Reasoning about knowledge

Epistemology, the study of knowledge, has a long and honorable tradition in philosophy, starting with the Greek philosophers. But what does it have to do with computer science? "Quite a lot,'' says Joe Halpern. "People intuitively use the word 'know' all the time when thinking about distributed computing. For example, we are apt to say 'this process can't terminate yet because it doesn't know that the other process has received the message.'"

In the 1980s, Halpern and his collaborators showed how reasoning about knowledge and common knowledge could help in understanding and verifying distributed algorithms. Consider the famous *coordinated attack problem*, introduced by Jim Gray, of Microsoft Research, as an abstraction of a commitment problem in distributed databases. Two generals must attack a village, but they want to attack together; if they attack alone, their armies will be wiped out.

There is a problem: they are miles apart, and the only way they can communicate is by means of a messenger, who might get lost or captured. General A sends a message to General B: "Attack at dawn!" General B gets the message, but will he attack? No, because he is worried that General A doesn't know that B got the message. So B sends an acknowledgement. General A gets the acknowledgement. Will he attack now? Again, no, because A is worried that B doesn't know that he (A) knows that B got the acknowledgement. So A acknowledges the acknowledgement. The messenger makes it again (although he's getting tired!). But this still doesn't work —General B is worried that A doesn't know that B knows that A sent the message.

There is never *common knowledge* that the message was sent, where common knowledge is the state in which everyone knows that everyone knows that everyone knows that ... . Indeed, as Halpern and his former student Yoram Moses showed, it is impossible to get common knowledge in a system where communication is not guaranteed. Moreover, as they also showed, common knowledge is a necessary and sufficient condition for coordination. Halpern and Moses were awarded the 1997 Gödel Prize for the paper describing this work, which included a general approach for using epistemic logic to analyze distributed systems.

The work on knowledge has found application in other areas. In AI, robots trying to coordinate have to reason about what one robot knows that the other knows. More recently, Halpern and his students Riccardo Pucella and Kevin O'Neill showed that knowledge is critical in reasoning about security. For example, what does it even mean to say that a message is sent anonymously? Roughly speaking, it should mean that no one knows who the sender is. This has typically been taken to mean that all possible senders are equally likely to have sent it. But it's a bit more subtle than that. Somehow, probability has to be involved. A contribution of $5,000,000 to Cornell may be anonymous, but not everyone is equally likely to have made it. Prior beliefs must be taken into account.

Knowledge and belief play a critical role in economics, particularly game theory. Here, the celebrated notion of *Nash equilibrium* is not always applicable because it requires that all agents be aware of all moves. Halpern, his student Leandro Rego, and Larry Blume and David Easley in Economics are exploring models for decision making and game theory that take this lack of awareness into account.

In 1986, Halpern organized a conference on Theoretical Aspects of Reasoning about Knowledge (later renamed Theoretical Aspects of Rationality and Knowledge) in order to bring together computer scientists, economists, and philosophers interested in issues of knowledge. Now, almost 20 years later, the conference is still going strong. The area continues to be active, and everyone *knows* that Cornell is at the forefront of these activities.

CS Professor Joe Halpern and his former student Yoram Moses received the 1997 Gödel Prize for their paper that showed that common knowledge is a necessary and sufficient condition for coordination.

Forty years of numerical analysis and scientific computing

Numerical analysis (NA) intersects applied math and computer science. NA has flourished at Cornell since the 1960s because it has always been a strong component of the faculties in Computer Science and Mathematics and because of easy participation in the Center for Applied Mathematics. For us, it is a best-of-both-worlds history, with advanced algorithmic ideas on the one hand and rigorous mathematical analysis and applications on the other.

Roland Sweet and Jim Bunch launched very productive research careers at Cornell. The Fast Poisson solvers of the early 1970s dramatically widened the class of solvable elliptical PDEs, and Sweet's work on cyclic reduction was a key factor. Bunch established our tradition in numerical linear algebra with fundamental work on symmetric indefinite systems. He also was one of four co-authors of LINPACK, a landmark in the history of numerical software and a precursor to LAPACK.

Faculty members Jorge More and John Dennis established a new way of thinking in the field of nonlinear equation solving and optimization. Their work on quasi-Newton methods showed just how well one could live with approximate Jacobians and Hessians. The sparse optimization research of Tom Coleman since the 1980s represents a continuation of this thread of NA. Incorporating ideas from graph theory and other areas, Coleman produced a steady stream of advanced numerical algorithms that could utilize fully the new generation of parallel machines. Coincident with his tenure as Director of the Cornell Theory Center, Coleman moved into computational finance and established the Financial Industry Solutions Center (FISC) in Manhattan. As an example of outreach to the applications community, FISC is one of the most remarkable realizations of Cornell's Land Grant Mission.

CS professor Steve Vavasis, with an interest in both differential equations and optimization, has brought advanced computer science ideas to several classical application areas. His book, *Nonlinear Optimization: Complexity Issues* (1991), reflects the department penchant for subareas-without-borders, in this case, NA and theory. The use of computational geometry and sophisticated data structures makes Vavasis's automatic mesh generator (QMG) one of the most important software tools for the solution of boundary value problems over irregular domains. The rigorous analysis of the matrix problems that arise in applications has been a hallmark of his work.

Similarly, former colleague Nick Trefethen's segue into the world of pseudo-spectra evolved from profound observations about the non-normal matrices that arose in various partial differential equation settings. This work is very much a tribute to the singular value decomposition (SVD), a matrix factorization that has come to dominate much of the numerical linear algebra scene. Charlie Van Loan's generalized SVD was used to solve several key real-time signal processing applications in the days of Reagan's "Star Wars" program, as were techniques developed by former colleague Frank Luk. Van Loan's SVD method with Gene Golub for total least squares created a new framework for doing least squares fitting when there are errors in the data, and his cache-friendly block Householder representation, developed with former PhD student Chris Bischof, is today the standard way of organizing orthogonal matrix computations like the SVD. According to Google Scholar, Van Loan's book with Golub, *Matrix Computations* (1983), is the most widely-cited text in the computing and mathematical sciences.

In the future, NA will play a critical role in the broader field of computational science and engineering. Typical of the new era of interdisciplinary research is an NSF-sponsored ITR project concerned with the simulation of turbulent combustion. Researchers in computer science, mechanical engineering, and math are combining their talents to address the problems associated with a high-dimensional numerical integration that is at the heart of the computation. The idea is to reduce the effective dimension through a table look-up scheme. To say something about the accuracy of the overall method requires physical modeling of the underlying phenomena (MechE), computational geometry (CS), and convex analysis (Math). The three-way collaboration makes it possible to track each table entry's domain of validity in "reaction space".

The intersection between NA and other departmental research areas is also greater now than ever before. For example, CS professor Uri Keich works on statistical and algorithmic problems that arise in biological sequence analysis. Researching the significance of matches of DNA sequences, Keich identified a problem in the extremely popular BLASTn program, a problem he is now working to fix. And let's not forget that PageRank is an eigenvector computation. The futures of NA and CS are coupled —and that is a good thing!

Uri Keich works on statistical and algorithmic problems that arise in biological sequence analysis.

Researcher Donna Bergmark receives the Best Paper Award for Collection Synthesis in the ACM Joint Conference on Digital Libraries.

Fred Schneider chairs the NSF ITR Program Review.

PhD '92 Daniela Rus and BA '93 Sendhil Mullainathan win MacArthur Genius Award.

Researcher Carl Lagoze, with three others, defines the Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH). The work has lead to renewed interest in shared metadata and increased ability to locate relevant digital assets regardless of geographical location.

Student Tim Roughgarden wins the "Danny Lewin Best Student Paper Award" at STOC 2002.

Ramin Zabih and student Vladimir Kolmogorov receive the Best Paper Award in the European Conference in Computer Vision. Their paper dealt with minimizing energy functions via graph cuts.

Uri Keich joins. David Gries becomes Assoc. Dean of Engineering.

2003

CS offers an undergrad Information Science major in Arts & Sciences.

Bob Constable is elected to the CRA Board.

Fred Schneider receives an honorary doctorate from the University of Newcastle upon Tyne, U.K.

Jon Kleinberg, Eva Tardos, and student David Kempe receive the Best Research Paper Award in the ACM SIGKDD Intl. Conference on Knowledge Discovery and Data Mining. Their paper, *Maximizing the spread of influence through a social network*, is one of a series of papers on social networks produced at Cornell.

The Cornell Game Design Initiative is formed, under the direction of Dave Schwartz.

Steve Marschner shares a Technical Achievement Award from the Academy of Motion Picture Arts and Science with Henrik Jensen and Pat Hanrahan for their model of subsurface scattering of light in translucent materials. The model has been used often, including for Gollum in *The Lord of the Rings* trilogy.

Lillian Lee's work with postdoc Regina Barzilay on a system that learns to paraphrase is featured in *The New York Times.*

The CS Programming Team wins honorable mention in the world finals in the ACM meeting at the Czech Republic.