Lab Notes

NEWS FROM AROUND LINCOLN LABORATORY

CRYPTOGRAPHY

Securing Data

A novel technology simplifies secure military communications and has the potential to be beneficial for a wide array of applications

The Department of Defense (DoD) military strategy relies in part on the development of advanced system technologies that can enable new capabilities for warfighters in the field. For example, imagine the advantages of a new drone or small satellite that can see through foliage and deliver high-resolution, tactical imagery. Because these new technologies are promising, developers focus on creating the systems' major subcomponents quickly, leaving important considerations like cyber security on the back burner. When developers eventually turn to incorporating security features, they often face several complications because they are so far along in the design lifecycle. At this stage, redesigning the system to add security features typically results in crippled system usability, major design delays, superficial security, and large cost overruns.

To address this problem, researchers at Lincoln Laboratory are developing new tools, including a software component known as the Lincoln Open Cryptographic Key Management Architecture (LOCKMA).[1] This software quickly and inexpensively simplifies the task of securing data and communication in a wide variety of systems and may even be employed during later stages of the design cycle.

Cryptography is a vital tool for passing sensitive information to intended recipients and keeping that same information from prying eyes. Through the encryption of data into an unintelligible sequence of characters called ciphertext, a sender scrambles a message with an algorithm and a cryptographic key, and the intended recipient decrypts, or unscrambles, the message by using a symmetric, i.e. used for both encryption and decryption of data, algorithm and key. Though several encryption solutions are widely used to secure data today, they all have one major shortfall: key management.

Key management is the process of managing cryptographic keys, that is, generating secure cryptographic keys, making them available to authorized users, and storing them. It is arguably the most difficult aspect of cryptography, says Daniil Utin of the Secure Resilient Systems and Technology Group at Lincoln Laboratory, because developing a new key management scheme may inadvertently introduce security vulnerabilities caused by system bugs and development oversights. "Developers make key management systems that combine low-level cryptographic functions into a secure design that supports high-level security functions; it is a complicated process. During the design process, developers can sometimes unintentionally create an insecure system. Even a small bug in the key management system, such as a biased random number generator that enemies can easily exploit, can create a big security vulnerability," says Utin.

Some key management solutions rely on manual key distribution. For example, if two military units plan to send encrypted radio messages to each other, they must first download cryptographic keys onto a Key Processor computer over secure phone lines by using the Electronic Key Management System (EKMS) or over a digital network by using the Key Management Infrastructure (KMI). The units must then manually program the keys into the radio of each communicating device. The key-loaded

[1] Lincoln Open Cryptographic Key Management Architecture (LOCKMA) is available for licensing through the MIT Technology Licensing Office (TLO) under MIT case number 16575L. For more information about LOCKMA and LOCKMA-related patents, contact the TLO at tlo-atto@mit.edu.

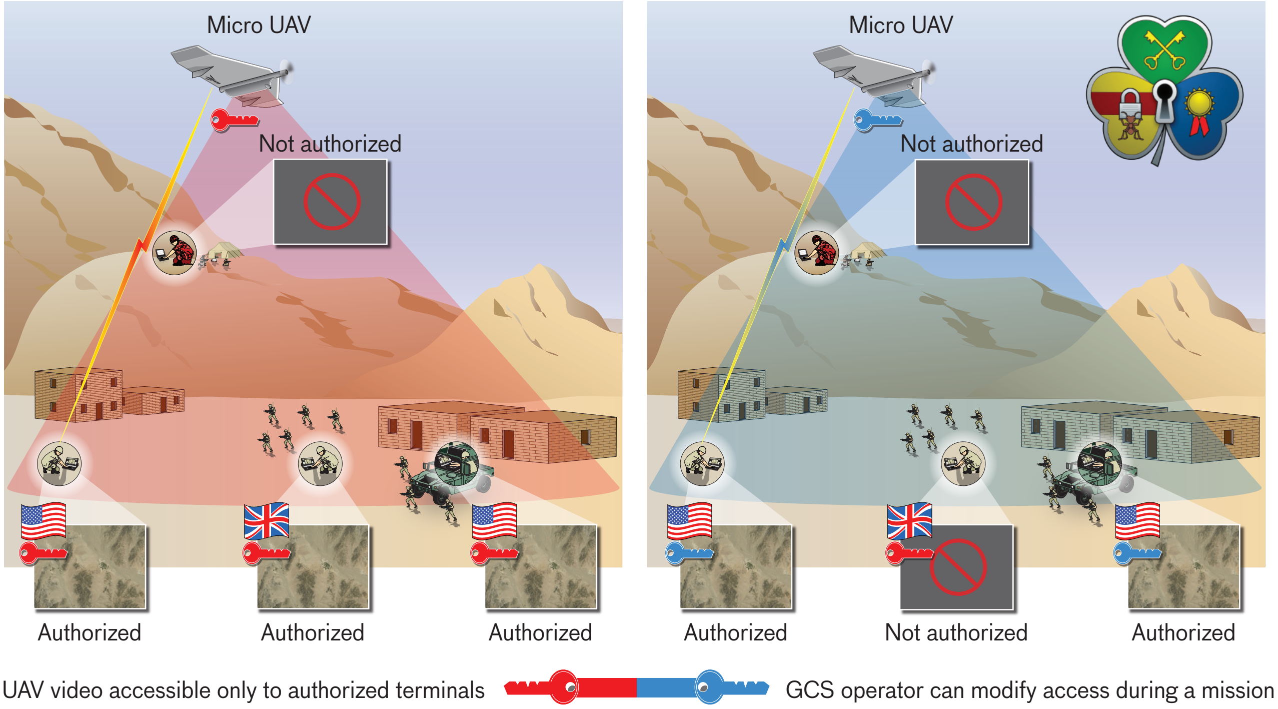

A LOCKMA user can transmit data from a ground control system (GCS) to intended recipients via an unmanned aerial vehicle (UAV) by providing keys to authorized users and can deny access to unauthorized users. The LOCKMA software transmits these authorizations to the intended recipients.

radios are used for just one mission; if the units need to send encrypted information during a future mission, they must download and install new keys into the radios. This key distribution process presents several risks, according to Benjamin Nahill of the Secure Resilient Systems and Technology Group. For instance, it may be difficult or impossible for units in the field to access EKMS or KMI. If an enemy captures a unit's radio, the enemy could eavesdrop on communication or impersonate the radio operator. All units involved in the communication must therefore obtain and manually program new keys into their radios. Says Nahill, "It is difficult to ensure that each unit has the correct key, so there is a need for dynamic key management."

Designed to reduce development errors and simplify the key management process, LOCKMA implements storage, organization, and management of key-related information, including key lifecycles, authorized users, and communication channels. Prior to field deployment, each LOCKMA-enabled device undergoes cryptographic provisioning during which LOCKMA generates private keys for each device that will be used in the field. The LOCKMA software uses a public key infrastructure (PKI) service to create certificates that cryptographically bind public keys to individual devices. If a unit wants to securely send a message, it can provide LOCKMA with each intended recipient's device certificate and then use LOCKMA to securely distribute symmetric mission keys and corresponding metadata for message decryption. Because device certificates typically have a lifetime of several years, units can use the same radios for consecutive missions, bypassing multiple journeys to distribution bases. If adversaries capture a radio, LOCKMA will enable the distribution of new mission keys to all authorized radios but the captured one, preventing the enemy from receiving any new communication. The certificate of the captured radio can later be revoked through PKI.

The military is increasingly relying on the use of unmanned aerial vehicles (UAV) to distribute tactical information. For example, a unit might deploy a UAV throughout mountainous landscape to locate enemy troops or scout tactical locations. The UAV's radio sends a signal back to the unit's ground station, displaying the live video feed from the UAV's camera. However, adversaries are gaining technology to intercept these feeds. To prevent unauthorized video access, LOCKMA can be integrated into the UAV's radio to encrypt the feed and create access restrictions. Using LOCKMA, a unit can identify intended video-feed recipients by their certificates and then send the symmetric key to a group of authorized recipients to decrypt the feed.

The Department of Defense is currently working with Draper Laboratory and the National Security Agency to integrate LOCKMA into devices that could benefit from its simplicity, focusing their research efforts on digital radios attached to small tactical devices like UAVs, according to Utin.

LOCKMA's key management messages are based on a cryptographic language standard called Cryptographic Message Syntax, which allows the software to understand and operate seamlessly with most cryptographic algorithms, modes, or key lengths. LOCKMA also works with many operating systems (e.g., Windows, Linux, Android, iOS) and independently, allowing application developers to integrate LOCKMA into devices with minimal changes to LOCKMA's application code.

"Overall, LOCKMA is much more flexible and easier to integrate than traditional key-management systems because, without its holistic approach to security, the entire cryptographic process, from key creation, to management, to delivery would be much more difficult and error prone," says Utin.

LOCKMA's application programming interface is easy for users to understand even without advanced cryptography knowledge. It hides all cryptographic complexities under the hood, allowing application developers to quickly integrate LOCKMA into devices. The technology saves time and costs compared to a custom key-management system that requires substantial expertise, time, and capital to develop, integrate, and test. Traditional custom-built solutions are also prone to security flaws that are often exploited by adversaries, resulting in substantial additional mitigation and repair costs.

The recipient of a 2012 R&D 100 Award, LOCKMA may become more commonplace as researchers look to integrate it in commercial applications. For example, an increasing number of homes are connecting to "smart" management applications that control energy use and security systems, e.g., Google Nest. Homeowners could use LOCKMA-enabled devices to secure communications within home networks and to thwart hackers. For government organizations, LOCKMA may be useful in tactical operations. "Consider the Boston Marathon bombings," says Utin. "After the attack took place, organizations, including the police and FBI, were communicating over standard shortwave radios. The radios gave

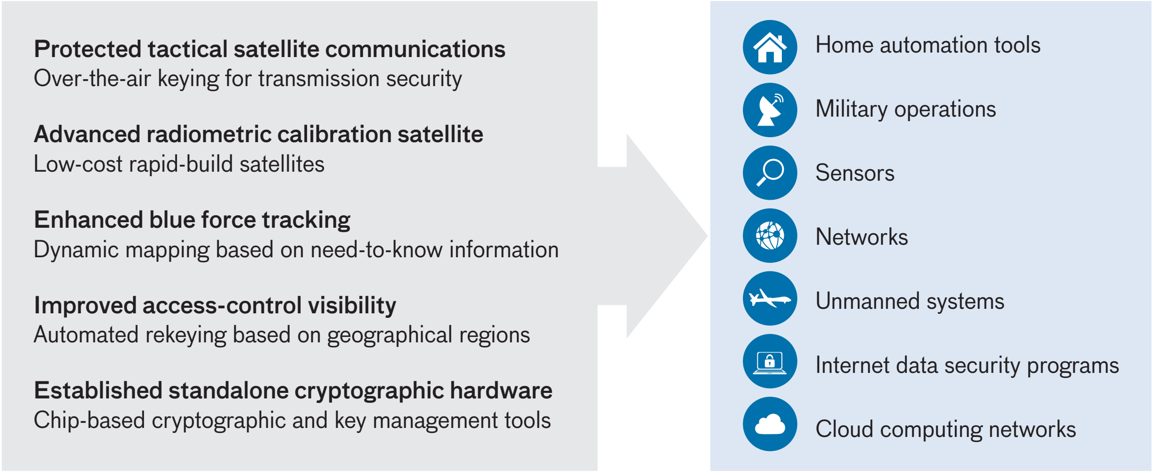

Protected tactical satellite communications
Over-the-air keying for transmission security

Advanced radiometric calibration satellite
Low-cost rapid-build satellites

Enhanced blue force tracking
Dynamic mapping based on need-to-know information

Improved access-control visibility
Automated rekeying based on geographical regions

Established standalone cryptographic hardware
Chip-based cryptographic and key management tools

Home automation tools
Military operations
Sensors
Networks
Unmanned systems
Internet data security programs
Cloud computing networks

The LOCKMA software can be applied to many applications that require cryptographic key management. Researchers are currently working to develop features (left) that enable LOCKMA to be used in commercial applications (right), such as home automation and cloud computing.

anyone, including the suspects, access to those communications. If organizations employ LOCKMA-enabled devices to protect their communications both online and in the field, they can securely provide necessary information to authorized recipients and dynamically accommodate access control changes in real time. LOCKMA really can make a huge difference in national security."

NETWORK SECURITY

Keeping an Eye on Cyber Threats

Researchers use real-time data from Lincoln Laboratory networks to monitor and develop countermeasures against cyber threats

It is a silent threat: as you read through your morning email and catch up on the news, hackers could be working to steal your passwords and sift through your files. Without warning, your private virtual world could become public. Many computer users know how devastating a cyber attack can be. But imagine the same thing happening to your office network. Now imagine it happening to an even larger network, such as that operated by a local government agency or financial institution. In 2014, threat became reality when researchers discovered an Internet security vulnerability named Heartbleed in the widely deployed network security library OpenSSL. Some analysts said Heartbleed had the potential to be the most catastrophic vulnerability ever found. It allowed hackers to probe web servers by simply sending a short command packet, i.e., a heartbeat packet, to the server. The packet would ask the server to echo information back to the user with extra data attached, leading previously secure websites to "bleed" information that could include private data, such as passwords and credit card numbers. Heartbleed affected more than 500,000 networks, including Lincoln Laboratory's, but the Laboratory's Security Services Department (SSD) and Information Services Department (ISD) had a unique cyber defense resource that allowed them to quickly access the data they needed to identify, assess, and neutralize the threat: the Lincoln Research Network Operations Center (LRNOC).

"When Heartbleed was released publicly, ISD and SSD requested the researchers at LRNOC to help determine its impact and ensure no critical information was leaked," says Tamara Yu of Lincoln Laboratory's Cyber Systems and Operations Group. "Using LRNOC network data, researchers were able to quickly use research tools to assist ISD and SSD in their investigation and assess potential impact. Fortunately, no sensitive information was leaked." Because the risk for a security breach and loss of sensitive information is high, researchers at LRNOC are working on ways to not only fight cyber threats like Heartbleed but also prevent them.

The LRNOC provides an environment in which cyber security researchers and analysts can use live network data to develop and test new techniques for defending Lincoln Laboratory's enterprise network.

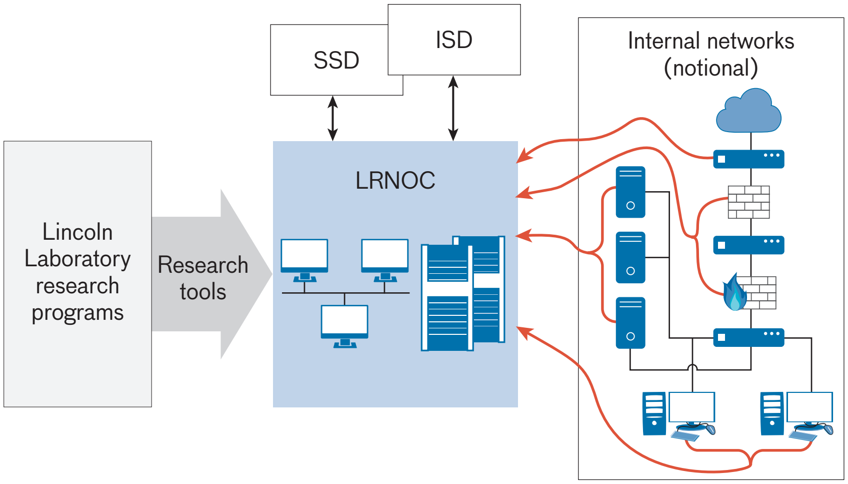
SSD
ISD
Internal networks (notional)
LRNOC
Lincoln Laboratory research programs
Research tools

The Lincoln Research Network Operations Center (LRNOC) is the hub of cyber traffic research at Lincoln Laboratory. It gathers information from internal networks (right) and shares important information (potential threats) with the Security Services Department (SSD) and Information Services Department (ISD). Research teams also use LRNOC network traffic and data to create cyber security tools.

For many research purposes, data produced by tools that utilize statistical models can be accurately labeled and can enable repeatable experiments. However, live data are "messy" and include many unusual and unexpected formats. Any algorithm developed to detect attacks needs to be tested out in a "real world" setting in order for developers to truly understand the algorithm's strengths and weaknesses.

"Researchers use LRNOC's high-quality, real-time traffic data to create security tools," says Jeffrey O'Connell of Lincoln Laboratory's Cyber System Assessments Group. "With real data, we're forced to push the limit of our tools, making them more capable, resilient, and adaptable. It is a very effective way to prepare for the next cyber attack."

Armed with standard and Laboratory-developed tools, analysts sift through troves of network data looking for suspicious anomalies that may warrant further investigation. LRNOC serves as an incubator and a proving ground for next-generation tools that government sponsors look to adopt to protect their own networks from cyber attacks.

To ensure that the live data are protected, the LRNOC is on a network that is separate from the main Laboratory network. Laboratory researchers are bound by a user agreement that is consistent with ethical practices and Laboratory security policies and that regulates data removal and research activities.

Because a vast amount of data, including system and application logs, network security appliance alerts, and raw traffic, are fed into LRNOC each day, researchers have created several tools to process the data. For example, one tool stores and aggregates each network packet entering and leaving the Laboratory network. The LRNOC infrastructure allows this tool to select packets on the basis of various features so analysts can define custom filters that capture and track communication with specific characteristics, such as heartbeat packets. Another tool, Scalable Cyber Analytic Processing Environment (SCAPE), works in a similar manner but creates feeds based on network information, such as logs and event data, rather than on raw network packets. SCAPE is a processing environment that monitors and correlates data feeds in real time to provide situational awareness about the state of the network.

Lincoln Laboratory researchers work in the Lincoln Research Network Operations Center (LRNOC) to identify and mitigate cyber threats.

These aggregator tools come in handy during attacks. Because LRNOC stores all network traffic data, researchers were able to carve out aggregated data from a period of time when the Heartbleed bug was active, gather pertinent data from that time period, and then investigate the selected data to determine abnormalities, such as a heartbeat packet. Researchers were then able to investigate the packets and work with SSD and ISD to determine if any sensitive information had been compromised.

Within LRNOC, research teams mainly work on two areas: monitoring Laboratory operations and developing or testing next-generation tools. Researchers collaborate with SSD and ISD to understand the challenges the network is up against, according to O'Connell. For example, if an LRNOC researcher finds suspicious traffic, he or she reports it to SSD and ISD staff, who potentially implement blocks on the network. If ISD finds a system on the network that might be compromised, they pass responsibility to SSD, who makes the final call on whether or not the system should be pulled from the network. If the system is pulled, SSD investigates it to find malicious components and may also isolate a malware sample and hand it back to the LRNOC for analysis. Working as a team, SSD, ISD, and LRNOC defend the Laboratory's network and support security protocols for sponsors.

"Our multifaceted, coordinated efforts in responding to not only daily incidents at the Laboratory but also several high-profile and potentially highly harmful exploits highlight the excellent collaboration among the various response teams, including LRNOC, ISD, and SSD," says Scott Mancini of SSD. "Their attention to detail and obvious dedication to protecting the Laboratory against daily threats are exceptional."

In addition to using LRNOC for defending the Laboratory's systems, research teams also use the network to develop security tools for government sponsors. Sponsors are eager to use programs created within LRNOC that can accurately protect their own networks from cyber attacks. For example, Laboratory researchers worked to enhance defender awareness of the specific types of scanning that adversaries conduct against the networks that they target. A Department of Defense sponsor requested and received a tool for fingerprinting five types of reconnaissance scans: Nmap, Strobe, Amap, Braa, Angry IP. This deliverable was developed and tested in the LRNOC by applying a detection algorithm on the real network data feeds in the LRNOC enclave.

Because LRNOC constantly evolves network security with each threat, its data are useful for developing current, accurate cyber defenses. "In just 30 minutes, LRNOC can see multiple reconnaissance activities probing the Laboratory's defenses," says O'Connell. "We keep seeing new threats and new ways to come out and be ahead of the curve."

CYBER TECHNOLOGY

Training the Cyber Defensive Line

A game-like competition is helping build experts in cyber "disaster response"

The number of attacks on computer networks is massive; for example, in 2013, the Pentagon reported getting 10 million attempted cyber intrusions a day.[1] These attacks are also growing in sophistication, primarily because cyber attackers are using combinations of techniques such as inserting malicious code (malware) or email phishing, and are adding complexity to the attack by involving multiple parties.[2] And, cyber intruders are breaching systems in just minutes.[2] Network operators, who are typically tasked with day-to-day maintenance of the computer systems, are hard-pressed, and often not trained, to address this flood of advanced, novel attacks.

In response to the proliferation and growing complexity of cyber threats, the U.S. Cyber Command (USCYBERCOM) over the last three years has created squads who will act as cyber "strike teams" in the field to protect the nation's networks. To help the Department of Defense (DoD) build such cyber protection teams, staff from Lincoln Laboratory's Cyber Security and Information Sciences Division, in collaboration with several other federally funded research and development centers (FFRDC) and university-affiliated research centers (UARC), developed and conducted a series of exercises designed to evaluate the capabilities of cyber defenders. Not exactly games, these exercises, collectively called Project C, pit a red team attacking the network against a blue team defending it. The red team plans an attack strategy, and the blue team develops countermeasures to thwart the attack. "The blue team needs to learn about the network and how best to defend it, locate any attacks, defeat them, and, finally, redefend the network," says Douglas Stetson, associate leader of the Laboratory's Cyber System Assessments Group.

Project C's primary goals are to assess and improve the performance of cyber teams and to advance technologies for cyber ranges (i.e., virtual environments for training cyber analysts and developing cyber defense tools). "Physical bodies are not the solution alone," according

[1] B. Fung, "How many cyberattacks hit the United States last year?" *National Journal*, 8 March 2013, available at http://www.nextgov.com/cybersecurity/2013/03/how-many-cyberattacks-hit-united-states-last-year/61775/.

[2] "Verizon 2015 Data Breach Investigation Report Finds Cyberthreats Are Increasing in Sophistication; Yet Many Cyberattacks Use Decades-Old Techniques," PRWire, 15 April 2015, available at http://www.prnewswire.com/news-releases/verizon-2015-data-breach-investigations-report-finds-cyberthreats-are-increasing-in-sophistication-yet-many-cyberattacks-use-decades-old-techniques-300066005.html.

to Lee Rossey, former leader of the Cyber System Assessments Group, who helped establish Project C. "You need the methodology and the tools." Rossey likens an effective cyber team to a football team: each player has a role, and they've all read the playbook and understand the team's offensive and defensive strategies. To develop a cyber playbook, Project C researchers investigated a number of questions: What makes one cyber team more successful than another? Why is one set of defenses more effective than another? How can we improve a team's capabilities? Answers to these questions will ultimately direct researchers to ways for improving subsequent rounds of training.

Project C sessions are conducted to help members of cyber protection teams be prepared to assist agencies undergoing serious cyber attack. How quickly a cyber team should be deployed to a site depends on two factors: the severity of the incident and the asset under attack. While an intrusion accomplished by a lone hacker most likely is handled expeditiously by an in-house computer security group, a coordinated assault by "well-armed" cyber adversaries requires highly trained, cyber security rapid responders. Because the DoD cannot constantly defend all data on its systems, the department has created a three-tiered Prioritized Defended Asset List for key missions and systems on a given network. Cyber teams are called in more quickly for higher-priority assets that are critical to the government's continued functioning than for lower-ranked systems. Rossey also notes that "just because a network goes down, it doesn't mean that you're under attack."

Blue team members analyze traffic and logs to determine whether an attack has occurred against their network. More than 60 personnel from active-duty, reserve, and guard units assigned to USCYBERCOM's cyber protection forces participated in the Project C exercises conducted at Lincoln Laboratory.

During the exercises, observers watch the cyber range activity and follow how blue team players respond to cyber incidents, gauging the effectiveness, creativity, and speed of the measures deployed to counter the red team attacks.

A Project C exercise is a multi-day event. At the start of each day's session, the staff members leading the exercise give participants a full briefing on the Project C format; the red team gets an additional briefing on their attack scenarios. The "battle" typically runs from 8:00 a.m. to 1:00 p.m. Before the red team begins its attacks, the blue team patches all known operating system (e.g., Windows 7) errors so that teams do not have to consider those errors when devising their stratagems. To make the exercise for the blue team as real as possible, the red team typically generates four to six different attacks derived from real-world threats detailed in Verizon's Data Breach Investigations Report (available on request from http://www.verizonenterprise.com/DBIR/). The blue team must ensure that their defensive actions preserve the integrity, confidentiality, and availability of all data. As the blue team works to mitigate threats, the red team is figuring out the blue team's strategies, and when one type of attack is defended, the red team tries another.

Noncombatant teams, called white teams, monitor the process, give advice if necessary, and score the results (number of successful and unsuccessful attacks, number of attacks identified by the blue team, mitigation results). Red team attack actions, blue team actions (even those not correlated to an attack), chat logs, network traffic, and other data are collected throughout the session, and a summary out-briefing is conducted in the afternoon. The five-hour multi-attack exercise, which in reality would be a situation spanning a few days, is fast-paced and stressful. In the out-briefing, blue team interactions resulting from the pressurized exercise (e.g., inadequate communication, heated discussions) are analyzed because the team dynamics are as important to the successful resolution of attacks as are the expertise and tools the team brings to the conflict.

One significant advantage Project C has over other serious gaming scenarios used in DoD cyber defense training is that it can simulate any of the various government networks and communication environments, such as ShoreNet for naval ships, warfighter communications in the field, power-grid-management networks, or command-and-control systems for the nation's missile defense systems. Project C allows cyber teams to work within a notional network-connected environment nearly identical to the real one they may be asked to defend. This virtual network environment, enabled by the Laboratory's LARIAT and KOALA tools,[3] includes all important elements, such as servers, users, and network activity.

Another advantage of Project C is that it is scalable and adaptable to different levels of attack severity and sophistication and to any type of network. The 2013 Defense Science Board (DSB) Task Force Report on Resilient Military Systems and the Advanced Cyber Threat classifies various types of cyber attackers. These cyber invaders range from individuals with commonplace equipment who simply employ malware developed by others to nation states that have the ability to execute cyber attacks that employ clever, new tactics. These classifications characterized in the DSB report also reflect the level of felonious intent of the perpetrators. Less malicious hackers break into networks for the challenge of doing so. Others invade systems seeking data that they can sell (e.g., the government's proprietary technical information). Critical threats to the United States are attackers targeting information that may give their nation states a military advantage. Project C's scalability and adaptability make it a valuable tool for improving the skills of respondents to all these types of attackers. It also provides an opportunity for participants to try out innovative cyber security technologies.

Experience gained by the researchers from Lincoln Laboratory, colleagues from FFRDCs and UARCs, and the Project C participants is being applied to the future strategy for training cyber protection teams. With guidance from the Laboratory's technical staff, the DoD held evaluation exercises last summer to compare the skills of three teams who had undergone a five-week Project C–type pilot training program in April and May to the skills of teams that had not engaged in such red team/blue team exercises. Analysis of the summer 2015 assessment sessions will be used to inform the direction of USCYBERCOM's cyber defense training. You might say Lincoln Laboratory is helping draft the playbook for the DoD's cyber protection defensive line.

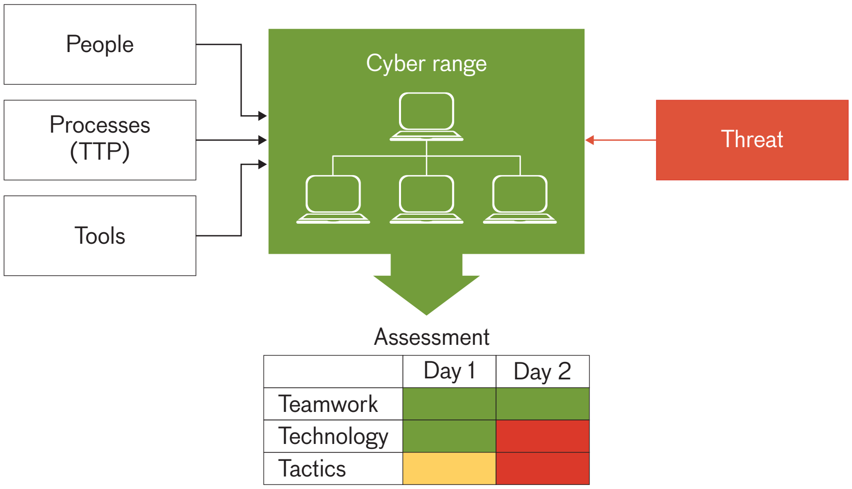

	Day 1	Day 2
Teamwork	🟩	🟩
Technology	🟩	🟥
Tactics	🟨	🟥

Project C sought to assess the people, processes, and tools of the cyber protection teams in a realistic environment with a realistic threat. The Project C format provides a day-by-day evaluation of how the team members interacted, how their technology worked, and how effective their tactics, techniques, and processes (TTP) were. In the "stoplight" evaluation chart, green indicates a highly successful performance; yellow, a satisfactory performance; and red, a breakdown or failure in performance.

[3] For more information about these tools, see the article "Advanced Tools for Cyber Ranges" in this issue of the *Lincoln Laboratory Journal*.

CYBER EDUCATION

Can a Game Teach Practical Cyber Security?

Lincoln Laboratory's Capture the Flag competition challenges college students to defend cyberspace

Thousands of teams around the world bearing names like the Plaid Parliament of Pwning, ghettohackers, or Shellphish compete each year in contests to infiltrate opponents' computer services while defending their own systems from cyber attacks. In these competitions, teams playing on networks of virtual machines earn points by breaching other teams' services to capture information that the contest administrators hide within the programming. Called a flag, the information is typically a lengthy string of random, hard-to-guess code. The first of these Capture the Flag (CTF) events was held at the 1996 DEF CON,[1] now one of the world's largest hacker conventions. Since then, CTF competitions have sprouted up in dozens of countries, often organized by university departments and technology companies seeking to improve students' and employees' skills in devising techniques and tools to ensure network security.

To investigate what educational benefits CTF competitions provide to participants and whether CTF play leads to the development of innovative strategies applicable to real-world cyber defense, researchers from Lincoln Laboratory developed a CTF event for college students.

Early on, the core team of technical staff members from the Cyber Security and Information Sciences Division—Joseph Werther, Michael Zhivich, Timothy Leek, and Andrew Davis—decided that their CTF competition would be structured as an attack-defend format. Some CTFs focus on either offensive or defensive actions. In contests in which only attacks earn points, competitors focus on techniques to breach security and forego protecting their systems. In defense-only matches, players employ functions to keep their services running despite assaults the CTF administrators have embedded in their virtual systems; these players do not face the pressure of devising defenses while also crafting attacks against others and foiling a steady barrage of onslaughts from other teams.

The dual format has significant advantages. Requiring success at both attack (capturing flags) and defense (securing services) to score points compels players to interact continually with opponents and their own systems. The attack-defend approach creates a dynamic, realistic environment in which defensive techniques must be developed under pressure and time constraints. To further simulate the demanding pace of real attack mitigations, the Lincoln Laboratory CTF also allowed teams to score points only if their services were operational. "Requiring that team services be up in order to score points either offensively or defensively provides a very strong incentive for every team to risk running services as soon as and as much as possible," says Davis. Finally, the attack-defend format is more challenging and, the Laboratory's organizers believe, more fun than a single-focus one.

Capture the Flag can provide a sandbox in which prototype technologies, both defensive and offensive, can be tested and evaluated.

Lincoln Laboratory's first CTF competition was held at MIT on 2 and 3 April 2011 and was open to Boston-area college students. Forty-five registered players from six schools showed up to spend 18 hours of their weekend attacking and defending a web application server. The virtual system was modeled as a Linux operating system running an Apache server and employing a MySQL database and Hypertext Preprocessor scripting language. So that students could experiment on an application whose code would be accessible, the open-source WordPress content management software for creating websites and blogs was chosen as the target. In addition, because the flexible architecture of WordPress allows plugins, the organizers could peri-

[1] DEF CON is an annual event that attracts not only computer hackers but also researchers from academia, industry, and government agencies.

Lincoln Laboratory's third Capture the Flag (CTF) competition drew 165 college students to MIT for a 48-hour marathon of attacking and defending Android services. Students came from MIT, Boston University, UMass-Boston, Northeastern, Brandeis, Wellesley, Worcester Polytechnic Institute, Rensselaer Polytechnic Institute, New York University Polytechnic School of Engineering, and Dartmouth College. Many of the participants and event organizers posed for a post-event photograph. The official MIT Lincoln Laboratory CTF flag is held by a participant in the back, while in front a member of one of the top three teams is holding a replica of a check for the team's prize money.

odically add new vulnerabilities for the students to mitigate.

Many universities and businesses that run CTF competitions do so online, with registered teams downloading the necessary software and instructions so that they can tackle the challenges made available on contest days. Lincoln Laboratory's CTF organizers chose to hold an onsite event. "The competition is so much more exciting live. The energy in the room is invigorating," says Leek. "There is a lot more interaction between team members."

The Laboratory's CTF development team, which in 2011 also included Nickolai Zeldovich from MIT's Computer Science and Artificial Intelligence Laboratory, found that the algorithm used for scoring the play is vital to the dynamic nature of the competition. Designing a scheme that rewards players for achieving the defensive goals of maintaining data confidentiality, availability, and integrity and that also awards points for offensive successes is a balancing act. After the first day of the 2011 event, the CTF organizers noticed that equally weighting offensive and defensive results encouraged teams to shut down their servers when they were planning their offense, thereby denying attackers access to the servers and increasing their own scores for maintaining data confidentiality and integrity. A revised weighting method to reward teams whose services were accessible created the motivation for them to focus efforts toward more defensive actions.

Patrick Hulin, a member of MIT's winning team in the 2012 CTF competition and now on staff in the Cyber System Assessments Group, credits his team's success to their emphasis on defense. "We narrowly focused on the essential tasks we had to complete in order to succeed under the scoring algorithm. It was more important to keep your services operating than to attack teams other than the leaders, so we wasted very little time working on the fun but not necessarily relevant offensive moves that took a lot of work for very little actual gain in the standings."

According to the surveys that participants filled out online after the competition, the 2011 CTF event was a success. The students appreciated the challenges presented by the game; most of them thought they had improved their skills; and many reported an increased interest in a career in cyber security (though these students did note a previous interest in such a career).

Era Vuksani, now a researcher in the Laboratory's Cyber Systems and Operations Group and formerly a member of Wellesley College's 2011 CTF team, says, "I learned a lot from being in that environment where you had to be very proactive in defending yourself from adversaries as well as be ready to wipe your machine and start over as need be. You had to be adaptable at a moment's notice."

The organizers applied lessons learned from the 2011 event to full-scale competitions held in 2012 and 2013 at MIT and a few practice sessions, or mini-CTFs, offered in 2014 at the MIT Lincoln Laboratory Beaver Works center near the MIT campus in Cambridge, Massachusetts.

In 2012, the students were tasked with sustaining the security of an enterprise web server, and the 2013 competition charged participants with supporting several apps for an Android platform and the corresponding backend services on Linux virtual machines. As word of the Laboratory's CTF spread, participation grew: in 2012, 62 students from six schools participated, and 165 players from 10 regional universities took on the 2013 Android challenge. The events also grew longer; in 2013, the students were in competition for 48 hours straight, eating while working and taking turns catching naps.

The 2013 CTF event also introduced a new element—evaluating an outside organization's technology. Employees from Raytheon BBN Technologies tested out their Advanced Adaptive Application (A3) Environment prototype by trying to defend the CTF's App Store against attacks from the competition teams. After a flaw identified in the A3 Environment software on the first day was remedied, the prototype was able to secure the App Store. The Laboratory's CTF organizers concluded that "CTFs can provide a sandbox in which prototype technologies, both defensive and offensive, can be tested and evaluated," but with the caveat that the technology developers need to be on hand to fix problems. From their participation in CTF, the A3 Environment researchers gained improvements to their code, validation of their defensive policies, and a corpus of attacks they could use in building later iterations of their technology.

Success at Capture the Flag depends a great deal on the teamwork exhibited by the competitors as they plan attack and defend strategies.

Creating a challenging, well-functioning CTF competition requires a significant investment in software development. The Laboratory researchers devoted a great deal of effort to conceiving the scenarios, cyber vulnerabilities, and scoring strategies, and then to building the software and interfaces that enabled these.

Lincoln Laboratory's CTF experience has been successful on a number of fronts. First, the researchers who worked on developing the competitions acquired some answers to their initial questions.

- Do CTF events help educate students in cyber security? The answer is a qualified yes. Students who are already inclined to engage in cyber defense, who have perhaps tried online CTF games, will strengthen their competencies in computer security. Zhivich likens the learning to that of athletes sharpening skills through practice: "Good ballplayers get better as they play more." It is harder to say how much CTF participation teaches students who do not have prior experience in computer security; in the Laboratory's CTF games, the less experienced students did not amass high scores but felt they took away new awareness of the cyber field. Although precompetition tutorials on cyber defense tactics, common cyber tools, and web applications were appreciated by the students who attended these sessions, the researchers cannot say definitively that the tutorials resulted in helping to “level the playing field” for the inexperienced CTF teams. “We believe CTF works as a kind of group self-guided, project-based instruction,” says the CTF research team’s 2014 paper chronicling their findings from hosting the events.[2]

- Can CTF events generate new ideas for real-world cyber defense tactics and tools? Again, the answer is not definitive. The researchers monitoring the competitions directed their attention to keeping the game running smoothly. They state in their 2014 paper that they would like to understand better what teams do to win CTF competitions and that “it may be possible to discover new advanced techniques for attack and defense by providing college students a safe [i.e., not incurring legal repercussions for hacking real networks] place to play.” However, the experience with the A3 Environment shows that a CTF event can be used as a test bed for new technology.

On another front, the collegiate CTF competition introduced the Laboratory and many talented young people to each other. Indeed, five staff members hired into the Cyber Security and Information Sciences Division had participated in one of the Laboratory’s CTF events. As the world becomes more dependent on computer networks to conduct all its activities, nations and private businesses are eager to find the best-qualified people to secure their network services. Hosting a CTF event could be one avenue for organizations to meet those people.

Finally, the CTF events resulted in personal successes for participants and organizers. In the postgame surveys, students cited not only increased understanding of cyber security but also improved teaming skills as takeaways from the competitions. Resolving the technical demands of crafting the scenarios and developing the automated scoring application were interesting projects for the Laboratory staff members. Furthermore, both students and the CTF staff had fun.

The researchers who conducted the CTF events under funding from the National Security Agency have completed their investigation into CTF’s role in enhancing education and development in cyber security techniques. Their published experiences can serve as a road map for future Lincoln Laboratory CTF events and for other organizations considering the establishment of CTF competitions.[2,3] In addition, the research team is looking to make their infrastructure available as an open-source codebase.

[2] A. Davis, T. Leek, M. Zhivich, K. Gwinnup, and W. Leonard, “The Fun and Future of CTF,” 2014 USENIX Summit on Gaming, Games, and Gamification in Security Education, available at https://www.usenix.org/node/184963.

[3] J. Werther, M. Zhivich, T. Leek, and N. Zeldovich, “Experiences in Cyber Security Education: The MIT Lincoln Laboratory Capture-the-Flag Exercise,” *Proceedings of the 4th Conference on Cyber Security Experimentation and Test*, 2011, available at http://www.ll.mit.edu/mission/cybersec/publications/publication-files/full_papers/2011_08_08_Werther_CSET_FP.pdf.

EDUCATIONAL OUTREACH

Recruiting the Next Generation of Cyber Security Specialists

Two Lincoln Laboratory outreach activities seek to steer high-school students toward careers in cyber security

Today’s cyber security specialists are too few in number and lack the skills needed to defend networks supporting the nation’s government agencies, financial institutions, power grids, and transportation systems. As cyber attacks escalate in frequency and sophistication, this shortage of adequately trained personnel will become even more acute, particularly within the U.S. government.

Lincoln Laboratory is trying to address one of the roots of the shortage in cyber security professionals: the lack of cyber security education in school curricula. Two programs designed for high-school students—CyberPatriot and LLCipher—have been a part of the Laboratory’s efforts to help fill this gap. By engaging these precollege students in activities that highlight the appeal of cyber security work, the Laboratory hopes they will be motivated to pursue undergraduate studies and eventually careers in the field.

Since 2011, Lincoln Laboratory has sponsored teams of high-school students participating in the CyberPatriot National Youth Cyber Defense Competition, a program initiated in 2009 by the U.S. Air Force Association to spark young students' interest in cyber security or other science, technology, engineering, and mathematics fields. A network defense competition, CyberPatriot challenges students to find vulnerabilities (e.g., malware, weak passwords, unnecessary services) within a set of virtual images that represent Windows or Linux operating systems while maintaining critical network services, such as email. Each image contains anywhere from 10 to 20 flaws; the teams that discover the most flaws within a six-hour time limit advance to subsequent rounds. Although the format of the rounds and the scoring system have evolved over the years to support the growing number of registered teams (eight to start and more than 2000 in the 2014–2015 season), the basic advancement process has remained the same, with teams competing at the state, regional, and national levels.

Helping a student prepare for the CyberPatriot competition, Robert Cunningham, leader of the Secure Resilient Systems and Technology Group, explains how to configure a Windows 7 system to ensure strong passwords.

In its first two years of participation in the CyberPatriot program, the Laboratory sponsored a single team; for the past two years, three teams have been sponsored. Teams typically consist of five to six students, many of whom compete in multiple CyberPatriot seasons. Veteran members are often paired with rookies, according to Chiamaka Agbasi-Porter of the Communications and Community Outreach Office, who coaches the teams and recruits Laboratory volunteers to serve as mentors. From September through March, the students and mentors meet once a week for two hours at the MIT Lincoln Laboratory Beaver Works facility near the MIT campus in Cambridge, Massachusetts. During these weekday sessions, students learn and practice the computer and teamwork skills they need to

For two years in a row, the first Laboratory-mentored CyberPatriot team, DoNut Hack Us, was one of 12 finalists selected to compete in the national championship held in Washington, D.C. More than 1000 teams entered the competition in each of those years. Seen above left are three of the five team members racing against the clock to detect vulnerabilities in the areas of policy, patch, configuration, and third-party management during the 2013 finals. After graduating high school, three Cyber-Patriot alumni from the team spent their summer interning in the Cyber Systems and Technology Group (above right). All three have chosen to pursue computer science in their undergraduate studies.

compete in CyberPatriot. Throughout the season, technical staff from the Laboratory give presentations on relevant topics, including cryptography, networking, Windows internals, and Linux security. On some weekends early in the season, all CyberPatriot teams participate in online qualifying rounds from their home base, finding vulnerabilities within virtual machine images downloaded onto laptops. These rounds could also include a Cisco Networking quiz or a Cisco Packet Tracer (a network simulation program for students to experiment with network behavior) challenge—one of the mechanisms through which teams can gain points beyond those acquired by fixing vulnerabilities. Points are also awarded for answering forensics questions about the steps taken to remediate the vulnerabilities. Teams lose points if they take any actions that make a system less secure (e.g., reintroducing a previously fixed vulnerability). Scores are automatically recorded by a centralized scoring system.

Jorge Coll, a technical staff member in the Secure Resilient Systems and Technology Group, is one of the CyberPatriot mentors. A previous Microsoft employee, Coll focuses on the Windows operating system, helping students identify misconfigured settings; configure their machines with policies, such as those for password restrictions; and ensure software patches are up to date. One of Coll's major contributions has been in the area of competition strategy: How can students maximize their time to gain as many points as possible?

CyberPatriot team members collaborate on finding malware and locking down a Windows virtual machine during one of the online weekend competitions.

"The two largest time sinks students struggle with during the competition are discovering what is wrong with any given system and applying security best practices to lock down their machines," explains Coll. To reduce the time spent on such tasks, Coll introduced the students to various automation tools, including Windows PowerShell (a command-line interface and scripting language), security policy templates, and techniques for recognizing configuration drift (i.e., changes to a system's hardware or software environments). "For example, with PowerShell, students can automatically query login records to see when the last time a particular user accessed his or her account, instead of having to manually sift through these records," says Coll.

The track record of the Laboratory teams has been impressive. For the 2011–2012 and 2012–2013 seasons, the one Laboratory-sponsored team advanced to the national competition in Washington, D.C., where they placed 7th among 11 finalist teams both times. At the end of the 2013 season, most of the team members graduated from high school. New team members were recruited for the following season (2013–2014), resulting in three teams, all of whom came very close to qualifying for the national finals. In 2014–2015, all three teams competed at the highest level in the statewide competition, and one went on to complete its season at the Northeast regional competition.

While CyberPatriot is at its core a competition, with scholarship money given to the top three teams, it is more than a game. "CyberPatriot gives students an early window into cyber security, a field that most students do not encounter until college," says Sophia Yakoubov, one of the mentors and a technical staff member in the Secure Resilient Systems and Technology Group. Yakoubov taught the team members about classical cryptography and cryptanalysis. "I showed them how, just by looking at an encrypted message or file, they can figure out which encryption scheme was used and then how to apply various techniques to crack it," she explains.

With the help of colleagues Emily Shen and David Wilson, Yakoubov served as the lead instructor for a new cyber security–focused outreach program, LLCipher, in summer 2015. Held at Beaver Works, this one-week cryptography workshop provides an introduction to modern cryptography—a math-based, theoretical approach to securing data. Lessons in abstract algebra, number theory, and complexity theory provide students with the foundational knowledge needed to understand theoretical cryptography. Students then construct provably secure encryption and digital signature schemes. On the last day, the students learn about two techniques that enable multiple entities to exchange data without disclosing to one another more data than necessary to perform a particular function: zero-knowledge proofs (proving a statement is true without revealing any information beyond the truth of the statement) and multiparty computation (computing a function over multiple parties' inputs while keeping the inputs private).

Workshop designer and lead instructor Sophia Yakoubov (standing) makes her way through the classroom as the students work on a physical secret communication challenge. Teams of three, an all-girls one of which is pictured above, assumed the roles of Alice, Bob, and Eve—common archetypes in the cryptography literature. The premise of the challenge is as follows: Alice is trying to securely communicate a secret to Bob; Eve is trying to eavesdrop. Alice and Bob are both given individual locks to affix to a writing notebook, which contains the secret, and corresponding keys. To solve the challenge, teams must figure out how the lock-key systems can be applied to the notebook so that Bob can read the secret but Eve cannot.

The idea for LLCipher came from Bradley Orchard, a technical staff member in the Advanced Sensor Systems and Test Beds Group and a part-time teacher at the Russian School of Mathematics in Lexington, Massachusetts. While teaching at this enrichment school for the past four years, Orchard encountered several remarkably bright students who were just entering high school yet were ready to take calculus—a course typically reserved for the senior-year curriculum. "These students are often two to three years ahead of their classmates in regular school," explains Orchard. Recognizing these students' need for learning opportunities beyond those offered in schools, Orchard set to work to design an introductory summer course for advanced students. With his academic training as a mathematician, he naturally thought theoretical cryptography would be the ideal subject matter for the course: "Theoretical cryptography combines beautiful mathematics with powerful, useful, and fun techniques and, most importantly, aspects of cryptography are very accessible to advanced students." Orchard proposed his idea to John Wilkinson, leader of the Cyber System Assessments Group, who reached out to cryptography experts within the Laboratory's Cyber Security and Information Sciences Division to help design and teach the course. Knowing how much she enjoyed teaching the CyberPatriot students about cryptography, Yakoubov was eager to get involved.

According to Yakoubov, the pilot program was a huge success: "The class was very interactive, with students asking questions that demonstrated they understood the material. The feedback we received from the students indicates they really enjoyed LLCipher and learned a lot." When asked about the most interesting thing he learned, one student replied, "Zero-knowledge proofs, as they seemed impossible. The idea of proving knowledge without sharing it is fascinating."

As Orchard had hoped, the subject matter of the course piqued student interest. "My favorite thing about this program was learning about cryptography, as it was different from traditional math and required a different type of thinking," another student commented. Among students, the most common suggestion was to extend the length of the program. On the basis of this feedback, the instructors will increase the sessions from two to eight hours per day next year.

CyberPatriot and LLCipher are two of the Laboratory's outreach programs dedicated to cyber security education. At the college level, a Capture the Flag competition based on an attack-defend approach seeks to equip students with the skills needed for real-world network security (see Lab Note titled "Can a Game Teach Practical Cyber Security?" for more information). The Laboratory's Science on Saturday demonstrations have made topics, such as computer authentication, accessible to the younger K–12 crowd.

By reaching out to students at different levels of their education, the Laboratory hopes to, at some point, incite their interest in cyber security—a field that will only expand in the coming years. "Every day, attackers break into computers holding sensitive information. The need to secure these data is great, but there is a shortage of people with the right knowledge and experience to meet this need. Currently, the Department of Defense is seeking to hire 6000 cyber security personnel but so far has only hired half of that," explains Robert Cunningham, one of the CyberPatriot mentors and leader of the Secure Resilient Systems and Technology Group. "Programs like CyberPatriot and LLCipher help grow the base of those who are knowledgeable about computer security while also teaching students about leadership and critical thinking."

Students in the LLCipher program gathered for class in the morning at Beaver Works. Here, Yakoubov provides a lesson on the ElGamal algorithm for public key encryption.