desire and ability is what the machine should provide,” he says. “In my mind, human ‘quality of life’ is the (ability) to live as independently as possible—to do what we can do for ourselves, rather than having it done by other people or machines.”

The practical and technical challenges associated with solving that equation are daunting, Kanade says. “One of the key components of quality of life technology is to understand what the user wants,” he says. “Because unless a machine understands what its user wants, it cannot help. The machine system and the human being must work together to make a better system. We call that a human–system symbiosis.”

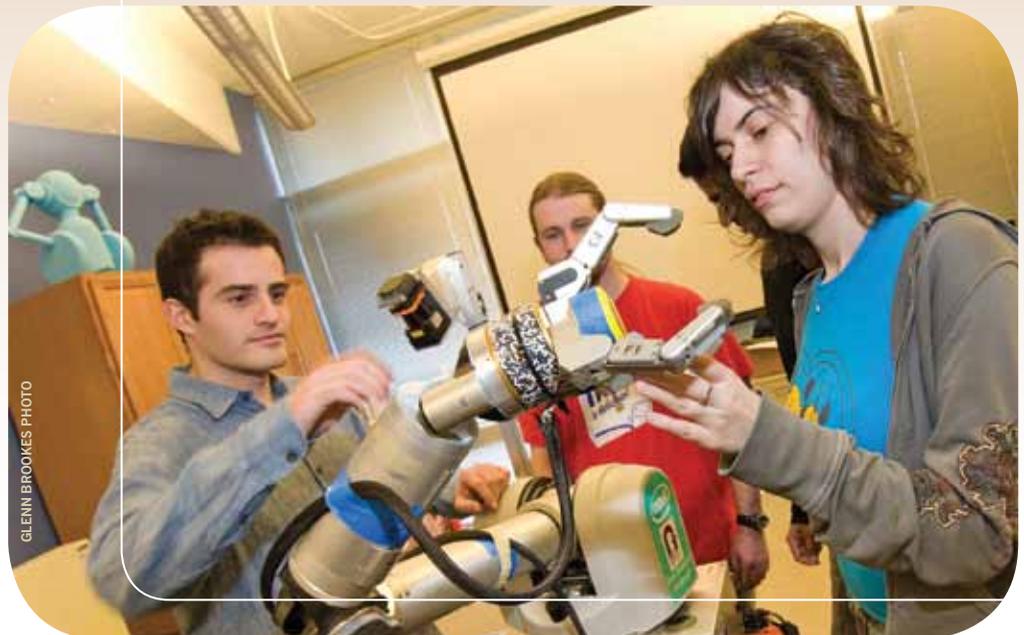
QoLT has reduced the problem of symbiosis between people and machines to four research thrusts:

- Living a Life: The Person and Society
- Knowing the Person: Perception and Awareness
- Using the Machine: The Human-System Interface
- Interacting With the World: Mobility and Manipulation

Rory Cooper, co-director of the QoLT and director of the Human Engineering Research Laboratories (a joint project of the University of Pittsburgh and U.S. Department of Veterans Affairs), says the technologies under development at the center must be able to address and mesh aspects of all four thrusts. For instance, the first area—“Living a Life”—requires researchers to assess the practical implications of machine assistance on a person’s needs. The second addresses how machines can best sense a user’s intentions and interpret them. The third and fourth areas focus on the ways that people can share information about their environments with an assistive technology, and then use that technology to interact with their environments and one another.

Some of the interfaces that users might deploy are as familiar as keypads, joysticks and touch-screen monitors. Others are more unusual. Take the visual recognition system called “First-Person Vision” being developed by a team led by Kanade and Robotics Institute professor Martial Hebert.

“First-Person Vision” uses a pair of tiny, eyeglass-mounted cameras to allow a computer to “see”



GLENN BROOKES PHOTO

Robotics Ph.D. students Alvaro Collet, Ross Knepper and Anca Dragan put HERB through its paces at the Intel Lab in the Collaborative Innovation Center on the Pittsburgh campus.

whatever the person wearing the glasses is looking at. The first camera uses a wide-angle lens to look at the full field of human vision, while the second uses a close-up lens to examine the user’s eye and pinpoint the exact area he or she is focusing on. From that composite result, the computer can determine what the user’s intentions are.

When a user is looking at a can of soup on a kitchen counter, for instance, the computer can guess that the user needs it. If the computer is controlling a robotic arm, it can reach out for that can and bring it back for the user.

“If a computer has access to a user’s field of view and area of interest, it will have a much easier time knowing what he or she is doing and thinking,” says Kanade, acknowledged as one of the world’s leading experts in computer vision—the science of enabling computers to see and interact with their environments. “First-Person Vision” is one of the biggest steps toward solving the quality of life equation, he says: “If the computer knows you want to reach something, and—if you reach it—knows that it doesn’t need to help you, that’s an effective quality of life system.”

Manipulating a physical object—that hypothetical can of soup, for instance—is part of the fourth research thrust. One of the technologies being developed at QoLT that Cooper points to with pride is the Personal Mobility and Manipulation Appliance Robot, or PerMMA, a two-armed intelligent wheelchair that can automatically transport a user around the house, open the refrigerator door, remove a snack from a shelf and put it to its user’s lips to eat.

It can map and retain local environments, memorize user activity and predict future behavior. It can even reach nine feet into the air to change a light bulb in a ceiling fixture, without using a

“Everybody wants to live independently, enjoy their life, feel secure. Robots have taken us into outer space, under the sea and into places too dangerous to go. QoLT gives us the opportunity to merge these technologies into rehabilitation and geriatrics—and at the same time spawn a new industry.”

Jim “Oz” Osborn (E’81, ’86)

ladder or breaking the bulb. Depending on a user’s individual requirements, PerMMA can be fitted with all manner of sensors, cameras, monitors, interfaces and tools. In order to facilitate squeezing through narrow halls and doorways, PerMMA’s robotic arms swing out of the way to the back of the device.

Currently being tested in a laboratory at the VA Pittsburgh Healthcare System’s hospital campus in the Highland Park section of the city, PerMMA isn’t quite ready for deployment to users. But it’s getting closer. And Cooper says the goal of QoLT is always to deliver its technology “into the hands of health care and rehab professionals and patients.”

One place where QoLT systems and sensors are already being used is in nearby McKeesport. The former mill town just south of Pittsburgh is home to the first of a series of “research cottages” built by a non-profit partner of the QoLT, Bluerroof Technologies.

Bluerroof’s research cottages are designed to be fully equipped and comfortable homes where senior citizens or others with limited mobility can live independent lives. They’re also >>>

test beds for what QoLT calls “active home” technologies. Equipped with more than 100 sensors, cameras and automation components, the research cottages allow a computer to log human activity inside.

“The sensors can detect how many people are in a room by the level of carbon dioxide,” says John Bertoty, chief executive officer of Blueroot. “They tell the system when the cabinet and refrigerator doors have been opened, so it knows if the occupants are eating. And if they’re not, it will send out a caregiver alert.”

The computer can also alert the residents to pending events like visitors, doctor’s appointments or emergencies, and remind them to take their medication, close doors or windows or shut off water accidentally left running.

Other active research areas include a “safe driving” effort, Cooper says, focused on applications

camera to automatically record and log pictures and sounds from a person’s daily life for future viewing and memory improvement exercises.

Besides their obvious benefit to mankind, these quality of life technologies also represent a chance for Carnegie Mellon and Pitt to aid the region’s economic development. Blueroot, for instance, is located in an economically depressed area where up to a quarter of the population lives below the poverty line. Bertoty says the non-profit is already putting local people to work building homes with assistive technologies and hopes to eventually create an assembly plant where it can manufacture complete pre-fabricated structures.

Stone, the director of the QoLT Foundry, predicts that projects spawned at the Quality of Life Technology Center will be responsible for more than 100 jobs in the Pittsburgh area over the next five years. “With the greater resources made possible by the NSF, we think we can spin out even more companies,” he says. QoLT is also preparing undergraduates and graduate students to develop new assistive technologies through education and research programs.



The Personal Mobility & Manipulation Appliance is a semi-autonomous robotic manipulator that also has a seat for a person. Shown testing PerMMA’s capabilities are JiJie Xu, a Pitt post-doc working in the Human Engineering Research Laboratories, and Benjamin Salatin, a graduate student in Pitt’s School of Health and Rehabilitation Sciences.

that help older adults and people with disabilities operate vehicles with confidence with the help of global-positioning systems and automated, driver-specific advice.

Another effort—Virtual Coaching—addresses not physical disabilities, but memory and cognitive disorders. MemExerciser, an assistive device for people with impaired memories, uses a wearable

Older Americans in 10 or 20 years won’t accept the sedentary lifestyle their grandparents’ generation might have tolerated, Osborn says. “The baby boomers represent unprecedented numbers, unprecedented wealth and unprecedented expectations” for maintaining a high quality of life as they age, he says. “That’s a whole different society than the world has ever known before. Quality of life technology represents a great commercial opportunity for companies that can provide those services.”

Tom Imerito is president of Pittsburgh-based Science Communications. A writer and technology consultant, his articles have appeared in Pittsburgh Quarterly, Research Penn State, the Journal of the Minerals, Metals and Materials Society and TEQ, the magazine of the Pittsburgh Technology Council. Email him at thomas@science-communications.com. Link Managing Editor Jason Togyer contributed to this story.

In the Loop

Lenore Blum

Distinguished Career Professor of Computer Science

Lenore Blum is Distinguished Career Professor of Computer Science at Carnegie Mellon University. As an undergraduate at Carnegie Tech, she took academic computing with Alan Perlis, and earned her doctorate in mathematics at Massachusetts Institute of Technology.

Before joining the SCS faculty, Blum founded the Department of Mathematics and Computer Science at California’s Mills College and served for 13 years as its head. She also held research positions at the University of California at Berkeley and the International Computer Science Institute, and served as deputy director of the Mathematical Sciences Research Institute.

Her research includes developing a theory of computation and complexity over the real and complex numbers, combining ideas from mathematics and computer science. A tireless advocate for increasing the participation of women in math and science, she was a founder and past president of the Association for Women in Mathematics and the founder of Women@SCS.

Blum recently spoke to Link Managing Editor Jason Togyer about her newest initiative—Project Olympus (olympus.cs.cmu.edu), which helps innovators turn their research into commercially marketable ideas.

What was your motive for founding Project Olympus?

There were actually two. One came out of my work as co-director of the ALADDIN Center. Our goal was to show that key computer science algorithms had applications across many different domains. It occurred to me some of these ideas could be commercialized to support our research and regional economic growth.

The other motive was to create an entrepreneurial environment that would encourage our students to stay in the area. When I looked at our student body a few years ago, less than five percent our SCS graduates were staying in Pittsburgh. Our competitors—MIT, Stanford, Berkeley—have huge technology infrastructures in their communities that are spinoffs of those institutions.



I wanted to create an environment where faculty and students would stay here and make a positive impact on the region's economy.

How does Project Olympus support innovators?

Olympus supports what we call "PROBEs"—Project-Oriented Business Explorations—which are teams of faculty and students working to explore the commercial potential of their research.

Some PROBEs need space, some need equipment and some need money. Everyone needs connections and visibility. We have great incubator space off Craig Street, and we provide connections with advisors and collaborators on and off campus.

We try to move quickly and we haven't created much bureaucracy. A lot of places have big long applications. We don't need that, because our students and faculty have already been vetted.

I mean, the faculty and students are brilliant. Why would I say no to them?

Don't you need to see a business plan first?

We certainly want to see a plan, and we have said no when things have looked very bad, but otherwise we rarely say no. We're more interested in seeing their enthusiasm, their excitement and their commitment to working on their project.

With faculty, for instance, I don't want them spending all of their time on business plans. We get teams of MBAs to provide market research and business plans for our faculty PROBEs.

And we do provide guidance. We are fortunate to have Babs Carryer as our embedded entrepreneur, and Kit Needham as our senior business advisor. One of our main goals is eventually to provide our faculty with "Innovation Fellows" who spend their time working on the business aspects.

If these innovations are so good, why doesn't the private sector support them?

Some of these ideas are at a very early stage. Whatever you think of as "early stage," we provide help even earlier than that.

Also, unlike Silicon Valley, the venture capital climate in Pittsburgh tends to be very risk-averse. And that's a concern. We're trying to change that climate, eliminate the roadblocks that keep innovators from receiving funding and show Pittsburgh that you've got to be less risk-averse.

What sorts of PROBEs are people working on right now?

I'm very excited about our SCS faculty PROBEs. A couple of our newer PROBEs have recently become companies. Safaba Translation Solutions, headed by Alon Lavie, is a spinoff of products developed at the Language Technologies Institute, while Jamie Callan's TellTale Information—which also comes out of LTI—allows government agencies to automatically manage all of the filings and documents they receive.

And then there is POW!—that's Mor Harchol-Balter of the Computer Science Department—which develops optimal algorithms for power management. Just starting up are Graphics Parallelism in the Cloud, which is Adrien Treuille of the Robotics Institute and CSD, which is poised to become the next revolution in the gaming industry, and GGideaLab—Carlos Guestrin from Machine Learning and CSD, and Seth Goldstein of CSD—which is trying to optimize the usefulness of social networks.

How many PROBEs do wind up becoming successful spinoffs?

So far, we've had about 45 PROBEs, and more than half have become companies. Our very first faculty PROBE was reCAPTCHA, which built on Luis von Ahn's work and has since been acquired by Google. One of our student PROBEs was Dynamics, which is developing a kind of fraud-resistant credit card, and that's received almost \$6 million in venture capital funding.

Does that mean Project Olympus is sustainable for the long run?

We've really been operating on a shoestring budget compared with other major university centers that do similar work with entrepreneurs. The Deshpande Center at MIT, for instance, was initially funded by a gift of \$20 million.

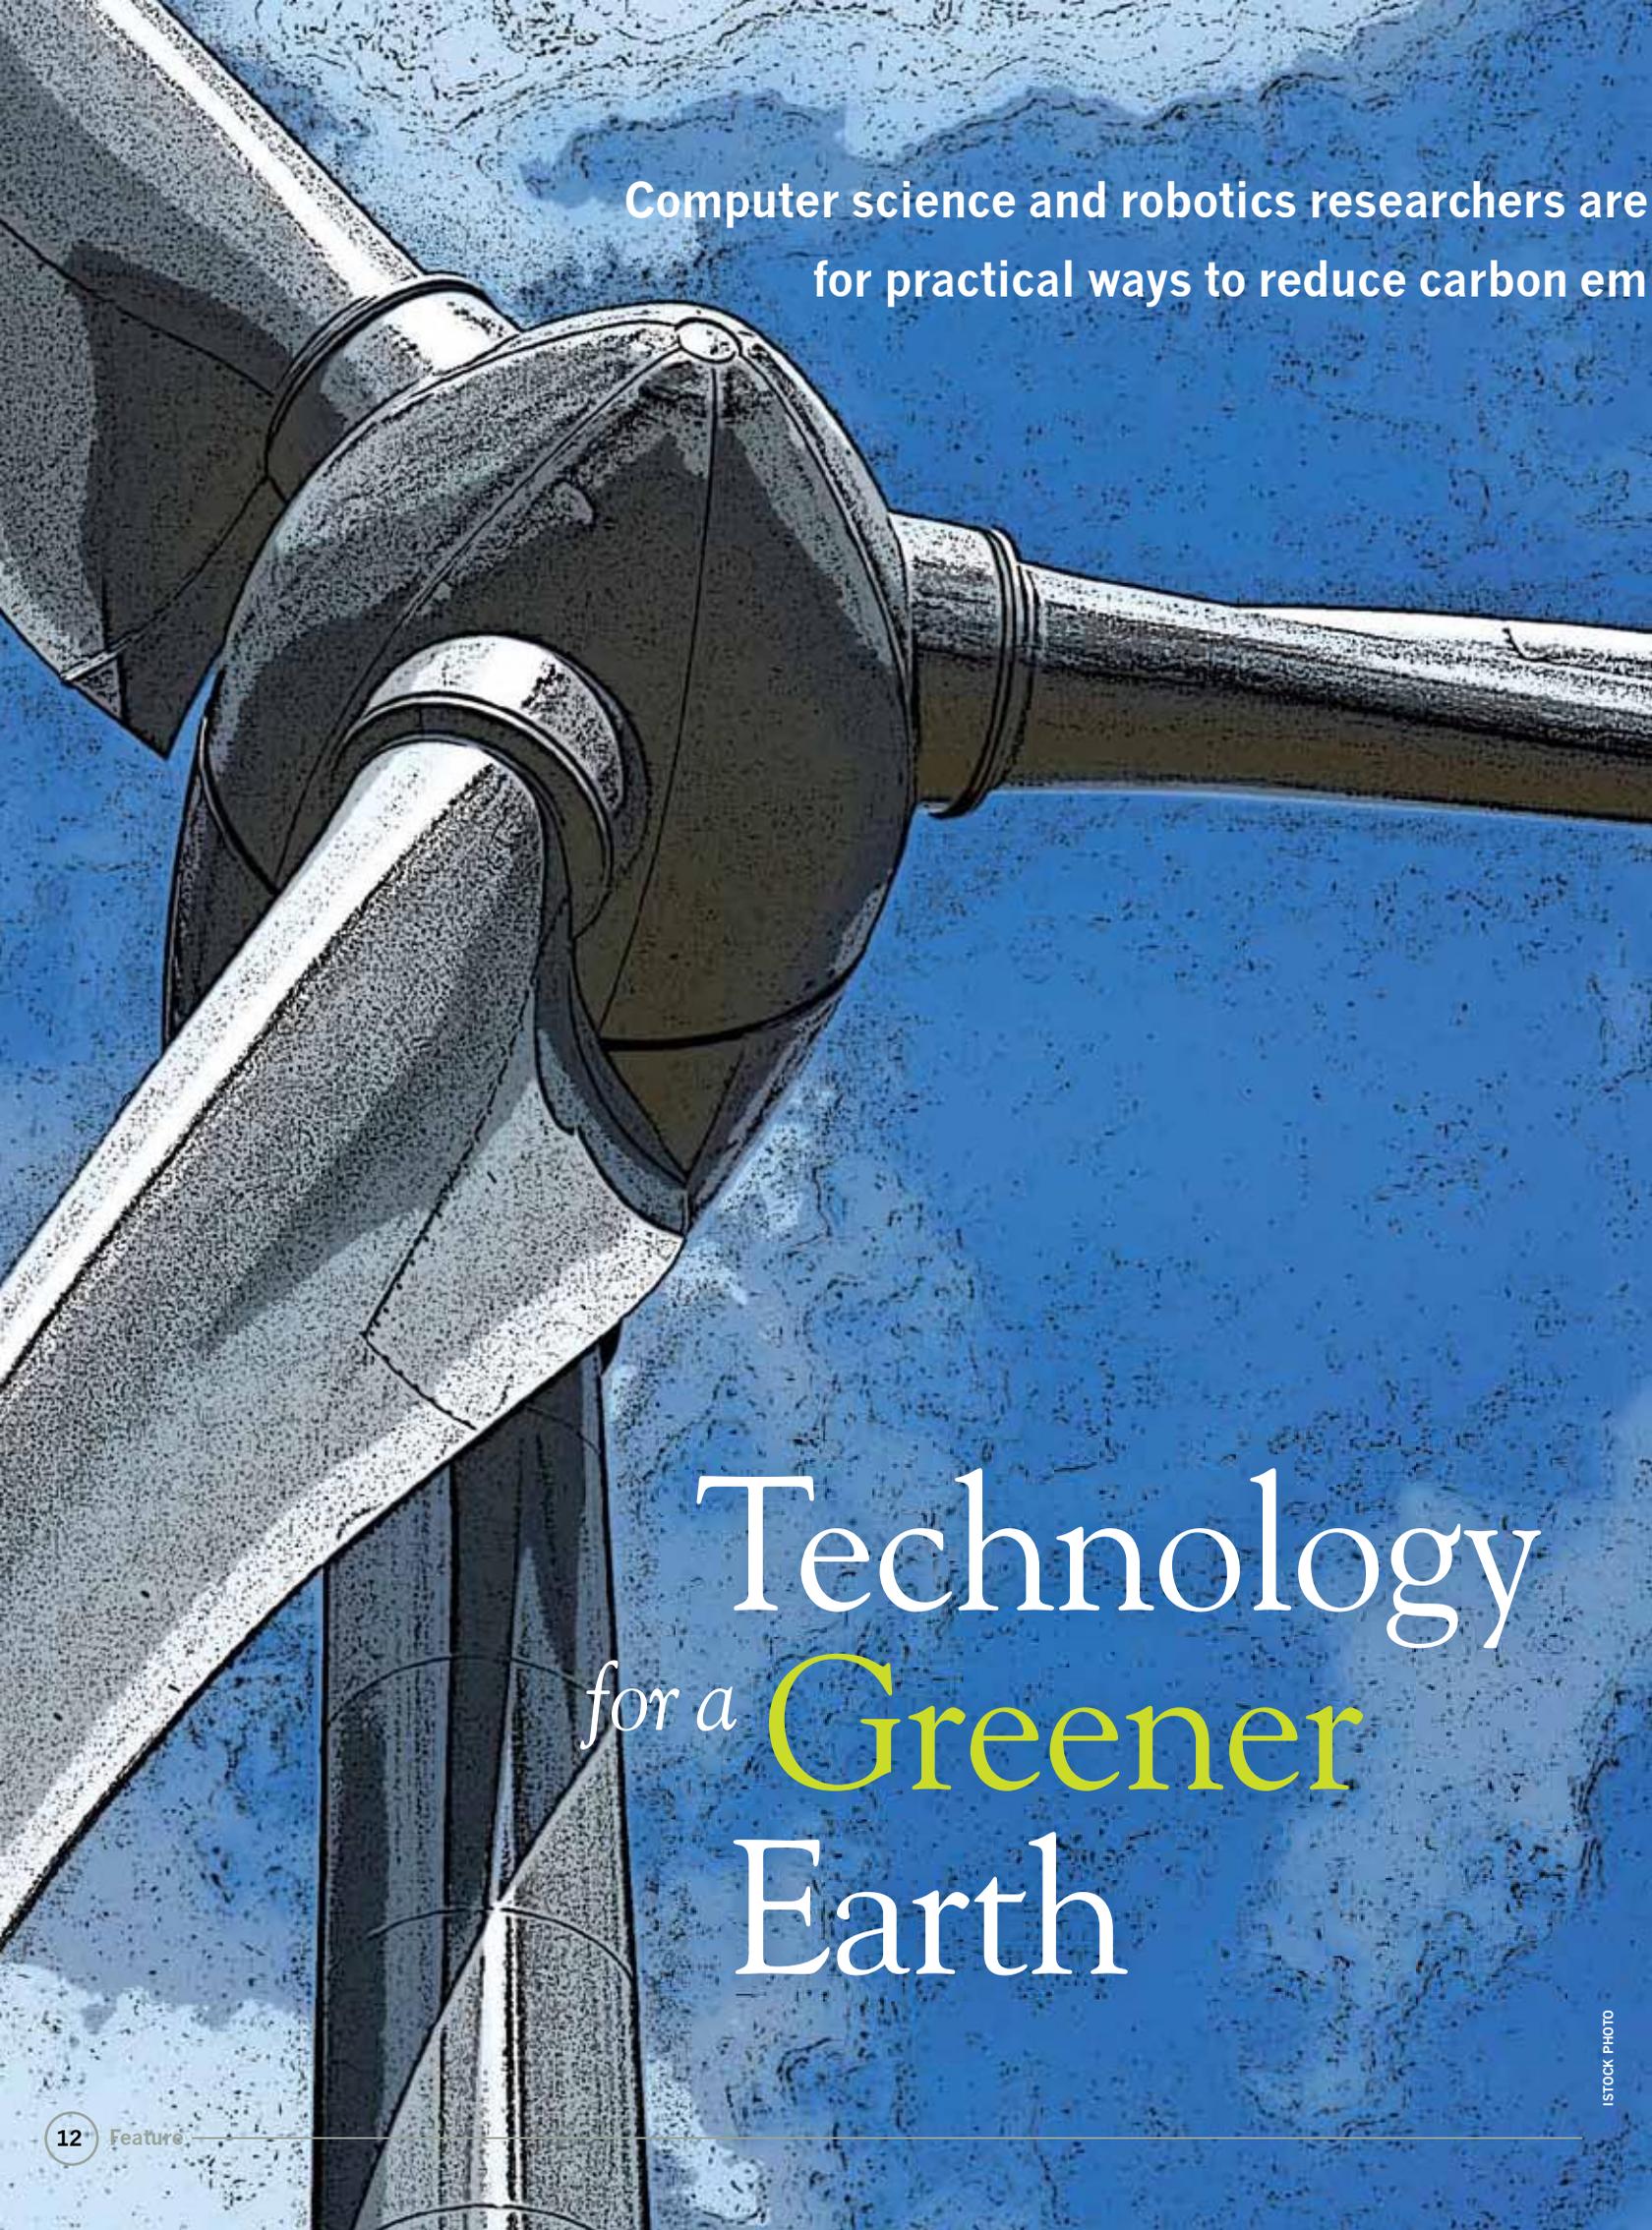
At the moment, we're relying on foundation grants. In our case, the Heinz Endowments provided initial funding of \$400,000, and then followed up with \$245,000. We've received additional support from various economic development agencies and the university, and we're seeking donor gifts.

Right now, we're exploring models for long-term stability, such as a kind of hybrid that would combine both for-profit and non-profit investment and equity ownership, as well as a membership-based model.

Could any of the PROBE successes contribute financially to Olympus?

We are part of the university—in fact we are an initiative of SCS—and so we can't directly invest in companies that come out of Olympus.

But we're not purely altruistic! We say to everybody: "When you make it big, remember us."



Computer science and robotics researchers are
for practical ways to reduce carbon em

Technology *for a* Greener Earth

searching missions and energy use

> By Jason Togyer

In a four-by-six-mile rectangle of a dusty mountain range about one hour east of Palo Alto, Calif., sits the Altamont Pass Wind Farm, a collection of nearly 5,000 power-generating electric wind turbines.

Since 1981, the Altamont Pass Wind Farm has been a landmark for tourists and—as the largest collection of wind turbines in the world—a symbol of the promise of harnessing the wind for renewable, clean electric power.

It's also very badly designed.

For one thing, many of the turbines are too low to the ground and too closely spaced. For another, air currents in the valley are unreliable, meaning the turbines regularly slow or halt, so the electricity they generate must be supplemented by natural gas power plants. And the turbines kill about 4,700 birds each year, including about 70 golden eagles and 1,300 other raptors, because the Altamont Pass Wind Farm was built smack in the middle of a major migratory path.

In short, Altamont Pass—as one environmentalist told the *San Francisco Chronicle*—was “the exact wrong place” to build a wind farm. But what's the “exact right place”? And what is the “exact right,” or optimized, configuration for the arrangement of the wind turbines?

Jaime Carbonell is no meteorologist or environmental scientist. Much of his research was focused on machine translation and language technologies. But a chance encounter with a utility-company executive during a long airplane flight inspired him to ask questions about wind energy.

It turns out that three decades after the construction of the Altamont Pass Wind Farm, the computer models used to design them remain crude at best, says Carbonell, head of the Language Technologies Institute in the School of Computer Science. He also was surprised to learn that many power companies lack in-house expertise and rely on turbine manufacturers to help them plan wind farms. “Each one does an ‘impartial’ study and concludes its products are the best,” Carbonell says.

Some preliminary investigation left him wondering if the optimization methods he was using on large passages of text could be used to optimize wind farm design. Now, Carbonell is working on just that in collaboration with Jay Apt, distinguished service professor in the Department of

One of the newest sources of wind-generated electric power in Pennsylvania is Iberdrola Renewables' Casselman Wind Power Farm, about 60 miles southeast of the Pittsburgh campus. Some of the same optimization techniques used in machine translation can be applied to development of a universal model for designing wind farms, says LTI Director Jaime Carbonell.



PHOTO COURTESY IBERDROLA RENEWABLES

Engineering and Public Policy and executive director of Carnegie Mellon's Electricity Industry Center.

It's a new area of research for Carbonell, but one he finds very attractive as a computer scientist. “Very often, the distance between ‘theory’ and ‘practice’ is very long,” he says. “This is a problem that matters. This is a case where some of the optimization techniques we're using can make a real difference.”

He's not the only person in SCS working on what might be called “green” technologies. Heightened awareness of global climate change—and the mounting evidence that human production of greenhouse gases is responsible—has spurred other researchers to see where computer science might help reduce our carbon footprint.

Some SCS faculty members, like Carbonell, are tackling environmental problems using computer models. Others, such as Jen Mankoff of the Human-Computer Interaction Institute and former SCS Dean Jim Morris, now working on the West Coast, are deploying web and mobile phone apps designed to encourage better behavior.

And a few, including Illah Nourbakhsh and Gregg Podnar of the Robotics Institute, are getting their hands dirty—literally—by launching a grassroots effort to convert gasoline-powered cars into sophisticated electric vehicles.

In some cases, green projects are showing benefits besides saving energy and cutting down on greenhouse gas emissions. Take David Andersen's effort to design less electricity-hungry servers using slower processors and low-power storage devices. Rewriting code to run on these so-called “wimpy nodes” not only saves electricity, the same techniques enable that code to run much more quickly on more powerful processors. “It was a very nice dividend,” says Andersen, an assistant professor of computer science.

In September, SCS launched a new monthly series of seminars on Sustainability and Computer Science as a forum for researchers to discuss ways in which computational thinking could be applied to energy conservation and the environment. The forums also provide opportunities for School of Computer Science personnel to share ideas with colleagues from other parts of the university, including the School of Architecture, the departments of Engineering and Public Policy and Civil and Environmental Engineering, and the Steinbrenner Institute for Environmental Education and Research.

Behind the seminars and the research activity is a new sense of urgency over climate trends that scientists say are both obvious and alarming. By examining ice samples dating back approximately 650,000 years, the National Oceanic and >>>

“I have two young kids, and a job that takes me away from them a lot.

*I always want to be able to tell my kids when I leave them that I’m
working on something important enough to leave them.”*

Jen Mankoff

Atmospheric Administration (NOAA) estimates that current atmospheric levels of carbon dioxide—a byproduct of burning fossil fuels and one of the main greenhouse gases—have never been as high as they are today. Over the past decade, NOAA says, ocean levels have risen at a rate twice that of the previous 100 years; polar ice caps have steadily shrunk in area and thickness; and storms have grown in intensity and volume.

Mankoff, an associate professor in HCII, remembers reading about those and other climate change “tipping points” and coming to a sudden realization “that I could have a bigger impact on these issues by bringing them into my research.”

With several colleagues, she developed a Web site called StepGreen.org, which suggests “green actions” that users can incorporate into their daily lives and measures the impact of their energy use. StepGreen also functions as an umbrella system for deploying new ideas, including sensors and mobile devices for helping its subscribers make better decisions about green choices.

“I have two young kids, and a job that takes me away from them a lot,” says Mankoff, who helped organize the SCS seminars on Sustainability and Computer Science. “I always want to be able to tell my kids when I leave them that I’m working on something important enough to leave them.”

There are flinty-eyed financial reasons for doing sustainability research as well. At a campus-wide workshop Feb. 24, Rick McCullough, CMU’s vice president of research, noted that the federal stimulus package alone has allocated \$4.4 billion to the U.S. Department of Energy for research into renewable energy. Other state and federal agencies are also funding research into sustainable power generation.

The private sector is interested as well. In 2006, the last year for which data is publicly available, servers and data centers in the United States consumed 61 billion kilowatt-hours of electricity at a cost of \$4.5 billion. Reducing that energy consumption even 10 percent would mean big cash savings. No wonder cloud-computing giants such as Google and major chipmakers such as Intel and AMD are funding research—at Carnegie Mellon and elsewhere—into ways to

reduce their energy consumption. In February, during the first-ever round of Google Focused Research Awards, Andersen and Mor Harchol-Balter, associate professor of computer science, received \$100,000 to study ways to build energy-efficient computing clusters.

“The bottom line is that there are opportunities in energy research, and they’re growing,” McCullough says. “We have some really fortunate strengths in the environment, solar power, electrical storage, carbon-capture technology, wind and water energy, policy issues—there are excellent, interesting things going on here at Carnegie Mellon.”

Wind energy represents one of the biggest research opportunities. A 2006 report by the U.S. Department of Energy suggested that 20 percent of the nation’s electric power needs could be met with wind by 2030. But under that scenario, new wind power installations have to come on line at a rate of 16,000-megawatts per year from 2018 through 2030. While engineers have learned from mistakes made at Altamont Pass and elsewhere, finding a location for a wind farm and then placing the turbines in a way to maximize the amount of energy generated isn’t a simple task.

Wind farm designers must be mindful of local and state environmental regulations, ground conditions, access to power lines—and of course, the prevailing winds, says Paul Copleman, a spokesman for Radnor, Pa., based Iberdrola Renewables, which operates 38 American wind farms, including the Casselman Wind Power Project in Somerset County. The 34.5-megawatt generating site on a former coal mine 60 miles southeast of Carnegie Mellon’s Pittsburgh campus was dedicated in 2008 and is one of Pennsylvania’s newest wind farms. “You obviously want to maximize the number of turbines you can fit on the land you’ve leased,” Copleman says. “By the same token, you can’t just clutter up the landscape with turbines.”

Understanding wind conditions, Copleman says, requires at least a year’s worth of monitoring and three or more sensor stations to provide continuous sampling of information that’s unavailable from public sources. “It’s very valuable data,” he says.

But those sensor stations aren’t cheap to operate, says Carbonell of the LTI, so while adding more monitors obtains more accurate data, it also increases the cost of the project and can delay construction. Building a computer model that approximates wind conditions using only a handful of monitoring sites would save both time and money, Carbonell says.

Such a computer model has to account for many other variables as well, he says. While the population of the United States tends to be concentrated along its coasts, the strongest, steadiest winds are found in the sparsely populated Dakotas. Wind farms need to be close to trunk power lines, but far enough from major cities that large tracts of property can be acquired affordably, away from nearby residential areas. And they need to be away from flight paths for both migratory birds and aircraft.

All of these variables can be optimized, and although wind farms have little in common with machine translation, designing a computer model employs many of the same techniques, Carbonell says. “In machine translation, we’re often trying to align words, phrases and texts in different languages,” he says. “You don’t align them in isolation, you align them in context.”

Wind farm design is a similar multi-variable optimization problem, Carbonell says. “You need to be able to provide the most power at the least cost. If the price of land acquisition goes up, it changes one part of the equation. If the distance to the power grid changes, it changes another part of the equation. So while these are completely different problems, the same tools can be applied.”

Tools—physical ones, such as wrenches and screwdrivers—are exactly what Nourbakhsh and Podnar have in mind at the Robotics Institute’s Community Robotics, Education and Technology Empowerment Lab. The CREATE Lab is home to a new open-source project to get both “code jockeys” and “gear heads” together to develop a practical process for converting gasoline-powered automobiles into electric vehicles.

Dropping an electric motor and some batteries into a second-hand car isn’t a new idea, says Nourbakhsh, associate professor of robotics and co-principal investigator of the Charge Car project with Podnar, an RI program manager. “There are tons of people who do it,” Nourbakhsh says. “And there are tons of people who do it—and fail.” Many of those projects flounder because adding batteries to increase the range of the vehicle also adds weight while subtracting passenger and cargo space.

Nourbakhsh has first-hand experience with such conversions—since 2001, he and his wife have been driving a Toyota RAV4 that runs on straight

electric power. But more important experience might have come while he worked on his bachelor's degree at Stanford, where he helped drive a solar-powered vehicle from Orlando to Detroit using only the sun for energy. "We were optimal controllers," he says. "We paid attention to motor control regulation, trying to predict where the hills were, and we paid attention to where the sun was on the highway."

Optimal control is a goal of Charge Car's prototype, built into a 2001 Scion. It uses an extra high-density capacitor—or "ultracapacitor"—to store electricity and lengthen the life of a vehicle's batteries by handling heavy current draws, such as when the Scion climbs a hill or accelerates away from an intersection. Charge Car is about to launch a monthly contest challenging developers to create intelligent, open-source power-management software that learns a person's driving habits and adjusts the charging and discharging of the ultracapacitor to maximize battery life and vehicle performance.

Using GPS devices, Charge Car is also collecting daily data on real commutes from all over the United States. That's already provided some surprising information—such as the fact that vehicles used on urban commutes spend up to 35 percent of their time decelerating, and not 3 percent, as older models assumed.

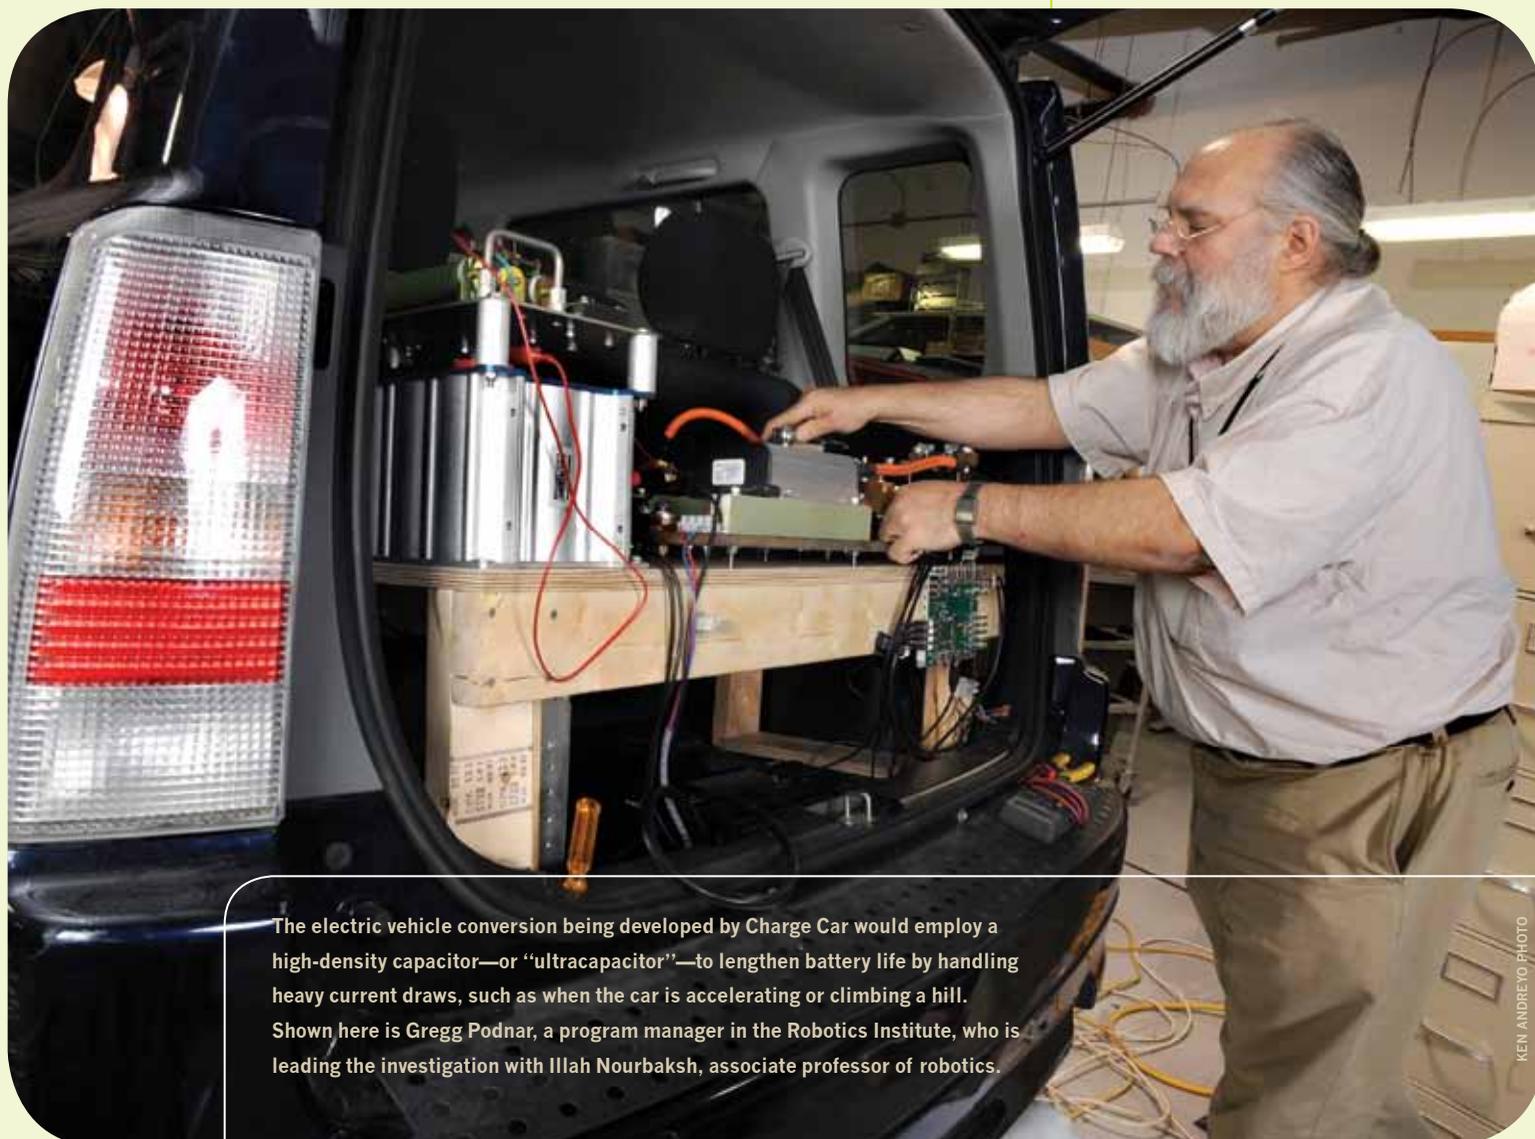
And it's not the only type of crowdsourcing envisioned by Charge Car. It's reaching out to Pittsburgh-area auto mechanics, vocational teachers and electronics and auto hobbyists, tapping their knowledge to create an optimum electric vehicle conversion kit that could be manufactured and installed in late-model cars for under \$6,000.

"Pittsburgh is perfect for this kind of project," Nourbakhsh says. "We have foundations who are interested in this, we have the university community and we have this incredible industrial background. And Pittsburgh has hills and different kinds of weather, so if Charge Car works here, it will probably work in 80 percent of the country."

On a smaller scale in terms of size—but with potential to make a big impact in computer science—are the so-called "wimpy nodes" being tested by Andersen. Servers in data centers use a tremendous amount of energy, he says, but much of that power is wasted. Speedy processors consume more power to perform operations than slower processors that can perform the same >>>

"The bottom line is that there are opportunities in energy research, and they're growing. We have some really fortunate strengths in the environment, solar power, electrical storage, carbon-capture technology, wind and water energy, policy issues—there are excellent, interesting things going on here at Carnegie Mellon."

Rick McCullough



The electric vehicle conversion being developed by Charge Car would employ a high-density capacitor—or "ultracapacitor"—to lengthen battery life by handling heavy current draws, such as when the car is accelerating or climbing a hill. Shown here is Gregg Podnar, a program manager in the Robotics Institute, who is leading the investigation with Illah Nourbakhsh, associate professor of robotics.

KEN ANDREYO PHOTO



JASON TOGYER PHOTO

David Andersen, an assistant professor of computer sciences, shows off one of his “wimpy nodes”—an older processor that can be reconfigured to run more efficiently in a large computer cluster than a newer processor. In preliminary tests, an array of these “wimpy nodes” sorted a gigabyte of data six times more efficiently than a conventional rack-mounted server.

work, and those speedy processors spend a lot of time waiting for memory, or worse yet, for conventional hard drives to access data, especially when they're engaged in jobs that require a lot of searching.

Yet a lower load on a processor doesn't directly equal lower power consumption. Even when a modern CPU is operating at 20 percent of its capacity, it's still drawing more than 50 percent of its peak power.

The architecture Andersen and his team are working on, called FAWN (for “a Fast Array of Wimpy Nodes”) uses less-powerful processors that are individually more efficient and reduces the amount of time those processors are idling. Instead of reading and writing data to a hard drive, FAWN caches part of the data in the memory of each “wimpy node,” where it can be retrieved in a fraction of the time and with a fraction of the energy that accessing a magnetic platter would require. FAWN further improves its energy efficiency by using solid-state Flash storage, which consumes about a tenth of the power of a hard drive and is up to 175 times faster at retrieving random blocks of data.

Tests have found that FAWN can sort a gigabyte of data six times more efficiently than a conventional rack-mounted server. The biggest FAWN cluster built so far had 25 nodes, and Andersen says the team is working on a 100-node cluster.

But getting existing software to run on FAWN isn't always easy. Many programs are tuned for speed, not for conserving memory, Andersen says. “Modern programmers expect a lot of memory to be available, and when you present them with a node that only has 256 megabytes of RAM, a lot

of things don't work,” he says. Researchers tried to run a popular virus-scanning program on a wimpy node, only to see it run out of memory when it got to its 1-millionth virus signature.

When the code was rewritten to allow it to run on a “wimpy,” it also performed two to three times faster on a conventional desktop machine, Andersen says, which was a happy discovery. But FAWN needs to be able to run existing software without requiring major modifications to the code, and that's the next hurdle for Andersen and his colleagues. “We need to be able to provide a black-box solution to current problems, and we're not there yet,” he says.

Any green technology that inconveniences developers or end-users will have a hard time finding acceptance, says Morris, currently on a sabbatical while working at Google in Mountain View, Calif. “Like many people, I've gotten deeply interested in climate change, but I've also done some philosophical inquiry,” he says. “I've asked myself, ‘What are people going to do to change their behavior to limit CO₂ emissions?’ And my conclusion was ‘not much.’”

Part of the problem, Morris says, is that it requires a sense of altruism that most of us can't manage. “The people who are going to benefit from me changing my behavior haven't been born yet,” he says.

It's better, perhaps, to avoid the problem by offering people green solutions that provide a clear, immediate benefit. One of Morris' professional and personal interests for more than 30 years has been encouraging people to carpool to work, but such efforts tend to stall, he says, because carpooling requires a high level of

personal commitment. Commuters don't want to be computer-matched with a stranger who's going to share their car forever, nor do they want to be tied to that person's schedule.

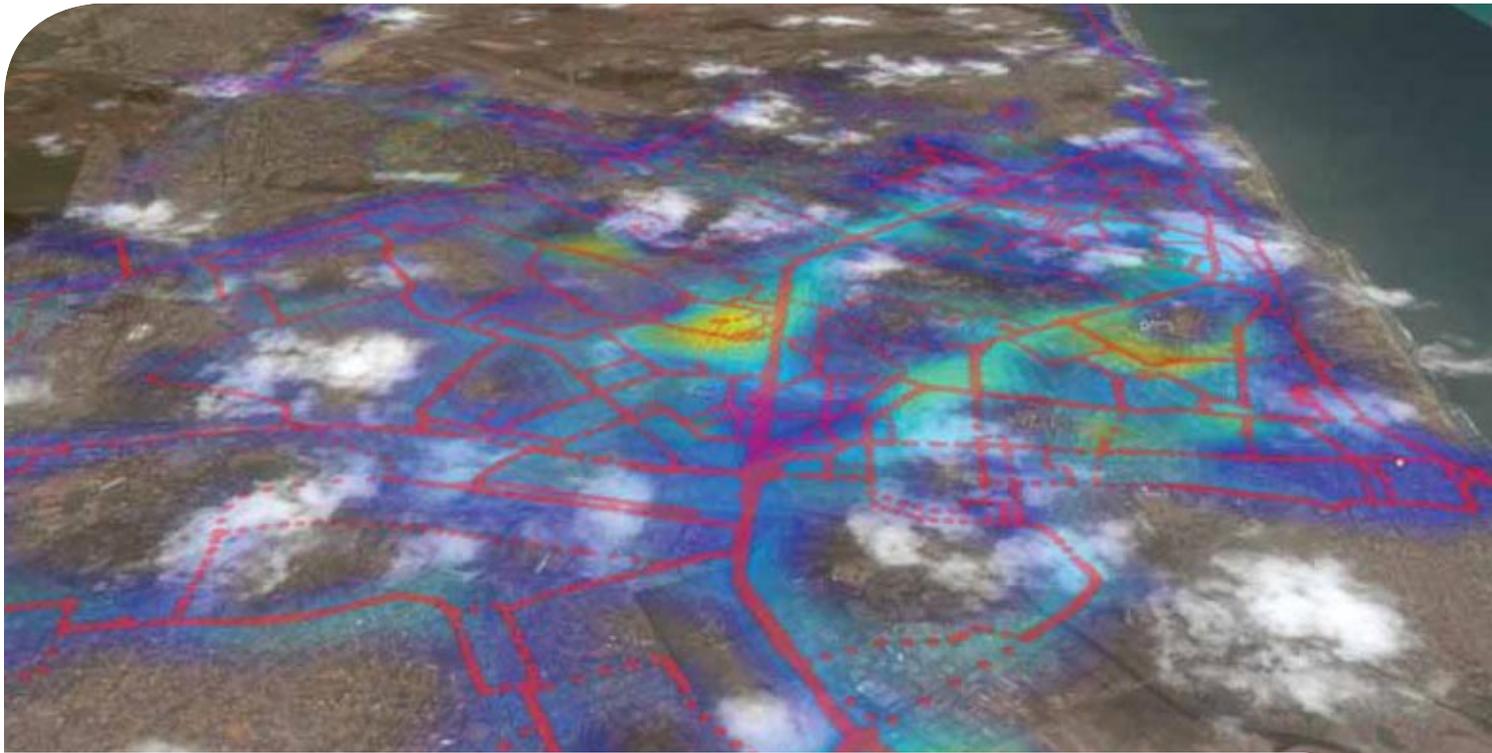
Morris has been working on one possible idea he calls “SafeRide.” Based on similar technologies to those that allow high school classmates to discover each other on Facebook and get directions from Google Maps, SafeRide would exploit email, cell phones and GPS to match passengers with drivers in real time. A driver might not be willing to share her car with a random stranger five days a week, Morris says. But she might be willing to share her car with someone once in a while if they personally reached out through a networking application.

With a web-capable cell phone, Morris says, SafeRide might even be able to match two people who are both at the airport and need to share a cab to a hotel. “We call it ‘high-tech hitchhiking,’” he says.

To Mankoff, those and other important applications demonstrate the different roles that computer scientists can play in reducing greenhouse gas emissions.

“There are all sorts of interesting, hard, technical problems that need to be solved,” she says. “The only way the country as a whole is going to address these problems is by getting everyone's help, along with the unique tools and knowledge that each of us brings to the table.”

Jason Togyer is managing editor of *The Link*. He wrote the cover story about the Gates and Hillman Centers for the Winter 2009-10 issue. Email him at jt3y@cs.cmu.edu.



The Rise of the Expert Amateur: Citizen Science and Neo-Volunteerism

If you cannot measure it, you cannot improve it. — Lord Kelvin

> By Eric Paulos

In the Living Environments Lab, we conduct research focused on societal problems. More specifically, we focus on the critical intersection of human life, our living planet and technology.

One direction of our work is shifting mobile phones from mere “communication tools” to “personal-super-computer-radio-stations-with-sensors” that can enable everyday people to become “citizen scientists” as they collect and share measurements of their environment. This new usage model for mobile phones will lead to important contributions in four research areas that:

- Improve the science literacy of everyday citizens through active participation in basic scientific data collection and use of scientific principles;
- Provide professional scientists and stakeholders with access to richer data sets for modeling, analyzing and advancing both fundamental and applied knowledge about people and ecosystems;

- Develop new usage models and user experiences for the mobile phone as a tool for promoting transparency and enabling grassroots participation in local community and civic government policy making; and
- Create a greater public awareness and understanding of the relationships between humans and the natural environment, particularly with regard to sustainability and environmental issues.

Through crowd-sourcing and new open-development platforms for mobile phones, we believe ordinary citizens around the world can become a central, driving element for collecting, reporting, interpreting and improving our health and our environment.

Challenges in citizen science

There are challenges that have to be overcome, of course. When collecting data, for instance, why

and how will people be motivated to participate? How and what type of data will be collected? When will samples be taken? How will problems of sensor accuracy, drift and calibration be addressed? What are reasonable types of sensors to use?

We must address ways the data is reported. How will we explain to our “citizen scientists” things like the range, accuracy and calibration of our sensors? Is the data better shared as raw numbers, or should it be interpreted and abstracted in a visualization such as a moving bar graph, meter or some other pictorial device?

How will the collected data be shared—on a personal mobile phone or in a public space? How will privacy be assured? How will our data be archived, preserved and authenticated, and what techniques will be used to insure that our data are valid?

Finally, what tools and techniques will facilitate change—productive debate and ultimate positive social benefit? How will people use the data >>>

to argue for and against various hypotheses?

We also have to study the best ways to invite and encourage active participation in both data collection and the use of data to develop real solutions to human problems.

Data collection at the grassroots

Ordinary citizens have long been agents of change, in grassroots movements ranging from neighborhood watch campaigns to political revolutions. Using citizens to collect scientific data is less well known, but there are precedents.

One successful, long-running citizen data-collection effort has been the National Audubon Society's Christmas Bird Count, which since 1900 has performed an annual census of birds in the Western Hemisphere.

Other movements in citizen data-collection indicate that many people have an intense interest in participating in scientific surveys. In the "Great World Wide Star Count," observers around the world are encouraged to go outside, look skywards after dark, and report the count of stars they see. In effect, they're measuring light pollution. Another is "Project BudBurst," where people submit time-stamped images of when flowers in their city bloom. This helps collect data on climate and pollen counts.

Our work builds on these kinds of existing citizen science projects but adds the ability to collect environmental data using sensor-equipped mobile phones.

Taxi Drivers in Ghana and the 'Prius Effect'

One of our early experiments was aimed at measuring air pollution from automobiles in Accra, the capital of Ghana. In 2007, we recruited taxi drivers and students to collect air quality data. The drivers and students were provided with sensors for the pollutants carbon monoxide, sulfur dioxide and nitrogen dioxide, along with a GPS unit (Figure 1). A set of visualizations from the sensor data as well as interviews resulted in a series of design recommendations and significant behavioral changes in our participants to improve air quality.

Perhaps the most surprising results from this study were not that air urban quality varies wildly over time and space, but that the participants recruited to collect this dataset began to experience what we referred to as "the Prius effect." Drivers of Toyota's Prius automobile get feedback on their real-time fuel mileage, and according to published reports, many have altered their driving behavior in an attempt to "score" 100 miles-per-gallon. Our participants in Ghana—exposed to real-time values for air quality—likewise developed strategies to minimize their exposure to poor air quality by altering their routes and times of travel across the city of Accra.

A Pocket-Sized Solution

After reviewing the results of our Ghana experiment, we integrated a series of air-quality sensors with a mobile phone (Figure 2). This system measured and reported carbon monoxide, nitrogen

dioxide, ozone, temperature and humidity and visualized the aggregate data on a public Web site.

While many research efforts in mobile sensing are focused on system-level challenges, such as integrating sensors and addressing hardware power issues, developing tools to share and visualize the data is a much larger challenge. Successful tools help persuade users to positively change their behavior while avoiding scare tactics.

A civic-citizen hybrid

Local governments and other agencies can play a role in data collection, and we obtained first-hand experience during a collaboration with the City of San Francisco to collect air quality readings using municipal street sweepers. We installed sensors to measure CO, NOx, ozone, temperature, relative humidity and speed on 12 street-sweepers. As the vehicles cleaned city streets, we were able to obtain real-time readings of environmental conditions across the city. (Figure 3)

That approach may seem to move away from using data collected by average citizens, but we think it's a realistic "hybrid" method where "citizen participation" includes citizens as well as civic agencies and public infrastructure. Such approaches are also great scaffolding techniques towards larger-scale public data collection.

Moving citizen science indoors

Citizen science can also be used to study indoor air quality. In one experiment, we designed a system integrated with Apple's iPhone to measure



Figure 1



Figure 2



Figure 3

and visualize particulate matter—the small particles of pollution found in air and a common problem in Pittsburgh. We also explored ways to use the visualizations to persuade people in a social network to modify behavior such as smoking.

Through a study of 14 households in Pittsburgh, we found that this system, called inAir, increased awareness of air quality, promoted behavioral changes that improved air quality, and demonstrated the persuasive power of sharing information and strengthening social bonds. (Figure 4) One of our participants, after viewing her air quality data compared with others, actually quit smoking! This is consistent with studies that have shown people are heavily influenced by the behaviors and actions they expect or know others to be performing.

Monitoring water conservation

Water is our most precious and most rapidly declining natural resource. More than 1 billion people do not have access to safe drinking water and more than 2.5 billion lack adequate sanitation. World health agencies estimate that more than 5,000 people—many of them children—die every day from water-related illnesses.

But having access to safe, plentiful water is influenced by a complex set of variables, including political, industrial and environmental constraints. And it's not just the developing world that faces water scarcity. Rapidly growing metropolitan areas in the western United States

also struggle to provide water to citizens with diminishing returns and steadily increasing costs.

We wanted to extend our environmental awareness research beyond air quality, so we applied similar techniques to monitoring water conservation using low-cost sensors (less than \$40) with persuasive displays.

Prototypes with two display styles—numeric and ambient—showing individual and average water consumption were installed in public and private spaces such as faucets and showers. The numeric display presents current water usage to the nearest tenth of a gallon while the visualization presents the same information in the form of a “traffic light.” (Figure 5)

Our work demonstrated that showing individual users the extent of their water usage helped promote conservation in real ways. It also inspired users to become curious about water usage in other contexts, such as dishwashing and lawn watering.

Areas for Future Research

“Citizen science” efforts have the potential to provide a vast and exciting new dataset for scientists and urban planners to use in improving public health and quality of life.

Involving non-experts in both data collection and community data analysis also has the potential to radically change public perceptions about urban planning and approaches to problem solving within neighborhoods and across cities.

Inviting public participation builds trust and encourages elected officials to make decisions that are transparent, and in the best public interest. It also allows policy makers to share their best practices with other communities.

Finally, successful citizen science projects could affect positive changes in society and help citizens of all ages participate in democracy and improve their understanding of the environment. The potential for grassroots efforts to emerge from such work and produce solutions to current local and global health and environmental issues is very real.

We hope our work will help motivate future research that empowers everyday people to learn, understand and improve their health and well-being, and broaden their awareness of their environment.

Eric Paulos is an assistant professor of human-computer interaction and joined the SCS faculty in 2008. Founder and director of CMU's Living Environments Lab, his research focuses on how technology intersects both human life and the living planet. He also serves as adjunct faculty in CMU's Entertainment Technology Center.

A native of California, Paulos earned his doctorate in electrical engineering and computer science in 2001 from the University of California at Berkeley. Prior to joining Carnegie Mellon, he served as senior research scientist at Intel Research in Berkeley, Calif., where he founded the Urban Atmospheres research group.

More information about citizen science and Paulos' other research is available at www.paulos.net.



Figure 4



Figure 5

“Citizen science” efforts have the potential to provide a vast and exciting new dataset for scientists and urban planners to use in improving public health and quality of life.

Stay Connected to SCS—In Your Own Backyard

It's spring again, and activity on campus seems to be reaching a fever pitch as we near commencement. And just as every spring is a season full of changes, there were changes this year to our Spring Carnival Weekend.

Beginning this year, all alumni reunions are being held during Spring Carnival Weekend. In addition, class reunions are being celebrated by decade. We hope these changes encourage more alumni to return to campus to reconnect with fellow alumni and the university.

There will still be special interest and school reunions held at various times throughout the year. In March, for instance, the Master of Software Engineering program marked its 20th anniversary. More than 120 alumni returned to campus to take part in reunion activities. In April, both the Robotics Institute and the Information Networking Institute celebrated anniversaries of their own.

On-campus reunions are just one way to reconnect with your fellow alumni and the school. We always encourage our alumni to return to campus to attend events, visit with faculty and staff, and see new spaces (such as the Gates and Hillman Centers). Still, given everyone's busy schedules, we realize that it can be challenging to try to travel back to Pittsburgh.

That's why we offer a number of ways for alumni to stay connected wherever they live or work. One example is the Carnegie Mellon Admission Council, an organization of 1000 alumni who volunteer with the Office of Admissions to help meet and coordinate activities for prospective students. Often, CMAC members help out by interviewing a prospective student in their area. This allows alumni to share their experiences and answer any questions the student may have about the university.

Another way to stay connected is by working with your regional Carnegie Mellon Alumni Chapter. The alumni chapters develop and organize regional events for alumni, students and parents. Some of the activities include picnics, sporting events, happy hours, faculty and entrepreneurship seminars, and community service projects. By serving as an alumni chapter volunteer, you help strengthen the alumni network and promote the university regionally. It also gives you an opportunity to make new professional and personal connections.

Carnegie Mellon has alumni chapters around the world, so visit www.cmu.edu/alumni and find out if there's one near you. And if you don't find one in your area, consider volunteering to help start one.

Serving as an alumni volunteer for the School of Computer Science is yet another way to stay connected. SCS alumni volunteers work on initiatives such as helping out with regional alumni events, talking with prospective graduate students, serving on our alumni advisory board and mentoring recent graduates. I'm always looking for volunteers to help out with our programs, and I appreciate your input and ideas about possible new activities.

The opportunities for alumni to stay connected are innumerable and we appreciate the chance to work and learn from our alumni. Alumni volunteers are invaluable resources who dedicate time, energy and expertise to many events, programs and ongoing initiatives. These activities would not be a success without the help of our alumni volunteers.

Remember, volunteering doesn't mean a big time commitment. Spending just an hour speaking with a prospective student or meeting a recent graduate for coffee are great ways that you can help, and both can make a tremendous impact on someone's life.



Tina M. Carr (HNZ'02)
SCS Director of Alumni Relations
tcarr@cs.cmu.edu

Still not sure you want to get involved? Why not just come to an event in your area to say hello to your fellow alumni? Planning for the annual SCS and ECE summer alumni gatherings is already underway. There are also many ongoing regional events for alumni from all of Carnegie Mellon's schools and colleges. Be sure to check both the SCS (cs.cmu.edu/alumni) and Carnegie Mellon alumni (cmu.edu/alumni) Web sites often.

And feel free to email me (tcarr@cs.cmu.edu) with your questions and comments. I look forward to hearing from you.

Cheers,

Tina M. Carr (HNZ'02)
Director of Alumni Relations
School of Computer Science

Bell to Address Grads at Diploma Ceremony

Entrepreneur and inventor C. Gordon Bell will address new graduates during the School of Computer Science's diploma ceremony at Carnegie Music Hall in Oakland at 1 p.m. May 16. In addition, during the university's commencement ceremony that morning, Gordon will receive an honorary doctorate in science and technology from Carnegie Mellon.

Gordon is a former member of the CS faculty who served as vice president of engineering at Digital Equipment Corp., overseeing the development of VAX. He later became founding assistant director of the Computing and Information Science and Engineering Directorate at the National Science Foundation.

Since 1995, Gordon has been with Microsoft,

most recently as principal researcher at Microsoft's Silicon Valley Research Group. He's also a founder of the Computer History Museum.

The diploma ceremony is open to degree recipients and their family and friends.

For details, see www.cs.cmu.edu/~commencement and www.cmu.edu/commencement.

Alumni Snapshots

Ben Ahroni

B.S., Computer Science, Carnegie Mellon University, 2004

What makes a good video game? “The short answer is, it’s fun,” says Ben Ahroni (CS’04), a producer at Seattle’s PopCap Games. The longer answer is that it depends on the audience and the player’s mood.

“Some games have a cooperative or competitive element, like the ones on Facebook,” Ahroni says. “There are other games that allow you to just sit on your couch and unwind. Some people want to escape for a while, but others want what’s almost an academic exploration—they want a puzzle to solve.”

Ahroni knows something about what makes a good video game. He’s been playing them since he was 3 or 4 years old, when his parents bought him the original 8-bit Nintendo Entertainment System. In middle school, he began taking programming classes and by high school was writing text-based games for fellow students to play on

their graphing calculators.

Even with the limited possibilities offered by a calculator’s hardware, Ahroni knew right away that he wanted to pursue a career in game development.

At Carnegie Mellon, he says, perhaps the most important courses he took involved writing software, designing systems and working with a team. That’s served Ahroni well at PopCap, where he’s had to understand how long it takes to develop a program, how inter-related software interacts, and what kinds of team members he might need to work with.

For anyone who wants to go into the video game industry, Ahroni has one piece of advice: “Learn everything you can. Don’t say, ‘I want to be a developer, so I don’t need to know what an artist does.’ Being a producer means being a generalist.



You have to understand everything—from how software works, to how what an artist does appears in your product, to understanding a little bit about the business of games.”

When he’s not working on games (or playing them) he’s liable to be out snowboarding. He also plays “quite a bit” of poker around Seattle, and lately, Ahroni has begun taking kung fu classes. “It’s both a mental and physical challenge,” he says. “The progression of learning new moves and steps and forms, and packing more elements into your repertoire is in some sense a video game progression. Maybe that’s why I found myself drawn to martial arts.”

—Jason Togyer (HS’96)

Eduardo G. Frias

B.S., Systems Analysis, Universidad de Buenos Aires

M.S.E., Software Engineering, Carnegie Mellon University, 1994

And no, it couldn’t be a “Band-Aid solution”—unless you mean a solution that reflected a 144-year-old Johnson & Johnson corporate culture that puts a very public emphasis on corporate responsibility to consumers, suppliers and employees.

Frias and his team realized they would have to work within the company’s culture, not against it. “We had to make it attractive to people, and we had to win their trust,” he says. “It was not something we could do just by changing the organizational chart. From Day One we had to have the processes and tools that would not only help us, but would also reassure people.”

Though some computer scientists were recruited from outside of Johnson & Johnson, he also was able to pull together many IT and CS professionals from the corporation’s worldwide operations. They understood its culture and the many challenges it faced, he says. “Because we needed a running start, we tried to use our existing

talent as much as possible,” Frias says.

Today, though spread across many continents, IT Global Services at Johnson & Johnson has a unified sense of purpose and can take advantage of its new sense of scale. It can harness the company’s worldwide buying power, Frias says, and also has a pool of people who can work together to apply their many different skills to problems both large and small. “We’ve created a software factory that can deliver services much faster—metrics, tools, applications—that enable us to do things better,” he says.

Lessons Frias learned at CMU’s master of software engineering program helped provide the confidence he needed to carry on, even as the Global Services organization department was doubling and tripling in size. “What I learned about process discipline, project management, architecture—I could go on and on and draw lessons from each one of my classes,” Frias says.

—JT



With everything under its corporate umbrella from Band-Aid bandages and Tylenol pain reliever to cosmetics, cancer therapies and surgical appliances, Johnson & Johnson has some complicated computing needs. Yet until recently, its 250 worldwide operating units and their own IT departments weren’t always sharing their needs or their know-how, and it wasn’t efficient.

Three years ago, Eduardo Frias (CS’94) was named vice president for information technology and together with his team was told to bring together those separate IT components to form one corporate-wide application services clearinghouse.

Giving Back: Harry Shum

> By Mark Dorgan

What do the following have in common?

- Directing research for Microsoft in Asia
- Leading development of search engine technology
- Rooting for the Pittsburgh Steelers

Well, they've all got Harry Shum, of course.

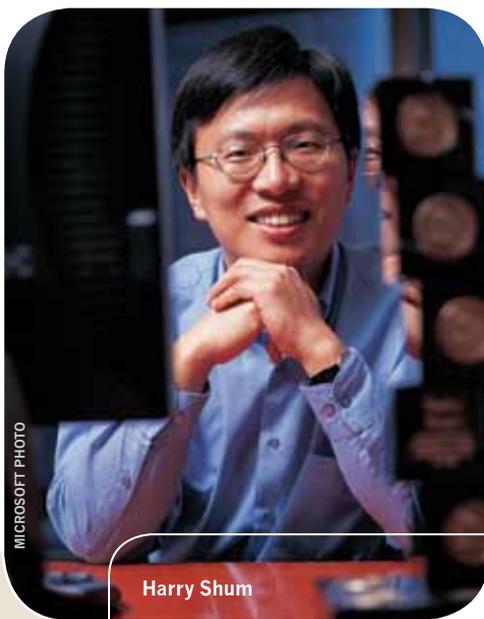
Harry (CS'96) has been with Microsoft since 1996 and currently serves as corporate vice president for core search development. Before that, he was managing director of Microsoft Research Asia. In his five years in the Ph.D. program at the Robotics Institute, Shum says he “learned to do research right,” and gained both knowledge and perspective from his professors and fellow graduate students.

As for rooting for the Steelers, Shum says that helped him get through the cold winters (or at least through football season anyway).

Harry says the strong sense of community at SCS is a primary reason he stays connected both philanthropically and professionally. In February, for instance, he made a recruiting visit to our Pittsburgh campus and delivered a “spotlight” lecture on Microsoft’s Bing search engine in the Rashid Auditorium.

According to Harry, he owes much of his career success to people at Carnegie Mellon, including his fellow students and the faculty. Harry fondly recalls hearing Herb Simon’s speech at the inauguration of Raj Reddy as dean of SCS; Allen Newell’s last lecture and his mathematical approach to artificial intelligence; Takeo Kanade’s use of Japanese cartoons to tell stories and inspire research papers; and Matt Mason’s glowing endorsement of the *International Journal of Robotics Research*. (Harry claims that Matt said the journal was “so good that I keep a copy in my bathroom.”)

Harry recently acknowledged the impact of one of those professors—Raj Reddy—by making a lead gift to name a conference room in the new Hillman Center for Future-Generation Technologies. He encouraged 10 other alumni to contribute to the Raj Reddy Conference Room and joined several of those donors for the official dedication in September.



Harry Shum

Harry credits Reddy with influencing him to both “think big” and to “think practically” when planning and evaluating research. He says Reddy’s students in particular “owe a great deal of gratitude and appreciation” for his influence on their lives and careers.

Harry’s gift was made as part of Carnegie Mellon’s “Inspire Innovation” campaign, which recently topped the \$628 million mark. To find how you can help the School of Computer Science through scholarships, fellowships, faculty support or gifts toward the Gates and Hillman Centers, please contact me at mdorgan@cmu.edu or call me at 412-268-8576. You can also learn more about the Inspire Innovation campaign by visiting www.cmu.edu/campaign.

Mark Dorgan is executive director of major gifts and development liaison for the School of Computer Science.

Discover, NSF Bring Robotics Panel to Pittsburgh



Shown from left are Murphy, Movellan, Powell, Brooks and Whittaker. To view the discussion, visit Discover Magazine’s website at www.discovermagazine.com/ events and click the link, “Grand Challenges of Science: Robotics.”

Discover Magazine and the National Science Foundation brought their “Grand Challenges” series to the Gates and Hillman Centers on Jan. 28 for a panel discussion among four eminent roboticists.

The forum, moderated by Discover Editor-in-Chief Corey Powell, asked how robots will transform industry, health care and warfare, and if they ever will equal humans.

Panelists included William “Red” Whittaker (E’75, ’79), the university’s Fredkin Professor

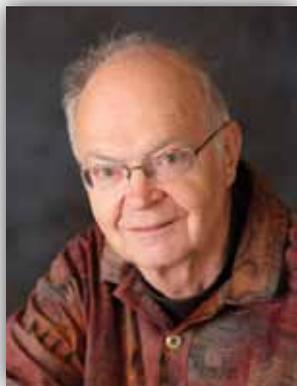
of Robotics; Javier Movellan, associate project professor at the Institute for Neural Computation at the University of California at San Diego; Rodney Brooks, director of the Computer Science and Artificial Intelligence Laboratory at the Massachusetts Institute of Technology; and Robin Murphy, Raytheon Professor of Computer Science and Engineering at Texas A&M University.

The discussion was held in the Rashid Auditorium of the Hillman Center for Future-Generation Technologies.

Stanford's Knuth, Cornell's Kleinberg honored with Katayanagi prizes

Stanford's Don Knuth and Cornell's Jon Kleinberg are the latest recipients of the Katayanagi Prizes in Computer Science.

Jointly awarded by Carnegie Mellon and the Tokyo University of Technology, the prizes were established in 2007 with a gift from entrepreneur and educator Koh Katayanagi. Katayanagi founded the predecessors of Tokyo University of Technology and following World War II helped reconstruct Japan's



Don Knuth

electronics industry, largely through engineering training and education.

Knuth, a 1974 Turing Award winner, is the creator of the TeX typesetting system and author of "The Art of Computer Programming," with three volumes in print and the fourth debuting this year. He has been called the father of algorithm analysis and presently serves as an emeritus professor at Stanford.



Jon Kleinberg

Kleinberg is the Tisch University Professor of Computer Science at Cornell. His

research focuses on the social and information networks that underpin the Web and other online media, and has helped shape the design of peer-to-peer file-sharing networks.

Knuth was recognized with the \$10,000 Research Excellence prize, while Kleinberg received the \$5,000 Emerging Leadership prize. Kleinberg delivered his Katayanagi lecture in the Rashid Auditorium on Jan. 21, while Knuth visits Carnegie Mellon to deliver his lecture and receive his award on May 5.



HCII director Cassell starts work Aug. 1

The new director of the Human-Computer Interaction Institute has a strong interest in understanding how people use both verbal and non-verbal forms of communication—and what that might mean for the next generation of intelligent tutoring systems.

Justine Cassell joins the Carnegie Mellon faculty on Aug. 1. Currently director of the Center for Technology & Social Behavior at Northwestern University in Chicago, she led development of the Embodied Conversational Agent, a virtual human capable of interacting with humans using spoken and written words as well as gestures.

Cassell's recent research includes investigating whether ECAs have a role to play as "virtual peers" for children using tutoring systems to learn language and social skills as well as science and math.

"We believe that she will expand the horizons of the institute while helping it continue as one of the world's outstanding centers for investigating how computers can better serve individuals and society," says Randy Bryant, SCS dean, who announced Cassell's appointment April 5.

She succeeds Dan Siewiorek, HCII director since 1998.

Cassell holds master's degrees from France's Université de Besançon and Scotland's University of Edinburgh, and a double Ph.D. in psychology and linguistics from the University of Chicago. Prior to joining Northwestern, she directed the Gesture and Narrative Language Research Group at Massachusetts Institute of Technology, where she won the Edgerton Faculty Achievement Award.

In 2006, Cassell was awarded the AT&T Research Chair at Northwestern, and in 2008 received the Women of Vision Award for Leadership from the Anita Borg Institute.

Siewiorek is resuming his duties as the university's Buhl Professor of Electrical and Computer Engineering and Computer Science; as a project leader for human system interaction within the Quality of Life Technology Center; and as leader for smart work technologies within the Heinz College's Center for the Future of Work.

Guinness record goes to RI-developed robot



Dragon Runner, a 20-pound "throwable" robot developed at the Robotics Institute, is now in the Guinness Book of World Records. The editors of the 2010 edition say it's "the world's most durable military robot."

Hagen Schempf, principal systems scientist at the Robotics Institute, led the development of Dragon Runner. The first 12 prototypes were delivered to the U.S. Marines in 2004.

Dragon Runner was commercialized by QinetiQ and can be used to make searches in areas too dangerous or small to send a human, such as underneath booby-trapped vehicles, or inside mineshafts, sewers and collapsed buildings.

It's just a little bit larger than a Pittsburgh phone book, yet can lift five to 10 pounds and be equipped with a wide variety of grippers, cameras and sensors. Hundreds of the rugged little robots are now in use by the U.S. military and its allies.

Ph.D. students cop Microsoft, IBM fellowships

SCS doctoral students claimed two of the 10 Microsoft Fellowships awarded for 2010. This year's Carnegie Mellon recipients were Chris Harrison, a Ph.D. student in the Human-Computer Interaction Institute, and Dafna Shahaf, a Ph.D. student in the Computer Science Department.

Meanwhile, another SCS doctoral student—Aruna Balakrishnan—is the newest recipient of the IBM Fran Allen Fellowship.

The highly competitive Microsoft Fellowships are awarded each year to doctoral students who are “advancing computing research in novel directions.” This year, 176 students nationwide were nominated, a Microsoft spokesman said.

The fellowships cover 100 percent of tuition and fees for two years and provide \$28,000 in living expenses and \$4,000 in travel expenses for professional conferences and seminars.

Shahaf is a graduate of Tel Aviv University and the University of Illinois at Urbana-Champaign whose research interests include artificial intelligence, learning theory, algorithms and handling uncertainty. Harrison is a graduate of

New York's Courant Institute of Mathematical Sciences whose research focuses on interface design and input technologies (see “Screenshot,” inside back cover).



Aruna Balakrishnan

The Fran Allen Fellowship was established in 2007 and is named for 2006 Turing Award winner Frances Allen, the first woman to be named an IBM Fellow—the computer maker's most prestigious honor for its computer scientists.

According to IBM, Fran Allen Fellowships are awarded to female Ph.D. research students based on their outstanding technical accomplishments and commitment to mentoring and community building.

Balakrishnan, a graduate of Harvard University, is a Ph.D. student in the HCII who studies different kinds of visualization techniques and their impact on the way that information is shared and understood.

She receives a \$30,000 grant from IBM and the assistance of a “career mentor” from IBM Research, and will be invited to present her work at an IBM campus and to collaborate with the company's researchers.

Three from CMU elected to NAE

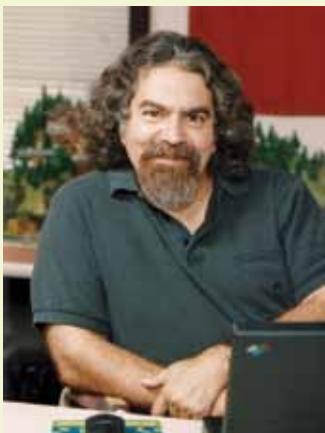
Tom Mitchell and Paul Nielsen are among the 68 new members and nine foreign associates elected to the National Academy of Engineering—one of the most prestigious honors accorded to engineers in the United States.

The elections of Mitchell, head of the Machine Learning Department and University Professor of Computer Science and Machine Learning, and Nielsen, director and CEO of the Software Engineering Institute, were announced Feb. 17.

Also elected from CMU was Jacobo Bielak, University Professor of Civil and Environmental Engineering.

Established in 1964, the NAE is a private, non-profit advisory board and a peer institution of the National Academy of Sciences, Institute of Medicine and National Research Council. Thirty-six members of the CMU faculty have now been honored by the NAE.

‘Tornado’ Marinelli pens memoir of ‘comet’ Pausch



The “tornado” tells all in a new book about the creation of the university's Entertainment Technology Center and his collaboration with the late Randy Pausch (CS'88).

The Comet and The Tornado is the first book by Don Marinelli, Carnegie Mellon professor of drama and arts management and executive producer of the ETC.

Published by Sterling, the book is subtitled “Reflections on the

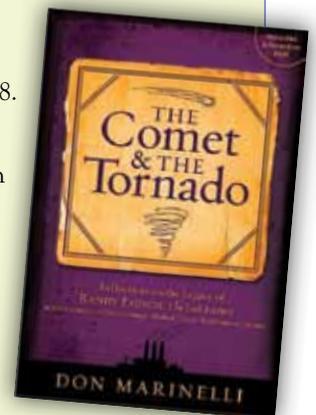
Legacy of Randy Pausch, The Last Lecture and the Creation of Our Carnegie Mellon Dream Fulfillment Factory.” The book centers around the six years Marinelli spent sharing an office with Pausch, professor of computer science, human-computer interaction and

design, who died of pancreatic cancer in July 2008.

Marinelli and Pausch were colleagues and co-founders of the ETC. In his Sept. 18, 2007 speech that came to be known as “the Last Lecture,” Pausch referred to Marinelli as the “Tornado”—a whirlwind of energy and creativity. The two had a yin-yang relationship, Pausch said, with Pausch serving as the logical, left-brain half and Marinelli as the creative right-brain thinker.

Marinelli calls Pausch a “comet” who “left millions of people stunned, amazed, happy, giddy and seeing light where there had been only darkness.”

The book includes a “Synthetic Interview” with Marinelli. Created by ETC and SCS researchers, Synthetic Interviews allow individuals to have a “conversation” with a character or persona as if that person were present in real-time.



SCREENSHOT

PHOTOS COURTESY CHRIS HARRISON



All You Really Need is Skin

Allan Sherman, the 1960s comedian best remembered for his Grammy-award winning “Hello Muddah, Hello Faddah (A Letter from Camp),” wrote many other parody songs, including a version of “You’ve Gotta Have Heart” called “You’ve Gotta Have Skin”:

It covers your nose,
And it’s wrapped around your toes.
And inside it you put lemon meringue,
And outside you hang your clothes.

But there’s another use for your skin that Sherman could have never envisioned—using it to dial a phone or send an email.

Chris Harrison, a third-year Ph.D. student in the Human-Computer Interaction Institute, is working on an application that makes it possible. Along with scientists Desney Tan (CS’04) and Dan Morris of Microsoft Research, Harrison has developed “Skinput,” which can turn your fingers, your forearms or any other part of your body into a giant touchscreen.

According to Harrison, Skinput requires a combination of some fairly simple sensors and some fairly complicated machine-learning algorithms. Users wear an armband containing a tiny video projector that creates the keypad image. When the “keys” are tapped, acoustic sensors (essentially, very sensitive microphones) pick up the different vibrations made in bones or soft tissues.

“There’s nothing super sophisticated about the sensor itself,” Harrison says, “but it does require some unusual processing.”

Skinput’s machine learning algorithms—which improve over time—analyze the audio captured by its sensors, examining some 186 different features to determine what area of the skin the user is tapping. In a test involving 20 people, the system was able to detect the correct location of “taps” with 88 percent accuracy overall. Skinput’s accuracy nears 96 percent when the sensors are located below the elbow, while finger “flicks” can be identified with 97 percent accuracy.

“We strap iPods and other devices onto our arms,” Harrison says. “Why not utilize all the external surface area that’s already with us? What’s great about skin, unlike tables, is that it travels with us.”

Instead of an iPod or mobile phone, imagine wirelessly controlling any application with the aid of a sensor pack no larger than a wrist-watch—and being able to move seamlessly from task to task with just a tap on your arms or a flick of your fingers.

The BBC and other international news outlets have already trumpeted the possibilities of Skinput. On April 12, Harrison presented a research paper on Skinput in Atlanta at the Association for Computer Machinery’s annual Conference on Human Factors in Computing Systems.

“Our skin is always with us, and makes the ultimate interactive touch surface,” Harrison says.

Just like Allan Sherman sang: “Ain’t you glad you’ve got skin?”

To see more, visit Harrison’s Web site at www.chrisharrison.net.

Office of the Dean
5000 Forbes Avenue
Pittsburgh, PA 15213

Then and Now

Though he wasn't a member of the computer science faculty, Carnegie Mellon Professor Clarence Zener had a profound effect on the design of digital computers and the algorithms they employ.

In 1934, while at Princeton, Zener discovered that certain electrical insulating materials—semiconductors—were able to regulate and limit voltage. That led to the invention 20 years later of Zener diodes, tiny solid-state voltage regulators that made reliable desktop computers a reality.

A physicist by training, Zener worked with Enrico Fermi and Edward Teller at the University of Chicago and taught there from 1945 to 1951, before becoming director of research for Pittsburgh's Westinghouse Electric Corp.

At Westinghouse, Zener became frustrated by engineers who tried to optimize their designs by setting arbitrary parameters and then using early digital computers to crunch numbers until they arrived at a solution. His more elegant approach—geometric programming—called for engineering problems to use adjustable

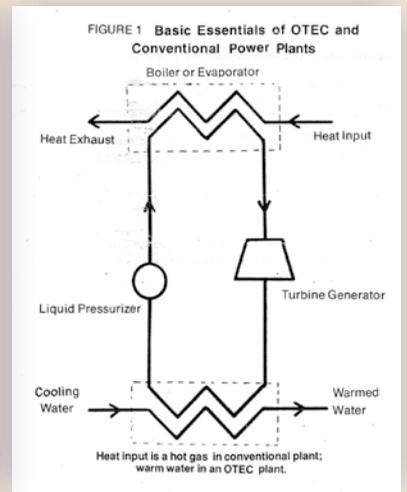


PHOTO COURTESY CUSHING MEMORIAL LIBRARY AND ARCHIVES, TEXAS A&M UNIVERSITY. OTHER MATERIAL COURTESY CARNEGIE MELLON UNIVERSITY ARCHIVES.

parameters that were defined by mathematical functions. Geometric programming remains a staple of optimization theory today.

In 1968, Zener was recruited to Carnegie Mellon, where he remained a University Professor until his death in 1993. During the energy crisis of the 1970s, Zener became intrigued with harnessing the thermal currents of the oceans as a clean, inexpensive source of electrical power. Using geometric programming, Zener created computer models to optimize the necessary heat exchangers and the locations where they could be deployed.

To Zener's disappointment, cheap oil and changing political priorities in the 1980s dried

up much of the research into ocean thermal energy conversion.

Today, concerns about manmade climate change are driving new research into “green” energy production and conservation at Carnegie Mellon’s School of Computer Science. And many of the optimization models being used carry echoes of the work of Clarence Zener.

There’s more about green technology at SCS in this issue, beginning on page 12. And there’s more about the life and work of Clarence Zener at our Web site, link.cs.cmu.edu.

—Jason Togyer