Algorithmic incentives: New twist on 30 year-old work could lead to better ways of structuring contracts

April 25 2012, by Larry Hardesty

Interactive proofs are a type of mathematical game, pioneered at MIT, in which one player — often called Arthur — tries to extract reliable information from an unreliable interlocutor — Merlin. In a new variation known as a rational proof, Merlin is still untrustworthy, but he's a rational actor, in the economic sense. *Image: Howard Pyle*

In 1993, MIT cryptography researchers Shafi Goldwasser and Silvio Micali shared in the first Gödel Prize for theoretical computer science for their work on interactive proofs — a type of mathematical game in which a player attempts to extract reliable information from an unreliable interlocutor.

In their groundbreaking 1985 paper on the topic, Goldwasser, Micali and the University of Toronto's Charles Rackoff '72, SM '72, PhD '74 proposed a particular kind of interactive proof, called a zero-knowledge

proof, in which a player can establish that he or she knows some secret information without actually revealing it. Today, zero-knowledge proofs are used to secure transactions between financial institutions, and several startups have been founded to commercialize them.

At the Association for Computing Machinery's Symposium on Theory of Computing in May, Micali, the Ford Professor of Engineering at MIT, and graduate student Pablo Azar will present a new type of mathematical game that they're calling a rational proof; it varies interactive proofs by giving them an economic component. Like interactive proofs, rational proofs may have implications for cryptography, but they could also suggest new ways to structure incentives in contracts.

"What this work is about is asymmetry of information," Micali adds. "In computer science, we think that valuable information is the output of a long computation, a computation I cannot do myself." But economists, Micali says, model knowledge as a probability distribution that accurately describes a state of nature. "It was very clear to me that both things had to converge," he says.

A classical interactive proof involves two players, sometimes designated Arthur and Merlin. Arthur has a complex problem he needs to solve, but his computational resources are limited; Merlin, on the other hand, has unlimited computational resources but is not trustworthy. An interactive proof is a procedure whereby Arthur asks Merlin a series of questions. At the end, even though Arthur can't solve his problem himself, he can tell whether the solution Merlin has given him is valid.

In a rational proof, Merlin is still untrustworthy, but he's a rational actor in the economic sense: When faced with a decision, he will always choose the option that maximizes his economic reward. "In the classical interactive proof, if you cheat, you get caught," Azar explains. "In this model, if you cheat, you get less money."

Complexity connection

Research on both interactive proofs and rational proofs falls under the rubric of computational-complexity theory, which classifies computational problems according to how hard they are to solve. The two best-known complexity classes are P and NP. Roughly speaking, P is a set of relatively easy problems, while NP contains some problems that, as far as anyone can tell, are very, very hard.

Problems in NP include the factoring of large numbers, the selection of an optimal route for a traveling salesman, and so-called satisfiability problems, in which one must find conditions that satisfy sets of logical restrictions. For instance, is it possible to contrive an attendance list for a party that satisfies the logical expression (Alice OR Bob AND Carol) AND (David AND Ernie AND NOT Alice)? (Yes: Bob, Carol, David and Ernie go to the party, but Alice doesn't.) In fact, the vast majority of the hard problems in NP can be recast as satisfiability problems.

To get a sense of how rational proofs work, consider the question of how many solutions a satisfiability problem has — an even harder problem than finding a single solution. Suppose that the satisfiability problem is a more complicated version of the party-list problem, one involving 20 invitees. With 20 invitees, there are 1,048,576 possibilities for the final composition of the party. How many of those satisfy the logical expression? Arthur doesn't have nearly enough time to test them all.

But what if Arthur instead auctions off a ticket in a lottery? He'll write down one perfectly random list of party attendees — Alice yes, Bob no, Carol yes and so on — and if it satisfies the expression, he'll give the ticketholder $1,048,576. How much will Merlin bid for the ticket?

Suppose that Merlin knows that there are exactly 300 solutions to the satisfiability problem. The chances that Arthur's party list is one of them

are thus 300 in 1,048,576. According to standard econometric analysis, a 300-in-1,048,576 shot at $1,048,576 is worth exactly $300. So if Merlin is a rational actor, he'll bid $300 for the ticket. From that information, Arthur can deduce the number of solutions.

First-round knockout

The details are more complicated than that, and of course, with very few exceptions, no one in the real world wants to be on the hook for a million dollars in order to learn the answer to a math problem. But the upshot of the researchers' paper is that with rational proofs, they can establish in one round of questioning — "What do you bid?" — what might require millions of rounds using classical interactive proofs. "Interaction, in practice, is costly," Azar says. "It's costly to send messages over a network. Reducing the interaction from a million rounds to one provides a significant savings in time."

"I think it's yet another case where we think we understand what's a proof, and there is a twist, and we get some unexpected results," says Moni Naor, the Judith Kleeman Professorial Chair in the Department of Computer Science and Applied Mathematics at Israel's Weizmann Institute of Science. "We've seen it in the past with interactive proofs, which turned out to be pretty powerful, much more powerful than you normally think of proofs that you write down and verify as being." With rational proofs, Naor says, "we have yet another twist, where, if you assign some game-theoretical rationality to the prover, then the proof is yet another thing that we didn't think of in the past."

Naor cautions that the work is "just at the beginning," and that it's hard to say when it will yield practical results, and what they might be. But "clearly, it's worth looking into," he says. "In general, the merging of the research in complexity, cryptography and game theory is a promising one."

Micali agrees. “I think of this as a good basis for further explorations,” he says. “Right now, we’ve developed it for problems that are very, very hard. But how about problems that are very, very simple?” Rational-proof systems that describe simple interactions could have an application in crowdsourcing, a technique whereby computational tasks that are easy for humans but hard for computers are farmed out over the Internet to armies of volunteers who receive small financial rewards for each task they complete. Micali imagines that they might even be used to characterize biological systems, in which individual organisms — or even cells — can be thought of as producers and consumers.

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Algorithmic incentives: New twist on 30 year-old work could lead to better ways of structuring contracts (2012, April 25) retrieved 3 July 2024 from https://phys.org/news/2012-04-algorithmic-incentives-year-old-ways.html