conduit!

Volume 8, Number 1 **Department of Computer Science** **Spring, 1999**
Brown University

GRAPH DRAWING and INFORMATION VISUALIZATION

Roberto Tamassia attending a workshop on graph drawing in Barbados

Information visualization is an emerging discipline that addresses the problem of communicating the structure of information space through diagrams. Quoting from an article by David Harel (*Comm. ACM* 31(5) 1988):

> "The intricate nature of a variety of computer-related systems and situations can, and in our opinion should, be represented via visual formalisms; visual because they are to be generated, comprehended, and communicated by humans; and formal, because they are to be manipulated, maintained, and analyzed by computers."

The benefits of information visualization include analysis through visual exploration, discovery of patterns and correlations, and abstraction and summarization. It is anticipated that in the next decade the use of information visualization techniques will be essential to the success of portals and other large information repositories on the Internet.

> *"...in the next decade the use of information visualization techniques will be essential to the success of portals and other large information repositories on the Internet"*

Traditionally, information visualization has focused on the display of ***quantitative information*** (e.g., bar charts, pie charts, function plots) and ***geographic information*** (e.g., road and subway maps), where a natural mapping exists between the data and their geometric location in the diagram. A great introduction to these types of visualizations is given in the books by Edward Tufte (*The Visual Display of Quantitative Information* and *Envisioning Information*).

From quantitative to relational

More recently, researchers have started addressing the problem of visualizing ***relational information***, where ***networks*** (also known as ***graphs***) model collections of objects and connections between those objects. Examples include:

- ***Web***: site maps, browsing history diagrams, presentation and refinement of query results, product catalogs.
- ***Software engineering***: UML class and state-transition diagrams, subroutine-call graphs, data-flow diagrams.
- ***Database systems***: entity-relationship diagrams.

Brown University, Box 1910, Providence, RI 02912, USA

- ***Real-time systems***: Petri nets and state-transition diagrams.
- ***Networking***: LAN diagrams.
- ***Enterprise and project management***: business process diagrams, organization charts, scheduling charts.
- ***Engineering***: circuit schematics.
- ***Artificial intelligence***: knowledge representation diagrams, belief and influence networks.

The fundamental problem in the visualization of relational information is the automatic layout of networks, which is the subject of the research area known as ***graph drawing***. Several pioneering commercial applications have begun to appear. For example, the AltaVista search engine by DEC/COMPAQ can visualize the results of a query with an automatically generated drawing of a graph whose vertices are relevant keywords, and supports the refinement of the query by a direct manipulation of the drawing. The Hyperbolic Tree (TM) technology for drawing trees by Inxight Software (part of the Xerox New Enterprise business initiative) is used in the Web

AltaVista Web search display

site of the **Wall Street Journal** (go to the "Money Tree") and is incorporated in a Web management tool by Microsoft. There is significant potential for enhancing electronic commerce Web sites with graph-drawing technology.

While the task of automatically producing a readable layout for a graph may appear simple to a nonexpert, it is actually computationally very hard. The cost of incorporating effective automatic network-layout capabilities into software systems is often grossly underestimated.

To complicate matters, some basic graph drawing problems for which theoretically fast algorithms are known turn out to be unwieldy to implement. Take, for instance, the problem of testing whether a graph is ***planar***, i.e., whether it can be drawn without crossings. While mathematical characterizations of planar graphs have been known since the 18th century, it was only in 1974 that John Hopcroft and Robert Tarjan published the first linear-time algorithm to test whether a graph is planar (*J. ACM* 21(4)). This algorithm was a major theoretical accomplishment and greatly contributed to their earning the prestigious Turing Award.

Coming up with a correct implementation of the algorithm was, however, a different matter. The intrinsic conceptual difficulty of the approach combined with data-structuring tricks and special cases defeated the efforts of many programmers for more than 20 years until finally in 1996 a research team led by Kurt Mehlhorn completed the development of a reliable implementation of the Hopcroft-Tarjan planarity testing algorithm that has been successfully tested on tens of thousands of graphs.

Graph drawing research at Brown

Since my first paper on automatic layout of entity-relationship diagrams was published in 1983, graph drawing has been one of my main research interests. Three of my six doctoral students to date have done their research on graph drawing:

Bob Cohen

Ashim Garg

Robert Cohen (Ph.D. 1992, now at Algomagic Technologies, Inc.), Ashim Garg (Ph.D. 1995, now at SUNY Buffalo), and Stina Bridgeman (current; a piece on her research appears in this issue).

My recent work on graph drawing has focused on:

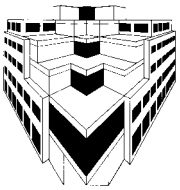

- interactive layout techniques
- Web-based graph-drawing systems
- software design patterns for graph drawing

Graph drawing on the Web

The Graph Drawing Server is a Web-based system that provides graph drawing services. It can be accessed in various ways:

- through an interactive graph editor implemented as a Java applet
- through an HTML form that allows the definition of the input graph in a variety of formats
- through a Java package that provides an API for involving the server

It supports various layout styles and drawing algorithms, including orthogonal layouts computed by the GIOTTO algorithm and hierarchical layouts computed by an algorithm that distributes vertices on horizontal layers, and has been successfully used by many researchers worldwide to experiment with graph-layout techniques.

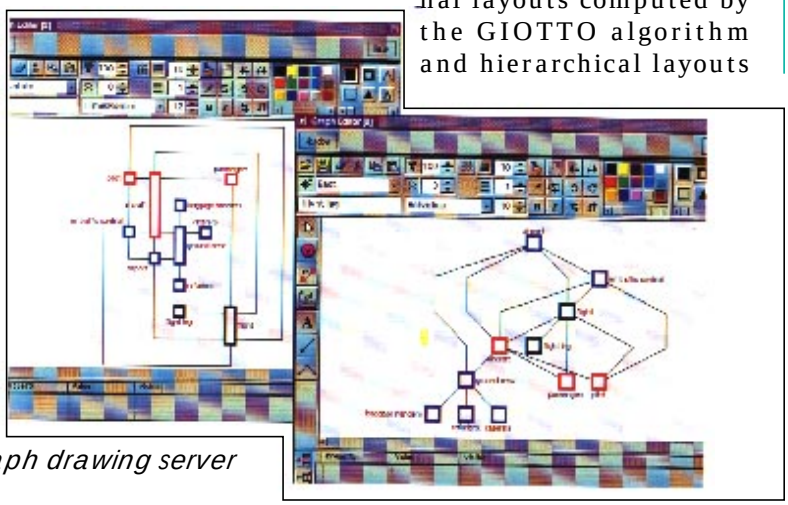

Graph drawing server

Graph drawing in Java

I am developing a library of reusable software components for graph drawing in Java. This work is based on algorithm engineering techniques that include the use of novel algorithmic patterns. The library can be used to incorporate automatic layout capability in various interfaces that make use of diagrams. I plan to show its applications especially to Web browsers and programming environments.

In addition to algorithmic and systems research, my future graph-drawing work includes the development and commercialization of a package of Java software components for graph layout in collaboration with Algomagic Technologies, a recent startup founded by Robert Cohen, Michael Goodrich from Johns Hopkins University and myself.

A book on graph drawing

I have recently published a book on the subject in collaboration with three other graph-drawing gurus: Giuseppe Di Battista from the University of Rome, Italy, Peter Eades from the University of Newcastle, Australia, and Ioannis Tollis from the University of Texas at Dallas. The rigorous treatment of the subject allows this book to be used as a text for graduate courses and as a reference for researchers, while the large number of examples and figures makes it also suitable for software practitioners.

Graph-drawing gurus. l to r: Ioannis Tollis, Roberto, Giuseppe Di Battista and Peter Eades

The latest issue of **Dr. Dobb's Journal** (June 1999, page 134) mentions the book as follows:

> "The final book this month is *GRAPH DRAWING: Algorithms for the Visualization of Graphs*, by Giuseppe Di Batista, Peter Eades, Roberto Tamassia, Ioannis G. Tollis. The title is an accurate summary of the book's contents, but doesn't do justice to its breadth. Section 5.1, for example, is devoted to angles in orthogonal drawings, while chapter 7 covers incremental construction technniques. The style is academic—there are a lot of references, and a lot of proofs and lemmas—but the book will be a rich mine of ideas for anyone who is trying to persuade a computer to turn data into dots, boxes, lines and arrows."

"HORIZON"

This imposing iron and azure glass monument by artist Costas Varotsos was commissioned by General and Mrs. Kanellakis in memory of their son Paris and his family. With majestic Mt. Parnassos as a backdrop and the town of Liya below, the sculpture is located on family-owned land—a favorite summer haunt where they'd gather each year to take family snapshots. The inscription reads:

DEDICATED TO OUR CHILDREN
PARIS - MATE - ALEXANDRA - STEPHANOS
THEIR PARENTS ELEFTHERIOS AND ROULA KANELLAKIS
20-12-95

Paris's legacy will endure via fellowships and awards established to honor his memory—the ACM's Kanellakis Award, Brown's Kanellakis Graduate Fellowship and MIT's Kanellakis Graduate Fellowship, to name a few. His parents have donated the land around the monument, aptly, to the SOS Children's Village International Parc. Said his mother, "It was their favorite summer place on earth and now their spirits shall dwell there—near to us—while those we cherish are young and together in eternity."

THE UNIQUENESS OF CS92

Roger Blumberg

The failure of contemporary technologies significantly to transform, much less improve, elementary and secondary education in the 20th century is by now well-known. Even if radio, television, film, and now computers have all contributed a great deal to changing what we teach and learn, they thus far have had remarkably little effect on *how* we teach and learn it in school.

In his insightful book, *Teachers and Machines* (Teachers College Press, 1986), the historian Larry Cuban presents data suggesting that, among other things, the failure of classroom technologies in American schools in this century has had much to do with a top-down method of implementation and promotion which has kept school teachers from using the technologies effectively and/or on a large scale. He writes that "claims predicting extraordinary changes in teacher practice and student learning, mixed with promotional tactics, dominated the literature in the initial wave of enthusiasm for each new technology. Seldom were these innovations initiated by teachers" (p. 4).

Cuban also described a cycle that has characterized the fate of classroom technology since 1920: exhilaration, followed by scientific credibility, then disappointment, and finally teacher-bashing. First published in 1986, Cuban's book had little to say, and no interesting data, about the use of computers, but anyone following even newspaper articles about computers in schools over the past few months will realize that all the stages of Cuban's cycle can already be identified in the short history of computers as a species of classroom technology.

In its preference for top-down design and promotion of educational technology, a good deal of university research has only contributed to the unfortunate story of

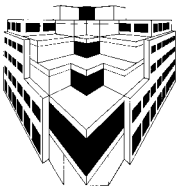

computers in school. The typical university-school collaboration model in computer science has been that a group of university researchers develops a tool, or set of tools and curricula, based on their own expert intuitions of good design, instructional need, and educational value, and then attempts to recommend the tool, and the vision implicit in its design, to working teachers. This model has met with remarkably little success, considering the tremendous expertise (technical and pedagogical) and resources (public and private) that have been devoted to it.

In light of the history of classroom technology, the difference in the approach to the study and use of computers at school embodied in The Educational Software Seminar at Brown (CS92/ED89) is clear. The Seminar begins not with a product, deduced from first principles, but with proposals gathered from working teachers. These teachers specify instructional objectives and technical constraints, as well as their speculations and aspirations concerning how the computer might improve student learning and their own teaching.

"The Seminar begins not with a product, deduced from first principles, but with proposals gathered from working teachers"

In the spring semester, Brown undergraduates enrolled in the Seminar read these proposals—which are solicited, gathered and negotiated in the fall—and choose those in which they are most interested. Then, working in teams of three or four, students in CS92 design, create, test and implement the programs proposed. In the process of building the software, the teams work closely with their sponsoring teacher and her/his students, and the production work (which includes the creation and presentation of storyboards as well as prototypes) is done in parallel with seminar discussions of literature drawn from computer science and the history and philosophy of education, as well as the cognitive and social sciences. As remarkable as the products of the Seminar have been—a phenomenon due in no small part to the excellent technical preparation the students receive in the CS Department—it is the process that, in my view, makes the Seminar a valuable, liberal arts (i.e. non-technical) undergraduate experience.

The Educational Software Seminar has been in existence for more than a decade, and it owes that existence to Andy van Dam. Orchestrating and teaching the Seminar with David Niguidula in the early and middle 1990s, when students built software in Hypercard for teachers in schools with Macintosh computers (small in size and number), Andy developed a curriculum for the course and his students produced a set of programs that remain interesting and useful in the study of educational software. Indeed, many of those programs are still popular with teachers and students—a great tribute both to Andy's and David's idea of the course and to how much their students were able to do even with very primitive authoring tools.

Last year I began teaching the Seminar, and decided first to expand the pool of possible projects to include proposals not only from teachers in K-12, but from Brown University faculty and educators in other community institutions as well. The range of possible platforms was similarly expanded to include not only the PC but the Web as well. With the kind and indispensable cooperation of the Multimedia Lab at Brown, CS92 students were able to choose from a variety of multimedia authoring tools, in addition to the programming languages available and taught in CS, and thus could design for a greater number of platforms, instructional goals, and technical constraints than in previous years. Finally, the readings for the course were revised to incorporate recent studies like Cuban's and the 1997 analysis of the Apple Classrooms of Tomorrow Project (Judith Sandholtz *et al.*, *Teaching with Technology* [Teachers College Press, 1997]), and to better reflect the interdisciplinary nature of the Seminar.

For all the changes, however, Andy's and David's idea of the Seminar remains, and this is what makes CS92 such a successful course and worthwhile experience for students as well as our clients. Last year

our students worked with elementary, secondary, university and community teachers, creating programs for settings as diverse as the Blessed Sacrament School (with Macintosh computers equipped with only 4MB of RAM) and the Brown School of Medicine; covering subjects as different as 4th grade social studies and undergraduate visual art; using tools as various as Hyperstudio and Java; and serving audiences as diverse as 7th grade science students at the Wheeler School and visitors to the City Streets exhibit at the Providence Children's Museum. All of the programs were completed and continue to be used, not only by their sponsoring teachers but by teachers and students who have downloaded them from the course Web site as well.

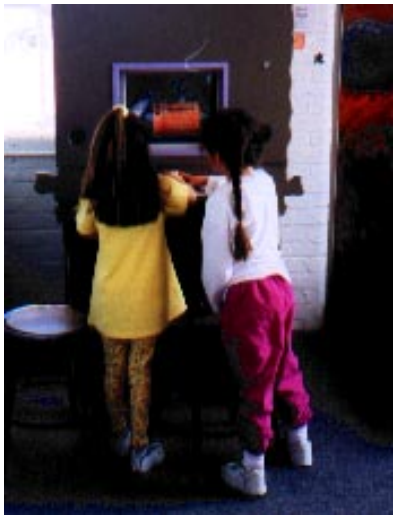

Visitors to the City Streets exhibit at the Providence Children's Museum

This year, we are again working with the Vartan Gregorian Elementary School and both undergraduate and post-graduate teachers at Brown, and we are excited to be working with the elementary school classroom at Hasbro Children's Hospital for the first time. The students in the Seminar are working with teachers and their students, documenting their activities on project pages, thinking carefully about pedagogy, learning and the design of electronic media, and are again doing remarkable work. This year I've been especially struck by the extent to which computer science students, some of them reading philosophy, cognitive science, and/or social science for the first time at Brown in CS92, are able to arrive at important, novel insights about the connections among questions concerning technology, pedagogy and cognition.

"...computer science students...are able to arrive at important, novel insights about the connections among questions concerning technology, pedagogy and cognition"

In his important book, *The Language of Education* (Charles Thomas, 1960), the philosopher Israel Scheffler wrote that "educational research must not be conceived as a single science, but rather as the common focus of many sciences with bearings on educational practice" (p. 73). While I won't raise local blood-pressures by making a similar claim about computer science research, the experiences of the students in CS92 clearly demonstrate how success in the creation of educational software necessarily involves ideas from the computer, cognitive and human sciences, as well as creative insights not clearly derivable from the first principles of any (known) science. This sort of non-algorithmic complexity makes the creation of genuinely effective educational software extremely difficult, and indeed the computer industry has largely retreated from education to entertainment and from the classroom to the home; but, of course, it is precisely this sort of complexity that makes educational software, and educational technology in general, such a challenging but potentially rewarding area of research.

The Educational Software Seminar provides not only a unique, productive model of university-school and university-community collaboration, but an important method of inquiry into the relation between technology and education as well. CS92 projects and programs, past and present, along with team project pages and our syllabus and bibliography, can all be found at the course Web site: *http://www.cs.brown.edu/courses/cs092/*.

Symposium speakers from l to r: Maurice Herlihy, Clare Rabinow, Lotus; Bob Morgan and Bill McKeeman, Compaq; Nikos Aneuris, AT&T Research; Ken Arnold, Sun; Bill Kayser, WorldStreet

THE 22nd IPP SYMPOSIUM

Host: Maurice Herlihy

The 22nd Industrial Partners Program technical symposium, held on November 12, 1998, was on "Realizing the Potential of Java," a topic selected for its relevance to the current interests of many Partners. The speakers included representatives from Partners Sun Microsystems, Lotus/IBM, and Compaq, as well as AT&T Research and WorldStreet. The topics covered fell into two broad categories: technology to support Java (Sun's Jini and Compaq's ahead-of-time compiler) and experience using Java for critical enterprise-wide applications (Jini, MARVEL, and eSuite).

Ken Arnold (Sun Microsystems) discussed the newly-released Jini system. Jini is described as "network plug-and-work": it is intended to ease the distinction between hardware and software by allowing spontaneous networking (connecting components on the fly using a simple common interface). Java itself provides a homogeneous network in the form of safe, portable object code and a single type system for the entire network. Jini extends this model by permitting a service provider to advertise its service across the network by registering itself with a Jini Lookup Service. A potential client can locate a server by matching a Java type, possibly augmented with additional attributes. Code moves from the service to the Lookup Service to the client. Code needed to use the service is dynamically loaded on demand from the client. Jini makes no distinction between hardware and software. Services can be local, remote, or a combination. It accommodates legacy services and languages. Jini is not an operating system, but eventually could replace the operating system with a ubiquitous and invisible network.

Clare Rabinow (Lotus), a Brown alumna (class of '72), described her group's experience implementing eSuite, a business productivity application that encompasses components for data presentation and editing and data access, as well as traditional business productivity applications for the network computer. The talk highlighted both the advantages and the pitfalls of using Java for an enterprise-level application. The principal advantages include the experience that it really is feasible to use Java to develop software for dissimilar platforms; it is relatively easy for customers to mix components from different vendors; and businesses can build applications that benefit from platform inde-

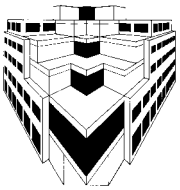

pendence and low-cost deployment without ever learning to program in Java. Among the pitfalls, Rabinow observed that every Java Virtual Machine has its own bugs and quirks, and every runtime environment its own schemes for security and for locating language resources, and that Browser/Java Virtual Machine deployment lags behind availability. The principal performance bottleneck turned out not to be execution speed, which is more than adequate for interactive applications, but delays in launching applications and downloading classes. The absence of pointers and explicit memory deallocation do eliminate many traditional sources of memory leaks, but leaks still occur in the form of references "buried" within tables and Abstract Window Kit peer resources.

Nikos Anerousis (AT&T Research) described MARVEL, an object-oriented Java-based toolkit for deploying scalable network management services. The architecture has two core components: an information model that allows network management information to be aggregated in spatial, temporal, and functional forms, and a presentation model that interacts with Web browser clients to launch Java applets to perform data display. Marvel is based on a client-server structure. Clients are simple Web browsers and servers are extensible agents that employ protocol adapters to interact with components using both standard and proprietary management protocols.

Bill Kayser (WorldStreet) described his company's experience building a productivity system providing real-time market information for securities trading. The tool described integrates real-time market status, a chat-style applet for group communication, and integrated access to database and enterprise information. Java provided substantial advantages in terms of programmer productivity, particularly due to exception handling, garbage collection, and the rich collection of class libraries. Experience with performance was mixed: WorldStreet encountered a wide variability in Java virtual machine implementations, and garbage collection introduced some scalability problems. The principal hazards encountered included spaghetti code introduced by careless use of exception handling, thread-related deadlock and performance problems, and unexpected behavior of static initializers. Development suffered from the lack of several kinds of tools: debugging, platform-specific development tools, and configuration management.

Robert Morgan and Bill McKeeman (Compaq) reported on Compaq's optimizing "ahead-of-time" Java compiler. They observed that generating good code for advanced processors is expensive, since one must deal with software pipelining and instruction scheduling, reg-

"I think that I shall never see
Such tortes as those for IPP.
Naught compares to cakes like that, oh
How I long for chocolate gateau!"

...eat your heart out, Melvin Belli

Besides distinguished speakers, IPP symposia are noted for an abundance of lavish desserts

ister allocation, loop transformations, and a host of similar optimizations critical for exploiting the power of modern processor architectures. These kinds of optimizations cannot be done effectively by a just-in-time (JIT) compiler because they take too much time. An "ahead-of-time" compiler requires advanced analysis to generate good code. Challenges include computing concrete types, optimizing dynamic method calls, placing data objects on the stack whenever possible, eliminating null reference and array bounds checks, and analyzing aliasing among references, as well as classical optimizations. They described a compiler structure that starts with Java byte codes and successively either computes information for later phases or moves the program closer to machine instructions. The final result is a "native" machine-code representation of the program, encompassing a level of optimization beyond what a just-in-time compiler could provide.

This IPP symposium gave an rare snapshot of how well Java is realizing (and sometimes not completely realizing) its potential to transform the computing world in essential ways.

INTERACTIVE GRAPH DRAWING and SIMILARITY METRICS

Stina Bridgeman

Being able to draw nice pictures of graph-structured information is a useful thing—data structure visualization, visual programming, database design, Web navigation, and a variety of other applications can benefit from clear and understandable drawings of graphs. A graph-drawing algorithm typically attempts to optimize certain aesthetic criteria while working within a given drawing paradigm. So, for example, an algorithm may attempt to minimize the number of bends in an orthogonal drawing, where vertices are placed at grid points and edges are chains of horizontal and vertical segments. There are, of course, tradeoffs inherent in this—requiring a straight-line drawing may increase the area requirements, minimizing the area may increase the number of edge bends, and so forth.

Interactive Graph Drawing

A great deal of work has been put into developing graph drawing algorithms and one can find an algorithm that does a good job of layout for a variety of applications. Many of these algorithms, however, were developed using a batch model in which the graph is redrawn from scratch each time. Such algorithms are not well suited for interactive applications, where the typical scenario is that the user requests an initial drawing of the graph, then makes some small changes and requests a new drawing, and then makes another small change or two and requests a third drawing, and then makes yet another change... These changes may be made because the user is explicitly editing the graph or because the graph represents some structure that is being updated, such as a map showing the user's navigation through a collection of Web pages or topics returned by a search engine.

In this interactive model, it is assumed that the user has gained some knowledge—built up a "mental map"—of the previous drawing, and so it is desirable to preserve the mental map by having small changes in the graph structure translate into small changes in the drawing, even at the expense of the other

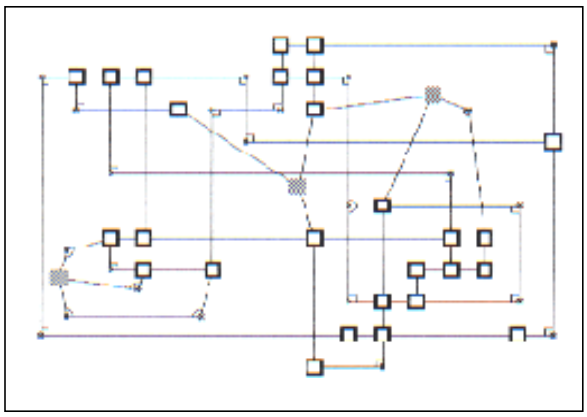

An orthogonal drawing, with some user changes

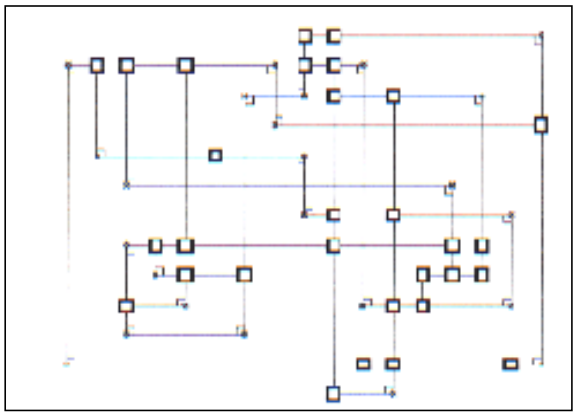

New drawing produced by InteractiveGiotto

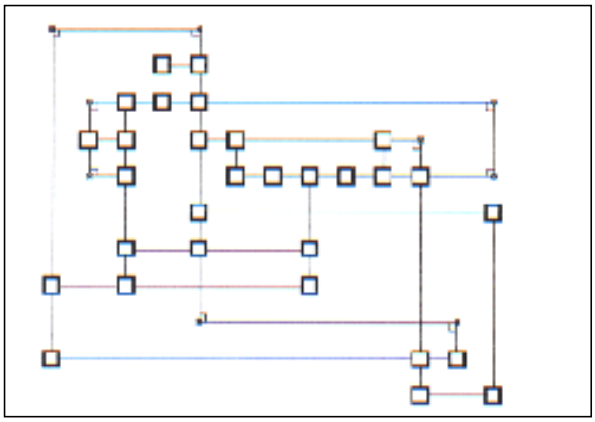

New drawing produced by Giotto

aesthetic criteria. Redrawing the graph from scratch at each step often causes large changes in the drawing, destroying the user's mental map and forcing her to spend considerable time refamiliarizing herself with the layout—imagine how disorienting it would be if the map of the Web pages you've explored were completely redrawn each time you visited a new page.

Of course, eventually the "at the expense of the other aesthetic criteria" bit adds up and it may be desirable to redraw the graph from scratch periodically, but a good deal of effort on the user's part can still be saved when this is done only periodically and not after every change.

Similarity

The goal of preserving the user's mental map brings up the idea of similarity—if the new drawing looks similar to the old, the user's familiarity with the old drawing will transfer to the new drawing with minimal effort. The question is, how does one measure similarity between drawings of graphs? Interactive drawing algorithms frequently seek to preserve the mental map by attempting to minimize the change between drawings, typically by allowing only very limited modifications (if any) to the position of vertices and edge bends in the existing drawing. Some layout adjustment algorithms (where the idea is to rearrange a drawing in order to improve some aesthetic criteria) use a notion of proximity to preserve the mental map, by requiring that a point's position in the new drawing be closer to its own old position than any other point. Orthogonal ordering has also been proposed as a factor in similarity—if the relative north/south/east/west relationship between pairs of vertices is preserved, the new drawing will tend to look more like the original than if these relationships are changed.

The goal of my current research is to define and evaluate similarity metrics for graph drawings, and in so doing, gain some insight into what features of the drawing are most important for visual similarity. Such insight will be useful in designing new interactive algorithms that better optimize the desired aesthetic criteria while preserving the mental map. Developing similarity metrics will also provide a basis of comparison between different drawing algorithms—currently the comparison is still done in terms of "traditional" optimization criteria like the area and the number of edge bends, and discussion of mental map preservation is limited to stating that the algorithm does a good job because it doesn't change the drawing much.

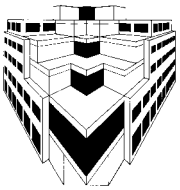

My focus has been on measuring similarity between two drawings of the same graph; determining how similar two different graphs are is another problem entirely, and one that is less immediately applicable to interactive graph drawing. The initial step has been to define metrics based on mental map models that have been proposed in the literature—vertex position, proximity, orthogonal ordering and edge shape. (For orthogonal drawings, edge shape is the sequence of left and right turns made along each edge.)

A successful similarity metric should pass three tests. It should *qualitatively* reflect human judgment so that ordering a set of new drawings with respect to the original based on the metric's value yields the same ordering a human would produce when asked to arrange the same drawings. It should also *quantitatively* reflect human judgment so that the value of metric is proportional to the perceived difference in similarity. And finally, it should *choose the correct rotation*; namely, given the original drawing and a set of new drawings identical to each other except for a rotation factor, the metric's value should be the lowest for the rotation that a human would say best matches the original drawing. This last condition is the easiest to satisfy, and was motivated by the "rotation problem" of InteractiveGiotto.

"A metric that chooses the correct rotation can be used to fix InteractiveGiotto's output, which is important because a large rotation is a very easy way to destroy the user's mental map"

InteractiveGiotto, largely developed by Jody Fanto '97, is essentially a front-end to the successful (batch-model) orthogonal drawing algorithm, Giotto. It seeks to preserve the look of the drawing by applying constraints fixing the angle between adjacent edges around vertices and the number and direction (left or right) of bends along edges. The rotation problem comes up because InteractiveGiotto throws away coordinate information when computing the new drawing, and thus may produce a drawing that is rotated by a multiple of 90 degrees with respect to the original. A metric that chooses the correct rotation can be used to fix InteractiveGiotto's output, which is important because a large rotation is a very easy way to destroy the user's mental map.

To evaluate the metrics it is necessary to obtain multiple drawings of the same graph. For orthogonal drawings, InteractiveGiotto is well suited for this purpose because its constraints can be relaxed on a vertex-by-vertex or edge-by-edge basis, making it possible to obtain multiple drawings of the same graph with varying degrees of similarity to the original drawing. (This is in contrast to most interactive drawing algorithms, which produce a single output drawing for a given input.) Having multiple drawings of the same graph is important because it is very difficult to judge whether one pair of drawings is more or less similar than another when the two pairs are drawings of different graphs.

So far I have focused primarily on evaluating the metrics based on their qualitative behavior and their ability to choose the correct rotation, because having a human assign meaningful numeric similarity values to drawings is very difficult. (One future project is to devise a way to get such values in as objective a manner as possible.) There have been some winners and some losers—the orthogonal ordering metrics tended to do best in both respects, with the proximity metrics faring the worst—but nothing so far has been an obvious "perfect metric." Most of the metrics are suitable for choosing the correct rotation, at least when the possible angles of rotation are widely separated (e.g., multiples of 90 degrees), and most of the metrics do an OK, but not stellar, job of ordering the drawings.

This may be because these metrics all consider the big picture, taking the view that all parts of the drawing are created equal and so a change has the same effect wherever it is applied. This is not necessarily the case, though—while

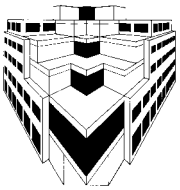

attempting to order the drawings for evaluating the metrics, I quickly realized that *landmarks* are important. One tends to focus quickly on sections of the drawing that are distinctive, and if these are present (and relatively unchanged) in the new drawing, it greatly increases the perceived similarity of the drawings. Also, landmarks vary in importance—a very distinctive structure in the graph is more noticeable, as is one that is relatively separate from other parts of the drawing. My current direction is to further investigate how to identify landmarks, how to measure their importance and how to measure how they change, in the hopes of improving the similarity metrics.

And if anyone has a desire to spend time staring at lots of drawings of graphs, come see me. :)

LETTERS & ALUMNI EMAIL

E. GORDON GEE, President

Dear Suzi—A note to tell you that I just read the latest issue of ***conduit!***. It is a first-rate publication and certainly represents the Department, as well as the University, very well. Would you please pass my delight in that product to others who contributed to this? Best personal wishes.

JASON LANGO AND NATE STAHL, ScBs '98

Hello friends and former colleagues! It would take a while to tell you all individually, so I thought that I would be lazy and spam you all. Nate and I have decided to leave Silicon Graphics and pursue career opportunities at a much smaller company named Network Appliance (*www.netapp.com*) working on systems software for the NetCache box.

We're unable to check our email at SGI, but the mail addresses *jal@rampartnt.org* and *nrs@rampant.org* will continue to reach us at our apartment.

ED LAZOWSKA, AB '72

The following was sent to Andy van Dam by David Salesin, ScB '83, under the subject line "Remarkable people:"

"Every Jan 1, the *Seattle Times* recognizes 'remarkable people...who added measurably to the quality of life in the Puget Sound region.' Among the half-dozen individuals who are recognized this year is our own Ed Lazowska. Read it online at *http://www.seattletimes.com/news/editorial/html98/bested_010199.html*, where they say "Ed Lazowska isn't exactly a household name, but at the University of Washington, where he is chair of Computer Science and Engineering, and among state education leaders, he is something of a hero. Lazowska has assisted Seattle Public Schools in a variety of ways, from Internet connections to technology training for teachers. He is a leader in the K-20 network that will someday link all the state's schools and colleges. Lazowska was a 1998 recipient of a UW Public Service Award. He is recognized for his efforts to bolster the state's economic future through technology-smart policies."

EVAN MAIR, BA '93

As the popular saying goes, you can take the Brown CS grad out of the Sunlab, but you can't core-dump the desire to code. During my years in medical school, I developed commercial medical software for the 3Com Palm series as well as helping the Boston University Medical Campus develop its Internet presence. I am now entering my residency in Diagnostic Radiology at Boston Medical Center, where I hope to continue my research on image compression and archiving systems for radiology services. I can be reached at *emair@bigfoot.com*.

LAURENT MICHEL, PhD '99

Hi! I am now working for Ilog S.A., a French company that specializes in the production of software libraries for combinatorial optimization problems. About a year ago they acquired CPLEX (the leader in linear programming). Basically, Ilog has a suite of products ranging fro libraries for LP/IP/MIP to constraint programming libraries (general-purpose CP) and even scheduling. They also have a

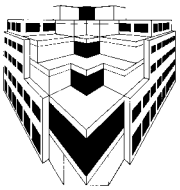

whole different sector of activity in computer graphics (high-performance 2D/3D rendering) and telecom.

I am working in the Optimization R&D team in the modeling language group. I am working from Belgium while Ilog headquarters are in Paris. You can find out more about them at *http://www-.ilog.com*. My work email address is *ldm@ilog.fr*, and I'm more likely to answer quickly on this one!

VISWANATH RAMACHANDRAN, PhD '98

After my Ph.D., I joined Netscape (now a unit of AOL), working on the Netscape Web server as an engineer. I've been focusing mostly on programming language work, developing language runtimes for Java and JavaScript. I am really excited about being in the forefront of the NetEconomy and electronic commerce. On the personal front, I got married to Shanthi in April 1998. Stay in touch folks, my email is *vishy@netscape .com*.

SRIDHAR RAMASWAMY, PhD '98

Hi everyone! My wife Seema (Tufts DMD '96) and I have moved from New Jersey to the Bay Area. I am leaving the comfortable and laid-back environment of Bell Labs for the thriving madness of Silicon Valley. I am joining a small company called Epiphany. We make ERM (enterprise relationship management) software. I would love to hear from friends. I can be reached at *sridhar@epiphany.com*.

MICHAEL RUBIN, ScB '99

Hello ***conduit!*** It's funny how only after I leave Brown do we get a chance to talk. Currently I am in California working at Network Appliance and having a great time.

After all the time and all the work, I think I am enjoying the fruits of my labors—a job where I sit in one room all day pushing buttons. Seriously, I am having a great time out here. It's amazing how much easier it is to work without sleep deprivation. It's hard to write one of these without sounding like all the hundreds of other recent grads who have nothing monumental to say. Work is better than I hoped. Love is better than I hoped. Life is good. Thanks to all those who helped make this happen, mhr. *mhr@netapp.com*.

fac.activities@cs.brown.edu

Eugene Charniak. This last year Eugene Charniak gave two invited talks at conferences on machine learning. He was pleased about these invitations because, while his research has always been in natural-language processing, it was only a few years ago that he began to apply machine-learning techniques to the problem. Thus he was happy to get this indication of "acceptance" in his new research community. Of these invitations, the more prestigious was from the NIPS (Neural Information Processing Systems) conference. However, the invitation that initially had him more excited was the Snowbird conference on machine learning, a conference famous because it is held in ski season at a ski resort, and afternoons are left open for skiing. Eugene always felt that clearly these folks knew how to throw a conference! Unfortunately, about two months before the conference Eugene had a skiing accident, and his physical therapist assured him that he would be crazy to try skiing given the state of his left knee. So, while his talk was very well received, the conference was less than a total success for him.

At the beginning of this last February Eugene (as well as Tom Dean) went to a conference center in Conway, CT, a noted ski center, for a retreat for students and professors connected with some new interdepartmental grants (Computer Science, Neuroscience, Cognitive Science, and Applied Math). Again, afternoons were free for skiing. With a year of knee exercises under his britches, Eugene ventured onto the slopes and is pleased to report that his knee held up just fine.

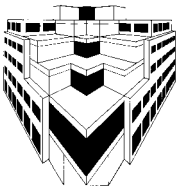

Tom Dean. Tom, program committee chair for IJCAI-99, attended a meeting of the IJCAI Inc. board of trustees in Paris during the World Cup Games.

John Hughes. Well...the graphics group sent six papers to SIGGRAPH—actually, something like a total of 12 items went out, from course and panel proposals to papers to who-knows-what. One cool thing was the paper-that-became-a-paper in the last three days before the deadline, with two undergraduates as lead authors. As for me personally, I've been working on the Andyfest (the May 27-28 gathering of Andy's friends, colleagues and ex-students worldwide to celebrate his 60th birthday), my courses, the five-year review of the grad program and being graduate advisor—all the stuff professors do.

Franco Preparata. Franco was appointed to the international committee that assigns the Gödel Prize in theoretical computer science. He was also appointed to the Scientific Board of ISTI (Istituto per le Scienze e la Tecnologia dell'Informazione) in Pisa, Italy, a unit of the Italian National Research Council. In the late fall of 1998, he successfully led the technical review of the project "Robust and Applicable Geometric Computing," which, funded by ARO, is the main support of our three-university Center for Geometric Computing. This review capped the initial three-year phase and has secured funding for the subsequent two-year option. At RECOMB99, the premier international conference in computational biology (held this year in Lyon, France), he presented the results of his work with Eli Upfal, which are the subject of a recent patent application.

John Savage. In April John served on a committee that wrote a report to the NSF entitled "Challenges for Theory of Computing." This report offers NSF advice on funding research in theoretical computer science. Its URL is *http://www.cse.buffalo.edu/~selman/report*.

Roberto Tamassia. Roberto cochaired the program committee of the Workshop on Algorithms and Data Structures (Vancouver) and served on the program committee of the Workshop on Algorithm Engineering and Experimentation (Baltimore). He also gave a keynote lecture at the Symposium on Algorithms and Computation (Taejon, South Korea).

Eli Upfal. Eli's financial computing research group visited Goldman Sachs in New York in December. He participated in an NSF review panel in January and was appointed associate editor for the *SIAM Journal on Discrete Mathematics*.

Andy van Dam. Andy has been awarded the 1999 IEEE James H. Mulligan, Jr. Educational Medal "for his field-defining textbooks, the introduction of innovative educational technology, and inspired undergraduate teaching."

Peter Wegner. Peter is retiring in June. This spring and summer he will be lecturing in Amsterdam, Lisbon and Tokyo. He was recently awarded the Austrian Medal of Honor in Science and Art.

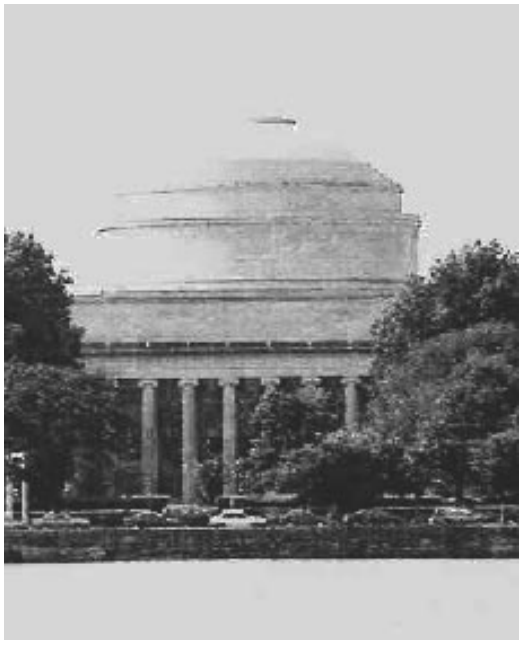

The Dome, MIT and The Bridge of Sighs, St. John's College

FROM CAMBRIDGE TO CAMBRIDGE

Peter Wegner

During the week of April 12-16, 1999, I attended the 35th anniversary of MIT's Laboratory for Computer Science and the 50th anniversary of the Cambridge (England) Computing Lab. The MIT event featured Bill Gates' $20-million gift and keynote talk about the future of computing, and the unveiling of MIT's Oxygen project for next-generation intelligent personal assistants. The British event, a gathering of about 500 mainly British computer scientists, featured a keynote talk by Maurice Wilkes, who developed the first stored-program computer in the late 1940s and retains his resilience at age 85. Since I cannot do justice, in this short article, either to the goals of Oxygen or to the British views of the past and future of computer science, I shall relate some anecdotes.

In talking to Bill Gates, I asked about Microsoft's goals for computer graphics. Bill seemed surprised at the question, as though he had not thought about it before. He spoke of Microsoft's wonderful researchers in the graphics area, and said that one of the goals of graphics was realism. A few days later, at the conference dinner in St. John's College in the other Cambridge, I sat next to Alvy Ray Smith, who designed the graphics effects in Star Wars at LucasFilm and Toy Story at Pixar and is now a researcher at Microsoft. When I told him I had asked Bill Gates about the goals of graphics, he cupped his ear and waited expectantly. On hearing that Gates had mentioned realism as a major goal, Alvy said there is much debate in the graphics community about the role of realism in graphics research. In fact, the Brown graphics group's project on non-photorealistic rendering specifically explores nonrealistic abstractions that focus on details of a visual image considered relevant while ignoring details considered irrelevant.

> *"Abstraction is the enemy of realism"*

Data visualization strays from realism even further, since its visualizations have no realistic image to which they correspond. Abstraction is the enemy of realism.

At dinner with Maurice Wilkes and his wife Nina, we revisited the early history of computing. Maurice discussed his trip to the US in the late 1940s to attend a course that enabled him to build the first stored-program computer in 1949, one year ahead of his US competitors. We talked about my days as a student in the

first ever computing course during 1953-54. I was later given a copy of the 1954 final exam, which I would certainly fail today because it included hardware questions as well as advanced numerical analysis.

On the final day of the conference I was invited to lunch at the Microsoft Research Lab by Roger Needham, who headed the Cambridge Computing Lab before Robin Milner took over three years ago. It occupies a very pleasant space in the center of Cambridge and has some impressive researchers, including Luca Cardelli and Tony Hoare, as well as a superb panorama overlooking one of the main thoroughfares.

These two anniversary celebrations allowed me to meet many old friends who had traveled on parallel tracks and were, like me, nearing the end of their academic career. These events marked not only a celebration of the past, but also a passing of the baton from the old guard to a new generation who must deal with new technologies of networking, collaboration, and electronic commerce. Oxygen is an ambitious effort to develop an integrated approach to technologies of the 21st century, but foundations for integration have yet to be developed.

The Microsoft connection at both Cambridge reunions reflects Microsoft's emergence as a ubiquitous presence in computing on both sides of the Atlantic. Microsoft's attention to the past could well prove useful in its bid to define the future. If Microsoft succeeds, historians will identify the shift from old to new technology with a shift from an IBM-dominated to a Microsoft-dominated computer industry.

FROM THE CHAIRMAN

Tom Dean

This short note reports some sad news, some news that is cause for celebration but makes us recall some sad events of the recent past, and some news about changes in the life of a friend and colleague that are cause for both reflection and celebration.

First, some very sad news. Markus Meister, in his second year of Ph.D. study in the department, died in a Christmas Eve automobile accident in British Columbia. Markus was well loved in our community. In addition to his research in graphics, Markus was prankster and proselytizer for classical music, mathematics for its own sake, and all things concerning Japanese culture. His whimsical side was obvious in his interest in *anime* videos, particularly those including the character Totoro, and in his fascination with all sorts of hardware, including sports cars well past their prime but nonetheless beautiful. More than anything else, we'll remember Markus for his generous and enthusiastic spirit and his infectious sense of fun.

The next piece of news concerns the 1998 ACM Kanellakis Theory and Practice Award. This was given to Randal E. Bryant, Edmund M. Clarke, Jr., E. Allen Emerson, and Kenneth L. McMillan for their invention of "symbolic model checking," a method of formally checking system designs that is widely used in the computer hardware industry and is beginning to show significant promise also in software verification and other areas. This award was established in memory of our friend and colleague Paris Kanellakis, whose tragic death in late 1995 cut short a distinguished research career. Paris and his family are frequently in our thoughts and we often hear from their many friends, students, and colleagues from around Brown and throughout the world.

Finally, on April 22 we had a party for Peter Wegner, who will be retiring at the end of this semester. We don't expect to see any less of Peter in his new role as *emeritus* professor, and indeed his retirement party was cause for a celebration of his scholarly achievements and his many contributions to the department. At the

Enjoying his retirement party at the Faculty Club. Clockwise from bottom left: Peter with President Gee, with long-time friends from the Sociology Department Marilyn and Dietrich Rueschmeyer, and with CS chairmen past and present

party President Gee thanked Peter for his long service to Brown and announced the news that Peter would receive the Medal of Honor for Arts and Sciences from the President of the Austrian Federation at a ceremony this summer. President Gee remarked on some of the extraordinary events in Peter's life, from his escape from Austria as a child on the "*Kindertransport*" to the award of the Austrian Medal of Honor almost exactly 60 years to the day later. There is certainly no reason for sadness that Peter has passed this milestone in his career; quite to the contrary, we look forward to following his scholarly achievements and comic shenanigans for decades to come. The sentiment on the brass plaque affixed to the classic spindle-back Brown chair that he was given at the party read, "To an unrepentant punster, indefatigable scholar, and generous friend."

BROWN ENTERS THE CAVE ERA

Rosemary Simpson

In the fall of 1997, six Brown departments—Applied Math, Chemistry, Cognitive Science, Computer Science, Geology, and Physics—were awarded a joint National Science Foundation MRI (Major Research Infrastructure) grant to develop a university-wide computing facility for collaborative research and education. The result is a new supercomputing and immersive virtual reality lab—the Technology Center for Advanced Scientific Computation and Visualization—located at 180 George St. This lab will be available as a research, computing, and education facility and is expected to foster scientific investigation and graduate and undergraduate instruction in a variety of fields, including chemistry, geology, cognitive and linguistic sciences, physics, computer science and applied mathematics. The lab contains a high-performance parallel computer, the IBM SP, and a cave.

What is a Cave?

Brown's cave is an eight-foot cubicle in which high-resolution stereo graphics are

Illustration of CAVE™ in use

projected onto three walls and the floor to create an immersive virtual reality experience. High-end workstations generate the 3D virtual world and create the sounds of the environment. Special hardware and software keep track of the positions and movements of a person entering that virtual environment, changing the images in the cave in a way that allows the visitor to feel immersed in the virtual space.

> *"...it is easier to understand how something works or is put together if you can hold it in your hand or walk around inside it"*

Our cave is based on the original CAVE™ developed at the Electronic Visualization Lab at the University of Illinois, Chicago, as part of long-term research on tools and applications for immersive virtual reality. It was first demonstrated at the SIGGRAPH '92 conference. Since then, dozens of caves have been installed worldwide with applications ranging from scientific visualization, industrial research, and training, to education, theatre, and the arts. (See *http://www.cs.brown.edu/research/graphics/research/vr.html* for cave-related resources and applications.)

Brown's Technology Center for Advanced Scientific Computation and Visualization

The Technology Center resources include an IBM RS/6000 SP parallel computer and the cave. The IBM computer will be used to simulate complex processes as diverse as the movement of the earth's mantle, subatomic particle interactions, and the function of the human heart. It will also drive the projectors that display those images in the cave. For example, a researcher studying a simulation of air flow around a space shuttle can walk around a model of the shuttle floating within the cave, even feeling compelled to duck while walking under it. The researcher could see patterns of airflow around the shuttle by positioning and moving colored streamers in the flow much like tails on a kite. "We anticipate that the cave will become a powerful tool for facilitating scientific insight," said

Samuel Fulcomer, Center director. "This is based on the idea that it is easier to understand how something works or is put together if you can hold it in your hand or walk around inside it."

People interacting with the cave's graphics wear tracking devices to monitor their movements, lightweight stereo eyewear to see objects in 3D, and slippers to protect the delicate polymer screens lining the cubicle. The projected virtual space appears seamless, without any intersections between floor and walls.

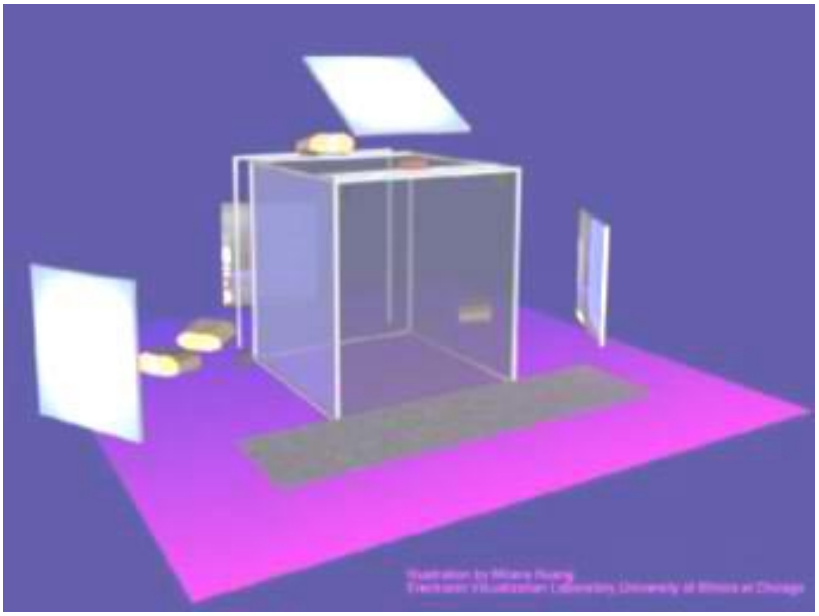

CAVE™ showing projection system

As viewers negotiate a virtual environment, objects may appear to hang in space in front of them, tangible and within reach. This is an effect of the stereo glasses, which create slightly offset images for the left and right eyes. The cave user may also wear special gloves that send the computer information about hand location and thus let him or her interact with the virtual environment using gestures. A viewer could change the lighting and artwork in a virtual gallery, precisely position an artificial valve in a virtual heart, or fly over a virtual Martian landscape.

Dedication Ceremony

The dedication ceremony will take place at 2pm on May 26, 1999, in the C.V. Starr Auditorium of MacMillan Hall, 167 Thayer St. and will feature Brown President E. Gordon Gee; Andries van Dam, the Thomas J. Watson Jr. University Professor of Technology and Education and Professor of Computer Science; George Karniadakis, Professor of Applied Mathematics; and Paul M. Horn, IBM Senior Vice President of Research. After the ceremony, there will be demos at the 180 George Street facility, among them a Mars topology demo by the Geology Department and 3D modeling and an architectural walkthrough of a virtual environment by the Computer Science Department.

Computer Science Application Possibilities

Like all the departments collaborating in the cave facility, Computer Science is developing research and educational applications that would either be impossible or much less effective without it. Some of these areas include:

Scientific visualization: In conjunction with NASA, the Graphics Group has been doing computational fluid dynamics studies that will benefit from the cave's environment. New studies of bloodflows are also being developed. Professor David Laidlaw is studying biomedical areas where he points out that graphically-based studies of deep structure and tissue in areas like the brain, the hand and the spine will enormously benefit from immersive 3D visualization.

Concept visualization: Professor Eli Upfal states that both research and teaching of theoretical areas would benefit from access to an immersive 3D visualization facility. For example, network analysis studies involve multiple dimensions and he feels that the development of intuition about these processes would be greatly accelerated by such a facility.

Algorithm visualization: Professor Roberto Tamassia currently uses extensive visualization techniques for algorithm teaching and research and feels that 3D immersive techniques would be highly beneficial. He pointed out, however, that the small size of the facility would severely limit its usefulness in teaching.

Behavior simulation: Professor Nancy Pollard and her students have been studying interactive teaching situations. One project, called Coach, is a 2D/3D electronic playbook for football. Coaches can develop and modify football plays and players can view animations of those plays within a 3D environment (e.g. from a first-person point of view). A second project explores animation of dance motions for archival and teaching purposes.

User interface research: Graphics Group members Andy Forsberg and Joe LaViola have been investigating multimodal user interfaces needed in computing environments where users no longer work solely in a desktop keyboard-screen-mouse mode. New approaches that incorporate new gestural, vocal, and visual vocabularies are needed to work in these environments.

Research and educational opportunities exist in all of these areas, as well as potentially in other domains such as systems analysis. For further information on research, support, and demos, contact: Director Sam Fulcomer at *sgf@cfm.brown.edu*:

OUR MOST RECENT PhDs

NAME	ADVISOR	THESIS	POSITION
Mitch Cherniack	Stan Zdonik	"Building Query Optimizers with Combinators"	Assistant professor at Brandeis
Hagit Shatkay	Leslie Kaelbling	"Learning Models for Robot Navigation"	Postdoctoral fellow at the National Center for Biotechnology Information. She is developing methods for faster and more accurate biomedical information retrieval and on applying machine-learning-type methods in bioinformatics
Laurent Michel	Pascal Van Hentenryck	"Localizer: A Modeling Language for Local Search"	Ilog S.A., Belgium
T.M. Murali	Jeffrey Vitter	"Efficient Hidden-Surface Removal in Theory and in Practice"	CS post-doc at Stanford working in the robotics and graphics labs with Professors Leo Guibas and Jean-Claude Latombe, among others

Laurent Michel and poultry after his defense and celebrating with his wife, Valerie

Hagit and Eadoh complete a construction project

Mitch Cherniack

CHARNIAK UNPLUGGED

Eugene Charniak

A few ***conduit!***s ago I wrote an article entitled "Bad Trips," about problems that have come up in faculty's trips to give talks, etc. The problems were mostly standard, from bad travel connections to bad weather. However, I just heard a new class of problem from one of my graduate students, Heidi Fox. She and a colleague were invited to give a talk at a conference on some work they had done jointly. They agreed that her colleague would give the talk, and thus Heidi was not planning on going at all. However, this was the end of last year, when there was a threat of a government shutdown, and as Heidi's colleague is a government employee, if the government did shut down, he could not go (or at least could not be reimbursed for travel expenses) because technically he would be unemployed. Thus Heidi had to plan on going and since the government eventually did not close down, they both were there. Having both authors of a paper show up at this conference seems to have been pretty unusual. Actually, just getting one author to show up was not that common either. For the four or so papers in the session, only three authors showed: Heidi, her colleague, and one other author. On the other hand, perhaps the missing authors knew something that Heidi did not—showing up was clearly

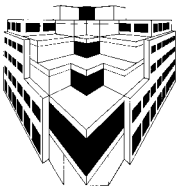

optional, since the speakers outnumbered the rest of the audience.

To understand the following bit of departmental goings-on, you need to know one piece of university/department trivia. The university distinguishes between very small purchases and bigger ones. Very small are handled from petty cash. While this is convenient for small purchases, the amount of paperwork per dollar is clearly higher than for larger ones. Thus when the department ran out of cookies last fall (they are served before departmental talks, along with tea and coffee), Jennet Kirschenbaum figured that we should stock up for a while and had a purchase order made out to the tune of $125. Then she and Mary Andrade walked through a local grocery store (the only one that takes Brown purchase orders) with calculator, trying to buy exactly $125 worth of cookies. This turned out to be a shopping cart full, which, as you might expect, attracted a lot of attention ("Can I come to the party?") The bill was $130. Jennet estimates that we use about $10 of cookies per event. Figuring about two events per week, the department goes through 130 dollars' worth of cookies in six to seven weeks. In the meantime, Jennet stores the cookies "on the 5th floor—locked up."

"...the department goes through130 dollars' worth of cookies in six to seven weeks"

On a related topic, as mentioned elsewhere in this ***conduit!***, our fall IPP symposium was on the Java programming language. Since parking is always tight around Brown, Suzi Howe had spaces reserved for the speakers near our building. However, in order to direct them to the spaces, and in order to make sure that they bunched up the cars so they could all fit in, Mary Andrade and Dawn Nicholaus took turns patrolling outside the building to intercept the speakers. As they had no idea what the speakers looked like and vice versa, they carried a sign (hastily improvised by attaching cardboard to a fly swatter) that advertised the conference. The sign said, with an admirable economy of words, "JAVA". But while the speakers knew how to interpret this, the average passerby, needless to say, did not, and this attracted even more attention than walking down supermarket aisles with $130 worth of cookies ("Where's the free coffee, lady?") The fly swatter is now retired to the CS archive.

Avi Silberschatz of operating systems fame was a speaker in the department's distinguished lecture series this year. In his talk he mentioned a common OS technique of keeping twice as much storage as one needs for an IO-bound task. One then reads into half of the storage while one is reading out of the other half. When the input half is all filled up, one just switches halves. In his talk Silberschatz mentioned that Stan Zdonik told him of someone who did this with dishes. I immediately knew what he was talking about because many years ago I decided that it would be a great idea to have two dishwashers. One puts dishes into one when they are dirty, then when it is full one runs it. In the meantime one is removing dishes from the other dishwasher. The point, of course, is that one never has the task of removing dishes to a cabinet. After the talk Stan told me that the person in question is Tom, one of the Click and Clack brothers. In researching this article I called Stan to ask whether Tom was Click or Clack. Stan said that this is a deep philosophical question about which there is a great deal of controversy. He did give me a bit of fast-breaking news however; namely, that MIT, after having U.N. Secretary-General Kofi Annan as its graduation speaker two years ago and President Clinton last year, is going to have Click and Clack this year (they are both MIT alums). Remember you heard it here first!

I learned about an even more ambitious home-improvement scheme a month or two later when I attended the Neural Information Processing conference to give an invited talk. I noticed that Mike Mozer, who graduated from Brown about

1980 or so and worked with me on an AI project during his senior year, was also speaking, and we got together to bring each other up to date on what was new in our lives. Mike is now a professor of computer science at the University of Colorado, and one of his projects is wiring up his entire house so that his home computer knows things like (a) whether he is in the house and if so, in which room, (b) what lights he has on, (c) what he has the heat set to, etc. As Mike is particularly interested in applications of machine learning (which is why we were both at this conference), his basic idea was to get the machine to learn his habits and preferences so that it could do things like turn off lights that were not in use, raise and lower the temperature to save on fuel, and even turn on lights when he is about to move to a different room (it notes motion in the room and has learned that at the particular hour in question Mike usually leaves that room to go to a second room). His house, as you might imagine, has been written up many times (both academically and in the popular press), and has its own Web page: *www.cs.colorado.edu/~mozer/adaptive-house*. One interesting result of this project is that Mike's water heater is the property of the University of Colorado.

Dawn and Julia

Finally, on the day of one of last winter's larger snowstorms, Dawn Nicholaus brought her five-year-old daughter Julia to work with her for a few hours. When I emerged from the elevator on the fourth floor, Julia was sitting there and greeted me with a cheerful "Hi, bro!" Not expecting a greeting at that point, and certainly not expecting that particular honorific, I managed only a weak "Hi" back. Later that day Dawn and I discussed whether a middle-aged white professor could be a "bro," and eventually decided that if Julia thought so, then yes. For the next few days the administrative staff called me "bro."

" ANDYFEST "

To celebrate Andy van Dam's 60th birthday, a symposium will be held May 27-28 entitled "The Computer, the Academy, and the World."

Andy was awarded the second US PhD in Computer Science in 1966. Since then, through his deep commitment to education and his boundless energy, his impact on the emerging world of computers has been enormous. He has broken new ground in the use of computers in education, educated a whole generation of experts in computer graphics (not to mention co-founding ACM SIGGRAPH, whose annual conference draws more than 30,000 attendees), helped with countless startup companies, and served on the advisory boards of many small and large companies, including Microsoft Research. He has also mentored a steady stream of undergraduates, many of them current or former chairs of the top-ranked departments in the country.

Among his awards are the ACM SIGGRAPH Steven A. Coons Award (1991), the ACM Karl V. Karlstrom Outstanding Educator Award (1994), and the IEEE James H. Mulligan, Jr. Education Medal (1999). In 1994 he also became an IEEE Fellow and an ACM Fellow. He has honorary PhDs from Darmstadt Technical University in Germany (1994) and Swarthmore College (1996). In 1996 he was inducted into the National Academy of Engineering.

Speakers at the two-day event will be **Ronen Barzel,** who joined Pixar in 1993 to work on Toy Story and has since worked on R&D of modeling, lighting and

animation tools; **Ingrid Carlbom ('80)**, Head of the Visual Communications Research Department in the Multimedia Communications Research Laboratory at Bell Labs; **Steven J. DeRose ('89)**, co-founder of Electronic Book Technologies and now Chief Scientist at Inso Electronic Publishing Systems; **Henry Fuchs**, Federico Gil Professor of Computer Science and Adjunct Professor of Radiation Oncology at the University of North Carolina and winner of ACM SIGGRAPH's 1992 Computer Graphics Achievement Award; **Dr. Alan Kay**, Disney Fellow and Vice President of Research and Development, The Walt Disney Company, who is best known for the idea of personal computing, the conception of the intimate laptop computer, and the inventions of the now ubiquitous overlapping-window interface and modern object-oriented programming and whose numerous honors include the ACM Software Systems Award; **Ed Lazowska ('72)**, Professor and Chair of the Department of Computer Science and Engineering at the University of Washington; **Raj Reddy**, Dean of the School of Computer Science at Carnegie Mellon University and the Herbert A. Simon University Professor of Computer Science and Robotics, who received the ACM Turing Award in 1995; and **David Salesin ('83)**, Senior Researcher at Microsoft Research and an Associate Professor in the Department of Computer Science and Engineering at the University of Washington.

http://www.cs.brown.edu/~andyfest

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