The science of networks

The Internet, pictured below, has inspired research in many areas of CS. Eva Tardos (left) and Jon Kleinberg have made important contributions.

Complex networks are ubiquitous, and the study of networks is now central to many scientific disciplines. The explosive growth of the Internet and the World Wide Web has led computer scientists to seek ways to manage their complexity and help users navigate their vast information content. Social scientists are confronted by social network data on a scale previously unimagined: datasets on communication within organizations, on collaboration in professional communities, and on relationships in financial markets. Biologists have discovered that network structure that defines pathways of cell metabolism can provide insight into fundamental biological processes. The drive to understand such phenomena has resulted in a new "science of networks" as they arise in the physical world, the virtual world, and society.

Cornell has occupied a central position in this emerging discipline since the beginning.

Much of the action in the science of networks is occurring at the boundary between computer science and the social sciences. Sociology has long been concerned with the processes by which networks form and with the social processes that they support. Game theory and other branches of economics have studied how people behave in network settings. The interaction of these fields with computer science is leading to deep and surprising insights. To illustrate, we examine results in the study of traffic and in network searching.

Commuters create traffic on the network of roads and highways, much as Web surfers create traffic on the Internet. If a central authority were allowed to specify everyone's route, it could make the average time in the network, accounting for slowdowns due to congestion, as small as possible. But neither commuting traffic nor the Internet has such a central authority, so how can the average time be reduced? Game theory suggests we consider "selfish routing", in which each person follows the path that is best for them, given what everyone else is doing. Selfish routing can be thought of as the only "stable" way for traffic to flow; it is the only kind of configuration in which no one has an incentive to change routes.

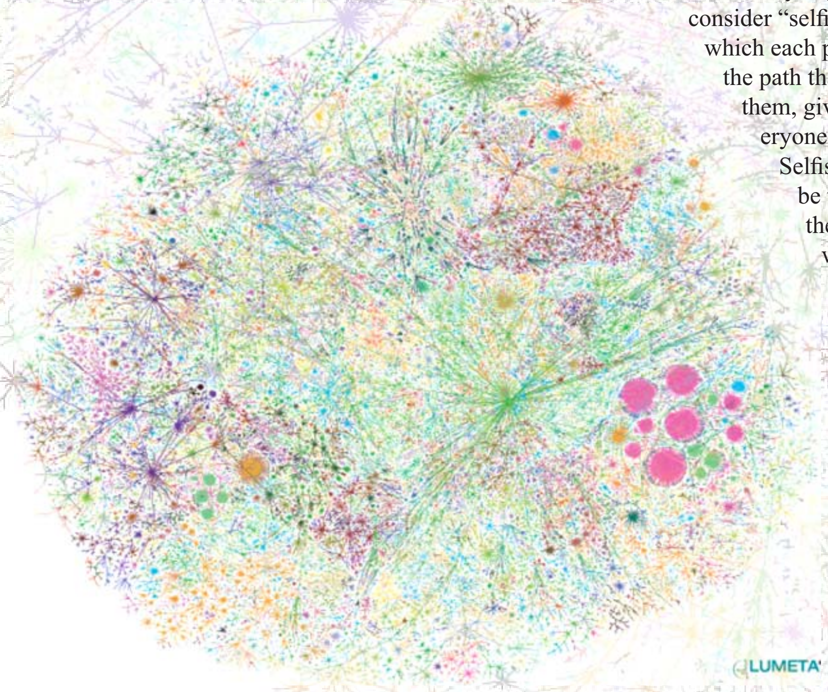

On the Internet, routers mimic "selfishness" by computing shortest paths, but congestion still occurs. To alleviate congestion, is it better to devise more complex routing protocols or just to throw more bandwidth at the problem? CS professor Eva Tardos and student Tim Roughgarden (now a faculty member at Stanford) proved the striking result that travel time under selfish routing is never more than the optimal travel time with twice the traffic. The Roughgarden-Tardos result says something fundamental: simply doubling the bandwidth is at least as beneficial as any amount of central control.

Different kinds of questions arise in social networks —the networks of relationships among people. For example, a popular belief is that people are connected by "six degrees of separation" —an average of six hops in the network that connects people who know each other on a first-name basis. The origins of this notion are rooted in experiments by social psychologist Stanley Milgram in the 1960s. He asked participants to forward a letter to a distant "target person", with each participant forwarding it to only a single acquaintance. The median length of the paths followed by the letters was six —the origin of the number six in this pop-cultural mantra. This striking experiment inspired researchers to develop network models designed to resemble real social networks in their profusion of very short paths.

CS professor Jon Kleinberg asked a natural question: how could the subjects in Milgram's experiments, lacking any global knowledge of the social network, be so effective at routing letters, essentially running a collective algorithm for finding shortest paths? In studying this issue, Kleinberg considered a network model developed by Cornell applied math professor Steve Strogatz and his student Duncan Watts. Kleinberg found that, in a variation on the Watts-Strogatz model, it was possible to perform highly efficient message-routing using only local information. The basic structural features of

Kleinberg's networks have since been found in the social structure of large organizations and online communities, and his message-routing approach has appeared in proposals for decentralized search methods in several peer-to-peer file-sharing systems.

Ongoing research at Cornell and elsewhere is seeking answers to other questions such as how new behaviors spread through large networks, how networks evolve over time, and how their structure shapes and is shaped by interactions among strategic agents. By grounding this research in models that are fundamentally computational, computer science is bringing a new perspective to long-standing questions in a broad array of fields.

Algorithms at Cornell

At Cornell, research in algorithms began in the early 1970s with a focus on the concept of asymptotic complexity and its role in graph algorithms, efficient data structures, and the abstraction of techniques such as divide-and-conquer and dynamic programming. A prime example of such work is the linear algorithm for testing planarity of a graph, developed by CS professor John Hopcroft and his Stanford PhD student Bob Tarjan (also a former CS professor), for which they received the 1986 ACM Turing Award.

Hopcroft also had a major influence on education and research in the field. His classic text with Al Aho and Jeff Ullman, *The Design and Analysis of Computer Algorithms* (1974), essentially defined the field, and nearly every CS department in the country developed a course around this text.

Cornell faculty members have made important contributions to a broad range of algorithmic domains. Dexter Kozen is widely known for his work on algorithms in computer algebra and symbolic computation and for his text *The Design and Analysis of Algorithms* (1991). Eva Tardos won the Fulkerson Prize in 1988 for her work on network flow algorithms, resolving long-standing open questions of Edmonds and Karp on the efficient solvability of minimum-cost flow problems. David Shmoys's work on approximation algorithms for scheduling and other problem domains was one of the key influences that gave approximability the prominent role it currently enjoys in the area of algorithms. Cornell professor David Williamson won the Fulkerson Prize in 2000 for his joint work with Michel Goemans on the use of semidefinite programming in the design of approximation algorithms.

Today, the field of algorithms has matured and broadened to connect with many other areas, both within computer science and beyond. Jon Kleinberg has led this transformation with his early work on search and information extraction in large networks like the Web; he received the 2001 Award for Initiatives in Research from the National Academy of Sciences for his introduction of link analysis techniques into Web search. Over the years, Kleinberg, Shmoys, Tardos, and Williamson, and more recently Hopcroft, have been leaders in this continuing evolution of the field, through their work on algorithmic issues in social and information networks, on approximation algorithms in discrete optimization, and on algorithmic and game-theoretic foundations for network routing. The new textbook *Algorithm Design* by Kleinberg and Tardos (2005) emphasizes the perspective of this emerging focus through its efforts to situate algorithms at the center of the field of computer science more generally.

Over the years, John Hopcroft (left) and Dexter Kozen have made important contributions in algorithms and complexity.

Algorithms are prominent in the research of Cornell faculty in many other areas of computer science. Ramin Zabih and Dan Huttenlocher have had a major impact in computer vision through their work using graph algorithms, especially using network flow techniques. These methods are also becoming increasingly influential in graphics. The work of Bart Selman and Carla Gomes has highlighted the key role of algorithms in artificial intelligence through an algorithmic approach to central problems in constraint satisfaction. Johannes Gehrke's work in data mining has established fundamental connections to algorithms. Work in computer security at Cornell includes algorithmic issues, led by Kozen's investigation of efficient code certification.

New faculty hire Bobby Kleinberg brings further strength in this area, with his already widely-recognized work on stochastic algorithms and learning-based models for networks and online economic systems.

Through its ongoing role in the evolution of algorithms as a discipline, Cornell moves forward with arguably the strongest algorithms group in the world.

CS undergrads do well on the Putnam Math Competition. The team of Kleinberg, Munoz, and Krosky places fifth out of 284, and Zhang places in the top ten individuals.

Juris Hartmanis is Chair of the NRC Committee that produces *Computing the Future* (National Academy Press). This influential report assesses academic computer science and engineering. It advocates a broader research and educational agenda that builds on the field's impressive accomplishments.

Charlie Van Loan publishes *Computational Frameworks for the Fast Fourier Transform* (SIAM).

Juris Hartmanis is elected to the American Academy of Arts and Sciences.

Monika Rauch Henzinger, Thorsten von Eicken join. Bob Constable becomes Chair.

1993

Juris Hartmanis shares the ACM Turing Award with Dick Stearns, "in recognition of their seminal paper, which established the foundations for the field of computational complexity theory" (see the entry for 1965).

Juris Hartmanis receives a Humboldt Foundation Award for Senior U.S. Scientists. This foundation, created by the German government in 1953, enables scholars to do research in Germany.

Researcher Yuying Li receives the 1993 Leslie Fox Prize in Numerical Analysis from the Institute of Mathematics and its Applications.

David Gries and Fred Schneider publish *A Logical Approach to Discrete Math* (Springer Verlag).

Stratus Computer acquires Ken Birman's Isis Distributed Systems, Inc. Isis technology is deployed in the NY and Swiss Stock Exchanges, the French Air Traffic Control System, and other places.

Brian Smith, Claire Cardie, Ramin Zabih join.

1994

Ken Birman becomes Editor-in-Chief of the *ACM Transactions on Computing Systems*.

David Gries receives the IEEE Computer Society Taylor Booth Education Award for his "commitment to education in CS and Engineering as demonstrated by a record of outstanding teaching and mentoring, writing of textbooks, curriculum development … ."