Protecting our energy infrastructure: New analysis targets cybersafety

Nancy W. Stauffer, MITEI

Working Paper CISL# 2019-12

May 2019

Cybersecurity Interdisciplinary Systems Laboratory (CISL)
Sloan School of Management, Room E62-422
Massachusetts Institute of Technology
Cambridge, MA 02142

Energy Futures

MIT ENERGY INITIATIVE MASSACHUSETTS INSTITUTE OF TECHNOLOGY SPRING 2019

Protecting our energy infrastructure from cyberattack p. 3

Removing CO_2 emissions from power plant exhaust p. 9

Technology and policy pathways to Paris emissions targets p. 17

MITEI launches online energy classes to meet global needs p. 27

Protecting our energy infrastructure

New analysis targets cybersafety

Nancy W. Stauffer, MITEI

IN BRIEF

Using a new, holistic approach called cybersafety, an MIT team has shown that today's energy systems are rife with vulnerabilities to cyberattack—often the result of increased complexity due to high interconnectivity between devices and greater use of software to control system operation. The methodology examines a spectrum of factors that influence system operation, from physical design to operator behavior to organizational and managerial actions, and then determines how interactions among those factors can affect system safety. The resulting analysis can point to specific steps a company can take to strengthen the cybersecurity of its facilities. In the past decade, cyberattacks on physical systems have demonstrated that traditional IT security measures are largely impotent in protecting critical infrastructure from advanced cyber adversaries. The researchers therefore stress the urgent need to identify and mitigate cyber vulnerabilities, as future cyberattacks could cause unimaginable disruptions such as interrupting the flow of fuels or shutting down the US electric grid.

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 number stolen can be alarming, 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 MIT Sloan School of Management, a professor of engineering systems at the MIT School of Engineering, and founding director of the Cybersecurity

Above Using their "cybersafety" methodology, Professor Stuart Madnick (left), graduate student Shaharyar Khan, and Professor James L. 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. Photo: Stuart Darsch

at MIT Sloan consortium. "But with infrastructure attacks, real physical damage can occur, and recovery can take weeks or months." Madnick and his colleagues are now releasing methods and analytical tools they've developed that organizations can use to examine their facilities for vulnerabilities to cyberattack—including those that may arise from organizational, managerial, or employee actions.

A few examples of attacks since the advent of cyber-physical infrastructures demonstrate the seriousness of 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 to do so. In the incidents just cited, all three factors aligned, and attackers effectively 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 lives," 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, such 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 disrupt 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 commonly 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 functions such as transferring files and updating software often require the use of USB flash drives, 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 of a system's 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.

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

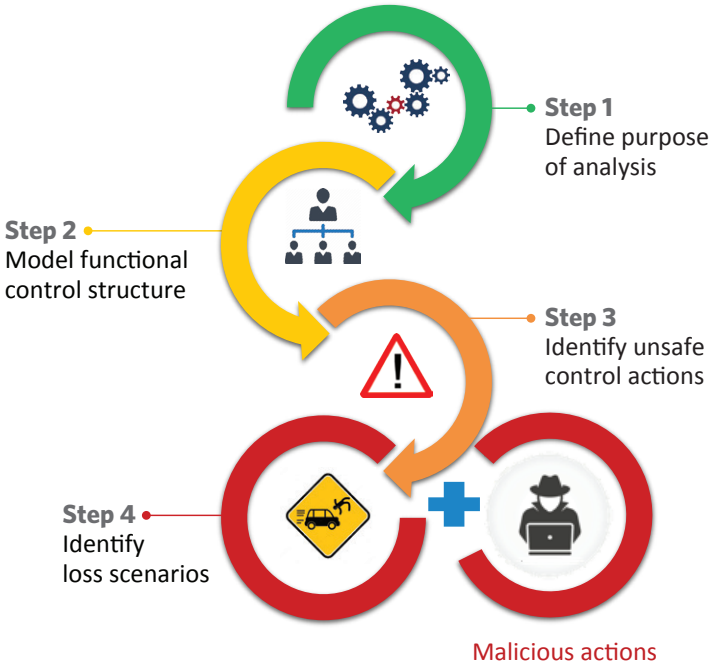

Overview of cybersafety analysis
This figure summarizes the steps an analyst takes in performing a cybersafety analysis.

Step 1: Define the purpose of the analysis by identifying unacceptable losses for the system as well as high-level system conditions that could be exploited to result in those worst-possible outcomes.

Step 2: Develop a model of the hierarchy of controllers and their interactions that together enforce safety and security constraints on system operation. (Controllers include human operators, automated systems, management, and even government and regulatory entities.)

Step 3: Identify control actions that could be unsafe and lead to system disruption or damage.

Step 4: Hypothesize how the controllers could interact to issue unsafe commands, given the malicious actions of an attacker. The analyst can now identify new requirements that would prevent the worst-possible outcomes named in Step 1.

electrical engineering, turned first to a methodology called System Theoretic Accident Model and Process (STAMP), which was developed more than 15 years ago by MIT Professor Nancy Leveson of aeronautics and astronautics. With that work as a foundation, they created "cybersafety," an analytical method specifically tailored to analyze the cybersecurity of complex industrial control systems (see the diagram above).

To apply the cybersafety procedure to a facility, an analyst begins by answering the following questions to define the task at hand.

- 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 thing 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. The diagram on page 6 shows a hierarchical control structure for one device in a power plant, illustrating the complexity of the factors and feedbacks involved.

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 at the wrong time or in the wrong order—any of which would lead to damage.

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—a limit needed 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."

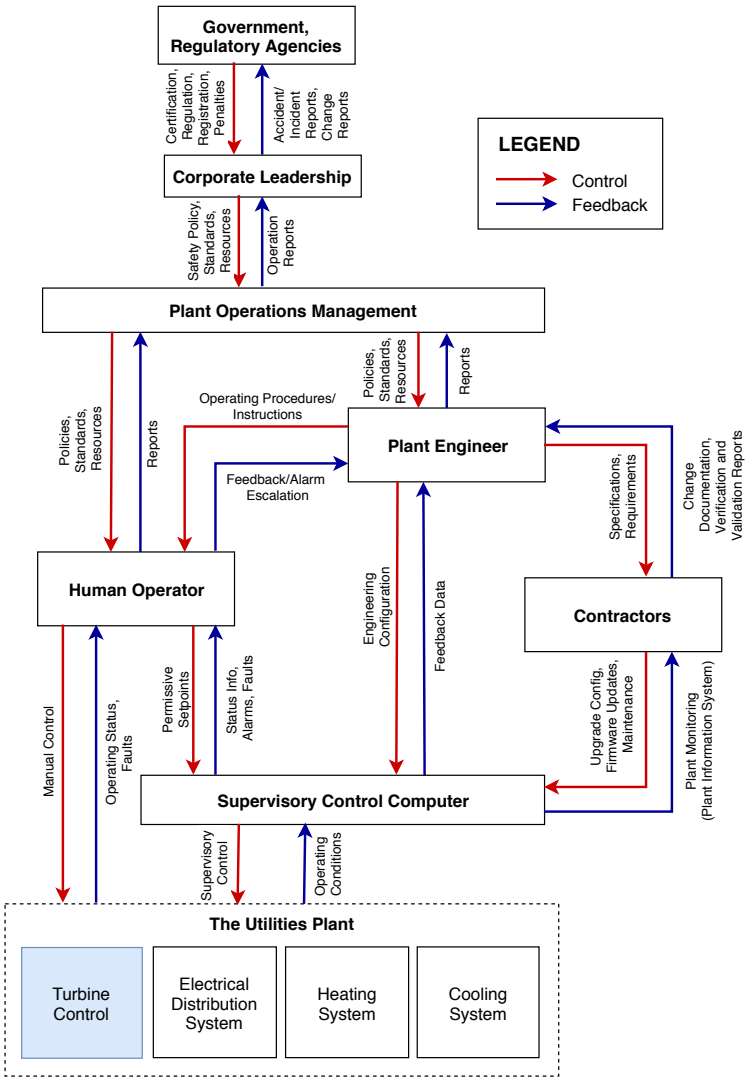

High-level hierarchical control diagram This illustration is a high-level diagram of the interplay among various actors in controlling a gas turbine in a typical power plant. Blue lines show information that each actor, or controller, receives in feedback. Red lines indicate actions that the controller can take in response to the received information.

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 revealed that attacks targeting software could produce several turbine-related loss scenarios, including 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 "overspeed" protection. Instead, overspeed 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 can subject the turbine to such extreme stresses that it sometimes fails. However, given the importance of cybersecurity, the researchers say it might be wise to bring back the mechanical bolt as a standalone safety device—or at least to consider standalone electronic overspeed 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; and 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—a practice that significantly increases the flexibility of operation and energy efficiency.

Like all motors, the motor driving this 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, which means that 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 the motor and the equipment connected to 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, PhD '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 computer code, they caused capacitors in the VFD to explode, sending smoke billowing out into the courtyard behind their MIT lab (photo at right). 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 systems. 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, which was publicly available on the internet. Other cases

In 2017, MIT researchers demonstrated a cyber vulnerability in a variable frequency drive (VFD) used to control the speed of a motor. For the demonstration, they used a lab-scale motor-development kit that had all the components found in a typical VFD. By altering a few lines in the VFD-control software, they caused a small-scale explosion that sent smoke billowing out into a nearby MIT courtyard, as shown above. Addressing the identified vulnerability is critical, as VFDs are ubiquitous in energy systems ranging from gas pipelines and oil wells to refineries and power plants.

involved operators plugging 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 more often such infractions were 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 put the plant at an even greater risk."

Fostering a new culture and mindset about cybersecurity requires a serious commitment from every level of the management chain, the researchers say. Mitigation strategies are likely to require some combination of reengineering the control system, buying new equipment, and making changes in processes and procedures, which might incur extra costs. Given what's at stake, the researchers argue that 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."

NOTES

This research was supported by the US Department of Energy, the MIT Energy Initiative Seed Fund Program, and members of the Cybersecurity at MIT Sloan consortium. More information can be found at cams.mit.edu and in:

M.G. Angle, S. Madnick, and J.L. Kirtley Jr. *Identifying and Mitigating Cyber Attacks that Could Cause Physical Damage to Industrial Control Systems*. Working paper CISL#2017-14, 2014. Online: bit.ly/madnick-cyber-attacks.

S. Khan, S. Madnick, and A. Moulton. *Cybersafety Analysis of a Central Utilities Plant (CUP) Gas Turbine Using STPA-Sec*. Working paper CISL#2018-12, October 2018. Online: bit.ly/madnick-turbine-analysis.

A. Nourian and S. Madnick. "A systems theoretic approach to the security threats in cyber physical systems applied to Stuxnet." *IEEE Transactions on Dependable and Secure Computing*, vol. 15, issue 1, 2015. Online: doi.org/10.1109/TDSC.2015.2509994.