Defending U.S. Critical Infrastructure from **NATION-STATE CYBERATTACKS**

Researchers combine cyberdefense expertise, network analysis, artificial intelligence, and collaborative-autonomy algorithms to defend the nation's industrial control systems.

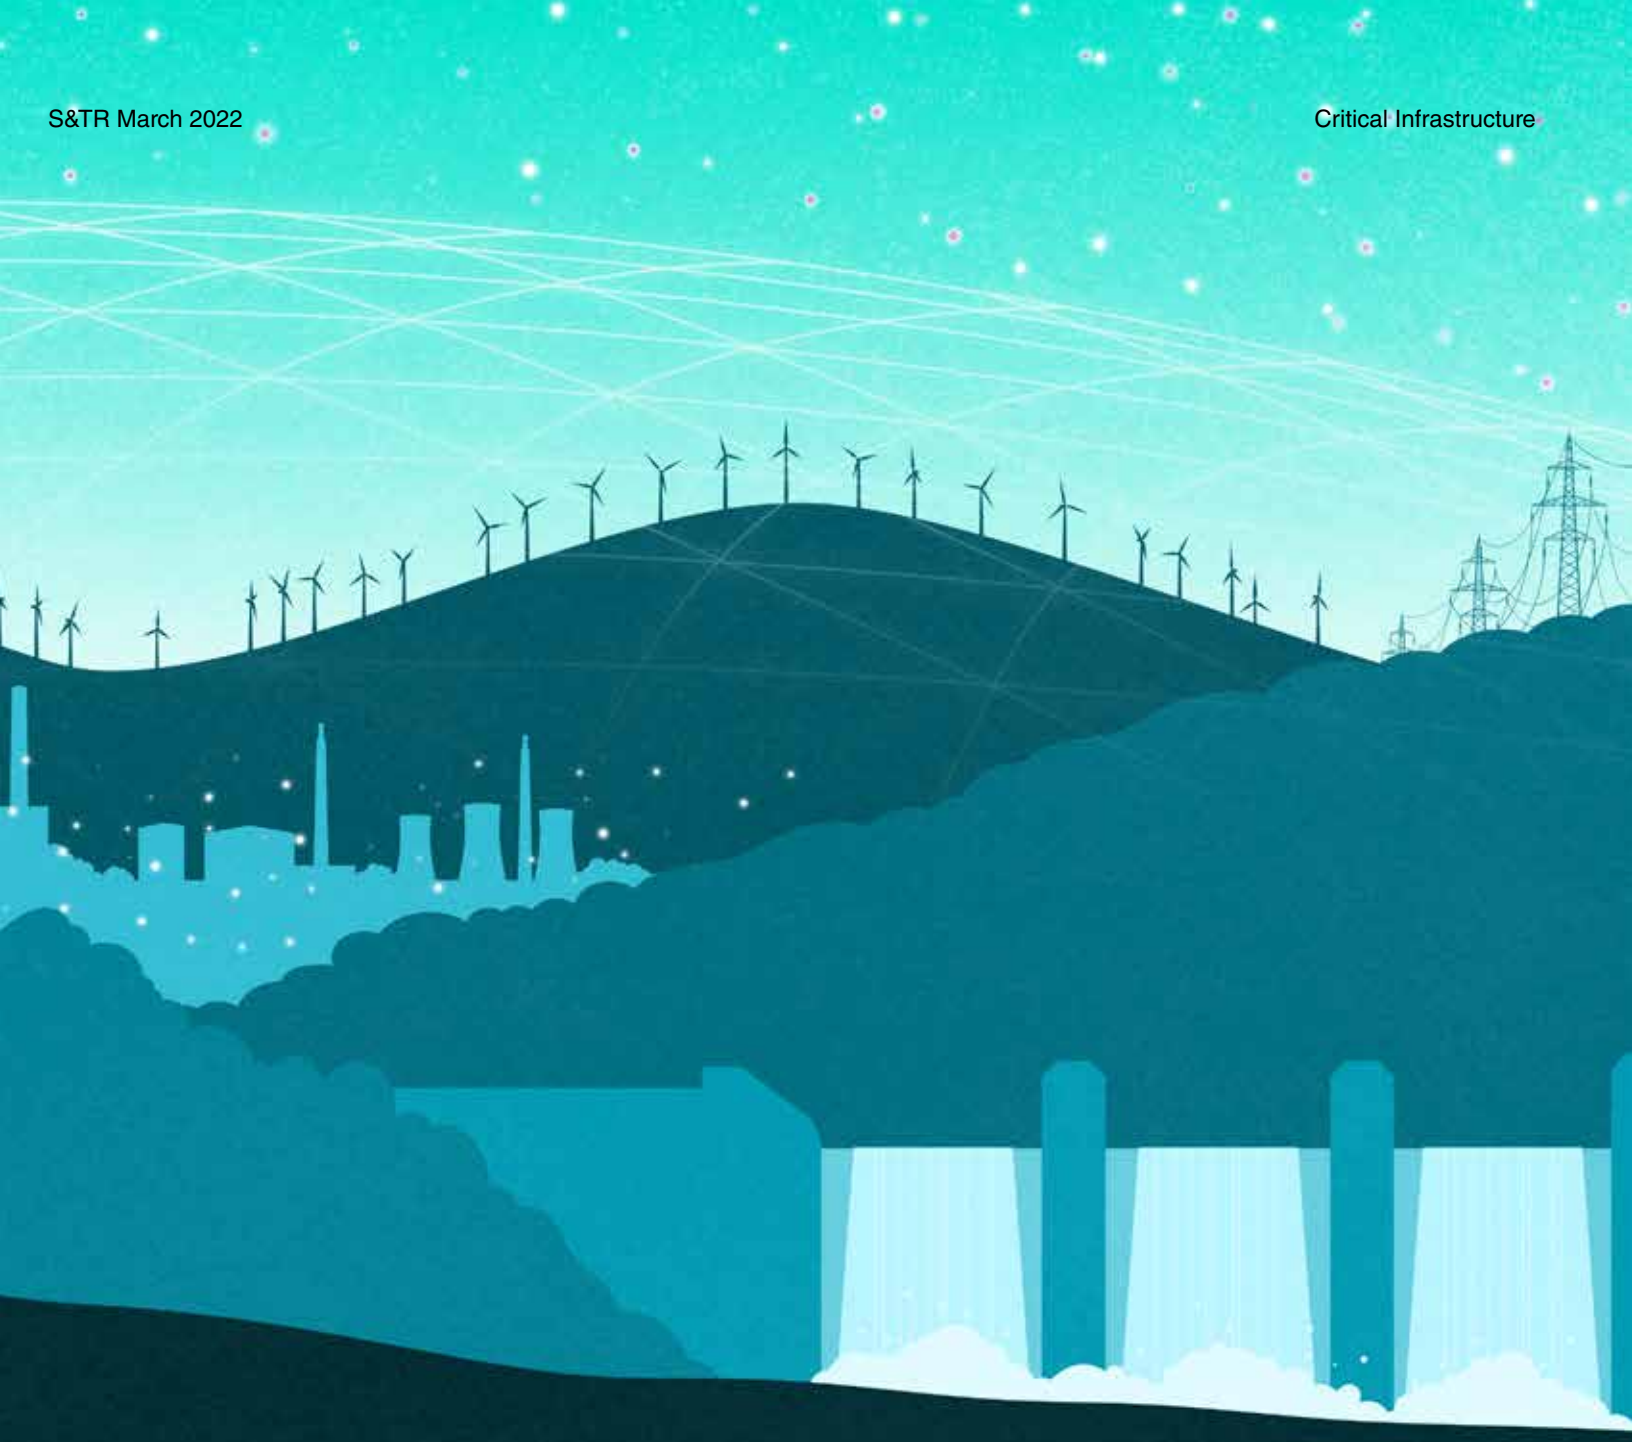

IN the news with increasing frequency are cyberattacks against critical infrastructure that supplies electricity, natural gas, water, transportation, and communication systems with the intent to disrupt these vital services. The first known cyberattack by a nation-state took place against Estonia in 2007. In response to the removal of a Soviet war monument in Tallinn, Russian-based attackers targeted state and commercial services, flooding them with junk digital traffic that rendered government, banking, and media websites inoperable. Then, in December 2015, hackers using malware compromised Ukraine's electric grid, disrupting electricity to hundreds of thousands of people. Subsequent attacks followed and in 2017, malware used against Ukraine spread globally, infecting banks, media outlets, and infrastructure in Germany, France, Italy, Poland, Russia, the United Kingdom, and Australia. Cyberattacks against electrical grids, natural gas pipelines, municipal railways, wastewater plants, and water utilities have also made headlines in the United States. Hackers have successfully extorted ransom money, collected intelligence, and disrupted critical infrastructure around the world, which, over the last several years, has become a digital front line in the conflict between nation-state adversaries.

For many years, Lawrence Livermore National Laboratory has been conducting

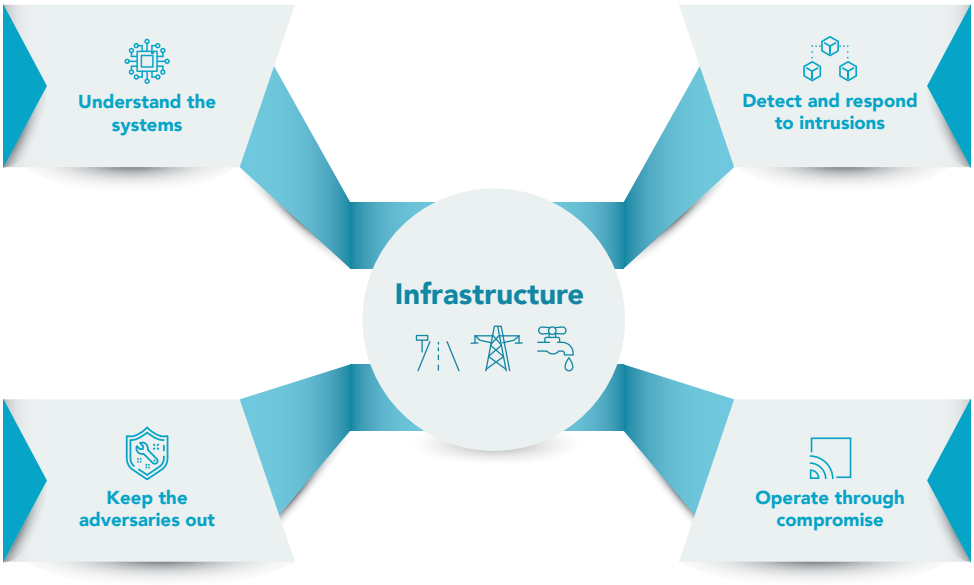

The Immune Infrastructure Program's vision is a strategic, layered defense that works to protect the nation's critical infrastructure from cyberattacks.

research on cybersecurity, as well as defending its systems and networks from cyberattacks. The Laboratory has developed an array of capabilities to detect and defend against cyberintruders targeting information technology (IT) networks and worked with government agencies and private-sector partners to share its cybersecurity knowledge to the wider cyberdefense community.

The proliferation of microprocessor-based industrial control systems (ICS) in recent decades has also exposed vulnerable points in critical infrastructure. In 2020, the U.S. Cybersecurity and Infrastructure Security Agency (CISA) identified 16 critical infrastructure sectors from the chemical industry to water and wastewater facilities, whose incapacitation would weaken national security and compromise public health and safety. With the rise of cyberattacks against critical infrastructure, Livermore has increased its focus on protecting the operational technology (OT) systems, which control and monitor utilities, transportation, communication, and manufacturing processes, as well as other industrial apparatus society relies on and requires.

Cyberthreats to critical infrastructure range from unsophisticated hobbyist hackers to organized crime syndicates and expert nation-state actors. Cyberattackers of any stripe may find and exploit well-known vulnerabilities, discover new ones, or create weaknesses where none previously existed. The least threatening adversary may be no more than teenagers with laptops breaking into an ICS to snoop around. Those looking for opportunities include organized criminal organizations using ransomware to extort money. The highest level of adversary can stealthily manipulate software, hardware, and networks using a full suite of capabilities to enter otherwise secure systems. Typically, these adversaries tend to be nation-states with deep expertise and resources whose motives are not necessarily financial—but intend to disrupt infrastructure or plot a future attack. Players at this level can manipulate the software supply chain by injecting malignant code into legitimate products or place insiders within companies and service providers to gain intelligence. "The vast majority of day-to-day threats are from lower-tier adversaries, and commercial industry does a decent job of defending against these threats," says Nate Gleason, program leader for the Laboratory's Cyber and Infrastructure Resilience Program. "We focus on protecting against national security-level threats to our critical systems, which generally means nation-state actors."

Layered Defense

The Laboratory has developed the immune infrastructure framework to protect critical infrastructure from these national security threats. Its goals are to create technologies and approaches that enable intelligent, self-healing, and resilient infrastructure by applying concepts from biological immunity to protect IT and OT networks from cyberattacks and physical disruptions. The Laboratory's design of an immune infrastructure framework departs from current cyberdefense strategies, which stress creating a robust, digitally secure perimeter to keep adversaries out—now an impractical defense against high-capability adversaries. "With lower level adversaries, we can keep the bad guys out of our systems," asserts Gleason. "When faced with more sophisticated actors, we have to assume that they will inevitably find a way to compromise our systems, so we must minimize the consequences of that intrusion and make it as difficult as possible for them to achieve their goals."

The immune infrastructure framework is comprised of four layers: understanding the systems, keeping the adversary out, detecting and responding to intrusions, and operating through compromise. Development of this approach leverages the Laboratory's ability to pull together multidisciplinary expertise in cybersecurity, data science

and machine learning, power-grid engineering, infrastructure, and systems analysis to address this predicament.

Understand the Systems

The Laboratory has developed a portfolio of capabilities that enable cyberdefenders to understand the asset inventory, which includes networks, hardware components, and software, as well as how an adversary might assess the system to successfully attack it. "Layer One focuses on understanding what's in networked systems in order to defend them," says Jovana Helms, associate program leader for Civilian Cybersecurity. Critical infrastructure such as the electricity grid or a water management system includes tens of thousands of networked, electronic devices that control a physical system and communicate with each other, forming a cyber–physical system. Some Layer One capabilities are designed to understand network software, hardware, and the connections between them; others are providing continuous network and device monitoring and creating digital twins of networks so that modeling and simulation can determine critical system nodes.

Livermore's Network Mapping System (NeMS) software is a longstanding capability that produces a comprehensive representation of Internet-based computer network environments. This tool, which has been licensed by public- and private-sector users, discovers and characterizes the devices, such as switches, routers, and hosts, and their connections on a network, allowing cyberdefenders to see how the system network is operating. It also determines where the IT network touches the OT network, so that cyberdefenders can monitor vulnerable points. "We don't want any unsupervised touch points between these two networks," says Helms.

The Laboratory Directed Research and Development (LDRD) Program has funded the development of a tool that goes one step further. The Industrial Control System Intelligent Device Characterization Tool (ICS ID ChaT) uses machine learning to analyze OT networks and characterize ICS devices on the network. The software identifies where the devices sit, who the manufacturers are, and what they do. "ICS ID ChaT's objective is to understand system behavior by analyzing network traffic to discover which components are on the network, their model numbers, and in the future, their firmware—the software permanently etched into a device by the manufacturer," says Helms. "It goes beyond existing network mapping by using machine learning to scan and infer what ICS devices are actually in these systems."

While ICS ID ChaT addresses existing cyber–physical systems, Livermore researchers are also working to secure the IT and ICS supply chain, a significant point of weakness in currently networked infrastructure systems. "A single firmware update can allow an adversary to compromise hundreds of devices," says Helms. Firmware updates to any device or system that manages an industrial process typically contain proprietary code written by the device vendor, as well as code from third-party software libraries from other companies, and sometimes,

Software Bill of Materials

Filesize 1183 kb
Executable Code

5 included components

Statically Linked Libraries 5		**55%**
libc v2.24	711 functions	45%
gcc v6.3.0	60 functions	4%
zlib v1.2.9	38 functions	2%
pcre v8.44	28 functions	2%
openssl v1.1.1d	27 functions	2%

Unidentified Code 45%

The Software Bill of Materials, much like the Nutrition Facts label on food packaging, provides cyberdefenders with information about the source of each piece of code in a software program, firmware, or automatic updates. Knowing exactly what code contains and where it came from will help cyberdefenders identify suspicious or malicious code.

even open-source code. Adversaries have successfully infiltrated and slipped malicious codes into firmware updates. A masterful example is the SolarWinds hack, discovered in 2020, which delivered a backdoor update to 18,000 corporate and institutional customers who use its network-monitoring software that gave hackers unfettered access to their networks. "We are developing tools that conduct automated software and firmware analysis, which helps us understand what's in software and crucially, firmware."

For the Valyrian Steel Project, a name drawn from the *Game of Thrones* television series, Laboratory researchers are developing a tool, Longclaw, which creates a software bill of materials (SBOM) or ingredients list for code that identifies where each component of code originated. The Department of Homeland Security–funded Longclaw tool extracts the code from a binary sequence and determines the libraries and functions that make up that code. "Knowing what libraries and functions are in software and firmware enables rapid and targeted response when a new vulnerability is discovered," says Helms. "Sometimes, the weak link is buried in the second or third layer of the software supply chain, so having SBOMs allows us to quickly recognize which systems have the vulnerability and how proliferated it is."

A third element of Layer One is modeling the physical process and the network controlling it to target vulnerable points. Sophisticated adversaries can launch attacks against dozens of components or points, while most utilities only have the computing power to model one or two components at best. Squirrel, part of the LDRD-funded Quantitative Intelligent Adversary Risk Assessment Project, is an algorithm that helps find these points by solving the inverse problem: For a given type of cyberattack and its consequence, such as shutting down components of an electrical transmission grid, Squirrel identifies the critical failures that lead to the consequence, and tells system operators which nodes of the network they should harden. The Laboratory is partnering with utilities to analyze their systems using Squirrel. One such run on a partner's grid revealed close to 200 weak nodes. Squirrel, however, also determined that protecting just 25 of those nodes would have eliminated all critical failures—a valuable insight that system operators can use to prioritize their limited resources for hardening their systems.

With funding from the Department of Homeland Security, Livermore has also launched the Night's Watch Project, bringing together several industry partners to demonstrate the Laboratory's infrastructure cyberdefense capabilities. Through efforts like these, the Laboratory is working to share these tools with electric and natural gas utilities and put them to use where they belong—defending the nation's critical infrastructure.

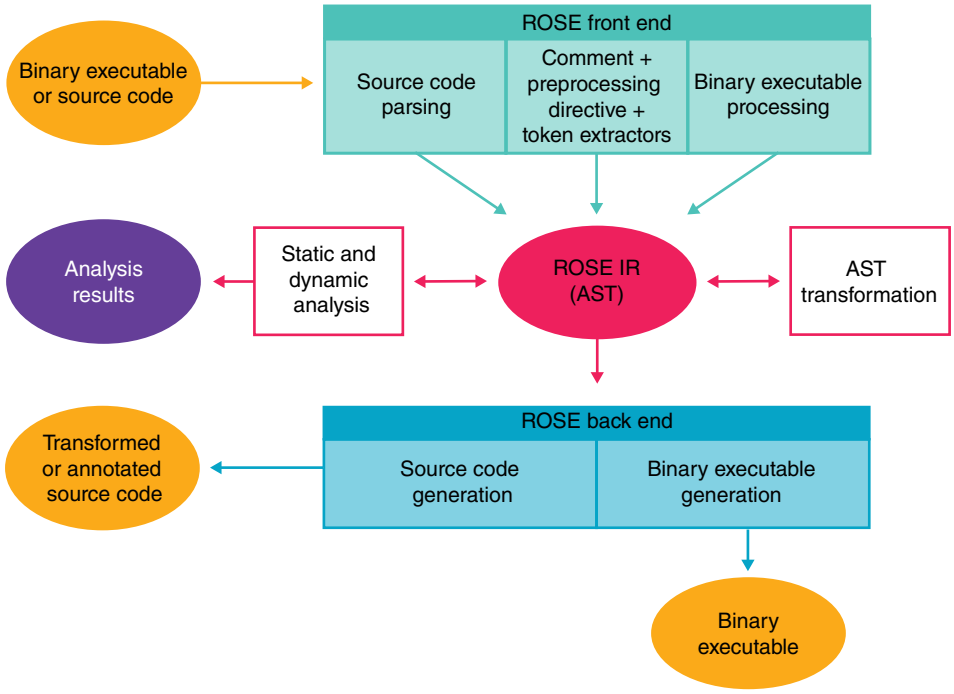

ROSE is a robust, open-source, compiler-based infrastructure for building source-to-source program transformation and analysis tools for static analysis, optimization, and other applications.

Keep the Adversary Out

Layer Two research provides tools to secure the hardware and software supply chain and keep the adversary out of systems. "Just because an adversary will inevitably find a way to penetrate a cybersecurity perimeter doesn't mean we should make it easy for them to do so. The idea with the layered approach is that each subsequent layer is there if the previous layer fails," says Helms. Researchers are building automated tools to increase the efficiency of device evaluation, enable cyberdefenders to verify the integrity of software and firmware updates, and develop self-verifying devices. "We have a software assurance focus, which tries to ensure that there are no flaws in the code the adversary can use to infiltrate the system," says Bob Hanson,

associate program leader for National Security Infrastructure at the Laboratory.

The Valyrian Steel Project also plays a key role in Layer Two. "It's very hard to assess large software programs even when they update monthly," explains Hanson. "Valyrian Steel focuses on automating software assurance analysis." Longclaw can also run analytics to detect suspicious features and evaluate the quality and security of the code, not just identify its contents.

ROSE, a Livermore-developed, open-source compiler for building source-to-source program analysis, is another essential Layer Two tool. With ROSE, cyberdefenders can perform automatic binary analyses of software updates in search of malicious code and scale their analysis to hundreds or thousands of pieces of software without requiring a high degree of user sophistication. FUNPAC (Firmware Updates Need Proof of Accompanying Code) is an LDRD-funded project to develop formal verification techniques for firmware. Using the ROSE compiler infrastructure, FUNPAC analyzes vendor-annotated binaries to determine if firmware conforms to security requirements. FUNPAC is designed so that grid devices such as processors, voltage converters, and regulators with limited computing power can verify that a firmware update does not contain malicious code before installing it.

Cybersecurity Testing for Resilient Industrial Control Systems (CyTRICS) is DOE's program for cybersecurity vulnerability testing, digital subcomponent enumeration, and forensic assessment. CyTRICS leverages best-in-class test facilities and analytic capabilities at six DOE national laboratories and strategic partnerships with key stakeholders including technology developers, manufacturers, asset owners and operators, and interagency partners. DOE's Office of Cybersecurity, Energy Security, and Emergency Response (CESER) funds CyTRICS. Lawrence Livermore is a participating laboratory in CyTRICS; performs testing on high-priority digital components in OT and ICS; brings its specialization in software analysis to the program; and leads efforts to automate testing, analysis, and SBOM generation for energy-grid devices.

Detect and Respond to Intrusions

In the Dragonglass Project, another *Game of Thrones*–themed name, Livermore researchers are building intelligent detection capabilities to automatically respond to unknown threats for the immune infrastructure framework's Layer Three. State-of-the-art detection algorithms look for signatures of network compromise. This approach works for detecting less-sophisticated adversaries, but not against the highest tier where adversaries deploy tactics that defenders are not likely to have seen before. To anticipate these unknown or unfamiliar tactics, researchers are using deep-reinforcement learning (DRL), a form of machine learning. By running a large number of transmission and distribution simulations, these algorithms can learn what a healthy system looks like and use that data to identify any behavior that could disrupt the health of the system.

"We're trying to teach the system to defend itself against attacks using deep-reinforcement learning," says Jean-Paul Watson, senior research scientist in Livermore's Center for

(a) Dragonglass uses a digital twin of a system to gather information, detect unusual behavior, isolate compromised components in a system, or counter suspicious commands.

(b) Longclaw creates a Software Bill of Materials by extracting code and determining the libraries and functions within it, as well as where it originated.

(c) FUNPAC (Firmware Updates Need Proof-Accompanying Code) provides formal verification techniques by analyzing vendor-annotated binaries to establish that firmware conforms to security policy before it is installed.

Applied Scientific Computing (CASC). "We are developing a simulator for each side of the system: one that captures the physics and one that looks at the control systems' points of contact with the physical systems." Using training sets of information about what normal behavior, as well as what attacks look like on the grid, this system acts like an immune system by learning how to recognize an attack, isolate or eliminate it, and return and restore normal system functioning.

To train the algorithms to identify abnormal activities, Dragonglass uses a digital twin of a system. When unusual behavior is detected, the algorithms gather information about the event and automatically respond. For example, by closing or opening relays or countering suspicious commands, the algorithms can isolate parts of the physical system that are compromised. The algorithm can also ignore normal actions, investigate abnormal actions, make corrections, and alert human operators. DRL algorithms developed for the Dragonglass Project also have climate-change relevance: they can be used to reroute power flows around sections of the grid affected by storms or wildfires. Elements of the Dragonglass Project, which address the electric grid, are funded by DOE's Grid Modernization Laboratory Consortium, and DOE CESER funds oil and natural gas applications.

The Laboratory's Skyfall Facility is a test bed that simulates an electrical grid where researchers can run Dragonglass's DRL and other detection algorithms to evaluate their effectiveness. (See *S&TR*, December 2018, pp. 16–19.) Livermore has also partnered with the University of Toledo, where a portion of the University's campus electrical grid has been heavily instrumented, and can provide real-time electric grid operational data to help train the software. Schweitzer Engineering Laboratory manufactures equipment for electric-grid controls and, as a partner in the Dragonglass Project, provides data on how their equipment behaves under normal operating conditions.

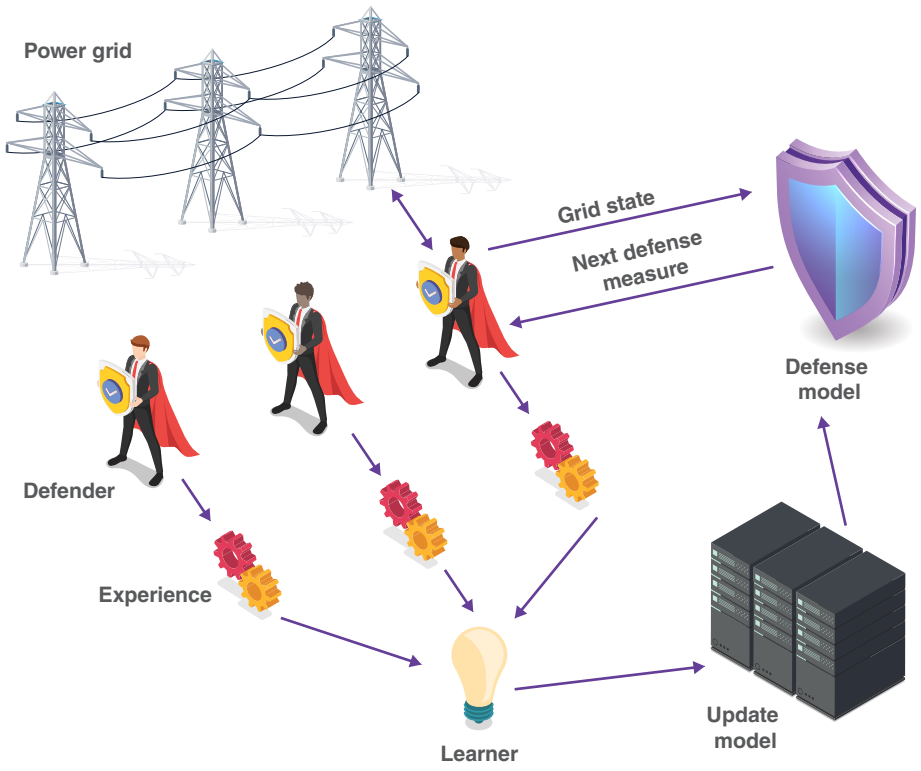

Dragonglass uses deep-reinforcement learning to monitor a control system, determine what normal and abnormal operations look like, build a model capable of recognizing a cyberattack or a natural catastrophe, take action to eliminate the threat, and keep the system running.

Operating through Compromise

A motivated, technically sophisticated cyberattacker will inevitably find a way into a cyber–physical system despite the preceding layers of protection. Layer Four research develops capabilities to facilitate infrastructure operations despite an attack on a part of the system. "Modern infrastructure is full of digital components," says Colin Ponce, a computational mathematician at CASC. "The nation's critical infrastructure is geographically diffuse, yet connected, and programmable. In these network-connected systems, all of the devices relay their data to a central control system, which performs analysis and sends out commands. The problem is that if a cyberattacker infiltrates the control center, they have the means to shut down the entire system."

To counter this vulnerability, Livermore researchers are deploying collaborative autonomy to decentralize control of physical systems. (See *S&TR*, June 2018, pp. 12–15.) Instead of all functions relying on a single, central control center or machine, low-power edge devices—such as solar inverters, smart meters, and vehicle chargers—distributed throughout the network can perform independent analysis, verify, and communicate with neighboring devices or nodes to reach consensus

on the next steps to take—the central axiom of collaborative autonomy. "Using collaborative autonomy, many devices can self-organize into a collective whole to reliably conduct monitoring and operations. No single device or control point can precipitate system or network failure," says Ponce. "We are leveraging the distributed nature of the system and using it to our advantage."

Livermore researchers are developing algorithms that allow distributed energy resources (DER) such as solar inverters and smart meters to share input data with their neighbors, and verify, for example, the voltages of different devices based on known data patterns to determine if they are normal. The underlying method uses "gossip" or update-and-share algorithms to propagate information about what the network is communicating from one component to another like a rumor. Each device then verifies its neighbors' computations mathematically.

The robust DERMS (Distributed Energy Resource Management System) Project is an effort to apply collaborative autonomy to keeping solar inverters, which convert the variable direct-current output of a photovoltaic (PV) solar panel into an alternating current that can be fed into an electrical grid, responsive to cyberattack. In a PV grid, multiple inverters are controlled centrally, so the adversary only has to compromise one device or central controller to gain control over multiple solar inverters. An adversary could destabilize the grid by telling all inverters to, for example, power down simultaneously. Funded by the Department of Energy's (DOE's) Office of Solar Technologies, Ponce's team has inserted a verification step into software that manages the inverters. When an operator (or adversary) sends a control command to DERMS, it is then relayed to the inverters, which uses collaborative autonomy to assess and decide whether to accept or reject it.

Managing DERs through collaborative autonomy also helps grids become more resilient to climate change. A grid composed of multiple DERs can be managed as a group, keeping itself running if one DER crashes during extreme weather events. DERs could also match variable loads with demand, for example, reconciling intermittent PV availability with charging batteries for electric vehicle fleets, to smooth out uneven supply and demand.

For electric utilities, a black start—restarting the power grid from a total blackout—is the most challenging responsibility. Typically, a black start requires a large power plant to anchor the restart of the rest of the grid. But if an adversary has compromised that larger plant, the grid cannot restart. Federally owned Plum Island, in waters near Long Island, New York, is the site of government facilities equipped with an instrumented electric grid researchers can utilize to test out cybersecurity infrastructure solutions. Livermore researchers have been conducting an experiment there, the Plum Island Blackstart Project, to execute a black start using collaborative autonomy via DERs using batteries on a PV system without a central controller. DOE's Grid Modernization Laboratory Consortium funds the project. "We installed software on the Plum Island system to collaboratively manage the DERs and were able to successfully execute a black start. The technology works," Ponce explains. "We can now use DERs as a control verification framework or as a redundant control system that kicks in if the central controller is compromised. These algorithms are applicable to industrial control systems, water, communications, manufacturing—anything that's too big to see in its entirety at a glance."

Transitioning to Operations

In the immune infrastructure roadmap, the first step to building and testing the components of the infrastructure protection framework will be followed by a proof-of-concept pilot at a cyber–physical demonstration site. Livermore has taken the first steps toward this goal at Site 300, a Lawrence Livermore experimental facility 24 kilometers from the Laboratory's main site. Site 300's core mission is assessing the operation of nonnuclear components of weapons systems, and its large open space provides an ideal place to build a pilot-scale ICS. Planning has begun. "Our vision is to bring together equipment vendors and asset owners so they can see how to incorporate immune infrastructure technology and this layered framework in their own devices," Gleason says. "They will receive firsthand experience incorporating these technologies while we identify the best solutions."

—Allan Chen

Key Words: Center for Applied Scientific Computing (CASC); collaborative autonomy; climate change; cyber–physical system; cybersecurity; Cybersecurity and Infrastructure Security Agency (CISA); Department of Energy (DOE) Grid Modernization Laboratory Consortium; distributed energy resources (DER); Distributed Energy Resource Management System (DERMS); DOE Office of Cybersecurity, Energy Security, and Emergency Response (CESER); Department of Homeland Security; electricity grid; industrial control system (ICS); immune infrastructure; information technology (IT); Laboratory Directed Research and Development (LDRD) Program; Network Mapping System (NEMS); operational technology (OT); photovoltaic (PV); resilience; Skyfall; software bill of materials (SBOM).

For further information contact Nate Gleason (925) 423-6278 (gleason6@llnl.gov).