# conduit!

Volume 1, Number 1 | Department of Computer Science | March, 1992
Brown University

## Development of the Industrial Partners Program

*John E. Savage*
*Chairman*
*1985-1991*

In March, 1989, ten years after the founding of the Department, we introduced our Industrial Partners Program (IPP). Those first ten years were challenging and rewarding. Our efforts earned us an excellent reputation, reflected in a gifted faculty, talented students, and superb administrative and technical staffs. During this period we also acquired superior computing equipment and excellent space in the handsome new Thomas J. Watson, Sr. Center for Information Technology. The time had arrived to build upon our good relations with industry and introduce the IPP, our vehicle for strengthening and expanding these relations.

### Our Relations with Industry

In 1982 we dedicated Foxboro Auditorium, the world’s first workstation-based classroom. By the beginning of the 1983-84 academic year this unique undergraduate environment was populated with 55 new Apollo DN300 desktop workstations. This impressive facility was visible evidence of the commitment of the Department and the University to the development and application of the new workstation technology, and Foxboro Auditorium attracted a great deal of attention to our research into the creative use of workstations. This auditorium, named by the Foxboro Corporation, is in Gould Lab, funded by the Gould Corporation, which adjoined Kassar House, named by the President of Atari Corporation for his father.

In 1981 the Department acquired an NSF equipment grant for instruction that allowed us to purchase our first workstations and at a time when this new technology was just emerging. In 1982 and 1988 we received NSF research infrastructure grants that let us provide workstations and staff support for faculty and student research. These grants helped us build our enviable reputation for the creative use of this technology. Our visibility in this area led to a number of multi-million dollar equipment grants and discounts from Apollo Computer, Digital Equipment Corporation, Hewlett-Packard, Sun Microsystems, and Xerox.

The Department moved to the new Thomas J. Watson, Sr. Center for Information Technology (CIT) in 1988. This facility, which provides office space for the Department, also houses the University's central computing staff, Modern Languages Laboratories, the mainframe, a computer store, and a computer repair facility. Its most distinctive features are its two lecture theatres and two classrooms outfitted with workstations, networked personal computers and video projection equipment. One of the lecture theatres is operated by the Department and has 65 color Sparcstations for undergraduate instruction. The others are generally available to the cam-

*Workstation Classroom*

pus. Major corporations, including Apple Computer, Claris Corporation, Dynatech, Hutchison Telecommunications, IBM, Motorola Corporation, the Novell Corporation, and Sun Microsystems as well as several individuals, have named space in the CIT.

## The New Program Takes Shape

To capitalize on our successful interactions with industry, Brad Clompus of the Development Office, Roy Bonner, a retired IBM manager in residence at Brown, and I began in the summer of 1988 to flesh out plans for our Industrial Partners Program. Joan Cerjanec joined the planning effort as it drew to a close and served as our first IPP Program Officer. In March, 1989, the IPP was introduced to senior executives of potential Partners with great success. At this event faculty members gave an introduction to the Program and its objectives as well as an overview of our research on artificial intelligence, computer graphics, databases, operating systems, multiparadigm design environments, theoretical computer science, and VLSI.

IPP PARTNERS 1992

Bellcore
Citicorp
Codex
DEC
GTech
IBM T.J. Watson Lab
Motorola
Siemens
Sun Microsystems

An important goal of our Program is to provide Partners with opportunities for pre-competitive cooperation. We do this through a series of technical symposia on topics of current interest to our Partners. Symposia are typically day-long events at which most of the speakers are Partner representatives or other non-academic visitors; usually only one or two speakers are Brown faculty or students. This arrangement encourages interaction between Partners and provides a program more likely to appeal to them.

*"An important goal of our Program is to provide Partners with opportunities for pre-competitive cooperation."*

For us at Brown, technical symposia provide an important window on Partner interests, helpful to faculty and students in understanding the concerns of the commercial world. Partners gain visibility in the Department, which is valuable for recruiting purposes, and have access to faculty and student research. The value of the latter is illustrated by the experience of Digital Equipment Corporation. At our first technical symposium representatives from Digital learned about FIELD, the software development environment produced by Professor Steve Reiss. When their need for such a product emerged, they revisited us and ultimately, in 1990, licensed FIELD; as FUSE, it is now one of Digital's product offerings.

## Content of Technical Symposia

Each of our eight technical symposia so far has had the goal of treating an important contemporary research or development topic from the perspectives of business and industry as well as academia. To insure that these perspectives are reflected in our symposia, most of our speakers are non-academics, as illustrated by the table opposite.

It is very exciting to see individuals from competing firms at our symposia engaged in lively discussions. Such interactions help all participants better to understand the decisions they are making, it increases confidence in these decisions, and provides an extremely healthy form of pre-competitive cooperation. We take seriously our role in shaping the agenda for symposia and providing a neutral ground for discussion, and we welcome suggestions for future technical symposia as well as for other types of events of value to our Partners.

## The Current Organization of the IPP

This year the Department has a new Chairman, Eugene Charniak, and a new Program Officer, Suzi Howe. Eugene, on the faculty since 1978, is an expert in artificial intelligence specializing in natural language understanding, and is co-author with Drew McDermott of the textbook *Artificial Intelligence*. Suzi joins us from Computing and Information Services, the group providing centralized computing support on campus. Eugene has overall responsibility for the Program and Suzi maintains day-to-day contact with Partners, helping them communicate with faculty and students and acquire information on the Department. She also coordinates technical symposia and is editor of this newsletter.

## Interaction with Faculty and Students

An important function of the IPP is facilitating Partner contact with faculty and students. Periodically Partners have job opportunities appropriate for our graduates. Suzi will circulate information on such opportunities through our undergraduate and graduate student communities, make introductions to students, forward resumes, and otherwise help to fill such positions. From time to time she will also contact Partners, especially at mid-year or late in the spring semester, to notify them of students who have just come on the job market.

Partners occasionally need advice or pointers to scientific literature. Here Suzi can also help by putting Partners in touch with appropriate faculty members or graduate students. In the past we have been able to find students and faculty members to offer lectures or short courses for Partners. In such situations, Partners will settle privately for the lectures and courses.

| DATE | TOPIC | SPEAKER ORGANIZATIONS | BROWN ORGANIZERS |
|---|---|---|---|
| 7/13/89 | Prototyping Environments | Object-Design, Inc.<br>Siemens<br>Brown IRIS<br>Brown CS (3 times) | Prof. Reiss<br>Prof. Zdonik |
| 2/21/90 | Scientific Visualization | Stardent, Sun, HP<br>Bellcore, Brown CS<br>Brown Applied Math | Prof. van Dam<br>Prof. Hughes |
| 5/3/90 | Experiences with the Object Paradigm | TI (3 times)<br>Siemens (twice)<br>Bellcore (twice)<br>Codex, Brown CS | Messrs.<br>Lejter, Kirman<br>Shewchuk<br>(Grad Students) |
| 7/12/90 | Robotic Systems Design | Denning Robotics<br>Transitions Research<br>Design Lab<br>Brown Linguistics,<br>Engineering & CS | Prof. Dean |
| 10/17/90 | OSF and UI Operating Systems | *Day One* -<br>Brown Tutorial<br>*Day Two* - OSF, Sun<br>Bellcore, DEC, IBM<br>Unix International<br>GTech | Prof. Doeppner |
| 3/14/91 | Parallel & Distributed Systems | DEC, Citibank, IBM<br>HP, Motorola, Transarc | Prof. Savage |
| 7/18/91 | Programming Techniques for Constraint Problems & Combinatorial Optimization | IBM, Siemens<br>Bellcore (twice)<br>Motorola, Brown CS<br>(twice) | Profs. Kanellakis,<br>Van Hentenryck |
| 11/7/91 | Privacy & Security | IBM (twice), Citibank<br>DEC (twice), Sun<br>Brown Philosophy | Prof. Doeppner |

The Department has approximately 50 Master's and 50 PhD students and seven technical staff who manage our complement of about 170 workstations. We award between 50 and 60 Bachelors degrees each year. We are well balanced in our research interests and represent just about every area of computer science. There are many ways in which a community of this size and diversity can be useful to our Partners. Eugene and Suzi are happy to discuss all such ways with Partners.

Ours is a young and exciting Program designed to meet the interests of our Partners. For this to continue we need your feedback. We would also like to expand the Program somewhat while keeping the number of Partners to at most approximately 12. Your suggestions for other Partners, especially those with whom you would like to interact, are most welcome.

## 1991 ACM SIGGRAPH Award to Andy van Dam

*Reprinted by kind permission of Professor Bertram Herzog of the University of Michigan. The article appeared in the July issue of Computer Graphics magazine.*

The 1991 Steven A. Coons Award for Outstanding Creative Contributions is presented to Dr. Andries van Dam for his unwavering pursuit of excellence in the field of computer graphics, his contributions to computer graphics education, and related fields. He is a stimulator and a leader, as shown by his key role in founding SIGGRAPH.

Van Dam entered the computing field in 1960 via pattern recognition and focused his doctoral dissertation on digital processing of pictorial data, inspired by Ivan Sutherland's seminal Sketchpad film. From the beginning, dealing with pictures has been central to his interests in computing. As a consequence of his early work, he always had a keen interest in the synthesis of the two areas: computer graphics and image processing. This led to his role in the establishment of the *Journal of Computer Graphics and Image Processing*, of which he was an editor from 1971 to 1981.

As one of the founders of the Computer Science Department at Brown University, and, as the first and two-term chairman, van Dam has been influential in that department's eminence. He always emphasized that computing is most effective through the crossing of disciplines and sub-disciplines. Thus, he and his graphics group worked in distributed graphics on multi-processors and networked computers in the early 1970s, nearly two decades before this topic became fashionable in mainstream computing.

> *"computing is most effective through the crossing of disciplines and sub-disciplines"*

In addition, van Dam was also an early proponent of, and contributor to, hypertext and hypermedia. It was through his presentation in the 1960s that many professionals were exposed to these concepts, long before they were recognized by the computing community at large. Under his leadership, nearly 20 years of hypertext and hypermedia graphics research was conducted before the first CM conference on this topic was held.

*Andries van Dam*

Professor van Dam played an important role in fashioning the professional status of computer graphics. Together with Sam Matsa of IBM, he presented the first ACM Professional Development Seminar in Computer Graphics, both in the United States and Europe. He and Matsa then founded the ACM Special Interest Committee in Computer Graphics, which evolved into SIGGRAPH.

In addition, van Dam was a prime mover in the 1976 launching of the SIGGRAPPH Core-Standards group that published, through SIGGRAPH, key specification documents. This work led to the formation of the ANSI X3H3 Technical Committee on graphics standards. Van Dam revisited that arena a decade later when he observed that the resulting GKS and PHIGS standards did not adequately support modern graphics workstations. He then cajoled a group of interested participants into producing a draft of PHIGS+, leaving it again to the formal committees to create a formal PHIGS+ standard.

Van Dam, with Brown University colleagues, introduced and established the concept of networks of graphics workstations for teaching and research, well before the term was coined in the computer graphics field.

Finally, this citation would be incomplete without mention of van Dam's literary achievements including, *Fundamentals of Computer Graphics*, co-authored with colleague James D. Foley in 1982, and the recent *Computer Graphics: Principles and Practice*, with J.D. Foley, S.K. Feiner, and J.F. Hughes.

SIGGRAPH has a tradition of recognizing individuals who have made major and long-term contributions to the field of computer graphics. This year's Coons award to Professor van Dam exemplifies that tradition. Professor van Dam is a tireless worker, an inspiration to students and a fine example to the industry as a whole. Andries van Dam has always had the right vision of what is important in computing, in computer graphics, and in related fields. Even as he was engaged in these and other technical activities, he found time to convince the computer graphics industry of needed improvements. He is truly one of the important people in our field — he has dedicated his talent, inspiration, and time to encourage everyone to excel.

### Previous Award Winners

| | |
|---|---|
| 1989 | David C. Evans |
| 1987 | Donald P. Greenberg |
| 1985 | Pierre Bézier |
| 1983 | Ivan E. Sutherland |

*Eugene Charniak*

# From the Chairman

We graduated a record number of PhDs this last academic year — eleven. They are:

| Name | Dissertation Title | Employer |
|---|---|---|
| Mark S. Boddy | "Solving Time-Dependent Problems" | Honeywell SRC |
| Robert Goldman | "A Probabilistic Approach to Language Understanding" | Tulane U |
| Eric J. Golin | "A Method for the Specification and Parsing of Visual Languages" | U Illinois Urbana |
| Cheryl L. Harkness | "An Approach to Uncertainty in VLSI Design" | HP Design Tech Ctr |
| Richard P. Hughey | "Programmable Systolic Arrays" | UCSC |
| Robert Ravenscroft | "Generating Function Algorithms for Symbolic Computation" | U Waterloo (Postdoc) |
| Peter Revesz | "Constraint Query Languages" | U Toronto (Postdoc) |
| Andrea Skarra | "A Model of Concurrency Control for Cooperating Transactions" | AT&T Bell Labs |
| Lynn Andrea Stein | "Resolving Ambiguity in Nonmonotonic Reasoning" | MIT |
| Markus G. Wloka | "Parallel VLSI Synthesis" | Motorola |
| Felix W. Yen | "CI2–A Logic for Plural Representation" | Space Telescope Science Institute |

It is interesting to note that their current positions are almost evenly split between industry (5) and academia (6). I suspect the size of this class is, in part, just a fluke, but it also reflects the increase in the size of our PhD program over the last few years. I should also mention that our record number of PhDs helped propel Brown as a whole to a record number of PhDs at the last Commencement.

We have a new addition to our faculty, Leslie Kaelbling, who received her PhD from Stanford in learning and robotics. We were looking for someone whose research was in the AI area, and we are very happy that we were able to attract her. The next issue of ***Conduit!*** will include an article about her.

In July of 1991 the Department received a $2.6 M grant from DARPA for research on high-performance design environments. This research starts from the observation that, while computers are becoming more powerful, programming them (or more generally, solving problems with them) is not becoming easier. In some cases, for instance on highly parallel machines, it is becoming harder. The goal is to create environments for design (and particularly program design) to alleviate this problem. We hope to do this in part by design aids particularly suited to the environment, and in part by leveraging the power of the machine itself to make more powerful (and computationally expensive) tools feasible.

Although it is hard to believe, it is now four years since the last time we set out to buy new departmental computing equipment. At that time we were replacing our aging Apollos, and we chose to do so with the Sun Sparcstation 1's we now use. Four years is a long time in this business, and it should come as no surprise that these machines are no longer state of the art. We purchased the Suns with funds from a five-year NSF "institutional infrastructure" grant. We are now nearing the end of that grant, and we have the final year's money available to replace our equipment again. Thus we are now in the process of finding out what is about to come on the market. My impressions so far are 1) what the chip makers can do is simply incredible, and 2) desktop workstations are in the process of becoming a relatively standardized commodity. It is definitely going to be a tough choice.

Philip Klein

# In Search of Something Short of Excellence

*The Department is very proud that Phil Klein has won a* ***Presidential Young Investigator*** *award. This award is given primarily for research accomplishments, but excellence in teaching also plays a role. The PYI award encourages contact with industry via its matching funds arrangements, and we would like to take this opportunity to thank Xerox and Thinking Machines for their grants to Phil under this program.*
*Eugene Charniak, Chairman*

‘‘Lower your standards,” Phil Klein advises his students. ‘‘Avoid perfectionism.” Klein isn't advocating that his students do anything less than high-quality work. He's encouraging them to work in the burgeoning research area of *approximation algorithms*.

> *“Lower your standards.... Avoid perfectionism”*

In 1975, Richard Karp of U.C. Berkeley made a discovery that rocked the world of algorithm designers, people who design ways for computers to solve complicated problems. ‘‘Karp demonstrated the ubiquity of intractability,” Klein explains. ‘‘He showed that very hard-to-solve problems were all around us, that many problems currently under attack were in fact unlikely to yield, and that the boundary between easy and difficult problems was narrow indeed.”

Klein cites the example of matching, a problem that might arise, e.g., in the airline industry. An airline must match up flight crews with scheduled flights in a way that satisfies as many constraints as possible (e.g., the flight crew should be located in the flight's originating city). Figure A (on Page 7) schematically represents such a problem. The problem of finding a best matching—one that matches the greatest possible number of flights to suitable crew—is a well-understood and easy-to-solve computational problem.

A slightly more sophisticated version of this problem, however, turns out to be intractable. Say that the airline needs to construct *triples*, each consisting of a scheduled flight, a flight crew, and an airplane, instead of mere pairs as in the previous version. Figure B (on Page 8) illustrates an example of this more complicated problem. In this case, there is no known computationally tractable method of finding the best matching—and Karp's work showed that none is likely to be discovered.

What's the airline to do, when faced with the news that finding the best triple matching is too hard? One option is to settle for less than the best. For example, one way to get a matching—albeit not necessarily the best—is simply to construct compatible triples, one by one, setting aside the matched elements, until no compatible triples remain. One can in fact prove that this simple approach never fails disastrously—that the number of compatible triples constructed is always at least one-third of the number in the *best* matching of triples.

A method whose output is always nearly as good as the best possible output is called an *approximation algorithm*. Well over a hundred approximation algorithms have been discovered and analyzed for problems in such diverse application areas as scheduling, VLSI design, facilities placement, transportation and communication, numerical analysis, and even biology. Most of these algorithms are more complicated than the simple technique outlined above, but many are surprisingly straightforward. The difficulty lies in analyzing the algorithm to determine how well it performs.

In many cases, however, an approximation algorithm performs far better than predicted. For example, Klein and a student, Ajit Agrawal (now graduated and working for DEC) studied a new approximation algorithm they had discovered in collaboration with another student, R. Ravi. The algorithm deals with solving *sparse* linear systems, those with only a few non-zero entries per row. Since these systems can have thousands of variables and thousands of inequalities, it is critical to take advantage of their sparsity by manipulating only the non-zero entries. In solving these systems by Gaussian elimination, however, each time another variable is eliminated, new non-zero entries are introduced. The order in which variables are eliminated thus affects how sparse the system remains, and thus how long the entire process takes.

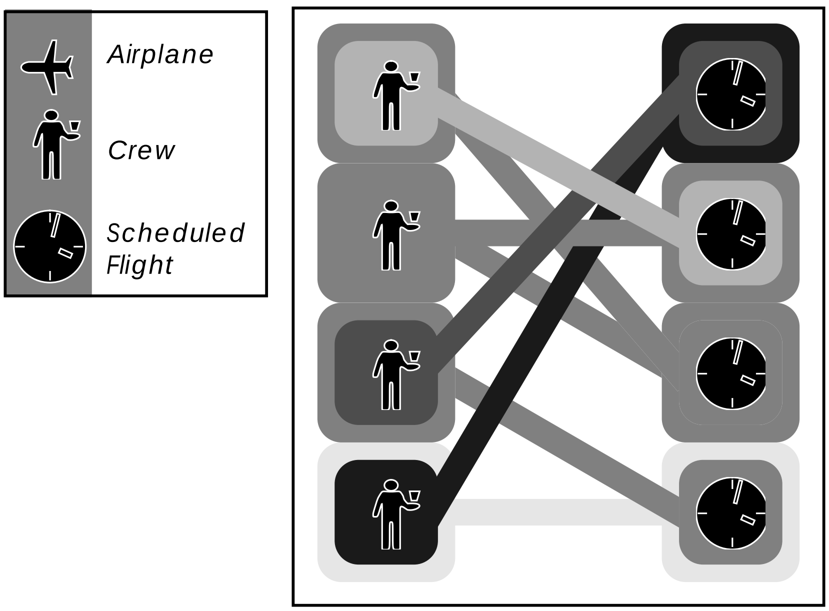


**Figure A**

*This figure represents a matching problem. Each of the compatible pairs (flight crew, scheduled flight) is indicated by a common crosshatch pattern. The goal is to pair up the flight crews with the scheduled flights in such a way as to maximize the number of pairs that are compatible.*

Finding the best ordering of variables is known to be intractable, and until recently no method was known for finding an ordering that was provably near-optimal. The newly discovered algorithm does just that, but the performance guarantee, while theoretically good, was not good enough to warrant use of the algorithm in practice.

Klein believed, however, that the new algorithm would perform much better than the analysis could show. He and Agrawal implemented the algorithm and tested it on a variety of sparse linear systems that arise in practice. They compared the results to those produced by a commonly used heuristic, Liu's minimum-degree code, a code that has been tuned over many years. The comparison revealed that the new algorithm, without any tuning, is comparable in performance to the minimum-degree heuristic and even beats it in time required for parallel execution of the elimination. Thus the algorithm's performance was indeed far better than analysis predicted. Klein attributes the discrepancy in part to the limitations of the analyzers and in part to the limitations of analysis. ‘‘We still have far to go in developing techniques for analyzing the performance of approximation algorithms. But even when we know exactly how badly the algorithm can perform, it often performs much better for most instances encountered. The instances where the algorithm is at its worst are fairly contrived.”

Klein doesn't see this as a reason to abandon analysis. Just the opposite. ‘‘If we keep in mind its limitations, theoretical performance analysis can be quite useful in guiding us toward promising algorithmic ideas and away from less promising ones. An algorithm whose performance can be analyzed and shown to be good should be investigated, even if the performance guarantee seems less than stellar, there's a good chance the algorithm will perform quite well. Basically, worst-case analysis tends to be over-pessimistic.”

Klein is also investigating heuristics that are not approximation algorithms themselves but make use of them as subroutines. One system currently being implemented finds a good layout for a circuit. Another project, carried out by a student in Klein's seminar, involves assigning tasks to processors of a parallel computer in order to minimize communication time.

‘‘Parallel processing is a special interest of mine,” says Klein, whose thesis work concerned new parallel algorithms that made efficient use of processors. ‘‘One of my most recent results combines the paradigms of parallel processing and approximation algorithms. A fundamental problem in optimization is finding the cheapest path through a network.

Unfortunately, we have no satisfactory way of solving this problem in parallel: the fast methods require far more processors than are practical. What I've done is to give a fast parallel algorithm for finding an approximately cheapest path."

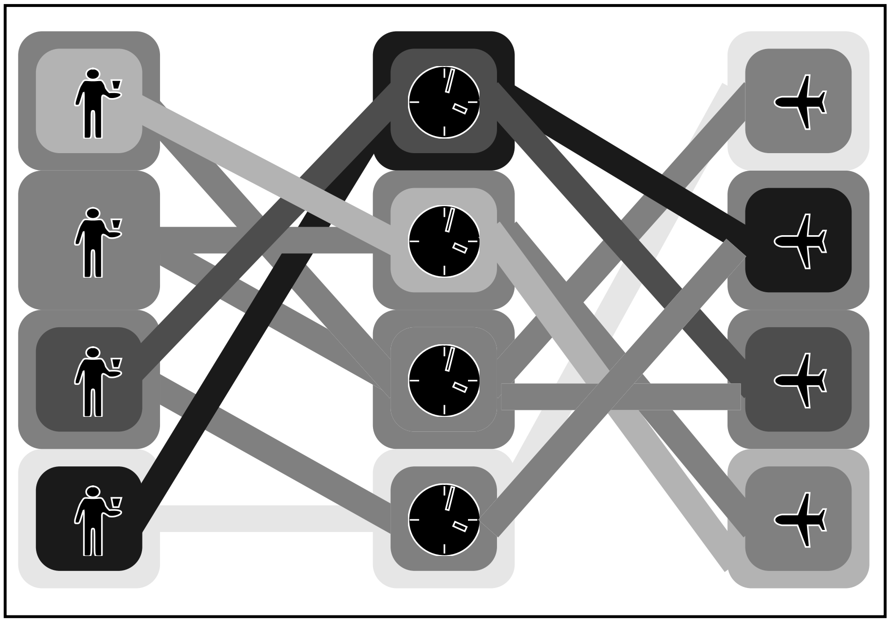

***Figure B***

*This figure represents a triple matching problem. Each of the compatible triples (flight crew, scheduled flight, airplane) is indicated by a common crosshatch pattern. The goal is to group the flight crews with the scheduled flights and the airplanes in triples so as to maximize the number of triples that are compatible.*

Klein's new algorithm can be used in solving certain multicommodity flow problems; these can be used in turn in further approximation algorithms. Thus, by attacking one fundamental problem, Klein has made possible the parallel implementation of algorithms for many others.

Klein continues to work in the two areas of approximation algorithms and parallel processing. "I'm focusing on problems in network design and analysis. It's a promising area because so many of the interesting optimization problems turn out to be intractable and yet really need to be solved—they're problems for which the quality of solution can make a big difference in time or money expended."


*conduit!*

A publication of the Computer Science Department
Brown University

Suzi Howe . . . . . . . . . . . . . . . Editor
Katrina Avery . . Editorial Assistant
Jeff Coady. . . . Technical Support
Dan Robbins . . . . Design Assistant



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



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

***conduit!***