Protecting our energy infrastructure from cyberattack

5 June 2019, by Nancy W. Stauffer

sing their “Cybersafety” methodology, Professor Stuart Madnick (left), graduate student Shaharyar Khan (right), and Professor James Kirtley Jr. (not pictured) identified several cyber vulnerabilities in a small power plant, including a system that poses a risk because it relies on software rather than mechanical safety devices to keep turbines from spinning out of control. Credit: Stuart Darsch

Almost every day, news headlines announce another security breach and the theft of credit card numbers and other personal information. While having one's credit card stolen can be annoying and unsettling, a far more significant, yet less recognized, concern is the security of physical infrastructure, including energy systems.

"With a credit card theft, you might have to pay $50 and get a new credit card," says Stuart Madnick, the John Norris Maguire Professor of Information Technologies at the Sloan School of Management, a professor of engineering systems at the School of Engineering, and founding director of the Cybersecurity at MIT Sloan consortium. "But with infrastructure attacks, real physical damage can occur, and recovery can take weeks or months."

A few examples demonstrate the threat. In 2008, an alleged cyberattack blew up an oil pipeline in Turkey, shutting it down for three weeks; in 2009, the malicious Stuxnet computer worm destroyed hundreds of Iranian centrifuges, disrupting that country's nuclear fuel enrichment program; and in 2015, an attack brought down a section of the Ukrainian power grid—for just six hours, but substations on the grid had to be operated manually for months.

According to Madnick, for adversaries to mount a successful attack, they must have the capability, the opportunity, and the motivation. In recent incidents, all three factors have aligned, and attackers have crippled major physical systems.

"The good news is that, at least in the United States, we haven't really experienced that yet," says Madnick. But he believes that "it's only motivation that's lacking." Given sufficient motivation, attackers anywhere in the world could, for example, bring down some or all of the nation's interconnected power grid or stop the flow of natural gas through the country's 2.4 million miles of pipeline. And while emergency facilities and fuel supplies may keep things running for a few days, it's likely to take far longer than that to repair systems that attackers have damaged or blown up.

"Those are massive impacts that would affect our day-to-day life," says Madnick. "And it's not on most people's radar. But just hoping that it won't happen is not exactly a safe way to go about life." He firmly believes that "the worst is yet to come."

The challenge for industry

Ensuring the cybersecurity of energy systems is a growing challenge. Why? Today's industrial facilities rely extensively on software for plant control, rather than on traditional electro-mechanical devices. In some cases, even functions critical for ensuring safety are almost

entirely implemented in software. In a typical industrial facility, dozens of programmable computing systems distributed throughout the plant provide local control of processes—for example, maintaining the water-level in a boiler at a certain setpoint. Those devices all interact with a higher-level "supervisory" system that enables operators to control the local systems and overall plant operation, either on-site or remotely. In most facilities, these programmable computing systems do not require any authentication for settings to be altered. Given this setup, a cyberattacker who gains access to the software in either the local or the supervisory system can cause damage or disruption of service.

The traditional approach used to protect critical control systems is to "air-gap" them—that is, separate them from the public internet so that intruders can't reach them. But in today's world of high connectivity, an air-gap no longer guarantees security. For example, companies often hire independent contractors or vendors to maintain and monitor specialized equipment in their facilities. To perform those tasks, the contractor or vendor needs access to real-time operational data—information that's generally transmitted over the internet. In addition, legitimate business needs, such as transferring files and updating software, require the use of USB sticks, which can inadvertently jeopardize the integrity of the air-gap, leaving a plant vulnerable to cyberattack.

Looking for vulnerabilities

Companies actively work to tighten up their security—but typically only after some incident has occurred. "So we tend to be looking through the rear-view mirror," says Madnick. He stresses the need to identify and mitigate the vulnerabilities of a system before a problem arises.

The traditional method of identifying cyber-vulnerabilities is to create an inventory of all the components, examine each one to identify any vulnerabilities, mitigate those vulnerabilities, and then aggregate the results to secure the overall system. But that approach relies on two key simplifying assumptions, says Shaharyar Khan, a fellow of the MIT System Design and Management program. It assumes that events always run in a single, linear direction, so one event causes another event, which causes another event, and so on, without feedback loops or interactions to complicate the sequence. And it assumes that understanding the behavior of each component in isolation is sufficient to predict the behavior of the overall system.

But those assumptions don't hold for complex systems—and modern control systems in energy facilities are extremely complex, software-intensive, and made up of highly coupled components that interact in many ways. As a result, says Khan, "the overall system exhibits behaviors that the individual components do not"—a property known in systems theory as emergence. "We consider safety and security to be emergent properties of systems," says Khan. The challenge is therefore to control the emergent behavior of the system by defining new constraints, a task that requires understanding how all the interacting factors at work—from people to equipment to external regulations and more—ultimately impact system safety.

To develop an analytical tool up to that challenge, Madnick, Khan, and James L. Kirtley Jr., a professor of electrical engineering, turned first to a methodology called System Theoretic Accident Model and Process, which was developed more than 15 years ago by MIT Professor Nancy Leveson of aeronautics and astronautics. With that work as a foundation, they developed "Cybersafety," an analytical method specifically tailored for cybersecurity analysis of complex industrial control systems.

To apply the Cybersafety procedure to a facility, an analyst begins by answering the following questions:

- What is the main purpose of the system being analyzed; that is, what do you need to protect? Answering that question may sound straightforward, but Madnick notes, "Surprisingly, when we ask companies what their 'crown jewels' are, they often have trouble identifying them."
- Given that main purpose, what's the worst that could happen to the system? Defining

the main purpose and the worst possible losses is key to understanding the goal of the analysis and the best allocation of resources for mitigation.

- What are key hazards that could lead to that loss? As a simple example, having wet stairs in a facility is a hazard; having someone fall down the stairs and break an ankle is a loss.
- Who or what controls that hazard? In the above example, the first step is to determine who or what controls the state of the stairs. The next step is to ask, Who or what controls that controller? And then, Who or what controls that controller? Answering that question recursively and mapping the feedback loops among the various controllers yields a hierarchical control structure responsible for maintaining the state of the stairs in an acceptable condition.

Given the full control structure, the next step is to ask: What control actions might be taken by a controller that would be unsafe given the state of the system? For example, if an attacker corrupts feedback from a key sensor, a controller will not know the actual state of the system and therefore may take an incorrect action, or may take the correct actions but at the wrong time or in the wrong order—any of which would lead to damage.

Overview of Cybersafety analysis: This figure summarizes the steps an analyst takes in performing a Cybersafety analysis. Credit: Massachusetts Institute of Technology

Based on the now-deeper understanding of the system, the analyst next hypothesizes a series of loss scenarios stemming from unsafe control actions and examines how the various controllers might interact to issue an unsafe command. "At each level of the analysis, we try to identify constraints on the process being controlled that, if violated, would result in the system moving into an unsafe state," says Khan. For example, one constraint could dictate that the steam pressure inside a boiler must not exceed a certain upper bound to prevent the boiler from bursting due to over-pressure.

"By continually refining those constraints as we progress through the analysis, we are able to define new requirements that will ensure the safety

and security of the overall system," he says. "Then we can identify practical steps for enforcing adherence to those constraints through system design, processes and procedures, or social controls such as company culture, regulatory requirements, or insurance incentives."

Case studies

To demonstrate the capabilities of Cybersafety analysis, Khan selected a 20-megawatt, gas turbine power plant—a small facility that has all the elements of a full-scale power plant on the grid. In one analysis, he examined the control system for the gas turbine, focusing in particular on how the software controlling the fuel-control valve could be altered to cause system-level losses.

Performing the Cybersafety analysis yielded several turbine-related loss scenarios involving fires or explosions, catastrophic equipment failure, and ultimately the inability to generate power.

For example, in one scenario, the attacker disables the turbine's digital protection system and alters the logic in the software that controls the fuel-control valve to keep the valve open when it should be closed, stopping fuel from flowing into the turbine. If the turbine is then suddenly disconnected from the grid, it will begin to spin faster than its design limit and will break apart, damaging nearby equipment and harming workers in the area.

The Cybersafety analysis uncovered the source of that vulnerability: An updated version of the control system had eliminated a backup mechanical bolt assembly that ensured turbine "over-speed" protection. Instead, over-speed protection was implemented entirely in software.

That change made sense from a business perspective. A mechanical device requires regular maintenance and testing, and those tests subject the turbine to such extreme stresses that it sometimes fails. However, given the importance of cybersecurity, it might be wise to bring back the mechanical bolt as a standalone safety device—or at least to consider standalone electronic over-speed protection schemes as a final line of defense.

Another case study focused on systems used to deliver chilled water and air conditioning to the buildings being served. Once again, the Cybersafety analysis revealed multiple loss scenarios; in this case, most had one cause in common: the use of variable frequency drives (VFDs) to adjust the speed of motors that drive water pumps and compressors.

Like all motors, the motor driving the chiller's compressor has certain critical speeds at which mechanical resonance occurs, causing excessive vibration. VFDs are typically programmed to skip over those critical speeds during motor startup. But some VFDs are programmable over the network. Thus, an attacker can query a VFD for the critical speed of the attached motor and then command it to drive the motor at that dangerous speed, permanently damaging it.

"This is a simple kind of an attack; it doesn't require a lot of sophistication," says Khan. "But it could be launched and could cause catastrophic damage." He cites earlier work performed by Matthew Angle '07, MEng '11, Ph.D. '16, in collaboration with Madnick and Kirtley. As part of a 2017 study of cyberattacks on industrial control systems, Angle built a lab-scale motor test kit equipped with a complete VFD with computer code familiar to the researchers. By simply altering a few key lines of code, they caused capacitors in the VFD to explode, sending smoke billowing out into the courtyard behind their MIT lab. In an industrial setting with full-sized VFDs, a similar cyberattack could cause significant structural damage and potentially harm personnel.

Given such possibilities, the research team recommends that companies carefully consider the "functionality" of the equipment in their system. Many times, plant personnel are not even aware of the capabilities that their equipment offers. For example, they may not realize that a VFD driving a motor in their plant can be made to operate in reverse direction by a small change in the computer code controlling it—a clear cyber-vulnerability. Removing that vulnerability would require using a VFD with less functionality. "Good engineering to remove such vulnerabilities can sometimes be mistakenly characterized as a move backwards, but

it may be necessary to improve a plant's security posture," says Khan. A full Cybersafety analysis of a system will not only highlight such issues, but also guide the strategic placement of analog sensors and other redundant feedback loops that will increase the resiliency of system operation.

Addressing the challenge

Throughout their cybersecurity research, Khan, Madnick, and their colleagues have found that vulnerabilities can often be traced to human behavior, as well as management decisions. In one case, a company had included the default passcode for its equipment in the operator's manual, publicly available on the internet. Other cases involved operators connecting USB drives and personal laptops directly into the plant network, thereby breaching the air-gap and even introducing malware into the plant control system.

In one case, an overnight worker downloaded movies onto a plant computer using a USB stick. But often such actions were taken as part of desperate attempts to get a currently shut-down plant back up and running. "In the grand scheme of priorities, I understand that focusing on getting the plant running again is part of the culture," says Madnick. "Unfortunately, the things people do in order to keep their plant running sometimes puts the plant at an even greater risk."

Enabling a new culture and mindset requires a serious commitment to cybersecurity up the management chain. Mitigation strategies are likely to call for reengineering the control system, buying new equipment, or making changes in processes and procedures that might incur extra costs. Given what's at stake, management must not only approve such investments, but also instill a sense of urgency in their organizations to identify vulnerabilities and eliminate or mitigate them.

Based on their studies, the researchers conclude that it's impossible to guarantee that an industrial control system will never have its network defenses breached. "Therefore, the system must be designed so that it's resilient against the effects of an attack," says Khan. "Cybersafety analysis is a powerful method because it generates a whole set of requirements—not just technical but also organizational, logistical, and procedural—that can improve the resilience of any complex energy system against a cyberattack."

More information: Identifying and Anticipating Cyber Attacks that could cause Physical Damage to Industrial Control Systems: web.mit.edu/smadnick/www/wp/2017-14.pdf

Cybersafety Analysis of Industrial Control System for Gas Turbines: web.mit.edu/smadnick/www/wp/2018-12.pdf

Arash Nourian et al. A Systems Theoretic Approach to the Security Threats in Cyber Physical Systems Applied to Stuxnet, *IEEE Transactions on Dependable and Secure Computing* (2015). DOI: 10.1109/TDSC.2015.2509994