From *SIAM News*, Volume 37, Number 5, June 2004

The Price of Anarchy

By Sara Robinson

Interstate 5 stretches through California's San Joaquin valley, straight and flat as far as the eye can see, giving drivers a unique opportunity to observe game theory in action.

With only two lanes in each direction, most cars stay in the left lane, zipping past the slow moving trucks, which keep to the right. Every so often, an impatient truck driver lumbers into the left lane, blockading the road for several long minutes as he edges ahead of his even more sluggish comrades. Meanwhile, a long line of cars quickly builds up.

Eventually, car drivers start to swerve into the right lane, zooming past the line of cars, only to duck back in at the front of the line. Each time a car cuts back in line, the cars behind brake, slowing traffic further. Soon, the right lane is clogged with impatient car drivers trying to advance, which, in turn, induces more truck drivers to try to pass on the left.

Although some of the scofflaw traffic does get to Los Angeles in less than the average travel time, the average time increases. In other words, the selfish behavior of some individuals increases the congestion of the highway.

This "highway game" is reminiscent of the well-known Prisoner's Dilemma from economic game theory. There, each of two arrested suspects is given the choice of confessing or remaining silent. Each suspect knows the rules: If both confess, both get light (five-year) sentences for the crime. If neither confesses, both get six-month sentences, on a minor charge. If one confesses while the other remains silent, the one who confesses walks free while the silent one gets the full (ten-year) sentence. These rules yield a strange paradox: Each player improves his own lot by confessing, regardless of what the other player does, yet the optimal outcome for the pair occurs when both suspects remain silent.

An instance of the game in which both suspects make the selfish choice is an example of a Nash equilibrium, a set of strategies such that even if each player knows what the other player's strategy is, he has no incentive to change his own. A Nash equilibrium is a simple model of what happens in a competitive market.

From a global perspective, as these examples show, competition often has a cost. Still, for most markets, achieving the global optimum is not a reasonable expectation, either because there is not enough information to compute it or because centralized control is undesirable for other reasons. In such situations, the Nash equilibrium may be a reasonable approximation to the optimum, and the best that can be hoped for. But how good an approximation is it? Computer scientists have found an answer.

Game Theory and Computer Science

Theoretical computer scientists attempting to model the Internet have necessarily turned their attention from engineered systems to the decentralized systems described by economic game theory. In the process, they are modifying traditional game theoretic analysis by introducing techniques of their own.

A new thread of results, stemming from techniques used to analyze approximation algorithms, is rapidly moving from computer science into other fields. The central idea is to quantify the loss of efficiency from competition for a game by bounding the ratio between the worst possible Nash solution and the social optimum. Because a ratio remains constant as the problem scales and is independent of the type of quantity being measured, economic efficiency loss is a fixed parameter for a given model.

Economists are well aware that selfish behavior can lead to inefficiency. The economics literature abounds with results giving the difference between Nash equilibria and the social optimum for a given instance of a game, or proving that a Nash equilibrium achieves a global optimum. Still, economists have not systematically bounded the ratio of the worst-case Nash equilibrium to the global optimum, says Éva Tardos, a professor of computer science at Cornell University. "Once it's not optimal, [economists] somehow feel you're done."

A March workshop at the Institute for Mathematics and its Applications, "Control and Pricing in Power and Communications Networks," brought together an eclectic group of electrical engineers, mathematicians, operations researchers, and theoretical computer scientists.

Hal R. Varian, a professor in the Haas School of Business, the economics department, and the School of Information Management and Systems at Berkeley, agrees that efficiency-loss ratios are new to economists. "This sort of "worst-case ratio" analysis has not been used much in economics," he says. "It's a standard technique in computer science, and people have used it in operations research as well, but I haven't seen it used in economics in the same way."

Economists have traditionally focused on situations in which a competitive market is usually the best that can be attained, says Christos Papadimitriou, a professor of computer science at the University of California, Berkeley. Computer scientists, by contrast, often work with engineered, centrally controlled systems. "Economists never embarked on a systematic investigation of efficiency loss only because the social optimum is not an option," he says, whereas computer scientists "come from a domain where the social optimum is the status quo."

Measuring the cost of competition as a ratio is in the spirit of long-standing efforts to quantify the effectiveness of approximation algorithms and of "on-line" algorithms, which must produce outputs based on partial information. Computer scientists use corresponding ratios to measure the effects of a lack of unbounded computational resources for approximation algorithms, and of a lack of information for on-line algorithms. For game theory, Papadimitriou says, the efficiency-loss ratio quantifies the effect of a system's "lack of coordination."

From its origins in computer science, the idea of quantifying efficiency loss is gradually spreading into operations research, information theory, and more recently, electric power research. Following Papadimitriou's example, researchers in these fields call the efficiency-loss ratio "the price of anarchy."

An Intriguing Question

On a Friday afternoon in the fall of 1999, Tardos, who was visiting Berkeley for a semester, joined a lively gathering of theoretical computer scientists in a ground-floor lounge in Soda Hall. The occasion was the Berkeley "theory lunch," a weekly gathering of computer science graduate students and faculty for lunch, followed by an informal, half-hour talk. Initiated by Papadimitriou, the lunch gives students and faculty a forum for discussing interesting research ideas or work in progress in a setting less formal than a seminar.

On this particular Friday, the speaker was Leonard Schulman, visiting from the Georgia Institute of Technology. After lunching on salads and sandwiches from a local Greek café, the group gathered for Schulman's discussion of an intriguing traffic paradox he'd read about in a recent issue of *American Scientist*.

Introduced in the 1950s by Dietrich Braess, a civil engineer, the paradox demonstrates that a new road built to alleviate traffic congestion can actually make the congestion worse, even when the quantity of traffic remains the same.

Braess's example is a simple directed graph with a source, a sink, and linear "delay functions" on the edges (see Figure 1). The network has a single Nash equilibrium, a state in which no driver can decrease his travel time by changing his route. Assuming one unit of traffic flow, the traffic delay is 1.5 at equilibrium. Braess then added a road through the middle, providing an attractive shortcut for drivers. As a result, the system eases toward a new Nash equilibrium with a delay of 2, though the optimal rate is still 1.5. Selfish utilization of the new road thus increases travel time by 33%.

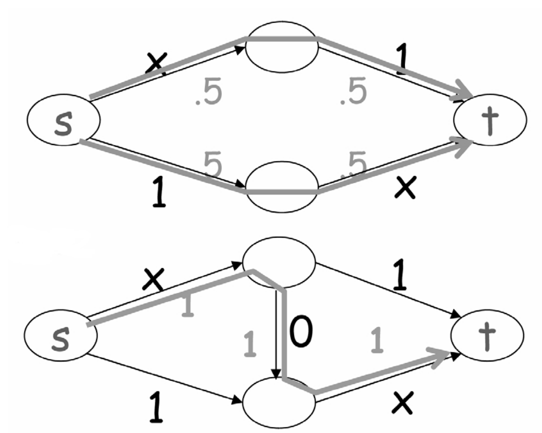

Figure 1. *Original network (top) in which the cost of Nash flow is 1.5. With an added edge (bottom), the cost of Nash flow becomes 2. From Éva Tardos's talk at the IMA workshop.*

This model is equally useful for thinking about packets flowing over the Internet, whose routes are controlled by the sending agents, said Schulman, who is now at the California Institute of Technology. He then posed a research problem with a computer science flavor: Given such a network, is there an efficient algorithm that improves the equilibrium travel times by removing edges from the graph?

Papadimitriou, sitting in the audience, had a different question, related to his work of the previous year with Elias Koutsoupias, a professor of computer science at the University of California, Los Angeles. The two had introduced the notion of using a ratio of the worst-case Nash equilibrium to the social optimum to measure the cost of competition, and then computed bounds on this ratio for sample problems about Internet routing.

With this paper in mind, Papadimitriou pointed out that in Braess's example, the ratio of the Nash flow rate to the optimal rate is 4/3, and the "price of anarchy" is thus an efficiency loss of 33%. "More generally," Papadimitriou wondered, "can you quantify the efficiency loss due to competition for all such networks?"

Intrigued, Tardos began to work on the problem, focusing first on graphs with linear delay functions. A few months later, she and Tim Roughgarden, then a student of hers, solved it. (Independently, Roughgarden, who will become an assistant professor at Stanford University in the fall, answered Schulman's question by showing that his problem is NP-complete.)

In their paper, first published in the proceedings of the 2000 IEEE Symposium on the Foundations of Computer Science, Tardos and Roughgarden define a mathematical model of a network consisting of a directed graph with source–sink pairs, flow rates $r_i \geq 0$ for traffic through the network, and latency functions $l_e(x)$ for each edge e. They go on to prove two theorems: The first, in essence, states that Braess's paradox is as bad as it gets. The researchers show that the general cost of a Nash flow is at most 4/3 of the minimum-latency flow for graphs with linear latency functions, a "price of anarchy" of 33%. (With general edge latency functions, they point out, the efficiency loss is unbounded, as in the example of a two-edge network with edge latency functions of x^N and 1 and a flow rate of 1.)

The second theorem states that for any network with continuous, nondecreasing latency functions, the cost of Nash flow with a given set of flow rates is no more than the cost of an optimally routed network with all the flow rates doubled. This suggests that, for some routing networks, the efficiency loss can be effectively dealt with simply by increasing edge capacity throughout the network.

Cornell University computer scientist Éva Tardos, named editor-in-chief of SIAM Journal on Computing *late last year, is a significant player in the "Braess paradox story" told here.*

The Word Gets Out

The Koutsoupias–Papadimitriou and Roughgarden–Tardos papers spawned a series of results, presented at theoretical computer science conferences, on the price of anarchy for other network games, including designing and creating networks and optimally locating resources. At the same time, as the larger trend of applying computer science techniques to game theory took shape, leaders in this effort began to convey the theoretical computer science results to broader audiences. The price of anarchy featured prominently in their talks.

As the word spread, Tardos says, she and Roughgarden were deluged with requests to speak about their paper. In the audience at a talk Tardos gave at the Massachusetts Institute of Technology in 2001 was Ramesh Johari, then a graduate student in the electrical engineering department. Johari, who characterizes his field as "the design and control of large-scale decentralized systems," remembers being intrigued by Tardos's work, although it would be more than a year before it influenced his own research.

Before arriving at MIT, Johari had worked with Frank Kelly, a professor of the mathematics of systems at Cambridge University. In 1997, Kelly had looked at the problem of optimally allocating a shared resource, such as bandwidth. He gave a simple market mechanism for pricing the resource and showed that it results in a socially optimal allocation, as long as individual agents act as "price takers," i.e., do not anticipate the effect of their actions on the price. (Johari wrote about this work, described in a talk Kelly gave at ICIAM 99, for *SIAM News* (www.siam.org/siamnews/03-00/congest.pdf.)

With his adviser John Tsitsiklis, a professor of electrical engineering and computer science and co-director of the Operations Research Center at MIT, Johari had intended to look at "peering relationships"—the rules governing the interaction of Internet service providers—comparing the efficiency of decentralized solutions with that of optimal peering. Operations researchers commonly compare decentralization with optimality, Johari says, although, like economists, they had not looked specifically at efficiency ratios.

A year later, stalled on his own thesis work, Johari was asked to speak at an MIT game theory seminar. Thinking back on Tardos's notion of efficiency loss as he prepared the talk, he wondered whether it could be applied to Kelly's work on resource allocation.

Efficiency Loss in Resource Allocation

In his 1997 paper, Kelly considered the situation of several users sharing a fixed supply of a single resource, such as bandwidth. Each user gets a nonzero share of the resource and has a concave, non-negative, strictly increasing utility function specifying how much he is willing to pay for a given amount of the resource. In this framework, the socially optimal allocation of the resource is obtained by maximizing the sum of the utility functions, subject to the constraint that the sum of the amounts given to each user is less than or equal to the supply.

To price the resource, Kelly supposes that each user submits a bid and receives an amount of the resource proportional to the ratio of his bid to the sum of all the bids. This is equivalent to allocating the resource to each user at a unit price, where the price is equal to the sum of the bids divided by the supply of the resource.

Given a set of bids, one for each user, suppose that an individual user is reconsidering his bid. The goal is to maximize his payoff—the difference between his utility and the bid. Assuming, as Kelly does, that the individual does not know the effect of his bid on the price, he uses the current price to determine the utility obtained from an alternate bid.

Under this model, Kelly proved that all competitive equilibria yield a socially optimal allocation of the resource. For any set of bids and the corresponding price, in other words, he showed that if each user's bid is optimal for him at that price, then the allocation is also socially optimal.

Kelly's result, however, depends on the assumption that the market participants do not anticipate their effect on the price. What happens if the players know the price-setting procedure and anticipate the effect of their bids? In this case, the price is no longer a fixed parameter in the payoff function, and the resource allocation problem becomes a game in which the setting of bids is strategic.

In 2002, Bruce Hajek, an electrical engineer at the University of Illinois, Urbana–Champaign, and his student Ganesh Gopalakrishnan had shown the existence of a unique Nash equilibrium for this game such that the resulting allocations are the unique socially optimal solution for modified utility functions. They presented their result at a meeting on stochastic networks at Stanford University to an audience that included Johari and Tsitsiklis.

Thinking back to Hajek's game, Johari, who will become an assistant professor at Stanford in the fall, wondered if it would be possible to compute the efficiency loss in this setting. He and Tsitsiklis began to work on the problem and were soon able to show that the price of anarchy is a 25% efficiency loss, which they measured by computing the ratio of the Nash equilibrium utility function to the socially optimal utility function and showing that it is at worst 3/4. They also extended their result to more general network allocation games in which users request service from multiple links, submitting bids to each one. Links then allocate traffic flow rates proportionally, and each user sends the maximum flow possible for the allocated rates. For the multilink game, the price of anarchy is also a 25% loss.

These price-of-anarchy results, Johari and Tsitsiklis note in their paper, offer hope that robust, engineered systems can be designed around competition, because selfish behavior does not arbitrarily degrade the mechanism's performance. "It's striking that over such a huge class of problems, the efficiency loss is so small," Hajek says.

After reading an early draft of the Johari–Tsitsiklis paper, Hajek and his students, who work in information theory and communications, were also intrigued by the notion of the price of anarchy. They have since produced a slew of related results.

Sujay Sanghavi and Hajek showed how to find the mechanism for a two-user resource allocation that has the lowest possible price of anarchy, given that the bids are equal to the payments. This mechanism, which utilizes nonuniform pricing, yields an efficiency ratio of 7/8, and thus an efficiency loss of only 12.5%. Hajek and Sanghavi also look at the general n-user case and demonstrate numerically that the price of anarchy is likely to be no worse than 12.65%.

Sichao Yang and Hajek showed that by dropping the condition that the payments equal the bids, and requiring merely that they be a function of the bids, it is possible to find nonuniform pricing mechanisms such that the allocations are 100% efficient. Thus, in this case, anarchy has no price. (Johari and Tsitsiklis showed recently that if bids do not equal payments and price is uniform, the efficiency loss is still 25%.)

Hajek and Yang also introduced a new network game model, similar to that of Johari and Tsitsiklis for multiple links but taking into account the rebalancing by other players when one player changes a bid. For this game, the price of anarchy approaches 100%.

Power Systems and Deregulation

In March, the Institute for Mathematics and its Applications at the University of Minnesota convened a workshop called "Control and Pricing in Power and Communications Networks." The idea was to encourage a cross-fertilization of ideas by bringing together researchers in electric power and communications networks.

About half of the 40 odd attendees were electrical engineers who work on resource pricing and control issues for the new, partially deregulated power networks. The other half—the "communications" side—consisted mainly of electrical engineers and mathematicians specializing in the mathematics of networks and operations research, like Tsitsiklis and Johari, or information theory/probability, like Hajek. Also present were Tardos and Roughgarden, the lone representatives of theoretical computer science. (Papadimitriou had been invited but had to cancel.)

The issue of efficiency loss was a major theme at the week-long workshop. In his talk, Tsitsiklis described recent work with Johari demonstrating that it is possible to limit the price of anarchy in a deregulated power market by constraining the bidding mechanism for electricity generators.

In their model, demand is inelastic, but generators compete to meet it: Each generator submits a supply function—which describes how much electricity the generator will supply for a given price—and a price is set to ensure that supply equals demand. The researchers show that, by constraining the class of supply functions, it is possible to guarantee that even if some of the N electricity generators have market power, the aggregate cost of electricity production is no more than a factor of $1 + 1/(N-2)$ higher than the socially optimal cost.

Chris DeMarco, a professor of electrical engineering at the University of Wisconsin, Madison, who specializes in power systems and who was one of the workshop organizers, expressed amazement that no one working on power systems had attempted to systematically quantify the effects of power market restructuring. This is the case, he said, even though the social optimum for power systems (which amounts to minimizing generation and transmission costs) is prac-tical to compute, and was routinely computed in the historic, regulated system.

"In computer and communication networks, where market approaches have only been implemented in small-scale experiments, we have a rich literature making comparisons of the performance of various market designs to a theoretical global optimum," he said. However, "in power networks, where we've jumped in with both feet to trying market approaches, we have no such literature."

For power systems, researchers have compared the social optimum with real market data, and with data from experiments in which paid participants emulate power generators, DeMarco conceded. Still, he said, there have been no systematic efforts to prove theorems about worst-case performance.

Other power researchers had reservations as to whether such analyses are meaningful to real-world settings. Sean Meyn, a professor of electrical engineering at the University of Illinois, Urbana–Champaign, explained that while it's easy to optimize a power system at equilibrium, complex dynamics over short time scales can lead to behavior that's quite difficult to model mathematically. "I agree that one needs simple models for analysis," he said, "but one can't claim to understand the physical reality completely based on such simple models. This is why I think one must be careful in interpreting the 'price of anarchy' conclusions."

Global models of power systems are almost always static in the economic sense as well, said Ross Baldick, a professor of electrical engineering at the University of Texas, Austin. "If the prices are high, then [new generators] may enter the market and drive the prices down," he said. "In fact, this dynamic may be the principal benefit of restructuring: encouraging innovation by competition."

Nevertheless, some of the communications researchers are hopeful that determining efficiency loss for simple models can contribute to an understanding of real-world power markets. "It's useful to understand the cost of strategic behavior," Hajek said.

Papers describing much of the work discussed in this article can be found at the IMA Web site. See http://www.ima.umn.edu/talks/workshops/3-8-13.2004/.

Sara Robinson is a freelance writer based in Pasadena, California.