conduit!

Volume 11, Number 2 | Department of Computer Science, Brown University | Fall, 2002

NEXT-GENERATION STREAM-BASED APPLICATIONS

Developing under Debian Linux, XEmacs, CVS, Java, and C++: *free*

Gas for a round-trip, two-hour commute from New Hampshire to Rhode Island: *$15*

Funding four professors, 14 graduate students, and four undergrads for six months: *$250,000*

Getting a data stream processing system to process 100 gigabytes in six months: ***priceless!***

Stream-based Applications

While rip-offs of MasterCard™ commercials are passé, continuous data-stream management is relatively new. Applications that deal with potentially unbounded, continuous streams of data are becoming increasingly popular due to a confluence of advances in real-time, wide-area data-dissemination technologies and the emergence of small-scale computing devices (such as GPSs and micro-sensors) that continually emit data obtained from the physical environment.

Consider a data-stream application that involves all MasterCard™ transactions in the United States. As transactions take place, they would be fed to a monitoring system as a continuous stream. The system would then run queries over this stream, in real time, in order to detect fraud or other unusual purchase patterns and take appropriate actions (such as temporarily canceling the credit card or placing an automatic phone call to the card's owner notifying him or her of the unusual activity). A similar financial application involves monitoring stock transactions: a company might be interested in tracking certain stock price trends and want to react by dynamically re-evaluating portfolios and reallocating funds in real time.

Another application domain is health care. In a hospital, the vital signals coming from critical patients would be very useful data streams. By processing such streams in real time, emergency situations can be detected and appropriate targets can be alerted in a timely fashion. Other sources of data streams include GPSs that track the positions of mobile entities, smart-spaces that monitor the status and location of objects (such as books in a library), and plant-monitoring systems that report environmental conditions (e.g., inside a nuclear power plant).

L to r, top to bottom: Adam Singer, Alex Rasin, Matt Hatoun, Anurag Maskey, Eddie Galvez, Jeong-Hyon Hwan, and Ying Xing; Christina Erwin, Christian Convey, Michael Stonebraker, Robin Yan, Stan Zdonik, Don Carney and Nesime Tatbul (not present: Ugur Cetintemel, Mitch Cherniak, Daniel Abadi)

All these applications require timely processing of large volumes of continuous, potentially rapid and asynchronous data streams. Moreover, in many cases, the "live" data stream must be compared or combined with stored "historical" data—for example, in order to detect an unusual situation, one has to know what the usual case is.

Databases Upside Down

Existing data-management solutions cannot meet these challenges. Traditional data-management architectures are de-

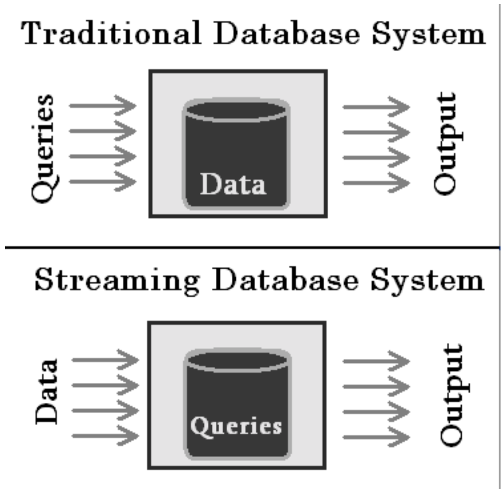

signed and optimized for processing large amounts of stored data and temporary, one-time queries initiated by humans. Because of the continuous-streaming nature of the data and the longevity of queries in stream-based applications, this classical database paradigm must shift dramatically—instead, incoming streams of data must be processed using persistent queries that produce results continually. Figure 1 (above) illustrates the high-level differences between the two paradigms. Traditional databases essentially store the data and process incoming one-time queries, whereas a stream-processing system must process incoming data using stored, persistent queries. In other words, the database architecture is essentially inverted. Addressing this fundamental difference requires a critical rethinking of many existing data-management and processing issues, and the development of new proactive data-processing concepts and techniques.

The Aurora Project

Aware of the strong emergence of this new class of applications and the fundamental limitations of the existing approaches, Brown University professor Stan Zdonik and MIT professor Michael Stonebraker initiated the Aurora research project during the summer of 2001. The goal of Aurora is to design and implement a general-purpose architecture for efficiently and effectively processing data streams. The project quickly gained significant mass and momentum, and is currently being undertaken by a group of academics from Brandeis University, Brown University, and MIT. As of October 2002, the active members of the group include four professors (Ugur Cetintemel and Stan Zdonik at Brown, Mitch Cherniak at Brandeis and Michael Stonebraker at MIT), eight PhD students, four Master's students, and several undergraduates. Even though the group is broadly distributed geographically and academically (we range from undergrads to full professors), the group has been able to develop a functional Aurora system prototype over the past six months.

Aurora from 100,000 Feet

In Aurora, the data streams are assumed to come from a variety of data sources such as computer programs or hardware sensor units, and to enter the system through a well-specified input interface (such as a socket). The basic job of Aurora is to process the incoming streams in the way defined by a network of operators:

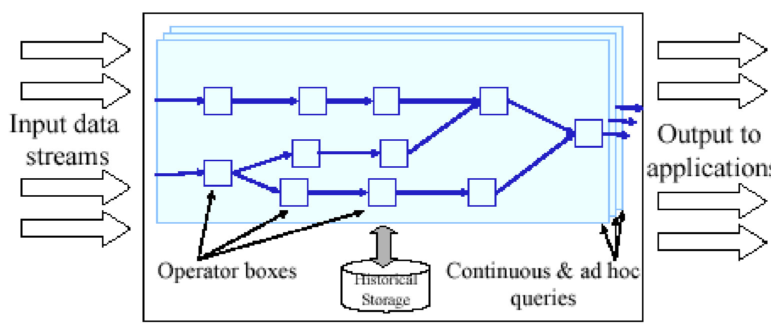

the data flow through a loop-free directed graph of operators, *boxes* in Aurora terminology, are processed along the way, and are ultimately presented to client applications. Besides providing support for processing data streams, Aurora also incorporates mechanisms for managing and processing stored data. Figure 2 (above) shows an overview of the basic system model in which multiple concurrent continuous queries process incoming data streams, potentially also accessing stored historical data. Each box has an input queue and an output queue, and the execution of a box entails removing tuples from its input queue, processing them, producing result tuples, and then writing them to its output queue.

Features

Aurora query networks are built using a well-defined set of operators. Some Aurora operators are essentially borrowed from relational algebra when they translate well into stream processing (e.g., the "select" or "union" operators). Others are novel operators specifically designed to meet the unique needs of data streams. For example, traditional database operators assume that their operands are finite data sets. This assumption clearly does not hold for data streams, which are potentially unbounded. In order to address this inherent difference, many Aurora operators define and exploit the concept of "windows" that effectively chop the infinite data sets into manageable, finite data sets over which the operators can be applied in finite time. Windows are typically defined over temporal attributes of the data (e.g., a window of all tuples that are generated within the last five minutes), but this is not a requirement. In addition, windows can be disjoint or can overlap. After finishing processing on a particular window, each windowed operator then starts processing the next window of tuples.

Another difference between stored data and streaming data is that while the former is always "there" and accessible, parts of streaming data may arrive late or, even worse, may be lost and never arrive. Because timely results are crucial, Aurora operators cannot afford to wait for all pieces of data to arrive, since in the worst case this might mean blocking indefinitely. Therefore, the operators incorporate simple but effective "timeout" mechanisms to deal with missing or stale data.

Every client application is associated with a query that defines its stream-processing requirements and a Quality-of-Service (QoS) specification that defines its performance expectations. Aurora's operational goal is to maximize the QoS delivered to its client applications. As a result, all resource allocation decisions (such as CPU scheduling and buffer management) in Aurora are driven by QoS specifications.

A key goal of Aurora is to attain extreme scalability and acceptable performance even in scenarios in which the incoming data stream rates and cost of the existing queries exceed available resources. Under unexpected "overload" situations, Aurora gracefully and systematically degrades the quality of its results. It achieves these "approximate" results using a process called "load shedding", which involves strategically eliminating some of the data from processing in order to gain the extra resources necessary to produce perfect results. Aurora also quantifies the degree of approximation it makes in order to help client applications reason about the quality of answers they receive.

"A key goal of Aurora is to attain extreme scalability and acceptable performance even in scenarios in which the incoming data stream rates and cost of the existing queries exceed available resources"

Finally, Aurora is being designed as a self-adaptive system. It incorporates introspection techniques that continually monitor system conditions, including the load and available resources, and uses this information to optimize its execution by retuning system parameters and reallocating resources dynamically.

Current Implementation and Future Plans

As of September 2002, the Aurora team has a working Aurora prototype that implements the operators and QoS-driven resource-allocation mechanisms described above. The prototype has a Java-based GUI that allows the construction and execution of arbitrary Aurora networks, specification of QoS requirements, stream-type inferencing, and zooming.

In Figure 3 (over), which is a snapshot of the Aurora GUI, a stream of input tuples appears at the two left input boxes. As the tuples arrive at the system, they are continually processed and pushed towards the single output box on the right. In this particular example, streams of credit card purchases and card-holder locations are processed in order to detect potentially fraudulent activity. On the top half, usually inactive credit cards (measured over one month) that suddenly experience excessive amounts of activity over one day are passed through a join operator. On the bottom half, the credit cards that experience activity over varied locations within a time window of five minutes are also passed through the bot-

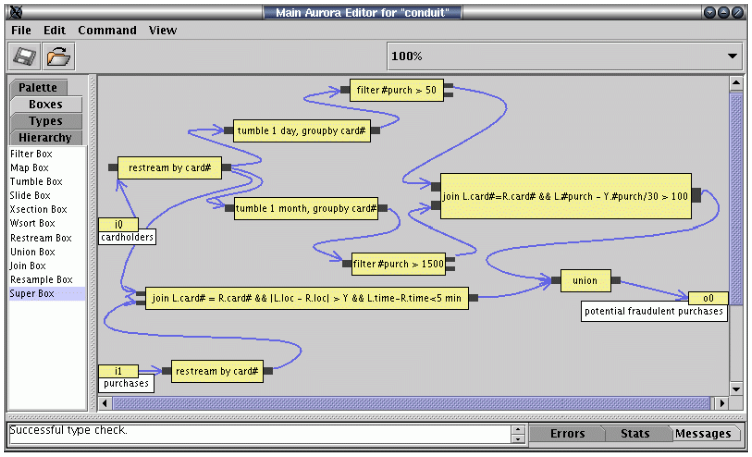

tom join operator. These two streams are then passed through a union operator and then passed through to an output application that expects tuples containing potentially fraudulent credit card activity.

In the GUI, a query can be created by dragging boxes from the left-side palette of various stream operators. Arcs can then be drawn by clicking and dragging from the output port of the source operator box to the input port of the destination operator box. The interface further lets the user recursively encapsulate multiple boxes in groups called a "superbox", effectively implementing a macro utility. As a result, predefined operator boxes or superboxes may be saved and copied from a palette, allowing cut-and-paste functionality. Last, the interface can type-check the types of tuples streaming into operator boxes and type-infer the tuples output by the operator boxes.

Future development plans include implementation of ad-hoc queries and user-defined functions to let a network administrator define scripts for processing special aggregates and other functions. In addition, the team is actively working on dynamic optimization and load-shedding techniques.

A distributed peer-to-peer version of Aurora named Aurora* is also in the works. Aurora* is being designed to achieve higher availability and scalability by distributing data processing across multiple distributed Aurora servers, each running a copy of the full single-node Aurora system.

Further information on the Aurora project can be found at:

http://www.cs.brown.edu/research/aurora.

Students from the Rhode Island School of Design (RISD) and Brown team up in CS237, a course on interdisciplinary scientific visualization

CS237 is a new course co-taught by David Laidlaw, Brown CS assistant professor and Fritz Drury, professor of illustration at RISD. Daniel Keefe, a CS graduate student, is TA for the effort. The course explores the design process involved in creating scientific applications for immersive virtual reality. Brown's virtual reality Cave is the medium for which students develop designs that address a range of scientific problems. Frequent group crits in class meetings (such as the one shown) give students the opportunity to evaluate and iterate proposed designs in a collaborative environment.

VAN DAM GIVES CONVOCATION ADDRESS

Professor van Dam gave this address at Convocation on September 4, 2002.

John Abromowski, Photographer, Public Affairs and University Relations

L to r, President Ruth Simmons, Provost Robert Zimmer, Andy van Dam, Chancellor Emeritus Artemis Joukowsky and Chancellor Stephen Robert

First, let me congratulate the class of 2006, as well as the incoming graduate and medical students, on launching the next phase of your lives at this world-class university. Besides being excited, many of you may be just a little apprehensive about the experience you're about to have here.

Some of you may wonder what you're doing here and even think that maybe you were accepted because someone in admissions missed the dart board altogether. Not to worry: our admission office never makes mistakes; you belong here! But even if you always knew you belonged at Brown, you may be feeling some apprehension about embarking on this brand-new part of your life. Let me reassure you that any such feeling is an absolutely normal response to such a big change.

Which brings me to my main theme—change. I will give you sneak previews of some changes you are likely to experience in the technology-driven world in which you will live, and some suggestions for how to cope with both those challenges and the more immediate ones that you face as students here at Brown. These will be drawn partially from my own personal experiences with change. I'm inspired to talk about change in part because of the great changes that I myself have lived through, especially as a kid. Before I was 14 I had lived in three different countries, experienced a world war, and become immersed in a new language.

I was born in the Netherlands in 1938. When I was less than a year old, my father moved us to Indonesia, then the Dutch East Indies, to head up a marine biology research institute. We were trapped there when WWII broke out in the Pacific and were interned for nearly four years in a Japanese concentration camp. After liberation in 1945, we went back to the Netherlands, and then when I was 13 we emigrated to the U.S.; that was in 1952, exactly 50 years ago next week. I can still remember seeing the Statue of Liberty as our ship approached New York City and feeling both exhilarated and apprehensive, knowing I would have to learn a new language and culture, deal with a new school, and make a new set of friends. I went to high school and then to Swarthmore College in the 1950s, in the relative calm of the leave-it-to-Beaver era of post-war prosperity. After graduate school at the University of Pennsylvania in the early '60s, I came directly to Brown in 1965, where I've been ever since.

Having told you briefly about some of the changes in my personal life, I'll now talk about the technological part of my theme—technological change and its impact on society. Consider the impact of relatively recent inventions such as the telephone, automobile, antibiotics, computers, and the Internet. The important thing is not just change but the increase in the rate of change. You can see this by looking at a rough calculation of how long each era of societal development has lasted. If we were to live through all those time periods in the space of a single year, the pre-agrarian age would take 11.5 months and the agrarian age two weeks; the industrial age is a single day, and the current information technology/post-industrial age was the last four hours. This dramatic increase in the rate of change affects all our lives.

What is the cause of this acceleration? For one thing, change begets change: current inventions make the next set of inventions occur more rapidly. For example, computers are increasingly vital in the design of next-generation computers. Another reason for this acceleration is that we live in a technology-based society. Amazingly, something like 90% of the scientists and engineers who ever lived are alive today. This means that the future is coming at you faster than ever. Like drivers on a race track who encounter ever steeper hills and sharper turns, you will experience more excitement and thrills but need more skill to keep the car under control. In your lives you will experience ever more dramatic changes in everything from how you work and recreate, to how your children will be born, to how you will die.

In a period of change like this, effective workers and citizens will need not just to understand some specific subject matter or learn a specific trade, but also to develop an unprecedented level of flexibility and adaptability. You are here to learn not only sets of facts, laws, and theories, but more importantly, to learn how to learn and problem-solve at a far deeper level than you've been challenged to in your prior schooling.

Now some of you, who are planning to major in the humanities may think, “eh, technology change does not apply to me.” But you're wrong. Technology is relevant to all of you. Let me give you a couple of examples from our own campus. We have a state-of-the-art virtual-reality facility, one of only two on the East Coast, that is part of a research collaboration with IBM. It is called the Cave because it's a small room made of four eight-foot display screens, three walls and the floor. You have the feeling of being fully immersed in a computer-generated 3D world. This facility is used not only by faculty and students from the sciences and engineering but also by those in the arts and humanities. For example, Professor of Old World Archeology and Art Martha Joukowsky's Cave representation of excavations of the Great Temple in Petra, Jordan features a virtual reality visualization of the dig and its finds for further exploration of the site. The *New York Times* last August featured English professor and novelist Robert Coover's creative writing class project in the Cave, a project called, appropriately enough, “cave writing.”

In your lifetime you will experience ever more revolutionary advances both in traditional fields and in fields that don't even exist today...these advances in science and technology will profoundly affect all your lives

No field today is untouched by technology. It is changing how we relate to society and to one another in our day-to-day lives and is even changing what it means to be human—the essence of what the humanities seek to explore.

Now let me attempt to give you some sneak previews of some of the really fascinating and high-impact potential breakthroughs that are on the horizon. Just for fun, here are some famously wrong predictions:

1) On flying: in 1901—two years before the Kitty Hawk flight—Wilbur Wright said to Orville Wright, “Man will not fly for 50 years.”

2) On computers: Thomas J. Watson, Sr., founder of IBM (and father of devoted alumnus Ambassador T. J. Watson, Jr.) in 1943 commissioned an economic forecast and decided not to enter the computer field, saying, “I think there is a world market for maybe five computers.” Fortunately for IBM and for all of us, he changed his mind after a few years.

In your lifetime you will experience ever more revolutionary advances both in traditional fields and in fields that don't even exist today. The fields that I think will show the most amazing progress are nanotechnology—the technology of the ultra-small—biotechnology, and information technology, especially in their combinations. These topics are not just of interest to techies—I repeat that these advances in science and technology will profoundly affect all your lives, and all of you therefore need to know what is coming down the pike right at you!

One of the most promising and potentially transformative technologies that is being developed now is microminiaturization “to the max”. Materials scientists at Brown in our Division of Engineering and elsewhere are learning to create microscale machines, including sensors and motors, in silicon. Even more exciting, there is good progress on combining organic tissue with silicon to create tiny implantable blood-chemistry monitors combined with insulin pumps that will free diabetics from the need for shots.

Beyond that, we're moving into the thousand-times-smaller molecular scale. Molecular nanotechnology will allow us to build machines literally an atom at a time. When we learn how to make them, you'll be able to fit several million nanotech machines on the head of a pin. It may even be possible in this century to build cellular-scale robots that can move through our bodies unfelt, destroying cancer cells or repairing the damage that comes from aging and allowing us to live longer, healthier lives.

Biotechnology and the life sciences are likely to be to the 21st century what physics was to the 20th. You've all heard of the mapping of the human genome and ongoing research to understand the biochemical implications of that digital code. Equally profound in its implications is cutting-edge research to understand how the brain works, how it senses information and controls the body. Brown's cross-disciplinary Brain Science Program, led by Professors John Donoghue and Leon Cooper, is pioneering the effort to build brain-machine interfaces to help patients whose thought processes are unimpaired but who can't control their muscles. The goal of the research is to enable the brain to control machines such as computers and wheelchairs purely by electrical impulses in the brain generated by thought.

Now I'm going to talk a bit about my own field, which is not only a rich intellectual discipline in its own right but also a major force in all other disciplines and our economy, indeed our society.

John Abromowski, Photographer, Public Affairs and University Relations

Andy and Ruth after the ceremony

Thanks to microminiaturization, today's electronic greeting card is more powerful than one of those early dinosaur computers that filled a large room, but it costs pennies, not millions of dollars. You all know about instant obsolescence of any computer you buy because next year's model will be twice as powerful and cost less. Nothing in nature or manmade has ever benefited from such an exponential improvement curve. You can't be blasé about an exponential, and what's more, we're just beginning to climb the curve—the best is yet to come! If, over the past 30 years, cars had improved at the same rate as computers, they would cost pennies, go 100,000 miles in an hour, use only a gallon of fuel—but would fit on the head of a pin.

So what may we expect? A profusion of computing devices of all sizes and shapes, from tiny to huge, all with amazing computational intelligence. Computers will be everywhere, not just in all our appliances and vehicles but in our work and play spaces and, through nanotechnology, even in our bodies.

How you will interact with computers will change dramatically. Instead of being restricted to the very limited, clumsy means we have today, mainly typing and mousing, we will interact much more as we do with an intelligent partner who knows our background, habits, tastes, and needs. Computers will finally be able to understand you when you speak to them, They will even react appropriately to our moods as a sensitive friend does. The interactive display surfaces that know you shown in this summer's science-fiction movie "The Minority Report" will come to pass. At Brown we're actually working on many aspects of this vision, especially using the Cave I mentioned earlier. For example, in the Geology Department's simulation of flying over the Martian north polar cap looking for potential landing sites, you can speak simple commands to the computer and point in the direction in which you want to fly.

Even further out, AI and robotics experts predict that not only will machine intelligence surpass human intelligence by the middle of this century—within your productive lifetimes—but that humans and robot evolution will start to converge as humans become more and more bionic while machines approach human levels of sophisticated and intuitive intelligence. Not all computer scientists agree with these aggressive projections, of course.

You may dismiss the further-out speculations as just science fiction, but remember, in the last century we've seen an amazing number of science fiction speculations actually happen, often in a shorter time frame than anyone predicted, things such as landing a man on the moon or wearing a computer on your wrist. As I mentioned earlier, advances in computer technology fuel almost all other technical advances as well as both social progress and social problems, such as identity theft or the potential terrorist attacks on our information infrastructure. To sum up my technology predictions, you will live a science-fiction future.

Now I want to switch gears and take advantage of this bully pulpit to leave you with some thoughts that may help you cope with change in a science-fiction future. First, I want to tell you about the value of persistence. Here's what former President Calvin Coolidge had to say:

> Nothing in the world can take the place of persistence. Talent will not; nothing is more common than unsuccessful men with talent. Genius will not; unrewarded genius is almost a proverb. Education will not; the world is full of educated derelicts. Persistence and determination alone are omnipotent.

Learn to handle Brown's pressure. Being persistent and working hard is at least as important as native ability. And just to personalize that advice, I nearly flunked out during my freshman year in college because, having by and large breezed through high school, I was overconfident and thought I could handle not only a full course load, but also late-night socializing in the dorm, JV soccer, band, theater, and a new girlfriend (who later became my wife).

I am extremely proud that among the academic leaders turned on by their ugrad TA and RA experience with me have been seven department chairs and well over two dozen faculty members...Furthermore, few of them got straight As but their persistence and passion enabled them to become creative and productive professionals

I also saw that many of my peers weren't struggling the way I was and had far more impressive backgrounds. I started feeling as if I might not belong. Thus I was painfully forced to reassess my abilities, cut back and regroup, get my act together; I then did OK. But many of the superstars I was intimidated by never lived up to their early promise. You don't have to be a star as an undergraduate to be very successful in your chosen career and in your life as a whole.

Right up there in importance with persistence is developing a passion. The thing that has sustained me when I felt scared or inadequate has been passion. If you find something that you love doing, you can do it even in the face of great obstacles.

I want to address a specific outlet for passion that has been very important to many, many hundreds of my former students here at Brown, and that applies especially to you undergraduates. This is the ability to be not merely a passive consumer of courses but to be actively, even passionately engaged in a field of study and involved in a department by becoming an undergraduate teaching assistant or writing fellow and/or a research assistant. Involving undergraduates this way was all but unheard of, even controversial, when I started doing it in 1965, but over the years it has become quite common, both here and in other schools.

TAing or being a research assistant is an experience that not only teaches you more about a specific subject than taking a course, but may well help you find your passion, your calling in life. And it provides terrific leadership training. I am extremely proud that among the academic leaders turned on by their ugrad TA and RA experience with me have been seven department chairs and well over two dozen faculty members. Other former ugrad TAs and RAs who are now leaders in industry include the person in charge of all the modern Intel microprocessor designs, the person in charge of Windows95 and Internet Explorer, and the people in charge of modeling and lighting and much of the software for Pixar's award-winning Toy Story movies. By the way, most of these students took my course as an experiment without any intent to major. Furthermore, few of them got straight As but their persistence and passion enabled them to become creative and productive professionals.

The context within which persistence and passion must operate is freedom to experiment. It is crucial that you experiment here, personally, intellectually, artistically, and socially. As to intellectual experimentation, explore widely to take advantage of being in a great liberal arts institution. I did the writing I am most proud of in a college course on Goethe's *Faust* in which I combined my fledgling knowledge of optics from physics and perception from psychology to examine Goethe's theories of color vision and the wonderful light imagery he used in that great verse drama. I was so enthralled that I flirted for several weeks with the idea of switching out of engineering sciences to some combination of German literature and perceptual psychology, and I still love teaching the basics of color theory and perception in my graphics course.

Experimentation includes false starts. Take the advice of your advisor seriously but don't obey slavishly. Define your own path. Find your own inner compass. As a graduate student I had to plead with my advisor to let me do my Ph.D. thesis on computer graphics, an area that was unknown in my school because at that time there had been only one other Ph.D. in the brand-new field of computer science as a whole.

Finally, back to the theme of change, the rapid technological change not only causes changes in society but also confronts us with brand-new ethical dilemmas. Our society is already being faced with hard choices confronting the implications, for example, of human cloning, stem cell research, and reproductive technology. Will you determine the gender of your children? How much do you really want to know about your biological destiny? Do you want the government (or your parents) to know your whereabouts within a meter day and night? I believe that it is no longer possible, as it seemed in my college days, to study any subject for its own sake without considering such questions. Roger Blumberg will consider provocative questions like these in his freshman seminar on Computers and Human Values. What is profoundly exciting about the technologies that will change our world is that many of you will play pivotal roles in their exploration and application. You will be the leaders who will make the scientific discoveries, lead society through adapting to these changes, and wrestle with the political, socioeconomic, and ethical issues. You have both the opportunity and the responsibility to confront these challenges. To finish Calvin Coolidge's remarks about persistence and determination alone being omnipotent: "The slogan 'press on' has solved and always will solve the problems of the human race."

Thanks very much for listening to me, and press on!

NEW FACULTY MEMBER

Our newest faculty member, Anna Lysyanskaya, was born and raised an only child in Kiev, USSR. Her father is a math Ph.D. and her mother an electrical engineer. In '95 her parents moved to the States and went to live in Colorado, where an aunt had already settled.

The educational system in the USSR was rigid and choosing one's own course of study was not an option. Fortunately, hers was a specialized high school with an intensive program in English, which was particularly exciting because her study there coincided with the fall of the Iron Curtain and increased contact with the West. Tourists and foreign visitors were often sent to visit the school as a showcase and because the students spoke English so well. These contacts fomented Anna's desire to go to college abroad. The U.S. educational system, and the idea of a liberal arts education, were especially appealing. Considerations such as the virtual impossibility of getting into a good university in Ukraine without connections in the right places also entered into the equation. Going to the U.S. was not an easy proposition because of fierce competition for full financial aid for foreign students. Luckily for us, Anna was successful and was accepted at several schools. She chose to attend Smith College, graduating *magna* in 1997.

Members of the WICS group (Women in Computer Science) enjoy their initial meeting of the semester, traditionally a bagel fest

Anna Lysyanskaya

Unsure what to study at Smith, she took an introductory course in programming and ended up TAing the course. Eventually she decided to pursue a combined Math and CS degree. She particularly enjoyed working with Joe O'Rourke, whose interest is computational geometry and the philosophy of artificial intelligence. Interestingly, during her years at Smith, Ruth Simmons was the president. Anna felt Ruth the perfect president for Smith and was very disappointed to learn a couple of years ago that she had left for Brown. But Ruth had anticipated such sentiments among alums and actually toured the country to talk to them and explain her decision. Anna considered Ruth a no-nonsense president, unafraid of not being PC and someone who spoke her mind. Said Anna, "At Smith the way many students dressed and behaved was quite frivolous and bordered on the indecent, but most of the faculty and administration were afraid to say anything for fear of not being politically correct. Ruth didn't like it, she said so openly and actually went on to do something about it. For a Smith administrator that was super-brave."

At graduate school at MIT (she also got into Brown, by the way) she decided to study something that had applications in practice but had a theoretical core—cryptography exactly fit the bill. Her advisor, Ron Rivest, was her inspiration. During the summers she preferred to escape the city; she spent two summers at Lucent doing research, two at IBM Zurich, and one at IBM New York.

Anna is a serious British mini-series fan, "Pride and Prejudice" and "Brideshead Revisited" being her current favorites. She has recently been convinced to start watching "Upstairs Downstairs"; however, to prove her affinity for American culture, Anna admits to having loved the movie "Clueless."

Recently she bought a house on the East Side and is being traumatized by the realities of home ownership. When asked whether she was entertaining yet, she reacted with horror and admitted that she has yet to buy furniture! She loves the Providence scene but has been too busy at Brown even to see a WaterFire or visit the RISD Museum—she has much to look forward to. To continue with Rhode Island lore, on a recent supermarket trip she came across a jar of Buddy's Marinara Sauce complete with photo of a grinning Buddy Cianci (Providence's ex-mayor) on the label. She was almost tempted to make the purchase but, recalling his recent trial and conviction, decided against it. Still, if anyone has an EMPTY Buddy's Marinara Sauce jar, Anna would very much like it...

AUTOMATA

Master's student Eduardo Hueso, author of this article, was the technical director for this very clever animated short film. Our thanks to Prof. Michael Black for giving us the tip!

In a dark place called Autopolis, automata live and operate. This is the story of an automaton called Dave whose control program aborts. Lost and confused, he is forced to make a decision and find his own way, and ultimately discovers the unexpected results of "free will".

Supervising Animator	Alberto Lara
Technical Director	Eduardo Hueso
Concept Art and Design	Octavio Villegas
Music and SFX	Diego Bauducco
Modeling and Animation	Pedro Omedas and Francisco Berrizbeitia
Narrated by	Derrick Bishop
Written, Directed and Produced by	Carl Zitelmann

Automata is a 12-minute independent short animated film made using 3D computer graphics and produced mainly in Caracas, Venezuela and partially all around the planet. A group of seven colleagues, friends and classmates worked on this project for a whole year, largely to pursue a dream. However, *Automata* has succeeded more than expected; by obtaining good critiques and being accepted for competition in international film festivals, it has become an important piece of our portfolios that we all proudly embrace.

After some years of playing with animation software and experimenting with digital visual effects, I started feeling that, although computer graphics is an amazingly interesting and challenging field, it is little more than a tool for art and communication, another brush in the artist's box. Graphics by itself, without a purpose or a message to deliver, is interesting only from a research point of view.

Knowing the tools, understanding the technical issues of CG, and being able to create interesting images is one side of the problem. Using all these tools and knowledge to communicate something meaningful, to evoke emotions and get a message across to an audience is normally treated as a separate problem and has traditionally been considered the role of the artist.

For me, these two aspects of computer graphics are two extremes of a continuum that are intimately related and depend on each other. The effectiveness of a scene depends on the technical viability of a certain visual effect, and in turn the effectiveness of a visual effect is determined by the power of the argument behind it. This intimate relation can be seen clearly in the work of the traditional painter who is at the same time the director who decides what color will best evoke an emotion and the technical master of the tool who actually creates the necessary brush-

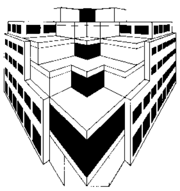

Eddy Hueso working on *Automata* in a Budapest café

strokes. In animation production the complexity of the tools and effects is such that both technical and art specialists must work together to achieve a unified result. This kind of hybrid area, where art and science come together, is in my opinion the role of computer graphics.

For some years I had been looking for the right story and people to start an independent animated film, and then Carl, a friend and former classmate at Universidad Simón Bolivar in Caracas, Venezuela, showed up in my office with a great idea. He brought along a rough script and a couple of sketches that, together with his narration of the story, captured the interest of everyone in the office. *Viva Carl!* Carl has always been a very creative person; his mind is full of pictures, scenes and sounds that he is able to describe in great detail.

The team got together immediately. Francisco and Pedro, both students of computer engineering at U.S.B. Caracas who were doing part-time research in human facial modeling and had been in the office when Carl showed up, couldn't resist this venture. Octavio, a very talented illustrator and designer, was the only formal artist in the group. Alberto, another computer engineer from USB who was working full time at a multimedia production company, joined as the main animator. Diego, the sound engineer, who graduated from Berklee College of Music in Boston, offered to take care of all the music composition and sound editing. And I, who had also graduated from U.S.B. in computer engineering, had had experience with various 3D packages, production tools, and CG technical issues.

The contract to join the team was very simple: no money, no office space, no computers, and no employee benefits. We would have to work in our spare time and use our personal computers. However, it was clearly a perfect opportunity to execute a dream we all shared and these conditions were just part of what made it exciting. I believe that working in a small independent film is the best way to gain experience in the whole production process and to be a creator in all aspects of it.

We all implicitly agreed on what I would say was one of the main goals of *Automata*: we wanted to make a balanced production in which the story, the argument, the art design and the technical execution were present in the right measures. Since most of us had technical backgrounds, we found it challenging and exciting to experiment with the whole creative process and find our own answers to questions like: how can we take the base storyline that came from one person's mind and use our team's collaboration to make it a richer narrative with graceful details? How can we communicate in order to give the whole team a clear vision of the desired look, atmosphere, and feel of the film? This issue was particularly important in the case of *Automata* because we worked independently with very little supervision and yet all the pieces had to have a unified look. We also got to experiment with cinematography issues like camera motions,

changelog

Andy van Dam, who has been at Brown since 1965, has been named Vice President for Research, *a new position.*

Andy continues to teach his accustomed CS courses and retains his CS office, but also has one in University Hall.

His appointment reflects President Ruth Simmons' wish to emphasize research as a university priority, providing more generous resources and additional faculty. Andy also oversees the Office of Research Administration, Brown's intellectual property policies and efforts to identify research with real-world applications.

Said Andy, "I know from long personal experience that the quality of research produced by the faculty at Brown is excellent, and that the potential for growth in the university's research effort is substantial...I'm very hopeful that framing, lighting, and the like that are key to the effective transmission of a coherent story. One of our difficulties here was the proper management of the absolute freedom of camera motion available in a digital medium. The absence of the physical constraints of traditional filmmaking can tempt the inexperienced to abuse this “advantage” of the medium. In *Automata* we made a big effort to remain “traditional” in our lighting and camera motions. We emphasized using cinematographic “clichés” rather than exploring the capabilities of the digital media in that respect.

Given our strict limitations in time and resources, technical sophistication quickly dropped to the bottom of the priority list. We had to design compelling but simple characters that could be animated easily and still be expressive. The environments needed to be simple and reusable. Doing lip-synch was totally out of the question, so having a narrator seemed the most appropriate from the beginning. Music and sound effects were key in achieving the desired atmosphere.

Taking these factors into account, we went through a design stage in which most of the characters, machines, and environments were roughly designed on paper. Pre-production started simultaneously: for about two months we met in cafés and on Carl’s balcony once or twice a week to discuss the story. The story was broken up into numbered acts, scenes and shots. We brainstormed on one or two acts per meeting. During the meetings, while we all deliberated on the creations of our right brains, Octavio drew amazing real-time pictures that miraculously depicted our ideas and later became the storyboard. *Viva Octavio!*

As a learning process we tried to keep our production stages well defined and aligned with the formal stages of the production of animated films, namely:

Storyboard: hand-drawn pictures that depict the main action nodes and overall look of the story. This was the first production stage but was extremely useful at all stages as a reference document and communication tool.

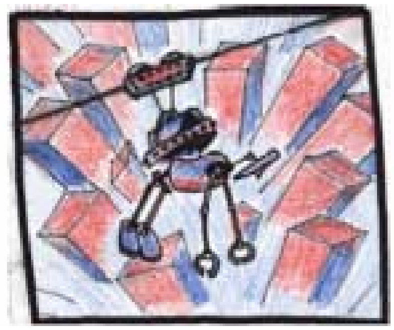

Animatics: a version of the film using the storyboard drawings as the video frames and a mock audio. The main goal of the animatics is to confirm the coherence of the story and create rough timing for each shot. The animatics is the step in which the story is first laid out on a time line—before this step it is difficult to tell how long the story will be.

Modeling: in this stage most of the characters, props and environments are modeled in 3D. For this we

made an inventory of models, split it among us and worked on them indepedently for 3-4 weeks. Posting our intermediate results in our private webpage helped to get feedback from the rest of the team and give things a consistent look.

Layout: a simple version of the film in which the 3D models are roughly placed in 3D scenes and only their positions within the scene and the cameras are animated. The layout was normally rendered in preview mode with very simple lighting and

changelog

we will, in five to ten years, take our place in the major leagues of the top-tier research universities.”

------8<------

“Activist and journalist” Dilip D’Souza, ScM ’84, has a newly-published book, The Narmada Dammed: An Inquiry into the Politics of Development. In an incisive analysis of the controversial judgment to resume construction of India’s Sardar Sarovar Dam, Dilip’s first-hand experience among affected people led to his conclusion that dams are being built less to solve the problems of water, flood and power, and more for the sake of politics. http://www.penguinbooksindia.com/books/aspBookDetail.asp?ID=4995.

------8<------

Peter Wegner celebrated his 70th birthday at a surprise reception in the Faculty Club last August. Feted by family and

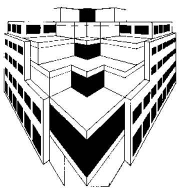

no textures or atmospheric effects. The idea of this stage is to fine-tune the timing, relations between models, camera framing and camera motion. The layout scenes were in many cases used as templates for the animation stage.

Animation: the stage in which the characters are brought to life and their interaction with the environment is animated. In *Automata* we used purely key-framed animation, and relieda lot on Alberto's talent as a character animator. *Viva Alberto!*

Texturing and lighting: the stage in which most of the textures, environment effects and lighting setups are created. Texturing and lighting are crucial in the achievement of the desired atmosphere and feel of the film. We used nearly 100 digital

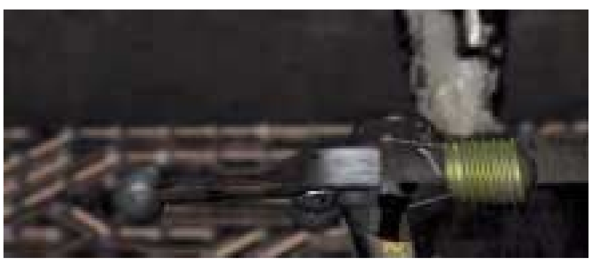

pictures of different surfaces to create digital materials that were applied to the 3D models. The lighting in most cases aimed for the dark and foggy environment that was conceived in the original story line.

Rendering: the stage in which the 3D scenes are converted to 2D images, 30 per second of animation. This was technically the most difficult part, where most unexpected flaws showed up and most time was wasted in repeating endless renderings. The renderings took 100% of our personal CPU cycles for long days. A typical 30-second scene took between 15 and 40 hours to render, and in many cases was not satisfactory the first time. We made a big effort to speed up renderings by compromising atmospheric effects, polygonal resolution and other effects, but even so the final 12 minutes of animation took roughly 400 CPU hours.

Audio production: the stage in which original music is composed, produced and mixed with the sound effects. The audio track had to match the final edited video, and for this reason was largely left for the very end and created under tight deadline stress. Diego was in charge of the audio and responsible for the music composition and execution. *Viva Diego!*

Postproduction: the final phase, in which the rendered clips are edited and put together, the rendered frames are manipulated to emphasize the desired look and compensate for the excessive crispness and perfection of the 3D rendering, the audio is mixed in and synchronized, and the final master is built. This phase happened almost a year after the beginning of the project and under the pressure of deadlines for submission to film and animation festivals. Not having access to specialized hardware for video mastering forced us to get help from a local studio, Nexus, where a good friend of mine, Alfredo, works. They got excited about our results and helped us create the masters without any charge. *Viva Nexus!*

The production of *Automata* took a whole year during which many important professional and personal events happened in the team members' "parallel" lives. In particular, my acceptance at Brown University at the beginning of the project made it seem like the right moment to take some time off to travel. I left my job three months before the beginning of classes and went to Europe with a backpack and *Automata* in a laptop.

My long-term absence implied that the office in Venezuela where Francisco, Pedro and I worked as a satellite group for our mother company in Los Angeles and Budapest needed to be closed. Given the circumstances, Francisco and Pedro had the chance to go to L.A. and later Budapest to continue their work on human modeling closer to the rest of the research team. Simultaneously Diego went to Miami to work in a recording studio, and only Carl, Alberto and Octavio remained at home.

changelog

friends, Peter was in fine form. The highpoint of the evening was Judith's witty paean to Peter sung to the nursery rhyme tune 'Old MacDonald Had a Farm'!

Peter is also the subject of a new 21-minute video entitled "Peter Wegner is Alive and Well in Providence." Based on his experiences after his near-fatal accident in the summer of '99, the video will be used as a teaching text for current issues in biomedical ethics. Professor Edward Beiser hosted the screening; Luke Walden created the video.

------8<------

Yet again, the all-graduate football team, Public Enema, won the 2002 High-Intensity Intramural Flag-Football Championship. Over half the team are members of CS. Both this year and last, the finals were against the Thete fraternity, and both games were won by one point! Go

The time we spent all together at the beginning of the project in the preproduction stages later proved to be of key importance. Working from different places of the world with very little communication was possible because we all had a clear picture of the project in our minds and we had worked on creating it together. Also, of course, the Internet was our main means of communication even when we were in the same city, so physical distance didn't make that much difference during the intermediate production stages when each of us had a set of well-defined tasks. Worth mentioning is Carl's great management job in motivating and coordinating the team's efforts and keeping everyone up to date.

During the final production stages most of the team was back home and I was in Brown's grad school. After Fedexing a number of CDs with my final rendered frames I had little to do but wait. Carl and Alberto took care of most of the postproduction and final touches back in Venezuela. Without having participated in the final production stages, I was absolutely shocked and surprised by the magnificent result. I think none of us was expecting the visual quality and story-telling effectiveness we achieved. This was the first serious independent production experience for all of us, and the biggest surprise we encountered was, happily, the good results of our work.

Currently *Automata* has been submitted to a number of international film and animation festivals around the world and this far we are proud to announce that it has been accepted for competition in three major ones, Savannah Film Festival, L.A. Shorts Fest and the International Short Film Festival Berlin. *www.automata-film.com* is still under construction but should be up by the time this is published. It will describe the whole process in detail and have lots of pictures, including parts of the storyboard and other production phases as well as the trailer and latest news on the status of the film. Unfortunately, *Automata* cannot be screened online while it's still participating in festivals; however, it will be screened on campus very soon. Send an email to *hueso@cantv.net* if you'd like an invitation for the screening of *Automata*.

John Guttag '71, ScM '72 (l) and Ed Lazowska '72 are on the International Advisory Board for the computer science program of the National College of Ireland. They were touring a new facility, nearing completion—both MIT, where John is and UW, where Ed works are also completing new facilities simultaneously—a prime opportunity to troll for ideas while touring the NCI construction site

changelog

Computer Science! See Page 25 for a picture of the victors.

------8<------

Luigi Di y los Gapasutras is the official latin-rock band of the CS Department. The all-cs-grad-student band boasts a repertoire of original songs (with quasi-poetic lyrics in Spanish) featuring a mixture of Spanish folk, rock, funk and an Afro-Peruvian style that is making headlines. Expect their first album by the end of 2003.

(l) Luigi Di (aka Luis Vega, MSc '01) and Dan Keefe (r) write the songs, sing and play guitar; Daniel Acevedo plays bass and Luis Ortiz (PhD '02) is the percussionist. They will be adding a latin flair to this year's holiday party.

ALUMNI @ SUN MAKE CONTACT

By Mike Shapiro and Bryan Cantrill '96


(l to r) Mike Shapiro and Bryan Cantrill outside the CIT in '94 and outside Sun this year!

When we met during our first-year orientation at Brown, we had been on campus for only a few days. We could hardly have guessed the journey we were beginning—over the years at Brown, we took CS classes together, TA'd together, head-TA'd together, did our CS32 final project together, head-TA'd together again and then again once more. So it will perhaps come as no surprise that after ten years, we're still working together: now (as since Brown) in Solaris Kernel development at Sun Microsystems. Less expected, perhaps, is the number of Brown graduates we have subsequently recruited into our department: in addition to the two of us, there are Stephen Hahn PhD '97 (physics), Dan Price '98, David Powell '99, Matt Ahrens '01, and Adam Leventhal '01. Among these seven samurai, we count five former CS169 head TAs, two former SPOCs and one former head consultant.

But we're not making offers to Brown students because we want to spend our professional lives reminiscing about late-night credit at the Gate. It has been our experience that Brown CS graduates are, consistently and by a wide margin, the best suited for our environment. Why does Brown routinely generate such high-caliber graduates? The answers are manifold, but they include at least Brown's undergraduate TA program and its balanced computer science curriculum.

While sophomores at Brown we recall reading a special "career" edition of *Computerworld* aimed at college undergraduates. While this was mostly a shameless advertising vehicle to target the highly prized 19-24-year-old male demographic (replete with "Mountain Dew" ads featuring testosterone-fueled BMXers), it contained an interesting article in which several industry leaders were asked to name the schools they recommend. The answers were surprisingly disjointed, with virtually no school appearing on multiple lists—with one notable exception. Brown appeared on nearly every list, and some contributors (notably Alan Kay, inventor of Smalltalk) came close to implying that anything other than a Brown education amounted merely to vocational training.

And certainly we had a hunch that the CS department was special: where else would a sophomore be allowed to participate in the education of a first-year? Or a junior be allowed to help redesign an upper-level course? Or a student be trusted with the root password to everyone else's data? Being given such responsibility was exhilarating—and frightening. While TA'ing or working for the tstaff, one quickly learned one of life's most important acquired skills: when to say "I don't know" and how to find the right answer. And when TA'ing a course, one was rarely explicitly told what to do—TA's were expected to figure out what needed to be done to keep the course on an even keel and to do it. Working as a TA or for tstaff also afforded endless opportunities to design and write new software that would actually be put to use immediately by students and faculty throughout the department.

And then there were the courses themselves. In the second semester of our first year, we went looking for an upper-level course to take in the first semester of our sophomore year. (Like an intrepid pair of war orphans, we had already decided that we would increase our chances of survival if we stuck together.) Specifically, we were trying to decide among CS141, CS123 and CS169. As we interviewed juniors and seniors, it became clear that while CS169 had a reputation of having incredibly difficult programming assignments, it had also engendered in its graduates a die-hard allegiance. We remember one senior excitedly exclaiming, “It’s really fun: you get to cast to function pointers!”, to which at least one of us remembers thinking “Great, what’s a function pointer?”

In the end, we agreed that we wouldn’t let the fear of the unknown deter us from taking CS169. The course *was* difficult: the lectures were thought-provoking (well, except the Kerberos lecture—yawn), the programming assignments were hard (a difficulty exacerbated by a massive transition in the course that particular year), and the bugs were absolutely brutal. We loved it. The euphoria we felt when nailing a killer bug or finally understanding an elusive concept is what ultimately led us both to Sun after graduation, where we have pioneered new solutions to problems we first met in CS169.

The euphoria we felt when nailing a killer bug or finally understanding an elusive concept is what ultimately led us both to Sun after graduation, where we have pioneered new solutions to problems we first met in CS169

The challenge offered by CS169 is emblematic of the best the department has to offer: a course that demands the student master both theory and implementation. Brown excels at maintaining this balance—and in our years of recruiting for Sun, we have been surprised to find that a number of allegedly top-flight schools fail miserably in this regard. These schools, which count two high-profile institutes of technology among their ranks, view implementation with a sort of disdain. They seem to take the implicit view that asking a computer scientist to write a program is like asking a civil engineer to dig a ditch—it is demanding a mundanity that is beneath the gentleman engineer. Of course, computer science is not civil engineering, and software systems have little in common with dams and superhighways. By depriving students of the details that would give them intuition for theory, these schools leave their graduates ill-equipped to solve real problems in any aspect of our field. And ironically, Brown’s balance of theory and implementation generates not only better implementors but substantially better *theoreticians* than any other school at which we have interviewed.

This balance may explain why Brown students are well prepared to thrive in a challenging and dynamic work environment, but why have so many graduates specifically chosen our environment? Like many other companies in Silicon Valley, we are solving problems that have not been solved before and that, once solved, will advance the state of the art. But unlike many other companies, we have an enormous number of such problems—a consequence of the innumerable markets in which Solaris must be competitive. Indeed, we have a seemingly infinite amount to do because Solaris now has what no operating system has ever had before it: substantial markets from the smallest to the largest possible general-purpose computers[1]. Each of these markets demands that the operating system scale linearly with available resources while not degrading exponentially with additional load; that applications built upon it can be easily developed, debugged, and deployed; that hardware faults induce the minimal possible system failure; that the operating system itself be completely reliable; and that the whole mess be managed with as little human labor as possible. These demands lead to a bevy of hard problems requiring creative solutions. For example, most traditional OS algorithms break

1. For you old fogeys: yes, System/360 spanned a similar range of machines (roughly two orders of magnitude in size), but OS/360 infamously could not and did not. And for you young punks: Linux starts coughing blood at about eight CPUs—far below the 128 CPUs to which Solaris scales.

down with a terabyte of main memory, a petabyte of storage, hundreds of CPUs, or millions of threads. And the sheer number of problems lets us give tremendous latitude to those who endeavor to solve them: if one problem domain becomes stale, there is always another to attack.

Finally, there is the sense of teamwork and camaraderie we remember well from our days and nights in the CIT. Whether it was having someone else around when you're working in the lab all night, having someone else to talk over your latest impossible bug with, having someone else appreciate the magic of your OS simulator coughing up its first shell prompt, or attending TA meetings that sometimes resembled Marx Brothers movies, we always found strength in the teamwork and shared philosophy developed at Brown. The creation of a little piece of that environment in our corner of Silicon Valley has proven successful for Sun and rewarding for us and the other Brown alumni we have recruited, and is strong evidence of the value of Brown's balanced program.

IPP SYMPOSIUM ON COMPUTER AND NETWORK SECURITY

Symposium host Tom Doeppner

The spring IPP symposium (the 29th) was held on April 25 and had an all-star cast. **Joe Pato of Hewlett-Packard**, Brown '81 and a former technical staff member in the CS department, led off with a highly informative talk on "The Cyber-Security Revolution: Protecting Critical Infrastructure with the Emperor's New Armor?" He began with an overview of the security problems facing industry today and their ramifications. He mentioned a number of industry-wide groups and efforts that have been put together to cope with these problems, in many of which Joe is an active participant, and gave us a summary of what every organization should be doing, from the CEO down to the private citizen.

The next speaker was **Stephen Kent of BBN Technologies** (now part of Verizon). He instructed all of us on the promise and problems of biometric technologies for authentication in his talk, "A System-Level Analysis of Biometric User Authentication." He described a number of fascinating technologies being used or envisioned for biometric use and explained the risks of each, concentrating on ease of spoofing. Interspersed with his slides were some great cartoons and a few beautiful photographs from his recent vacations.

The third speaker, holding our interest in spite of coming just before another bountiful and tasty lunch organized by IPP manager Suzi Howe, was **Steven Carmody,** Brown '71, now of **Brown's Computing and Information Services**. Steve is part of the national Internet 2 project, working on their middleware architecture. He talked about an important aspect of this—Shibboleth, which is "an initiative to develop an architecture, policy framework, and practical technologies to support inter-institutional sharing of resources." A big concern is to provide "attribute-based" authorization, in which access to resources is controlled using attributes that don't identify the users. Thus authorization can be achieved without degrading privacy. Steve outlined the technical details of the project as well as progress to date.

Following lunch we were entertained and enlightened by an amusing yet informative talk by **Radia Perlman of Sun** on "How to Build an Insecure System Out of Perfectly Good Cryptography". She described a number of important basic concepts in security and cryptography and showed how much common use of them is misguided. She went through some current work in the security area, discussing some important protocols, and showed us how much of this can hardly be called advances.

Next up was **Christopher Spirito of EMC** speaking on "System Security Methodology: Protecting your ASSets." Christopher continued on the theme Joe Pato established earlier: what an organization must do to protect itself. After characterizing the threats and the adversaries, he described a system security methodology in which not only the networking people and the OS people but also the applications people all work together towards protecting the organization's important assets: its data.

The penultimate talk was by **Roberto Tamassia, a Brown CS faculty member**, on "Distributed Data Authentica-

"MODERN TIMES" COMES TO CS009

Brown's First-Year Seminar program, inaugurated this year, features a fall course in the CS Department created by Roger Blumberg entitled "Computers and Human Values" (CS009). This course discusses a number of provocative ideas in contemporary computer science in the context of traditional liberal arts debates about humanity and society. The first unit of the course explored developments in contemporary robotics and the questions they raise about the nature of our humanity, with the students reading works by Hans Moravec, Hannah Arendt, and N. Katherine Hayles. In this photo, Tom Dean (3rd from left) joins a session on Moravec's *Mind Children* to discuss contemporary "strong AI" as well as his own views about humans and/as machines. The group is watching the opening scene from Charlie Chaplin's classic *Modern Times* to illuminate a discussion of the relationship between technology and human beings in industrial and post-industrial society. Most of the students had never seen Chaplin's movie (!), and the experience evoked many interesting insights about the different ways technology can challenge our sense of what it means to be human. The syllabus for CS009 can

tion," an area in which he not only does research, but also has a company commercializing his ideas. The concern he addressed was how to authenticate large volumes of data and transactions in a not necessarily trustworthy environment. He showed how a trusted source can replicate its data on non-trusted "responders" that nevertheless provide authenticated answers to queries posed by users.

Mirek Kula, GTECH, at the post-symposium reception

Last, **Mirek Kula of GTECH** spoke on the intriguing topic of "A Commentary on Lottery System Security in the Internet Era." In particular, how does one go about running lotteries over the Internet? GTECH, which in the past has run its business over relatively secure proprietary networks, now must do all of this over the public Internet. He outlined their requirements and assumptions, the most important of which is that 100% secure systems are impossible. Thus a careful risk assessment is absolutely necessary. He discussed what GTECH is doing about all of this, which, in summary, is to shift from "reactive" security to "proactive" security, to spread the "security mindset" across the company, and to implement a number of specific security measures including risk assessment, security design reviews, and security testing.

Symposium speakers: top, l to r: Steve Carmody, Brown CIS; Radia Perlman, Sun Labs; host Tom Doeppner; Stephen Kent, BBN. Bottom l to r: Roberto Tamassia, Brown CS; Joe Pato, Hewlett-Packard; Christopher Spirito, EMC

Michael Black. This spring and summer Michael served as an area chair for the European Conference on Computer Vision (ECCV), on the papers committee for SIGGRAPH, and as a program committee member for the International Conference on Pattern Recognition and the International Conference on Automatic Face and Gesture Recognition. Michael was excited to join the Brain Sciences Executive Committee this summer at a time when the Brain Sciences Program (BSP) is still growing at Brown. In a BSP collaboration with John Donoghue and Elie Bienenstock, Michael and his colleagues received a grant from NIH-NINDS to study the motor cortex and its reorganization. He had a busy summer of European travel, including a trip to Copenhagen, where he and his collaborators had four papers at ECCV. In August he attended a workshop on motor control in humans and robots in Edinburgh, where he had a co-authored paper. A program committee meeting in Sweden and a trip to Grenoble for a thesis defense and workshop rounded out the travel schedule. In an endless summer of travel, Michael found a little time to go surfing and turn 40!

Ugur Cetintemel. Ugur served on the program committees for two conferences: Information and Knowledge Management 2002 and Applications and the Internet 2003. He also served on the organization committees for Distributed Objects and Applications 2002 and Cooperative Information Systems 2002. He co-authored papers presented at Principles of Distributed Computing 2002 and Very Large Data Bases 2002 in Hong Kong. Microsoft provided wireless pocket PCs to support Ugur's data dissemination research. He is currently using these in his pervasive computing class for mobile computing projects.

.

Amy (...make my day) Greenwald on her recent trip to Africa. The gun belonged to a hitchhiker they picked up!

John Hughes. In her list of topics for the *fac.activities* column, Suzi included 'Interesting trips abroad'...Said Spike, "Well....it's not actually abroad, but I went to San Antonio for SIGGRAPH in July. The conference was great. Because I was the Papers Chair, my part of the work was mostly done before we got there. But there was one fun thing that we did onsite: working from a suggestion David Laidlaw made, we planned something called the "Papers FastForward." The idea was to give folks a chance to see what all the papers would be about before any papers sessions began, and to do it really quickly.

"With 67 papers, we gave each presenter 52 seconds and two PowerPoint slides to make a pitch telling folks why they should come to the talk. I was the MC. In the week before, I put together a huge PowerPoint deck that had, for each speaker, the two slides followed by one with an alarm clock and a ringing sound. At the event itself, the A/V folks were told to cut off the microphone when the bell rang. I got all the speakers lined up in order along one wall near the speakers' platform, and since I'd gotten slides from all of them, I knew we were ready to go.

"Everything went beautifully—the 2500-person room was standing-room-only, and the speakers did a great job: we had a rap, a presentation in the form of a sonnet, and a speaker who put up quotes from two of his reviews: "This paper thrilled me" and "The results are not good enough to merit publication at SIGGRAPH," and then told the audience to come and decide for themselves.

"The only problem arose after about ten speakers. The previous speaker finished up, the bell rang, and the next slide showed up, but the next speaker didn't come up to the platform. I waved him up, but he said in a stage-whisper, "Not my slides!" Well, as I said, I had this autoadvance PowerPoint deck running, which meant we were going to have 52 seconds of "dead air." So I got up and did a quick presentation by reading from the slides and doing my best to fill in with things I remembered hearing during the papers committee meeting. It's true that I hadn't actually read the paper, but I figured I couldn't really go too far wrong. So I spoke for 40 seconds and finished with "and if you want to hear the real authors present this, come here at 4 o'clock Thursday." It happened twice more (same group of authors—apparently there'd been some problem with their plane reservations from China), but other than that, the whole thing was a lot of fun.

"Various people were impressed that I knew every single paper so well, but I had to admit to them that it was just the basic academic skill of being able to sound as if you know what you're talking about on short notice. It's odd—I've never had those dreams about taking an exam in a class I haven't ever been to, etc., but here actually I got to live the experience of giving a talk on a subject that I didn't know—and I got to do it three times. That's enough for one entire life...".

The TeachScheme! 2002 participants—high school and college educators—outside the CIT building with Kathi Fisler (center front) and Shriram Krishnamurthi (far right, 2nd row from top), who jointly ran the program

Thomas Hofmann. Thomas and his colleagues received the best-paper award at the European Conference for Machine Learning in Helsinki. "Helsinki in summer?—wonderful! People are out on the street all night, walking, talking, dancing and yes, drinking!" He continues to work as chief scientist for RecomMind, a company he founded in 2000. Thomas attended Snowbird, co-organized the Neural Information Processing Conference in Vancouver and organized a workshop on learning from text and images. In September he gave an invited talk at a workshop on Language Models in Information Retrieval.

Shriram Krishnamurthi. Shriram led another successful TeachScheme! workshop (with Kathi Fisler) during the summer. With Matthias Felleisen, he also inaugurated an advanced TeachScheme! workshop for those who'd survived the first installment. It was even more fun than he had dared hope. Participants culminated the week writing a program that had three computers playing roles from Hamlet (the eponym, Rosencrantz and Guildenstern), trading their lines (from "My honoured lord!" to "Man delights me not...") with one another.

Meanwhile, Shriram served on the program committees for Component-Based Software Engineering, the Workshop on Generative Programming, IBM Watson's PL Day and Programming Language Technologies for XML. He also jointly organized the Functional and Declarative Programming in Education workshop, part of PLI 2002.

He is serving as freshman advisor this year and learning a good deal about Brown in the process. He has also begun tenure as a Faculty Fellow in the Royce Fellowship Program. He is mastering the art of saying no, but clearly not quickly enough!

Finally, provoked and then inspired by the beguiling Eugene Charniak, Shriram is working on a theme involving bonobos. But you'll just have to wait until the next ***conduit!*** to hear more about that.

Nancy Pollard. Nancy received an NSF medium ITR grant on "Providing Intuitive Access to Human Motion Databases" with Jessica Hodgins, Christos Faloutsos and Randy Pausch at CMU and Amy Bruckman at Georgia Tech. She spent two months at ATR in Japan this summer, teaching a humanoid robot to tumble large and heavy boxes (but in a computationally efficient manner and with a guaranteed upper bound on contact forces). This work will appear in the Workshop on the Algorithmic Foundations of Robotics this December and videos can be seen at: *http://www.cs.brown.edu/people/nsp/papers/WAFR02.html*.

Visitors from Industrial Partner Sun Microsystems view an Exploratories project demo by Jean Laleuf, which Sun has been supporting. (l to r) Emil Sarpa, Sun; Jean, Anne Morgan Spalter, CS Graphics; Wendy King, Development; Andy van Dam, and John Hale, Sun

Franco Preparata. Franco spent the first part of his sabbatical leave at the National University of Singapore, where he held the Kwan Im Thong visiting chair professorship. He was also awarded a special mention by the InfoComm Development Authority of Singapore. While in Singapore he lectured in the International Workshop in Computational Biology. He spent the remainder of his leave at the University of Padova. Recently, he participated as a board member in the review of the Mathematics and Computer Science Division of Argonne National Laboratory and presented a lecture at the University of Chcago.

Roberto Tamassia. Roberto served on the program committees of the Symposium on Algorithms and Computation and of the Symposium on Graph Drawing. He gave an invited lecture at the NSF/CBMS Regional Research Conference in Mathematical Sciences on Geometric Graph Theory, at the University of North Texas. His information security project in collaboration with Michael Goodrich and Robert Cohen received a technology transition award from DARPA's Information Assurance Program. His latest textbook, *Algorithm Design: Foundations, Analysis, and Internet Examples* (coauthored by Michael Goodrich and published by Wiley), presents traditional algorithmic topics from a new Internet viewpoint and covers additional algorithms derived from Internet applications

Eli Upfal. Eli's grad student Gopal Pandurangan passed his defense in May and is now an assistant professor at Purdue. Eli participated in a June workshop on random graphs in a wonderful castle in Bertinoro, Italy, He was so impressed with the location (and food) that he agreed to organize another workshop there next summer. As of July 1, Eli replaced Tom Dean as the chair of the Department.

ROBOT 'LEWIS' TAKES PIX

By Bill Smart '02.

Here at WUSTL (Washington University in St. Louis), Cindy Grimm and I are working on a robot (Lewis), a direct descendant of Ramona, that wanders around an environment taking candid, well-composed pictures of the people in it—something like a photographer at a wedding or a conference reception. The idea is that people don't pose for the robot but ignore it as much as possible, so that we can get natural talking-to-your-friends sorts of shots.

TECHNOLOGY

The robot has a video camera mounted on a pan/tilt head. It's constantly classifying the pixels (about 5 frames/sec) as being skin or not skin. It clusters the skin into regions that are candidate faces. It then correlates the laser rangefinder data with these candidate faces. Faces that don't have legs associated with them are thrown out. Skin patches that are the wrong size or are not roughly face-shaped are also discarded. The remaining candidates are considered to be faces and are used to work out a good composition for the picture.

We've encoded some very simple composition rules into the system. For example, if there's one face in the picture, it should be in the middle third of the image horizontally and a little above the center of the image vertically. Using these simple rules and a knowledge of where the faces are in the image and in the world, we can work out a good framing for the shot. We then pan, tilt and zoom to get the correct framing, and take the picture.

Lewis currently uses a Sony DFW-VL500 1394 camera

All of these images are sent over a wireless link to a workstation, where people can browse through them and have them printed out or emailed to themselves.

Bill Smart PhD '02 and Cindy Grimm PhD '96 at their wedding in Brechin, Scotland, earlier this year

SCIENCE

Like most real robot systems, the hard part of making this all work was integrating a number of different systems. None of these are that complex on their own, but getting them all to work together is tough.

Face Finding. There are lots of good face-finding algorithms out there. Unfortunately, many of them are quite slow. Since we're trying to find faces as fast as we can, we need as close to 30 frames/sec as possible. To do this, we combine a very simple solution with other sensors.

It turns out that all skin tones cluster fairly tightly in YUV color space. In fact, they cluster well in the UV components of this space. So we developed a system that classifies pixels as skin or not depending on their U and V components. We have a GUI where you load in a number of example images, and paint the skin in them. These pixels are then converted into a table of U and V values (after a little processing), so that we can classify a pixel with just a table lookup. This is not the best way to classify skin, but it does a reasonable job and is really fast. We can

Who are these CSers?
(answer on back page!)

cover the deficiencies in our approach by fusing in data from other sensors.

Once we have the skin patches in the camera image, we do some simple filtering to turn them into connected regions. The problem with this simple approach is that a lot of things get labeled as skin. Posters, lights, and even carpet can look surprisingly skin-like. To work around this, we look at the readings returned by the laser range-finder. Most faces have legs somewhere below them. If we can't detect a set of legs underneath a skin patch, we throw it out.

If there are legs, then we have some idea how far away the skin blob is likely to be and we can work out how big it is. If it's not face-sized, we throw it out. If it's not face-shaped, we throw it out. If it's not between four and seven feet off the ground, we throw it out. What's left is a surprisingly robust set of detected faces.

Localization. Since the robot odometry isn't very accurate over periods of time, we've developed a simple localization mechanism: a set of painted Chinese lanterns that the robot can turn on and off remotely. When it needs to localize itself, it turns on these lanterns and finds out where they are and how large they are in the image. Then, using some simple trigonometry, it can localize itself to within about 10 cm (over the areas in which we've deployed the system).

We also give the robot a basic map of the world, annotated with the boundaries it should stay within. It then generates fake sensor readings, as if these boundaries were real walls. This makes the obstacle-avoidance and planning problems much easier.

Taking Good Pictures. The small set of basic photography rules we've built into the system seems to yield pretty good pictures. And we've noticed that people tend to be quite forgiving of the pictures we actually end up taking. We've also noticed that, in some situations, it's hard not to take a good picture: with a lot of people in the room, a lot of "bad" pictures (where the robot has gotten something wrong) still end up being good, since they catch some incidental action in the background.

HISTORY

The original purpose of this project was to get the undergrads in the lab excited about playing with robots. After working on it for a while, we decided to submit it to SIGGRAPH's Emerging Technologies program and take it to San Antonio last July. Two undergrads worked on it through the spring semester and during the summer before the conference to turn it from vaporware into an actual working system.

Once we got to San Antonio, it took us two or three days to get the lighting conditions to the point that all of the code did the right thing—we finally got that under control about an hour before the Emerging Technologies exhibit opened for the first day. We ran the robot about nine hours a day for four days, and another four on the fifth day, more or less continuously, with only a ten-minute stop for a change of batteries once a day. During all this time (40 hours!) the robot never collided with anything or anyone, and the code ran perfectly (with one minor exception). We guess that about 5,000 people passed through the exhibit, and we have a collection of about 2,000 pictures that the system considered good enough to keep.

Unlike anything either of us has done in the past, this project seems very press-friendly. We made it onto the BBC World Service, Tech TV, slashdot.org, and a bunch of other web sites. We were even interviewed by Japan's NHK, but were eventually left on the cutting-room floor.

Another two engagements are booked for the robot. One is a science writers' convention at Washington University, where we'll be taking pictures at the reception. The other is the wedding of one of our system guys. Yes, he knows what he's letting himself in for. Yes, so does his wife-to-be!

Public Enema, the victors! (top row, l to r), Cory Redd, Rob Hunter, Sam Heath*, Danny Acevedo*, Dave Tucker*, Don Carney*, Mike Harrison, Josh Reineke. Bottom row (l to r), Pete Morello, Dan Keefe*, Joe LaViola*, Andrew Huebner. (Not present: Russ Bent*, Sam Brenner) Names with an asterisk are CSers*

THE FUTURE

Now that the basic system is working, we'd like to push the actual science behind it in a number of directions:

Face Finding. Right now people in red shirts are sufficiently like skin in UV space that they get classified as such. We're currently working on techniques to segment out skin-like clothing from skin and get a more accurate face detector.

Another problem stems from using the laser range-finder to infer distance to a face. All we have is the distance to (what is probably) their feet. We have to assume that they're standing upright and that their face is about as far from us as their feet. We're currently working on adding stereo vision to the system. This is a pretty constrained problem, since we can do correspondence between the skin blobs and get a depth estimate from that.

Navigation. As always, we're working on more intelligent navigation problems. We typically rate a group of candidate shots according to how good we think they'll be. The goal is to plan a path to as many candidates as possible quickly enough that the opportunities don't vanish before we get there, and still be reactive enough to take opportunistic shots on the way. We think this can be modeled using Markov decision processes, and have started looking at getting this to work on the robot in a realistic amount of time.

Interaction. People like having a model of the internal state of things, especially of "intelligent" things. How can we give people a clue about the internal state of the robot? Will this make them more comfortable around it? Can we use this to persuade them to do things for us? For example, it's been shown that people are more likely to comply with a request if the requesting voice sounds annoyed. How can we use human-human interaction ideas like the notion of personal space in path-planning?

Composition. What other rules should we be using to compose good pictures? Can we create rules, based on human expert guidance, for what constitutes a pleasing picture?

People: Bill Smart, Cindy Grimm, Zach Byers, Michael Dixon, Kevin Goodier, Patrick Vaillancourt, Jacob Cynamon, Michal Bryc, Hui Zhang.

For Lewis's project webpage see: *http://www.sc.wustl.edu/~lewis.*

CHARNIAK UNPLUGGED

You may recall that in my last column I discussed the gun laws in Utah and how they are affecting the University there, as mentioned in the NYT, how I could not resist kidding a Utah colleague of mine about what sort of gun she takes to class, how she answered that another colleague of hers received email asking advice on what sort of gun he should take to Utah on his next research visit, and how all of this made me realize that both I and the unidentified colleague were really just EECLs (effete east-coast liberals) toeing the party line. Well, within a day of my column's publication a third EECL fessed up:

From: Shriram Krishnamurthi <sk@cs.brown.edu>

To: ec@cs.brown.edu

CC: kha@cs.brown.edu, sjh@cs.brown.edu

Subject: EECL

Date: Tue, 14 May 2002 14:16:08 -0400 (EDT)

The person you refer to in your Unplugged article *may* be me -- when I saw the news article, I dashed off a message to my own friend on the Utah faculty.

Of course, I suspect *everyone* on the Utah faculty was getting bombarded by their EECL acquaintances. I imagine that at their department retreat, they award a secret prize to the author of the stupidest comment they got from the east coast (and end up with a three-way tie).

Natteringly and nabobishly yours, Shriram

What really impressed me, though, was that Shriram's closing identified my "EECL" term as harking back to Spiro Agnew and his "nattering nabobs of negativism". As I figured it, Shriram must have been about -2 when Spiro Agnew resigned. When I asked him, Shriram said he was actually +2, but that he reads a lot of political magazines.

This past year I was program co-chair of the 40th annual meeting of the Association for Computational Linguistics, the major conference in my research area, held this year at the University of Pennsylvania. One of my duties was to write the program chair's preface to the proceedings. Having never read a preface to any of the many conferences I have attended, I went back and looked them over, only to discover that I had had good reason to give them the skip. I decided then to write the preface not in my (somewhat) formal academic style but rather with my "unplugged" hat on. In the preface I tweaked the executive committee of the organization (and the overall conference chair) for taking a week to decide that the official acronym for the conference was ACL-02 and not ACL'02, or ACL02, or ACL2002, or ACL 2002, etc. I also figured that nobody would read this preface either, so in the penultimate line I offered $5.00 to the first person who did (and told me so).

In my delight at my cleverness I forgot that the overall conference chair had to approve the thing. I have occasionally observed a lack of humor in some of my academic colleagues, but fortunately the chair, Pierre Isabelle, a noted European computational linguist, has a very robust sense of humor—or at any rate, he liked my preface. I was also pleased that two days before the conference officially started (during a pre-conference workshop) someone came to say that she had read my line, and was she first? I pulled out my wallet but she refused the money, telling me instead to give it to some of the homeless around the University of Pennsylvania campus, which I did.

Another conference task was inviting two keynote speakers. One person high on my list was Peter Norvig (Brown '78), who is now 'director of search quality' at Google. While I have known Peter for many years now, he has been at Google for only the last year or two and I did not have his email address. Naturally, I went to Google and typed in 'Peter Norvig Google'. In doing this I was reminded of the halting problem and humorously thought about Google crashing. (If this does not make any sense to you, review your CS51 notes.) I also mentioned this to Peter in my subsequent mail to him. His response:

> From: Peter Norvig <pnorvig@google.com>
>
> To: Eugene Charniak <ec@bohr.cs.brown.edu>
>
> Subject: Re:
>
> Date: Fri, 26 Apr 2002 15:09:21 -0700
>
> Yes, I would like to give a talk at ACL. I'm confident I can find enough interesting things to say without giving away Google secrets. We can talk more about the contents of the talk in a few weeks.
>
> You mentioned the circularity of searching for "Peter Norvig Google" on Google. Here's another example: you may know that we cache web pages, and offer them up for use when the original server is down or slow. We preface such pages with the disclaimer "Google is not affiliated with the authors of this page nor responsible for its content." Several people have pointed out that if you follow the cached link for Google.com, you still get this disclaimer. Our response is that rather than changing this behavior, we instead write back to people who complain and say "you obviously understand the issues of recursion and self-reference; are you interested in applying for a job at Google?".
>
> Best, -Peter

At the same time, a very cute paper was accepted for presentation at the conference that related to search engines. One problem in computational linguistics is to recognize the difference between 'simple' and 'simplistic', which mean sort of the same thing but with very different connotations. The basic idea of this paper was that if an article uses the word 'simplistic' it is more likely to be describing something as 'bad' than as 'good'; thus Google searches on 'simplistic bad' should get more hits than 'simplistic good', or some such (I have forgotten the exact ideas and I refuse to do research to write this column, so just remember that PERSONS attempting to find accuracy in this narrative will be prosecuted; persons attempting to find a moral in it will be banished; persons attempting to find unattributed quotations in it will be shot. BY ORDER OF THE AUTHOR).

At any rate, I mentioned this idea to my colleague Mark Johnson (a professor in Cognitive Science with a joint appointment in CS), and he suggested that this idea has wider applications. In particular, we could dispense with all the complications of tenure decisions at Brown. Rather than sending out letters asking if, say, Eugene Charniak should be given tenure, we just type in 'Eugene Charniak stupid' and 'Eugene Charniak smart' into Google, and see which gets more hits. I just tried this and was relieved to find that the latter outnumbered the former, 394 to 98. Less good is that I tried the 'simplistic' example and got the WRONG result, 271,000 to 170,000. Probably the fact that I tried this at all says something more relevant about my smarts.

Lastly, walking by Tom Doeppner's office recently I noticed that he has a collection of (well, four) ceramic elephants. For a long time he had had one such elephant, and Tom explained to me how he had come by the other three. The stories were interesting, but nothing compared to how he got the first one. This story dates back to pre-***conduit!*** times, but is such a classic that it deserves to be told. One day Tom received a package, clearly ad-

l to r: Trina Avery, Andy van Dam and Tom Doeppner, 50ft down off the coast of Bonaire in the Caribbean

dressed to him, in which there was a very large ceramic elephant. The return address was Digital Equipment Corporation, or some such—at any rate, some address completely unrelated to the contents. Not having any idea of who sent it to him or why but liking the elephant very much, Tom put it in a prominent place in his office.

A little while later, Roger Simon wandered by Tom's office. Roger has his PhD in mathematics and a Master's in CS from Brown, and is a professor of CS at Rhode Island College. More to the point, Roger's wife is a ceramist and one of her specialties is, you guessed it, ceramic elephants. His wife had just had a major disaster when one of her large elephants was not delivered to the person who had commissioned it, and Roger was dumbfounded to see it in Tom's office.

The explanation turned out to be quite simple. Roger had found a box in the department trash and had taken it home for his wife to use. She pasted over the address on the box (which happened to be Tom's) with a new label, but, it seems, the new label came off, so it was delivered to Tom! Tom, of course, returned the elephant, but got a smaller, but still quite impressive, elephant as a make-up present.

THIS JUST IN ...

Takeo Igarashi, who was here as a postdoc with the graphics group, now at Univ. of Tokyo, reports from UIST 2002 in Paris that his paper "Clothing Manipulation" (co-authored with John Hughes) was selected for the Best Paper award!

conduit!

A publication of
The Computer Science Department
Brown University

Inquiries to: ***conduit!***
Department of Computer Science
Box 1910, Brown University
Providence, RI 02912
FAX: 401-863-7657
PHONE: 401-863-7610
EMAIL: sjh@cs.brown.edu
WWW: http://www.cs.brown.edu/publications/conduit/

Suzi Howe
Editor-in-Chief

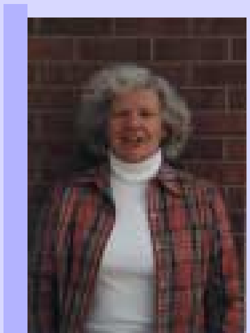

Katrina Avery
Editor

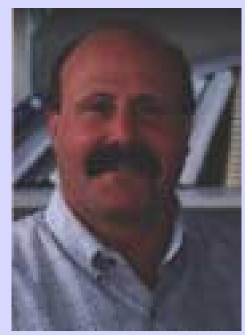

Jeff Coady
Technical Support

Department of Computer Science
Brown University
Box 1910, Providence, RI 02912, USA

conduit!

NON-PROFIT
U.S. Postage
PAID
Providence, RI
Permit #202

Printed on recyled paper

Address changes welcomed

conduit!

A publication of
The Computer Science Department
Brown University

Inquiries to: ***conduit!***
Department of Computer Science
Box 1910, Brown University
Providence, RI 02912
FAX: 401-863-7657
PHONE: 401-863-7610
EMAIL: sjh@cs.brown.edu
WWW: http://www.cs.brown.edu/publications/conduit/

Suzi Howe
Editor-in-Chief

Katrina Avery
Editor

Jeff Coady
Technical Support

Department of Computer Science
Brown University
Box 1910, Providence, RI 02912, USA

conduit!

NON-PROFIT
U.S. Postage
PAID
Providence, RI
Permit #202

Printed on recyled paper

Address changes welcomed