conduit!

Volume 4, Number 1 **Department of Computer Science** **Spring, 1995**
Brown University

COMPUTING FOR NONMAJORS

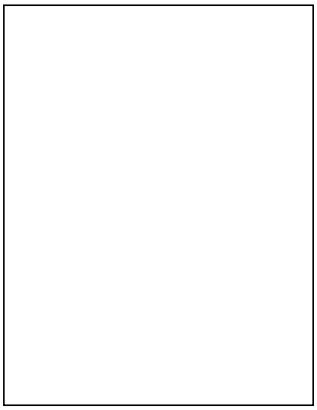

Peter Wegner

Brown's introductory course *Concepts and Challenges in Computer Science* (CS2), which has grown in size from 45 students two years ago to 85 last year and 180 this year, combines conceptual computer literacy with practical proficiency in the creative use of application packages. The MacPaint competition, assigned during the first week and judged by our team of undergraduate teaching assistants (UTAs), brings out interesting artistic talent. This year's 12 prizewinners (in the categories best artistic, best technical, funniest, and most original) and nine honorable mention entries are exhibited on the second floor of the CIT building. Two prizewinning pictures are shown here—last year's funniest picture (which turned out to be the star attraction at a number of public lectures on CS2) and this year's best technical prizewinner.

The course is assignment-driven with a simple Hypercard résumé assignment in the second week, a 12-hypercard assignment on "How Computers Execute Programs" in the third week, and a network assignment in the fourth week that includes a network treasure hunt, a simple HTML home page, and an essay on network architecture. By the fourth week students are familiar with MacPaint, MS Word, Hypercard, e-mail, and network surfing using a simple viewer like Netscape. They are encouraged to access the course home page, which contains information about the course syllabus and UTAs as well as a "Message of the Day" (MOTD) providing up-to-date information about assignments, exams, etc.

The artificial intelligence assignment "Can Machines Think?", given in the fifth week, is based on Turing's seminal paper, ongoing debate among scholars like Searle and Penrose, and Isaac Asimov's *Bicentennial Man*. Its "home card" (see Figure 1, page 2) has buttons for accessing background information, facts, "yes" arguments, and "no" arguments as well as content, help, and opinion cards. A "personal opinion" card from one of this year's assignments (Figure 2, page 2) shows the subtle reasoning this assignment can elicit. Later assignments include a spreadsheet in Excel, a coffeeshop cash-register program in Hypertalk that involves both interface design and simple programming, and an essay on the social impact of computers that this year allowed students to focus on the spring-1995 special issue of *Time* on cyberspace published just a week before the assignment.

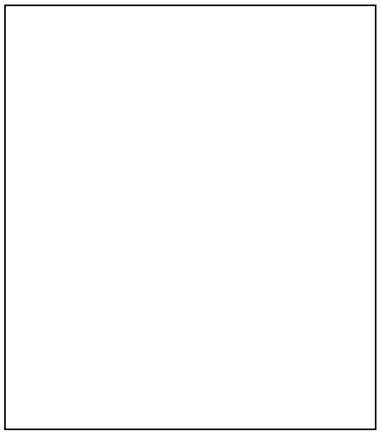

Funniest, 1994 "Bobbitt"

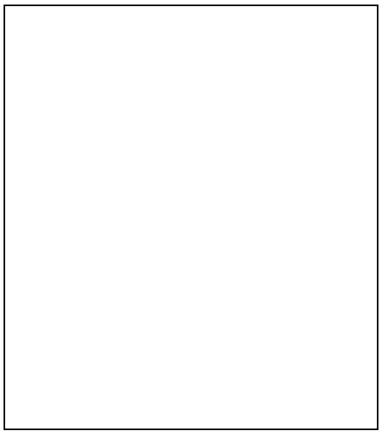

Best Technical, 1995

The month-long final project requires students to develop a hypertext on a topic in which they are interested, such as a hobby, a course, or an artistic or athletic interest. They are encouraged to develop a multimedia document that includes text, static and dynamic illustrations, and some audio. A "design" developed in the first two weeks must be approved before the implementation can go forward to ensure that it is neither too ambitious nor trivial. Last year's projects included hypertexts on the archaeology of Mesopotamia with artifacts and historical discussion of each period, a tutorial on modeling with clay with modern examples of pottery, adventure games based on Star Trek and other themes, a New Testament hypertext on Matthew illustrated with Leonardo's *Last Supper*, a tourist guide to New York emphasizing night spots, and a music hypertext on the band U2 with audio clips (see below).

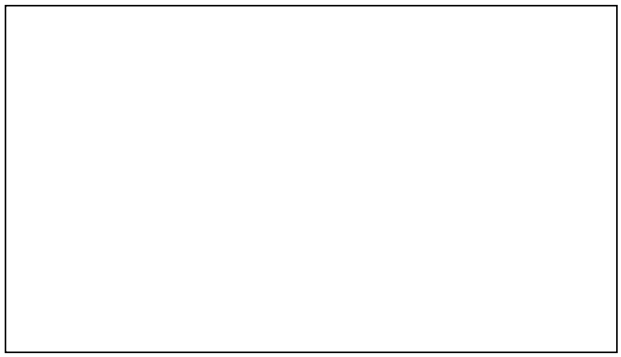

Final project topics in 1994

Students become "computer literate" in the first half of the course and deepen their ability to create multimedia documents in the second half. The midterm exam tests computer literacy as defined by a list of about 300 terms that students are expected to understand. Questions have the form: What are the similarities and differences between instructions and data? Binary and Roman numerals? Compilers and operating systems? The vocabulary list is available prior to the exam and evening help sessions devoted to answering questions about the "literacy vocabulary" are generally well attended.

One of the features of CS2, as of our other introductory CS courses, is its intensive use of UTAs. The 19 UTAs of CS2 provide 25 heavily used hours per week of consulting without which the intensive schedule of assignments would not be possible. During the final project period each UTA works directly with ten students on their final projects. The department's investment in 50 to 100 UTAs per semester not only benefits students but also develops a departmental *esprit de corps* and a sense of responsibility and belonging. The team of UTAs is generally quite diverse, including Indian, Pakistani, African-American, and Asian-American women and men. This year the CS2 UTA team includes the presidents of the Brown Islamic and Jewish societies. Over half of our majors serve as a UTA for at least one course.

CS2 has attracted attention both in industry and at other universities as a model for computer literacy. I was invited to talk on CS2 at Bellcore's "Electronic Document Delivery Conference (EDD-94)," at Apple Computer, and in Europe. My talk at a workshop on introductory courses in January 1995 at Harvard sparked interest at several colleges in using the approach and materials of CS2. We hope to develop an exportable version of the course materials this summer.

CS2 differs from first courses for majors in focusing on documents rather than programs. Hypercard, Excel, HTML, and MS Word may be viewed as tools for document management, while programs may be viewed as specialized documents. Students learn not only a set of general-purpose tools and concepts but also develop the ability to express themselves in a new medium. CS2 teaches technical writing and design skills for substantive multimedia documents that could not be written effectively

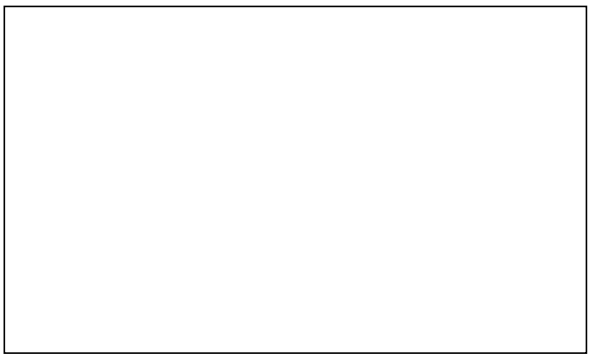

Figure 1—"home card"

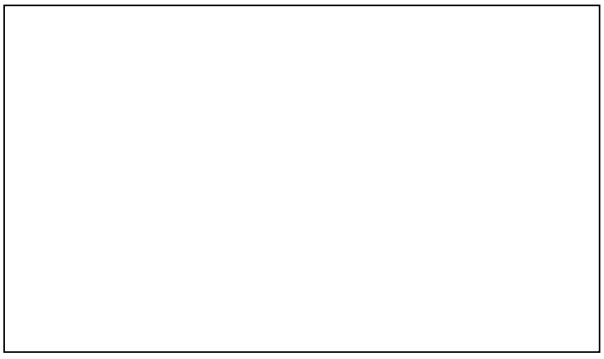

Figure 2—A "personal opinion" card

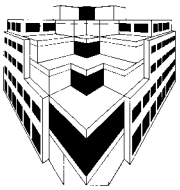

without computers. Methods of document design in CS2 are similar to those of program design, but are applied to the domain of documents rather than programs. The final "capstone" project requires students to apply their writing, design, and document management skills to a new domain about which they are already knowledgeable.

CS2 has a conceptual and technical coherence, focusing on the technology of *document engineering* (the creation and management of documents), which has many parallels with the established technology of software engineering. Large programs and large multimedia documents have a similar structure and similar management problems. Both are linked structures of components: links associated with buttons are introduced much earlier in hypertext-based courses than the corresponding concept of pointers in Pascal-based or C-based programming courses. Large hypertext documents are easier to create in the time span of a single course than large programs: CS2 can explore problems of largeness without many of the technical details that arise in programming. Moreover, large documents relate more directly to everyday experience than large programs and more effectively motivate students to explore substantive large applications.

As computer science becomes more application-driven and outward-looking, first courses in computer science may evolve from their current emphasis on programming to an emphasis on document engineering. It may well be that ten years from now first courses on computing for majors will be closer in content to CS2 than to programming courses in Pascal, C, or C++. First courses in programming are becoming more design-oriented and less preoccupied with low-level algorithms and control structures. The gap between courses for majors and nonmajors is likely to narrow as the technology of personal computers matures and a conceptual framework that spans both programming and document engineering is developed.

NEXAL COMPUTING

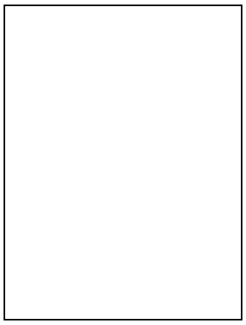

John Savage

On November 14, 1994, under the banner "Nexal Computing," the Department hosted its 14th Industrial Partners Program Technical Symposium. The symposium brought together five speakers from the Corporation for National Research Initiatives (CNRI), Bellcore, Sun Microsystems, DEC Cambridge Research Labs (CRL), and IBM Research Laboratories to discuss the technology and sociology of international communication networks. The title "Nexal Computing" was chosen to emphasize the importance of the nexus or connection point in such networks. The audience was made up of attendees from Cadre Technologies, Cambridge Technology Partners, BBN, DEC, Fleet Services Corporation, Ford Motor Company, GTech Corporation, Mitsubishi Electric Research Labs, and Sun Microsystems.

Dick Binder of CNRI led off the program with a discussion of gigabit networking. Dick ran the Gigabit Testbed for CNRI under contract to Federal agencies. Five separate testbeds were created around the US, with about $16M of Federal support and $100M of industrial support, to study the challenges and opportunities offered by gigabit-per-second communication speeds in order to assess the benefits of very high-speed networks as well as to advance the state of the art in networking technologies. The testbeds involved multiple industrial and academic partners who collaborated at a precompetitive level. Dick described some of the demonstration projects that hint at the potential of this new technology. They include remote collaborative supercomputing as well as metacomputing in which specialized supercomputers at differents sites are combined to solve those parts of a problem for which they are uniquely suited, providing in some cases superlinear speedups. Dick also reviewed the the technologies that are needed for gigabit-per-second speeds and pointed to some of the impediments to the full realization of this technology.

Mike Lesk of Bellcore gave the audience a glimpse of the potential, politics, economics and sociology of the national information infrastructure. For example, he cited data showing that in 1991 the American Chemical Society earned $88M selling paper and $57M selling bytes! Mike also reviewed the cost of libraries, noting that American university libraries spend $900M buying materials or about $120/student, and observing that it costs about $50 to scan a book, $20 to store it on a shelf but only $10 to store it in a computer. He

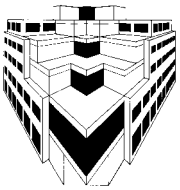

observed that many library buildings cost more to build than it would cost to scan their contents and record them on disks. These facts suggest that it would pay universities to band together to share their scanned books, thereby greatly reducing their acquisition and storage costs. As he observed, many issues are raised by the electronic storage of information, including how to protect the intellectual property rights of the owners of the original documents.

James Gosling of Sun described WebRunner, a new Web browser he developed on top of his new programming language Oak. WebRunner extends the features of other browsers by "adding dynamic behavior that transforms static documents into living applications.'' Applications such as interactive scientific experiments, electronic shopping applications, customized newspapers and many others can be added easily to WebRunner. The new browser also makes it possible to access software transparently over the network. If an application requires software not resident at a site, it need not be installed there because under WebRunner it migrates automatically to the site. Oak, the language used to implement these features, is "a simple, object-oriented, distributed, interpreted, robust, secure, architecture-neutral, portable, high-performance, multithreaded, dynamic language." Oak is based on C++ and has many of its features, although it avoids many of its problems, such as memory management for which it provides garbage collection. It supports distributed applications because pointers to remote objects are translated into RPCs.

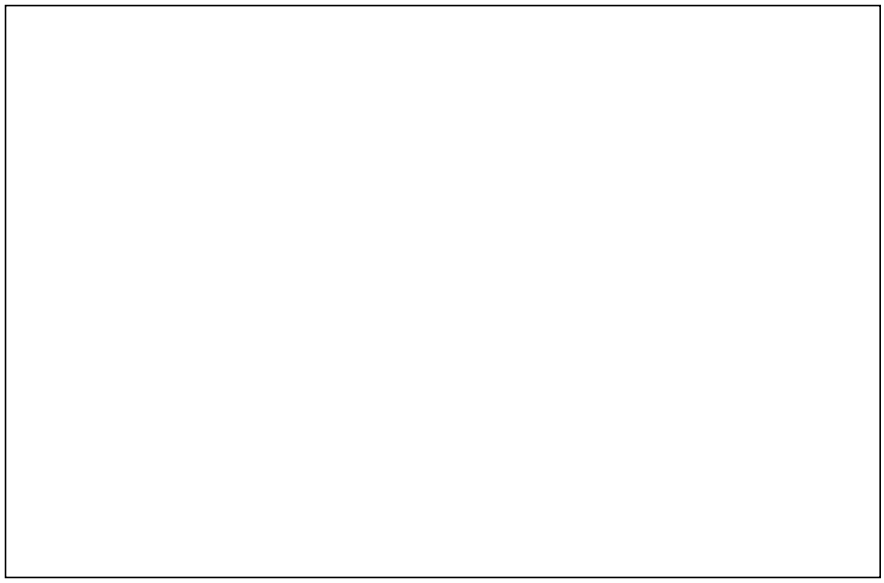

Symposium speakers: l to r, John Savage, James Gosling, Ed Balkovich, Dick Binder, Mike Lesk, Colin Harrison

After hearing from Dick, Mike and James we stopped for lunch only to have our lunch break interrupted by a fire alarm just as the Caesar salad was being tossed. We had barely returned to the buffet line when a second alarm went off and out we went again! Fortunately, the weather was brisk, the walks up and down invigorating and the interruptions brief. After a delicious meal we resumed the afternoon program.

Ed Balkovich of DEC chose as his topic "High-Tech Computer—A 21st-Century Oxymoron." His thesis is that in most future applications the computer will be invisible and low tech; it will appear to the user as just another communication appliance. Ed outlined a number of issues that will increase in importance with the advent of global communication. They include the globalization of commerce, commodity pricing for network access, and the advent of high-margin information services. Ed also cited a number of strategic problems that must be addressed before the new technology has matured. They include the development of billing systems to charge for network-based services, secure financial transactions, effective indexing and searching tools, and software for network management. Ed reported the startling fact that AT&T spends more on billing than it does on its transmission lines! The implications of observations of this kind for the development of technology are staggering. As Ed notes, the strategic problems facing the development of new technologies are often non-technological.

"AT&T spends more on billing than it does on its transmission lines!"

Colin Harrison of IBM, the clean-up hitter, described a large project underway at IBM to provide intelligent communication services. The goal of the project is to provide users with "integrated access to existing voice, e-mail, fax, pager and information services by voice and/or data with a single interface to multiple services using multiple access methods.'' A variety of abstractions based on a global object framework are being developed to simplify communication. They include proxies for sub-

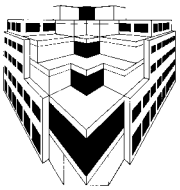

scribers, devices and services, an *alter ego* to provide static information on users and agents to control service behavior. An ICS is envisioned as a national (and possibly global) system providing users with access via a single phone number, priority access control and general management of information. The large scale of the system introduces many challenging technical issues, such as distributed large-scale directories.

The day ended with a discussion with a panel of the speakers. After a stimulating interaction we retired to whet our whistles with libations as well as grace our palates with the delightful delectables provided by Brown Food Services. All in all, it was a very enjoyable day.

COMPUTERS IN THE CS MUSEUM

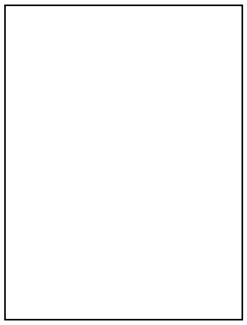

Tom Doeppner

Summary

The faculty members who in 1979 became the CS department had been involved with computers at Brown for well over a decade. Andy van Dam's early work with undergraduates on hypertext, producing first the *Hypertext Editing System* and then *FRESS*, ran on the university's IBM mainframe—a 360 Model 67 owned by the university for an incredibly long time (from the late '60s till the late '70s). Our first significant computer system was *BUGS*, the Brown University Graphics System, and our first general-purpose computer was a VAX-11/780 (named *Nancy* when "the VAX" no longer sufficed to identify it) that arrived in early '79. General-purpose time sharing prospered here with the addition of another VAX-11/780 in 1982 (*Sluggo*) and finally a 12-processor Encore Multimax in 1987 (*Zaphod*). Another VAX, a 750 (*Skyler*), was acquired in '85 and used as our mail server. Nancy now rests in the department's computer museum.

Brown was a pioneer in using workstations for CS education as well as research. Our first Apollo workstation (model DN400) arrived in '81. By that fall we possessed seventeen DN400s and used them in the classroom for our introductory programming and algorithms courses. The DN400s were supplemented with the next-generation Apollo, the DN300, in mid-'83: sixty were acquired for instruction, fifteen for research. The instructional machines arrived in the first two weeks of the '83 fall semester, just after Jeff Coady, newly hired to administer them. Jeff, who had never seen a DN400 before, soon had to cope with running what was perhaps the largest collection of Apollos outside of Apollo. The Apollos were joined by a Sun 1 workstation in late '82 (*Fritzi* now rests in the computer museum). We acquired a couple of Sun 2s in mid-'84 (one of which, *Munin*, is now in the computer museum) and some Sun 3s in another couple of years. Our first Sun 4 arrived in mid-'87.

Our original workstations weren't all that exciting to our AI folks. They investigated various Lisp packages for both the VAX and the Apollos, but finally decided that they would be best off with Symbolics Lisp Machines. Fortunately, money was found for these and five were acquired in '85 and '86—*Babar*, *Bimbo*, *Clyde*, *Dumbo*, and *Horton*. They served us well and were retired in '90: one of them now rests in the computer museum, the others were donated to Brown's Division of Engineering.

Late in the spring of 1988 we moved to our present quarters, in the newly constructed CIT building. We had hoped to install our recently ordered SPARCstation 1s in time for the fall semester, but instead, Sun leased us a number of Sun 3s and we used Zaphod (the Encore Multimax), originally a research

The leftmost display and box is an Apollo DN300, acquired in late '82. Next to it is Fritzi, a Sun 2 acquired in mid-'84. The rightmost display and box is one of our first Apollos (the fifth machine shipped by Apollo), a DN400. Behind and above the Sun 2 is a Sun 3/60.

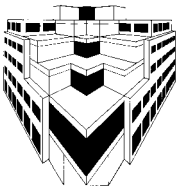

machine, as our central facility. We chose not to continue using the VAXes, but sold Sluggo and Skyler; Nancy we kept (at least its primary cabinet—it had grown over the years into two large cabinets holding an impressive 10MB of primary storage, three good-sized disk drives, and two tape drives. All these latter items were disposed of). Nancy became the basis of our computer museum (at a time when VAX-11/780s were still being used at a number of other places). The SPARCstation 1s finally started to arrive in late winter of ’89. By the summer we had enough of them that we had no further use for Zaphod, which we sold to Dick Bulterman, late of Brown and then (and now) of CWI in Amsterdam. By ’92 our SPARCstation 1s had become a bit dated, so we replaced them with SPARCstation 10s, which now form the bulk of our computer holdings.

The Early Days

The initial configuration of BUGS, built by van Dam’s graphics group, became operational in mid-’71. It consisted of a pair of Digital Scientific Meta4 processors and a Vector General vector-graphics display and was augmented with Simale (Super-Integral Microprogrammable Arithmetic and Logic Expediter) in ’75. Simale, designed and built by former undergraduate Harold Webber, had a four-processor 18-bit SIMD architecture—each processor had a 38-nanosecond cycle time for an effective peak performance of 105 MIPS. It supported real-time 3D and 4D vector graphics with matrix transformations, clipping, and dynamic level-of-detail management, and it was distinguished by never having a hardware failure in its seven-year lifetime—it was taken down only to replace light bulbs. Simale currently rests in the department’s computer museum.

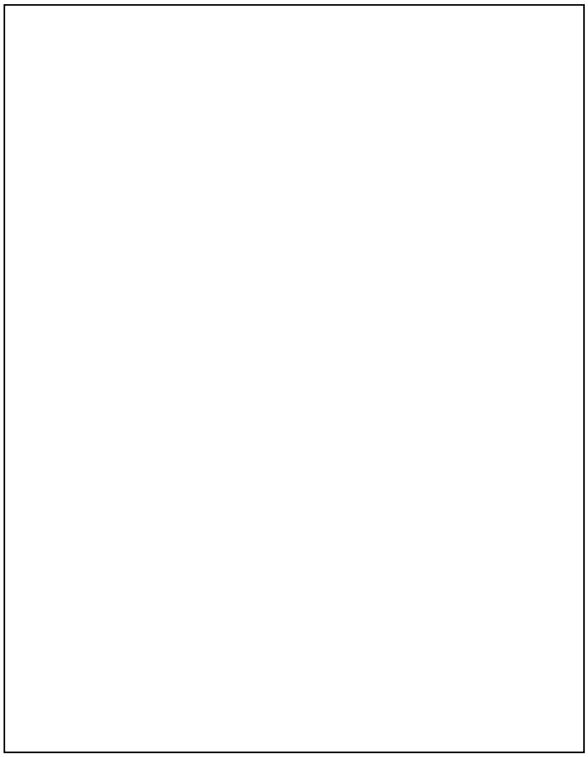

The machine on the left is Simale, 1975-82. On top of it is a VT100; on the right is Nancy, 1979-88. On top of Nancy is SSTARC (Stein, Stabler, Turrentine Automatic Relay Computer), familiar to AM100/CS100 students in the ’70s and early ’80s, where it was used to introduce logic design.

BUGS was originally installed in the University Computer Center at 180 George St. It was connected to the 360/67 via RPC used for dynamic division of labor experiments between the mainframe host and the graphics satellite. We believe this was the first published use of RPC. When we moved to the new building (Kassar House at 151 Thayer Street) in May ’79, BUGS moved to the basement, along with Nancy. Its tenth birthday was celebrated in the summer of ’81 (those dealing with the VAX were explicitly not invited). It was decommissioned in early ’82 when it became clear that there was no future in vector graphics and when work on the extension to 151 Thayer Street (Gould Lab) made part of the basement unusable.

UNIX Comes to Brown

In 1977 DEC announced the VAX-11/780. It was clear to us that this would be an ideal machine on which to run UNIX (an obscure research OS at the time); and it was also clear to us that the Program in CS (not yet a department) needed its own time-sharing system. NSF’s new equipment program for CS departments granted us a bit over $100K to purchase our time-sharing system. This was about $100K less than we needed, but, with DEC’s help, we were able to buy a VAX-11/780, configured with 512K of memory, one 67MB disk drive, and an amazingly slow tape drive. We considered purchasing a Prime 750 but, fortunately as it turned out, we stuck with our plans to get a VAX.

We intended to receive the VAX in fall ’78. However, though renovation of our building at 151 Thayer Street had begun, it was in no shape to house a computer. We knew that there was “plenty” of space in the Barus-Holley and

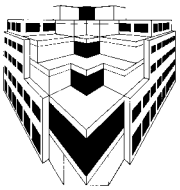

Prince Lab buildings—all we had to do was to get someone to part with some (temporarily as we hastened to point out). This turned out not to be easy. John Savage (acting director of the program at the time) and I had numerous conversations with our colleagues in Engineering and Physics. The room we thought was lined up fell through at the last minute (I placed a panic call to the loading dock of the DEC VAX factory one morning and convinced them not to ship our computer as it was about to be put on a truck.) By December we finally got a room (in Prince Lab) and the VAX arrived on January 8, 1979.

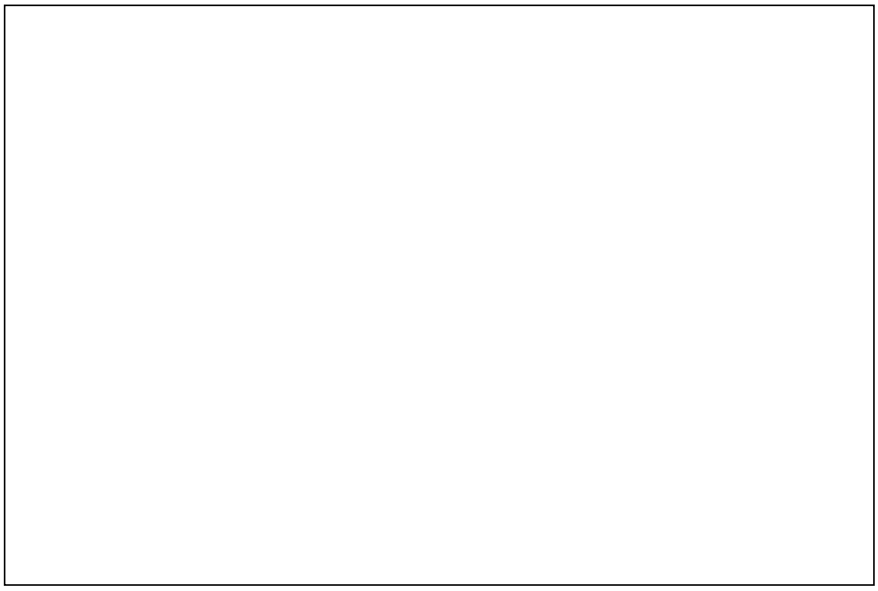

The display on the left is part of our Ramtek 9400 system, acquired in late '79. The rest of the Ramtek is the leftmost box on the floor. The middle display and box is Fritzi, a Sun 1 acquired in late '82. The rightmost display is Dumbo, a Symbolics Lisp Machine (model 3640) acquired in early '85: to the right of the display and on the floor is the rest of it.

UNIX was not quite ready for the VAX at this time (Bob Sedgewick, Steve Reiss, and I visited Bell Labs in summer '78 to check on its progress and were assured that it would be ready by early '79). So we ran VMS release 0.9, which came with no compilers and nothing of interest except for Adventure and a Scrabble game. (The student we hired to administer the system, Eric Albert, was a champion Scrabble player and enjoyed the game immensely.) In desperate need of a compiler, we became a beta site for DEC's Pascal compiler.

In May '79 the renovation of Kassar House was complete and the people of the department moved in, along with BUGS and the VAX. The building had been wired with RS232 cables and we were ready to put terminals in all faculty and student offices. Except we only owned four terminals. We had ordered four of DEC's new VT100s, but they were in short supply and we were allocated one; the others were due to arrive "soon." Even had they been readily available, they were too expensive for us to acquire in large numbers. So we settled for a cheaper alternative, the Zenith Z19 (also known as the Heathkit H19; despite our poverty, we did not acquire any in kit form). In the meantime, we had gone back to NSF and were awarded additional money. We used this to purchase another 67MB disk drive and 512K more memory. DEC helped out by granting us five VT100s, all of which arrived at the beginning of the '79 fall semester. One of these now rests in our computer museum.

In June '79 UNIX finally arrived. It was release 32V which, as advertised, ran on the VAX, but did not support virtual memory. But it had a compiler and all sorts of other nifty tools and the department finally entered the computer age. (We held no grudges and invited the graphics people to Nancy's fifth birthday party in '84.)

Once we had a C compiler, Steve Reiss wrote *b*, the first generation of the Brown editor. This was quickly adopted by most of the department and its successor, *bb*, is still used by a few diehards (including me).

With UNIX came email. Initially it was only for use in the department—we had no network connections. But in fall '79 I established our first email link to another computer—to "research" at Bell Labs via *uucp*, the Unix-to-Unix Copy program. This was a poor-man's approach to networks—point-to-point connections via phone lines. But it worked and eventually gave us world-wide (if slow and unreliable) email connections. In '82 we joined CSNET and had substantially improved mail service and in '86 we became connected to the Internet.

Late in '79 Berkeley UNIX was introduced and we were one of the first recipients of 3BSD, the first version of UNIX to support virtual memory. It was notable for being impressively slow—it compared unfavorably with VMS in many benchmarks, but even so, few people wanted to run VMS. Less than a year later we installed 4BSD, a much faster version of Berkeley UNIX, followed by 4.1BSD, which we installed as soon as it was available.

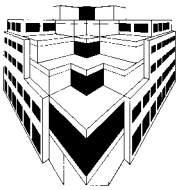

In ’82, as part of research collaboration with DEC, we were granted another VAX-11/780 (*Sluggo*) for support of graphics work. This was installed on October 6, 1982, just in time for the dedication of Gould Lab that evening. It now became important to get into the UNIX networking business, so we bought two ethernet boards and some cable and acquired an experimental version of Berkeley UNIX, 4.1aBSD (installed in November ’82), which added networking support to 4.1BSD.

In August ’83 we installed 4.1cBSD, which fixed a number of longstanding problems with UNIX, such as its file system. This was the biggest change since we started with UNIX. Despite a number of warnings, many people were caught off guard and had to do a lot of last-minute scrambling to get their code working again. However, the switch to the next official release, 4.2BSD, was made in late fall and hardly anyone noticed. 4.3BSD was introduced a few years later and again no one noticed.

Graphics, IBM, and Construction

With the demise of BUGS, the graphics group entered the worlds of raster graphics and UNIX. A special graphics room was built in the basement of Kassar and called “BURGER”—Brown University Raster Graphics Experimentation Room (constructed by graduate student and master carpenter Bill Smith). In it was installed in December ’79 our first (color) raster display, a Ramtek 9400, currently in our computer museum. A number of notable software projects used this display, including the Interactive Graphical Documents project, BRUWIN (the Brown University Window Manager), and the 4D animation project. Eventually the Ramtek was joined by a couple of Lexidata raster-graphics displays.

In ’80 and ’81 we began collaborative work with IBM. We needed IBM hardware for this work, so we began thinking about where to put an IBM computer. Serious thought was given to installing the computer in the Kassar House garage (despite the objections of those of us who parked our bicycles there). We eventually decided that the garage would go away to make room for Gould Lab, so the acquisition of IBM hardware was postponed.

When construction of Gould Lab began, a number of changes had to be made to let the construction workers use portions of the basement. BUGS was demolished and Nancy was moved into its place. (By this time Nancy had innumerable terminal connections, etc., so moving it was no easy chore.) BURGER became the construction crew’s office, so a small corner room of the basement was taken over for the graphics lab and christened “microBURGER.” Running the computers while construction was going on was interesting. Amazing amounts of dust were kicked up, so the computer areas of the basement were sealed off with plastic sheets. When construction work was particularly heavy, the machines had to be taken down so that the disk drives wouldn’t be damaged by the vibrations. Circuit boards were frequently reseated.

We were still thinking about an IBM installation and suddenly realized that it would put major demands on the basement air conditioning. We had to up the requirements for the air conditioner, which produced a considerable increase in the size of the air conditioner—so much so that the air conditioner required would not fit through any of the openings into the basement. So the construction crew removed a number of stones from the basement walls to make an opening (just) big enough to put the air conditioner through. We hired riggers (Zavota Brothers) to slide the air conditioner through the hole and set it up in the basement. This was spectacular to watch. They brought in some impressive equipment and some incredibly strong people and got the job done in seemingly no time at all, without a scratch to either building or air conditioner.

“an IBM 4381 was installed and the department had its first (and so far, only) IBM mainframe”

Finally, pretty much at the last minute, everything was cleaned up in time for Gould Lab’s dedication on October 6, 1982. There was now room in the basement, so an IBM 4381 was installed and the department had its first (and so far, only) IBM mainframe, which was used for research on text processing. It was removed a year or so later.

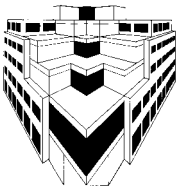

Graphics moved out of the basement and into a spiffy lab within Gould Lab. BURGER was history. The Ramtek and Lexidatas were moved in and were joined by high-end Apollos and a top-of-the-line Evans and Sutherland PS-300 vector display (which now rests in the computer museum).

Workstations

We became intrigued with the idea of workstations in the late '70s when we heard about what was going on at Xerox PARC. Finally in 1980 Three Rivers Computer announced the Perq workstation (but didn't deliver it until much later). This at least made it clear that workstations were about to become commercially available. One of the things that we wanted to use workstations for was instruction, so we applied to NSF's CAUSE program (Comprehensive Assistance to Undergraduate Science Education—a program that got the axe under the Reagan administration) and were awarded $150K. Workstations back then were being quoted for ~$35K each (they weren't being delivered yet), so this was not a whole lot of money. But it was something and we started searching for a workstation vendor.

DN400 of '82-'83 vintage in Foxboro Auditorium. On-screen is a very early version of BALSA

It was clear to us that we had to make ourselves look exciting so that we could get some assistance (i.e., attractive discounts) from vendors. We put together a brochure describing our needs and our vision for instructional computing. We commissioned an artist to draw a picture of our proposed lab for the brochure and we designed Gould Lab to feature a computerized classroom (holding up to sixty workstations), which became known as the Foxboro Auditorium. We sent the brochure to a number of prospective vendors and donors.

We narrowed things down to three serious potential vendors: Three Rivers, Xerox, and Apollo. Three Rivers was the early favorite, since they had actually announced a commercial product. Xerox, unlike anyone else, had actually produced workstations. Apollo was run by people who were already successful in the computer business (a number had come from Prime and had unsuccessfully attempted to sell us a Prime 750 a few years earlier). We had pretty well decided upon Apollo, but then new developments occurred at Xerox, so we delayed our order. This cost us the honor of receiving the first Apollos shipped. Things became clearer at Xerox in a few days so we put in a firm order to Apollo for seventeen workstations, two with disk drives (33MB each). We received the fifth and sixth machines shipped, in March '81. One of these, "node C," now rests in our computer museum (it was retired from active duty in May 1988). The other machines trickled in and students began to use them in the '81/'82 academic year. One of the more popular first applications written was "PACman," a copy of a then-popular computer game[1].

"we put in a firm order to Apollo for seventeen workstations, two with disk drives (33MB each)"

We became big proponents of workstations and Apollos, giving numerous demos for Apollo's potential customers. Marc Brown, a grad student/staff member, founded the Apollo Users

1. Kassar House was named after Ray Kassar, then the president of Atari, the owner of the rights to PACman, rights that it fought aggressively to maintain. An inadvertent poor move on my part was that someone was playing our bootleg PACman on the Apollo in my office when Kassar stopped by my office during the dedication of the newly named Kassar House in May '82.

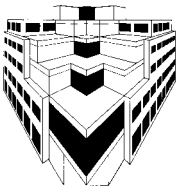

Group and organized its first meeting, held at the Biltmore Hotel in downtown Providence in the summer of ’82.

Great progress was made in instructional software and in Bob Sedgewick’s and Marc Brown’s research on algorithm animation (resulting in *BALSA*). This pioneering work helped us immensely in obtaining our first major equipment grant—$3 million over five years from NSF’s CER (coordinated experimental research) program, awarded in May 1982. We used this grant to purchase additional Apollo workstations, used for research, as well as additional memory and disk drives, etc., for Nancy (the first VAX).

The first generation of Apollo workstation was rather bulky. When the Foxboro auditorium was ready in September ’82, though we only had seventeen machines, they occupied a fair amount of space. The next generation of Apollo workstation, the DN300, was considerably smaller[1]. With the help of a generous grant from the Exxon Education Foundation, we were able to purchase enough DN300s to populate the Foxboro Auditorium. As mentioned earlier, these arrived all at once at the beginning of the fall semester in ’83. One of these machines, “Node 87C,” remains, resting in the computer museum. Rather than the two file servers serving fifteen diskless machines in our first-generation classroom, we now had sixty diskless instructional machines (along with fifteen research machines, each with a 33MB disk) served by sixteen file servers. We finally had the instructional facility we had described in our brochure three-and-a-half years earlier.

Foxboro stocked with DN300s running a later version of BALSA

The Apollos lasted us for the remainder of our days in Kassar House and Gould Lab. But by 1987 it was clear that our Apollos were rapidly becoming obsolete. So we embarked on another round of choosing a workstation vendor. We talked to a number of vendors and things finally worked out into a contest between DEC, NeXT, and Sun (Apollo dropped out, since they didn’t have a machine that met our needs at our price). DEC proposed a VAX-based workstation. This seemed to be a safe choice—it would run the same operating system we were running on the big VAXes. However, the workstation VAX was pretty slow compared to the machine Sun was proposing. (Little did we know that DEC had just started a project to produce a MIPS-based workstation. They couldn’t tell us about this until several months later, but even so, we would never have believed that it would be shipping in nine months.) NeXT seemed pretty exciting, and we were the subject of a very impressive sales pitch by Steve Jobs. Sun was proposing some exciting hardware, but they didn’t have it ready to demonstrate for us. After a lot of discussion and a few benchmarks, we decided in April ’88 to go with Sun.

None of the old Apollos made the move to the new CIT building, except for the two that went to the museum. Unfortunately the Sun SPARCstation 1s weren’t ready until March ’89, so the new teaching lab, christened the “Sun Lab,” was filled with Sun3s, leased to us through Sun. But by summer ’89 we had a full complement of SPARCstation 1s.

Our most recent round of vendor selection started in ’91 and finished in early ’92 by choosing Sun again. Our current SPARCstation 10s are aging rapidly and will soon be considerably slower than the PCs owned by many of our undergraduates. We have a recapitalization plan in place; the next chapter in the workstation story will begin to unfold within the next couple of years.

1. Bill Poduska and Dave Nelson, two of the top management at Apollo, came to visit us in the spring of ’83 and brought with them an early DN300 (in the back of Nelson’s Mercedes convertible), tied up with a red ribbon, as a gift for Andy van Dam, then chairman of the department.

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

One year ago, the department adopted the World-Wide Web as the medium for information delivery locally and to the Internet. Previously a student-run effort, the local Web has grown by hundreds of megabytes since then. Today it is home to department publications (including this one, http://www.cs.brown.edu/publications/conduit), research, course information, documentation and much more. Over the past year, a common refrain has been "put it on the Web!"

Computer Science's home page, http://www.cs.brown.edu/, presents information about the department in broad categories. The iconography of graphic-artist-in-residence Dan Robbins associates an image with each category and helps webnauts navigate the Web's hypertext space. With oversight from technical staffer and Webmaster John Bazik, the Web is constantly growing and changing—the product of dozens of authors.

John Bazik

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

http://www.cs.brown.edu/people/faculty.html

Course TAs have made the most effective use of the Web. Nearly all CS courses now offer syllabi, missives, handouts, notes and assignments via the Web. A few supply all lecture slides shown to date, and some use the Web as an interface to launch program demos. In an experiment this semester, CS4's Web offers not only the slides from each lecture, but also, at the click of a button, the lecture itself. Recorded live during class, audio is cut into slide-length bites and placed alongside the respective slide images in the Web. Students now have the opportunity to hear again the difficult passages lost to momentary distraction or somnolence.

Several research groups have established home pages, organizing information about projects, published papers and events. The robotics group and the object-oriented data-

base group are both active. The graphics group's Web supports their research collaboration with other schools, so staffers have made it a substantial document reflecting their work.

Some of the most useful Web services are those that are generated automatically. All techreports published by the department are submitted and distributed electronically. Their home page (http://www.cs.brown.edu/publications/techreports/) is formatted conveniently for browsing and downloading.

Every member of the faculty and staff, all graduate students and many undergrads have home pages. There are occasional gems to be found in these personal public spaces, and for those who know better than to take themselves too seriously, the Web o' Fun (http://www.cs.brown.edu/fun/) provides a home for the amusing and the strange.

http://www.cs.brown.edu/research/graphics/research/cmg.html

activities@cs.brown.edu

Thomas Dean. This spring Tom was lecturing at USC and the Santa Fe Institute. He is now on the Board of the Computing Research as the representative from the AAAI.

Thomas Doeppner. Tom gave an invited talk at the DCE Developers Conference in Boston in August.

Maurice Herlihy. Maurice gave an invited talk entitled "Algebraic Topology and Distributed Computing" at the Distinguished Lecture Series at SUNY, Stony Brook.

Franco Preparata. Franco presented two invited lectures on models and techniques of parallel computation at the PARLE94 Symposium in Athens, Greece. Subsequently, with M. Rabin, K. Mehlhorn, and A. Rosenberg, he lectured for a week at an international school on Parallel Computation in Lipari, Italy. In the fall of '94 he was named a Fellow of the Japan Society for the Promotion of Science (JSPS) and was invited to spend his sabbatical leave at the Department of Information Science of Kyoto University. While in Japan he presented lectures at several universities and was the selected speaker at the Tokyo IEEE and ACM yearly meetings. He also visited and lectured at the University of Singapore at the request of the JSPS as an associate member of a Japanese scientific delegation. Recently, he was elected a Fellow of the Association for Computing Machinery for "significant research contributions in Computational Geometry, Parallel Algorithms, Theory of VLSI Layouts, Fault Diagnosis in Computer Systems, and Algebraic Coding Theory."

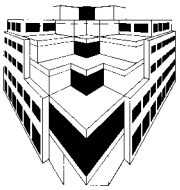

John Savage. In March, John was a panelist at the ACM 1995 Computer Science Conference. The panel discussed "Is Computer Science Obsolete?" In April he gave an invited talk at the NECUSE Forum on Parallel Computing Curricula at Wellesley College. John has again taken the reins as Director of our Industrial Partners Program, which he founded in 1989.

Andries van Dam. Andy has co-authored a book with graduate students D. Brookshire Conner and David Niguidula called *Object-Oriented Programmin in Pascal, A Graphical Approach*. It will be published in April, 1995. Other activities are detailed in the chairman's column.

Pascal Van Hentenryck. This spring, Pascal is giving invited talks at LIX in Paris and at CCP '95 in Venice. His edited book *Principles and Practice of Constraint Programming*, based on the April 1993 Newport workshop co-organized with Paris Kanellakis, is due to appear in April from MIT Press. Pascal is also on the program committee of CP '95, the international conference on constraint programming.

Peter Wegner. In February Peter was elected a Fellow of the ACM. He assumed his duties as Editor in Chief of *Computing Surveys* in January and encourages ***conduit!*** readers to submit contributions (pw@csbrown.edu).

As a member of the ACM Publications Board, he is involved in planning the transition to electronic publication (*CACM*, April, 1995). He has been invited to talk in Sweden in May, in Denmark in August, and in Germany in September.

WHAT'S IN A NAME?

"That which we call a rose By any other name would smell as sweet," but can that be said of a computer? "A good name is rather to be chosen than great riches," or so it would seem to the CS cognoscenti. Not only do computer names reflect the personalities and interests of their users, they serve to amuse and confuse—read on!

Cartoon characters are popular choices—*rocky, natasha* and *boris* are file servers, *igor* the firewall machine, *bullwinkle* the mail server, and *wilma* the anonymous ftp server. If you recall the article about faculty hobbies in the spring '93 issue of ***conduit!***, the following machine names will make good sense: Tom Dean's is *klee*; among Stan Zdonik's are *mandolin, lute, fiddle*, and *dulcimer*. Rob Netzer's machine is *bugs*—the name inspired by his research area, debugging, not by Bugs Bunny, as is usually supposed. Peter Wegner's machine, *geneva*, was so named because he had just spent a summer at the University of Geneva. When John Savage was chairman he decided to name his computer *icarus* to remind himself to keep out of the sun; he called his other one *daedalus*. One of Steve Reiss's machines is named for his son, Fred; however, it was with abject horror that his message "My Sun Fred is dead" was once received by the technical staffer on duty! Eugene remembers thinking the message too blasé even for Steve; there had to be another interpretation.

Peter Galvin: Max and I were setting up the system for a new faculty member (Philip Klein) and needed a name. We decided to welcome this new person to Rhode Island by naming his system after a famous RI tradition—wieners! Of course, we knew it wasn't an overly affectionate name, either. Upon this new faculty member's arrival we learned that he disliked the name and had asked for it to be changed to *remington*. It was then that we got nervous. It was only a few days later, when we learned Philip was a typewriter enthusiast and was naming his computer after an old model, that I started sleeping more soundly.

Philip Klein: My first workstation was originally named *wiener* when I arrived. (I suspect this was not a friendly act, but you'd have to check with Peter Galvin for the ori-

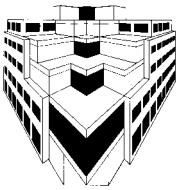

gin.) I requested it be renamed *remington*. Peter inferred that I must be a gun nut. I'm not—the machine was named for my old Remington Noiseless (ha!) mechanical typewriter; another one is called *underwood*. You might want to mention my Digicomp computer—made of plastic, has three bits of memory—I call it *ug*.

Tom Doeppner: *karla* is the name of the Russian master spy who appeared in a number of John Le Carre novels, and who was convinced to surrender to the British in *Smiley's People*. *botrytis* (a research machine), is otherwise known as the "noble rot"—the fungus that attacks grapes that have been left on the vine for longer than usual, shrinking and concentrating the grape and producing a very sweet grape that makes expensive and wonderful dessert wine.

John Hughes: Mine, *euclid*, reflects the graphics-group rule of naming all machines after streets in Providence (Euclid is the little street just past Store 24 on Thayer), combined with my desire to have a name that reflects my own interest in geometry. If there were a Riemann street in Providence, perhaps I'd have used that. Interestingly, Andy van Dam's is *power* (no need to elaborate).

Eugene Charniak: My machine, *bohr*, is, of course, named after the physicist. My undergraduate degree is in physics, I have an ongoing interest in the history of science, and the great scientists have always been heroes to me. I selected Bohr over others first because it is a short name, and second because of a funny story about the name. Many years ago Robert Brustein, formerly head of Yale Repertory Theater and now of American Repertory Theater in Cambridge (theater is another interest of mine), wrote an article about the relevance of quantum mechanics and things like the uncertainty principle to modern theater. The article was named "Theater in the Age of Einstein." A letter to the editor from a scientist pointed out that Einstein had little to do with either quantum mechanics or the uncertainty principle, and suggested that from that point of view a better name for the article would have been "Theater in the Age of Bohr"—though the scientist could understand why that title might not be appropriate either.

CS FOOTBALL: GEEKS NO MORE

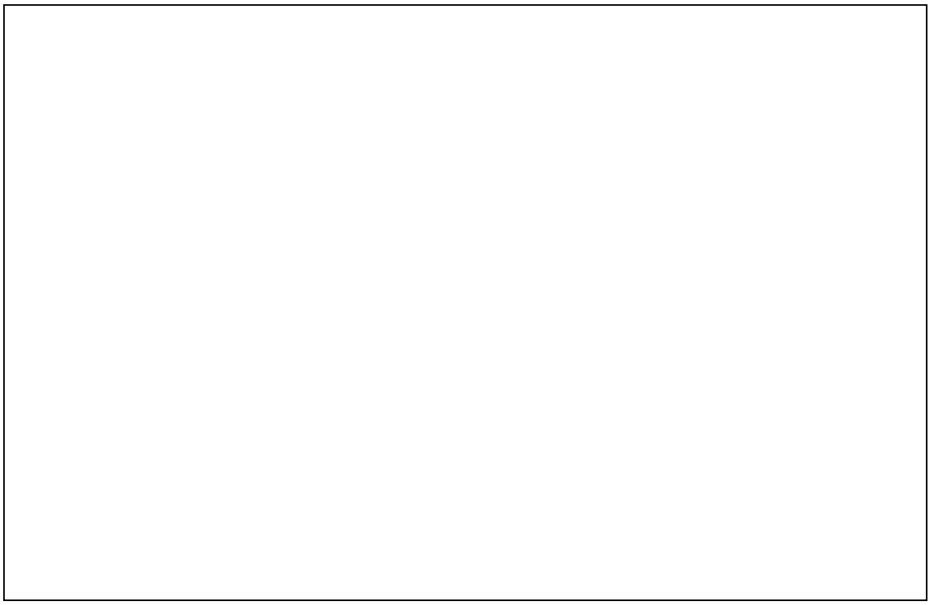

Front row: l-r, Kevin Regan, Bryan Cantrill, Bob Zeleznik, Matt Ayers
Back row: l-r, Kathy Davidson, Michael Littman, Lowell Kaplan

This year the CS intramural football team took the low-intensity league by storm, going undefeated in the regular season and racking up such high-flying scores as 42-6 over Marcy's Mercenaries and 48-14 over Alpha Delta Phi. Thanks to the protection of offensive line Kevin Regan and Matt Ayers, hacker-cum-quarterback Bob Zeleznik had plenty of time to look deep for the likes of Phil "Twinkletoes" Shen, Bryan "Crazylegs" Cantrill, William "Let's run the halfback fake" Etienne or Charlie "Tennis comes first" Hoecker. When the long bomb didn't work out, then a quick strike to playmakers Ben Boer, Michael Littman or Kathy Davidson was always guaranteed significant yardage. And when even that wasn't open, "Coach Z" would find Brown varsity football recruit Tim Brennan running a post for a cheap touchdown. During those rare times when opponents had the ball, defensive phenom Lowell Kaplan became a one-man

scoring sensation, running innumerable interceptions the length of the field for touchdowns.

The sweetest moment of each victory, however, occurred after the game, when the inevitable question was asked: "So which frat are you guys from?" "We're no frat," the response came, "we're the CS department." Then, usually tossed in, "Don't call us 'geeks' anymore and we won't have to tell your friends. Okay?" Uh, by the way, we lost a very close playoff game to close out the season—just wait till '95 though...

FROM THE CHAIRMAN, Eugene Charniak

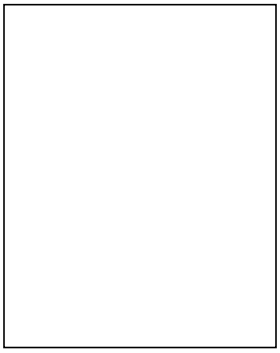

Eugene Charniak

Let me start out by noting that the e-mail from alums has slowed down. As I mentioned when we first started the e-mail from alumni column, I expected that most of you would either feel that the events in your lives were too uninteresting to warrant inclusion, or that you were too busy to write in. For several issues now you have proven me wrong, but things have slowed down a lot these last six months. Do send in contributions! We did, however, get a good reaction to the "fashion clone" picture from the last issue—the e-mail and my response:

Date: Mon, 21 Nov 1994

From: Ed Lazowska, lazowska@cs.washington.edu

So, listen Eugene, we all got the latest issue of ***conduit!*** *today, and we think the department should be incredibly embarrassed that an old geezer like van Dam has the best looking set of legs in the building. What are you going to do about this?*

No problem. I have a call in to Tonya.

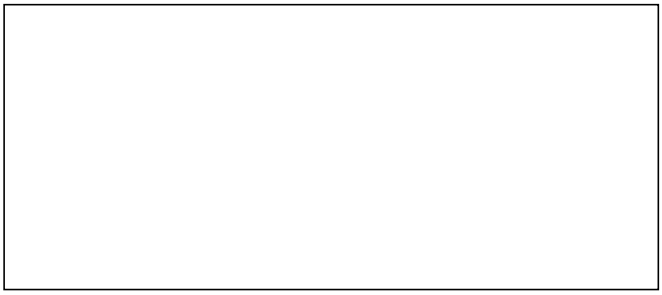

The legs in question are third from the left!

And while we are on the topic of Andy van Dam, I am happy to announce he has had his chair switched out from under him, so to speak. For several years now Andy has been the L. Herbert Ballou University Professor of Computer Science. Just this year IBM was kind enough to sponsor a new endowed professorship at Brown, in memory of T.J. Watson Jr., a Brown alum. Thus Andy is now the first T.J. Watson Jr. University Professor of Technology and Education. I might also note that, effective July 1, Andy will become the Director of the so-called "Gang of Five,'' more formally the five-university NSF/ARPA Science and Technology Center for Computer Graphics and Scientific Visualization.

The face in question is third from the left!

Other notable faculty awards in the last few months include Peter Wegner's and Franco Preparata's election as Fellows of the ACM. They join Andy van Dam to make three out of our 16 faculty members ACM Fellows. Franco's citation is detailed in the "activities@cs.brown.edu" column; Peter does not remember a particular citation, but expects he was cited for his work in the area of programming-language semantics and object-oriented programming.

Three Ph.D. students have finished up since the last issue of ***conduit!***— Glenn A. Carroll, whose thesis was entitled "Learning Probabilistic Grammars for Language Modeling,'' Gerd Hillebrand, who did work on "Finite Model Theory in the Simply Typed Lambda Calculus'' and Sridhar Ramaswamy, with the thesis "Indexing for Data Models with Classes and Constraints.'' Congratulations to all three.

Normally I only note when students pass their Ph.D. Defense, but we did have a nota-

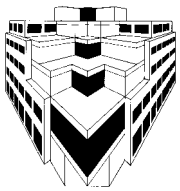

ble Ph.D. proposal a few months ago. As I mentioned in my column a few ***conduits*** ago, Jeff Vitter left our department to become the chair of Duke's Department of Computer Science. Naturally he had several students who were at various stages of progress toward their Ph.D. Those students who had passed their comprehensive exams here on one hand did not want to start over again to pass exams for a Duke degree, but at the same time wanted to be with their advisor. There are methods of handling this, and the students were accommodated. Recently one of the students, P. Krishnan, was ready to make his Ph.D. proposal presentation, and Jeff wondered if it would be possible to do the presentation and exam by video link. This was feasible because the members of the Gang of Five graphics consortium have a 24-hour-a-day video hookup among them and University of North Carolina (UNC), ten miles from Duke, is one of the group. Thus the Duke contingent went to UNC and the Brown contingent went to our Lubrano conference room, which is hooked up for this. The good news is that the technology worked flawlessly. The bad news was that many felt the normal dynamics of the exam situation were badly thrown off by the technology, so it was not a complete success. Nevertheless, another step, however halting, in the communications revolution!

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
Writer & Editor

Jeff Coady
Technical Support

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

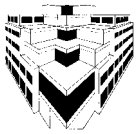

conduit!

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

Printed on recyled paper

Address changes welcomed