Game theory harnessed for cybersecurity of large-scale nets

November 16 2017, by Emil Venere

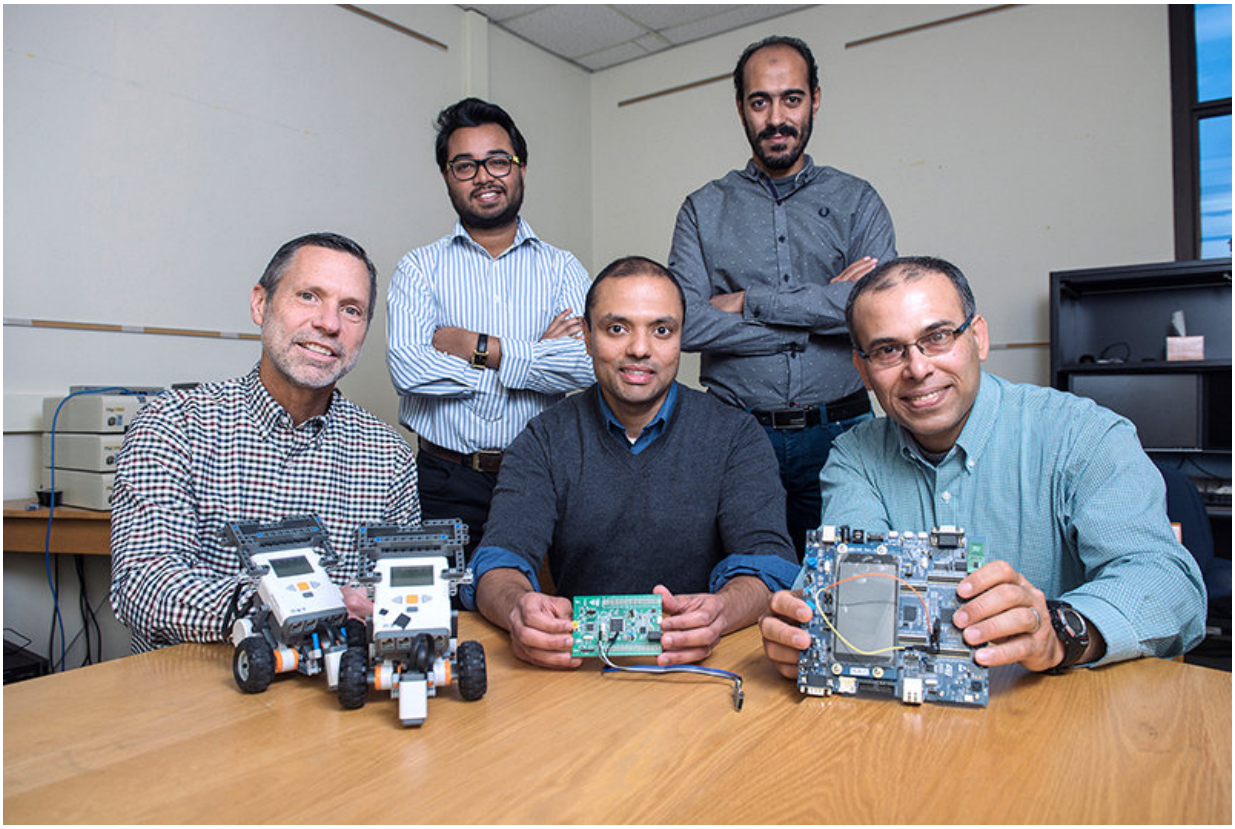

A Purdue University team displays hardware related to research to improve cybersecurity for large-scale systems like the power grid and autonomous military defense networks. From left are Timothy Cason, a professor of economics in the Krannert School of Management, graduate student Aritra Mitra, Shreyas Sundaram, an assistant professor in Purdue's School of Electrical and Computer Engineering, graduate student Mustafa Abdallah and Saurabh Bagchi, also a professor in the School of Electrical and Computer Engineering and Department of Computer Science. Credit: Purdue University image/Charles

Researchers have laid the groundwork for a method to improve cybersecurity for large-scale systems like the power grid and autonomous military defense networks by harnessing game theory and creating new intelligent algorithms.

Purdue University is leading the research, working with counterparts at Sandia National Laboratories. The work is funded with grants totaling about $700,000 from the National Science Foundation and Sandia.

The project harnesses the Nash equilibrium, developed by Nobel laureate John Nash, whose life was chronicled in the film "A Beautiful Mind." The work also applies "prospect theory," which describes how people make decisions when there is uncertainty and risk, decisions that are often "only partly rational," said Shreyas Sundaram, an assistant professor in Purdue's School of Electrical and Computer Engineering.

"The research will lead to a more complete understanding of the vulnerabilities that arise in large-scale interconnected systems and guide us to the design of more secure systems, with corresponding societal benefits," he said.

Sundaram leads the NSF-funded part of the project, working with co-principal investigators Saurabh Bagchi, a professor in the School of Electrical and Computer Engineering and Department of Computer Science, and Timothy Cason, a professor of economics in the Krannert School of Management.

"The two projects tackle the complexity of protecting today's large-scale systems," Bagchi said.

The NSF portion of the project started in August, while the Sandia portion began in October.

"In total, the project will provide new insights into the types of decisions that humans make when faced with security threats, via a comprehensive approach spanning theory and experiments," Cason said.

In conjunction with the project, Sandia researcher Abraham A. Clements has been working with the team while pursuing a doctoral degree at Purdue.

Their research is detailed in a chapter included in a new monograph titled "Game Theory for Security Risk Management – From Theory to Practice," to be published in early 2018 as part of Springer/Birkhauser's series on "Static & Dynamic Game Theory: Foundations and Applications." The chapter, "A Game-Theoretic Framework for Decentralized Defense of Interdependent Assets in Large-Scale Networks," was co-authored by Purdue doctoral student Ashish R. Hota, with Clements, Bagchi and Sundaram.

"Any large-scale system, like a power grid, an industrial control system or a consumer credit reporting agency, contains many parts, both cyber and physical, that you may have to secure," Bagchi said. "Human decision makers need to make pragmatic decisions about what to secure and to what extent. Our work seeks to put that kind of decision on a sound footing."

Complicating matters is that many large-scale systems involve "multiple owners" and subsidiary companies that are interdependent and may be vulnerable to attack, representing weak links in the network. However, budgetary constraints place limits on security measures. The project will tackle the complex workings of human decision making and the fact that various stakeholders could be acting in their own self-interest, Bagchi

said.

"You could use cooperative models where everybody decides to do what's best for the system, but that may not be realistic in practice," he said. "People may not want to reveal their own commercial secrets. There may be scenarios where I may actually choose to under-invest in security because I'm relying on somebody else upstream securing their assets. So, if I know that all attacks have to tunnel through you to get to me and I see that you've already invested a lot in security, I can save a lot of money by not protecting my assets because you've done the job for me."

The researchers have shown how to formulate an "optimization problem" that makes it possible to efficiently calculate how much a given stakeholder will decide to invest.

"We're not quite there yet, but our first steps are going to be to understand how the decision makers are going to make these security investments, to say, 'Stakeholder one is going to put in X amounts of dollars for protecting assets A, B, and C, and stakeholder two is going to do this and this,'" Bagchi said.

The team also will harness "moving-target defense."

"The basic premise is that the way most systems are set up is that they're sitting ducks. The attacker can try, try, and try again and only has to succeed once, whereas the defender has to succeed every time, every attack, 100 percent," Bagchi said. "Moving-target defense tries to change this dynamic and says that the system that I want to protect is not a sitting duck. It is going to reconfigure itself; it is going to change some aspect of itself from time to time so that if the attacker keeps trying the same thing, their chances of succeeding are not going to go up because the target that's being attacked changes somewhat."

Early research results in this aspect of the project were presented during the IEEE Security and Privacy conference in May 2017, in a paper co-authored by Bagchi and Clements, together with Matthias Payer, a Purdue assistant professor of computer science.

The project will include behavioral economics experiments using human subjects to evaluate the theoretical predictions and potentially yield new models of decision making.

"There is going to be a heavy component where we involve experiments using humans to actually see what kinds of security investments they make and use," Cason said. "The aim is to corroborate our models and to drive new research."

The NSF-funded part of the project is the game-theoretic decision -making side. The Sandia-funded part focuses on systematic ways to build more secure systems, particularly those made up of multiple interacting components, a goal with applications in the power grid and military defense systems. More secure systems are needed for "distributed platforms" using autonomous technologies in reconnaissance, surveillance and intelligence gathering.

"You will have not only human decision makers but also robots and UAVs that have been deployed, and there's also a sort of independent decision making that's happening by these autonomous systems," said Alex Roesler, a deputy director in the Integrated Military Systems Center at Sandia National Labs. "So, how do you come up with the algorithms to make these multi-UAV systems more secure?"

Future autonomous military technologies such as UAVs and land vehicles will be capable of persistent intelligence, surveillance, and reconnaissance, or PISR.

"The idea is that you deploy a fleet of UAVs or other mobile systems, where they're monitoring an area constantly," Sundaram said. "That's the persistent aspect of it: gathering intelligence, surveilling, doing reconnaissance."

When communications with base stations is compromised, the vehicles will be able to operate in a "self-organizing manner."

"If some of the UAVs are damaged, the others must be able to continue functioning together in a self-organizing network that needs to be smart enough to avoid these malfunctioning or compromised nodes," he said.

The team will build on recent work in developing secure distributed algorithms for situational awareness by Sundaram and his doctoral student Aritra Mitra, findings that were published in the 2016 *IEEE Conference on Decision and Control*.

Provided by Purdue University