Fall 2019

CYBERSECURITY

The BRIDGE

LINKING ENGINEERING AND SOCIETY

NATIONAL ACADEMY OF ENGINEERING

The mission of the National Academy of Engineering is to advance the well-being of the nation by promoting a vibrant engineering profession and by marshalling the expertise and insights of eminent engineers to provide independent advice to the federal government on matters involving engineering and technology.

The
BRIDGE

NATIONAL ACADEMY OF ENGINEERING

The Bridge (ISSN 0737-6278) is published quarterly by the National Academy of Engineering, 2101 Constitution Avenue NW, Washington, DC 20418. Periodicals postage paid at Washington, DC.

Vol. 49, No. 3, Fall 2019

Postmaster: Send address changes to *The Bridge*, 2101 Constitution Avenue NW, Washington, DC 20418.

Changes of address or requests to unsubscribe should be sent to PGibbs@nae.edu.

Mission Statement of *The Bridge*

The Bridge publishes articles on engineering research, education, and practice; science and technology policy; and the interface between engineering and technology and society. The intent is to stimulate debate and dialogue both among members of the National Academy of Engineering (NAE) and in the broader community of policymakers, educators, business leaders, and other interested individuals. *The Bridge* relies on its editor in chief, NAE members, and staff to identify potential issue topics and guest editors. Invited guest editors, who have expertise in a given issue's theme, are asked to select authors and topics, and independent experts are enlisted to assess articles for publication. The quarterly has a distribution of about 7,000, including NAE members, members of Congress, libraries, universities, and interested individuals all over the country and the world. Issues are freely accessible at *www.nae.edu/Publications/Bridge.aspx*.

A complete copy of *The Bridge* is available in PDF format at *www.nae.edu/TheBridge*. Some of the articles in this issue are also available as HTML documents and may contain links to related sources of information, multimedia files, or other content.

The BRIDGE
LINKING ENGINEERING AND SOCIETY
Volume 49, Number 3 • Fall 2019

Editors' Note

Features

(continued on next page)

The National Academies of
SCIENCES • ENGINEERING • MEDICINE

The **National Academy of Sciences** was established in 1863 by an Act of Congress, signed by President Lincoln, as a private, nongovernmental institution to advise the nation on issues related to science and technology. Members are elected by their peers for outstanding contributions to research. Dr. Marcia McNutt is president.

The **National Academy of Engineering** was established in 1964 under the charter of the National Academy of Sciences to bring the practices of engineering to advising the nation. Members are elected by their peers for extraordinary contributions to engineering. Dr. John L. Anderson is president.

The **National Academy of Medicine** (formerly the Institute of Medicine) was established in 1970 under the charter of the National Academy of Sciences to advise the nation on medical and health issues. Members are elected by their peers for distinguished contributions to medicine and health. Dr. Victor J. Dzau is president.

The three Academies work together as the **National Academies of Sciences, Engineering, and Medicine** to provide independent, objective analysis and advice to the nation and conduct other activities to solve complex problems and inform public policy decisions. The Academies also encourage education and research, recognize outstanding contributions to knowledge, and increase public understanding in matters of science, engineering, and medicine.

Learn more about the National Academies of Sciences, Engineering, and Medicine at **www.nationalacademies.org**.

Editors' Note

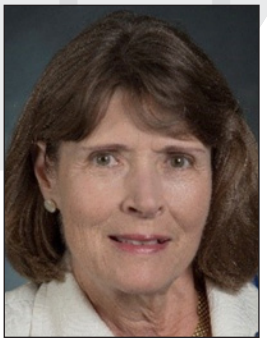

Ruth A. David (NAE) is retired president and chief executive officer of ANSER.

Robert F. Sproull (NAE) is retired vice president and director of Oracle Labs and now an adjunct professor of computer science at the University of Massachusetts at Amherst.

Cybersecurity: A Growing Challenge for Engineers and Operators

In today's increasingly connected and interdependent world, cybersecurity is an issue that touches virtually every individual, organization, and institutional entity—governmental and nongovernmental alike. According to *ITSP Magazine*, "There are three types of people in the world: those who have been attacked, those who will be attacked, and those who are being attacked right now and just don't know it yet."[1] The same could be said of organizations or institutions. All have a role to play—both in safeguarding personal devices and in contributing to the protection of the systems to which these devices are connected.

As computers shrink in size, computational capabilities—hardware and software—are increasingly embedded in everyday objects, from personal devices such as smart phones and watches to personal vehicles and even dwellings (e.g., smart home systems). Similarly, complex infrastructures on which people depend for transportation, energy, communications, food production, water distribution, and healthcare delivery are increasingly computerized, as are the manufacturing and design processes that will create the next generation of systems and infrastructures.

To be sure, the engineers who design and develop these systems and infrastructures play a vital role in addressing security concerns, but they cannot anticipate every conceivable threat that may manifest during operational use of their product—particularly since new attack surfaces are introduced when individual systems are connected to networks.

Cybersecurity is defined by the National Institute of Standards and Technology (NIST 2011, p. B-3) as "the ability to protect or defend the use of cyberspace from cyber attacks," and cyberspace is defined as "a global domain within the information environment consisting of the interdependent network of information systems infrastructures including the Internet, telecommunications networks, computer systems, and embedded processors and controllers." Thus, cybersecurity focuses on thwarting attack vectors that exploit network connections, but those connections typically occur in individual systems put into use by individual, organizational, or institutional operators who therefore play an equally vital role in addressing security concerns.

Origins and Evolution of Cybersecurity

Cybersecurity traces its roots to 1971, when Bob Thomas unleashed what is considered the first computer worm (Townsend 2019). The worm was not malicious but was capable of jumping from one computer to another, a behavior not previously observed. Years later, in 1989, the first recognized denial of service (DoS) attack was attributed to Robert Morris (Townsend 2019). Once again, the intent was not malicious—Morris intended to highlight security flaws—but this time the damages were real. In the three decades since, attacks have grown in both sophistication and variety, and motivations now include monetary gain, identity theft, espionage, and operational disruption. Cybersecurity has become an arms race—and unfortunately, the odds favor the attackers.

[1] https://www.itspmagazine.com/cybersecurity-quotes

The market for cybersecurity products and services is large and continues to grow rapidly, primarily because of increasing cybercrime activities. Cybersecurity Ventures reports that the global cybersecurity market was worth only $3.5 billion in 2004 and was expected to exceed $120 billion by 2017, and predicts that global cybersecurity spending will exceed $1 trillion cumulatively during the period from 2017 to 2021 (Morgan 2019a).

In spite of substantial growth in cybersecurity spending, damages attributable to security breaches are also growing. In its *Annual Cybercrime Report*, Cybersecurity Ventures identified cybercrime as one of the biggest challenges confronting humanity and predicted that related costs will double—from $3 trillion in 2015 to $6 trillion annually by 2021 (Morgan 2019b). Notably, the growing cost estimates are based not only on significant projected growth in adversarial attacks but also on increasing attack opportunities stemming from a cyberattack surface projected to be an order of magnitude greater by 2021 than it is today (Morgan 2019b).

Cybersecurity: Why Is It So Hard?

Attackers have the edge in the cybersecurity arms race. Reasons for this include the following:

- Attackers need to find only a single exploitable vulnerability, while defenders must identify and eliminate or mitigate all vulnerabilities.
- Because many copies of products or software are deployed, a single vulnerability can be widely exploited.
- A stealthy attack may not be detected while it's underway, leaving little opportunity for real-time defense.
- Geographic borders do not impede traffic in cyberspace, enabling stand-off attacks that often are hard to attribute.
- Cybersecurity is not just a technical problem—system operators and users also provide attack surfaces.
- System builders often sacrifice security measures for operational convenience or market share motivations.
- International norms, cybersecurity law, policy, and practice are not yet mature.

These and related factors contribute to the complexities inherent to cybersecurity.

In This Issue

This issue of *The Bridge* cannot hope to cover the breadth or depth of important cybersecurity issues. Rather, we have selected papers intended to provide a basis for understanding the evolving nature of cybersecurity threats, for learning from past incidents and best practices, and for anticipating the engineering challenges in an increasingly connected world.

In the first article, **David Clark** takes an internet-centric view of cybersecurity and parses the problem into four parts to provide insights about measures to improve the situation. He observes that "perfect security is not possible." Viable mitigation measures are context-specific.

Nicole Beebe and **Frederick Chang** make the case for expanding the traditional definition of the insider threat to include both unwitting human agents and technology that act as trusted agents. They note that the complexity of this issue cannot be effectively addressed absent a true systems-engineered solution.

Josephine Wolff uses perpetrator motivations—financial theft, espionage, public humiliation—as a framework for examining past cybersecurity incidents to distill recurring themes and lessons. She argues that improving cybersecurity best practices will require consideration of the entire security ecosystem, which extends well beyond a single entity under attack.

Christian Hamer offers an experience-based perspective from Harvard University in which he addresses two types of cyberattacks: unauthorized access and business disruption. He lays out the key elements of a risk-based program and describes *solutions that work* in a large and very diverse organization, while acknowledging the need to anticipate ways in which future threats may evolve.

Fred Schneider and Lyn Millett use case studies to examine the role of nontechnical issues in what appear to be technical matters, exploring the policy dimensions of cybersecurity engineering. Their paper draws on discussions hosted by the Forum on Cyber Resilience, a roundtable of the National Academies to facilitate the exchange of ideas among scientists, practitioners, and policymakers concerned with the resilience of the nation's computing and communications systems.

John Stankovic and Jack Davidson anticipate the cybersecurity challenges inherent in the rapidly growing Internet of Things, which presents new attack surfaces as well as new types of consequences for successful

breaches. They note the dangers that may arise if speed to market and low-cost competition drive design trade-offs to sacrifice security considerations during product development.

Ronnie Chowdhury, Mhafuzul Islam, and Zadid Khan address the rapid evolution of transportation systems based on advances in connected and automated vehicle technologies. Their article provides a deeper dive into one type of cyberphysical system that will connect to the growing Internet of Things.

In the final invited paper, Tom Longstaff and Noelle Allon provide a perspective on *what every engineer should know about cybersecurity*, summarizing some key points from the other papers and putting them into the systems engineering context. While acknowledging that every engineer will not become a cybersecurity expert, given the likelihood that virtually every newly engineered system will contain some computational elements and will be connected to other systems in its environment in some way, every engineer should have at least a basic understanding of cybersecurity principles.

Looking to the Future

This issue offers a limited window into current and expanding cybersecurity challenges confronting designers, developers, operators, and users of connected systems. The articles illustrate the need to address today's vulnerabilities while anticipating increasingly sophisticated attackers. They also acknowledge that the increasing propensity to interconnect systems yields new benefits while simultaneously increasing both the quantity and variety of attack surfaces. This environment continues to favor the attackers.

A key question is, therefore: *What is needed to change the cost-benefit equation for the attacker?* Answers will undoubtedly have policy implications and impact design choices. They could also result in different cybersecurity investment profiles for operators and may degrade user convenience, particularly for legacy systems. Options could include offensive as well as defensive measures.

While it is important to continue to learn from past incidents and to effectively implement relevant best practices, these measures are insufficient in an era of rapid growth in both the quantity and sophistication of attackers together with rapid expansion in both the quantity and variety of attack surfaces. Robust exploration of proactive ways to change the cost-benefit equation for attackers is needed.

Acknowledgments

We offer many thanks to the authors of these papers, who worked hard to deliver important messages to a broad audience in short articles. In addition, *The Bridge* asks outside readers to comment on papers and suggest improvements; we are grateful to Chris Bronk, Azim Eskandarian, Kevin Fu, **Paul Kocher**, **Steve Lipner**, Keith Miller, Bill Scherlis, David Sherry, and Mary Ellen Zurko, who generously dedicated their time to evaluate the papers in this issue.

References

Morgan S. 2019a. Global cybersecurity spending predicted to exceed $1 trillion from 2017 to 2021. Cybersecurity Ventures' 2019 Cybersecurity Market Report, Jun 10.

Morgan S. 2019b. Cybercrime damages $6 trillion by 2021. 2019 Official Annual Cybercrime Report. Northport NY: Cybersecurity Ventures and Toronto: Herjavec Group.

NIST [National Institute of Standards and Technology]. 2011. Managing Information Security Risk: Organization, Mission, and Information System View. NIST Special Publication 800-39. Gaithersburg MD.

Townsend C. 2019. A brief and incomplete history of cybersecurity. United States Cybersecurity Magazine, Jan 18.

An internet-centric view of cybersecurity challenges sheds light on ways to address them.

A Framework to Understand Cybersecurity

David Clark (NAE) is a senior research scientist in the Computer Science and Artificial Intelligence Lab at the Massachusetts Institute of Technology.

David D. Clark

Better cybersecurity is an admirable aspiration. But aspirations, as such, are not actionable. Calling for better cybersecurity does not give any hint of what actions should be taken, and by whom, to improve the situation.

The goal of this paper is to break the challenge of improved cybersecurity into parts that are potentially actionable, provide a roadmap to better security, and illustrate why the challenge of better security is so vexing. Since the internet is at the heart of cyberspace—without the internet or something similar there would just be a bunch of disconnected computers—I take an internet-centric view of security in this paper. There are, of course, other opinions about how to structure the thinking about security, but I find that this framework can provide good insights about which actors have to do what to improve the situation.

I consider four types of internet security challenges: (a) third-party attacks on mutually trusting users trying to communicate, (b) attacks by one user on another, (c) attacks on the network itself, and (d) denial of service attacks.

Third-Party Attacks on Mutually Trusting Users

In this case, two (or more) users are attempting to have a mutually desired communication and a hostile third party attacks it. This situation affects *information security*, which is traditionally described as having three subcomponents: confidentiality, integrity, and availability.

- The goal of confidentiality is to ensure that a hostile third party cannot observe what is being communicated. The motivation of that third party might range from demographic profiling for targeted advertising to content-based censorship to national security surveillance. A strong mechanism for confidentiality will ensure that these goals are thwarted.
- The goal of integrity is that a third party cannot modify what is being sent by the communicants.
- The goal of availability is that communication can still succeed despite attempts by a third party to disrupt it.

With respect to confidentiality and integrity, the current approach in the internet is end-to-end encryption. If a transmission is encrypted, a third party cannot observe it, and while encryption cannot prevent modification, the modification will always be detected, so the outcome may be a failure of availability but the attacker cannot actually modify what is being sent.

As tools for encryption are embedded in the most common applications (such as the Web) the fraction of encrypted traffic is increasing rapidly. This outcome is the result of efforts over many years by researchers, standards setters, and implementers.

The Limits of Encryption

Encryption is not magic. Just using it does not make all the problems go away. Its mathematics may be beautiful, but the actual encryption is embedded in a larger system that has the task of managing the encryption keys, the information that must be shared among the parties so that the encryption can work. The weak part of most encryption schemes is not the algorithm but how the keys are managed. If an attacker can steal the keys (which may just be stored on a user's personal computer—not the most secure device on the internet), or otherwise manipulate the key management system, they can completely thwart the goals of confidentiality and integrity.

Availability also cannot be improved using encryption. If an attacker is in a position to block the traffic between communicants, the only benefit of encryption is to blunt the attacker's instrument—if the attacker cannot see what is being sent, the blocking cannot be discriminating: it becomes an all-or-nothing attack. There are some contexts in which a third party (commonly a state actor or an internet service provider, or ISP, acting as an agent of the state) has blocked all encrypted traffic to try to force the communicants to send "in the clear," so that the third party can see what is being sent.

Ignored Warnings

For most users, fears about confidentiality are real—concerns about privacy include the fear that ISPs are observing what users send and selling that information for targeted advertising. However, loss of availability is also a problem, often arising from the key management system put in place for confidentiality and integrity.

Encryption keys are usually distributed in what is called a certificate, a signed attestation that binds the name of an entity to a particular key. For security reasons, certificates are valid for a limited time and need to be updated periodically by the owner of the key. If the owner neglects to do so (or makes other errors), software (e.g., the user's browser) will present warning messages that are both dire and sometimes obscure, saying that the user should not proceed because it is possible that the communication is under attack.

Encryption is not magic. Just using it does not make all the problems go away.

Even worse, the warning may not give the user a backup method to continue working—the choice is to quit (a total failure of availability) or take the risk and proceed. Since users need to get their work done, and most of these failures are not malicious, users conclude (correctly, in my opinion) that the pragmatic option is to ignore the error, which has the effect of undoing the benefit of encryption in the case of an actual attack. In my view, parts of the security community have somewhat ignored the goal of availability in their pursuit of confidentiality and integrity.

User Attacks on Each Other

While information security among communicating users has received a lot of attention from the research community, many of the problems that plague users do not fit here. Instead, problems arise because one of the parties to the communication is not trustworthy.

Most users understand that email may be spam, may contain malware in attachments, or may pretend to be

from someone other than the apparent sender (so-called *phishing* email falls in this category). Some users may understand that connecting to a website can lead to the download of malware onto their computer. Other applications suffer from different consequences of untrustworthy actors, including false reviews of restaurants, abusive comments on sites that allow posting, and fake news.

To make applications more secure, the challenge is not to fix implementation bugs but to change their design.

While attacks by one user on another can exploit many aspects of the internet architecture, I think the most common attacks exploit applications—spam and phishing via email, bogus restaurant reviews on Yelp, and so on. Unfortunately, while it would be nice if we could give the job of fixing all these attacks to one sort of actor (e.g., ISPs), we cannot do that. ISPs are not responsible for internet apps. Attacks have to be fixed one application at a time.

Built-in Risks

Why do applications allow these unwelcome behaviors? With respect to email, an application designed in the early days of the internet, the designers (including me) did not appreciate the many ways it could be abused. But with respect to more recent applications, risky modalities have sometimes been knowingly included because they provide powerful features. The Web-based attack in which malicious code is downloaded to a user's computer would not be possible if Web protocols did not include the ability to download and execute code in the first place. When that feature was proposed, security experts were clear on the risks; the feature was added anyway because of its utility. So to make applications more secure, the challenge is not to fix implementation bugs but to change their design. That challenge is a big one, because it may require changes to a highly distributed system where the parts are under control of many actors.

Despite these risks, people continue to use applications on the internet—the benefits outweigh the risks, and the benefits are substantial. For many people, the Web and email are essentially indispensable. So how do we deal with these risks?

In the larger world, when actors interact but don't trust each other, they depend on constraints on the interaction and trusted third parties to provide protection. When we buy or sell a house, we depend on registries of deeds, escrow agents, and the like to make sure the transaction is trustworthy, even if one of the actors is not. And the role of credit card processors in facilitating interactions between buyers and sellers (and essentially insuring the transaction by reimbursing fraud) allows buyers and sellers that have no knowledge of each other to complete a transaction in a trustworthy fashion.

On the other hand, people may engage in potentially risky activities because they know and trust each other (I might lend money to a friend, but not a stranger). People tend to treat friends differently from strangers.

Assessing Identity

How might an application be designed to provide both the advantages of powerful but risky modalities and protection from attack? Applications could be designed to modify their behavior depending on the degree of trust among the participants—an email client might allow attachments from a known party to be downloaded, but block them from an unknown sender. However, for this approach to work, it would be necessary for the application to "know" who the other parties are to a sufficient degree to determine the right degree of trust. It makes no sense to talk about whether someone is trustworthy if you have no idea who they are. This line of reasoning suggests that some concept of *identity* needs to be a part of internet security.

Calling for better identity verification could be taken to mean that all actions on the internet should be traceable to known persons. Calls for an "accountable internet" seem to suggest that what is needed is universal strong identity associated with every action. I think that would be a very bad idea. On the one hand, it would preclude any sort of anonymous action, which is often desirable, and on the other hand, it would not work well. How confident should we be about identity credentials issued by nation-states with interests adverse to ours? What sort of recourse would we have if one of those actors did something unwelcome?

In my view, identity cues need to be designed differently for different applications. Your bank really does

need to know who you are; Yelp does not. A site offering medical advice about sensitive issues may commit *not* to identify you. For email, what may matter is not that the recipient is given proof of who the sender actually is, but that the sender is the same person as in prior messages. Over a sequence of messages, a receiver can build up a model of whether a sender is trustworthy (this version of identity is sometimes called continuity identity). But this approach to managing identity again pushes the responsibility for improved security back to the application layer.

Trade-offs?

Depending on the security threat, the response may differ. Consider old-fashioned physical mail. Most people would probably say that their mail should not be opened by a third party while it is in transit—mail is private. But if the envelope contains anthrax dust, it would be very good if it were opened by someone else—properly trained and in a hazmat suit. In the internet context, it might be nice if incoming mail were inspected to see if it was spam or contained malware, but that task would be impossible if the email was encrypted, although encryption is the best way to preserve confidentiality.

Attacks on the Network

It is possible for a region of the internet itself to be attacked, by either another region or a user. As a packet carriage infrastructure, the internet is a global system, and it should not be surprising that some regions of it are malicious or have interests that are adverse to each other. Attacks often target a key service of the internet; the three most common services are the global routing system (the Border Gateway Protocol, BGP), domain name system (DNS), and certificate authority (CA) system.

BGP Vulnerabilities

Using BGP, each region of the internet tells other regions which addresses are located in that region so that packets can be sent there. If a malicious region announces that it is a good route to some addresses that don't belong to it, traffic may flow to that region instead of the legitimate destination and the information may be examined, blocked, or manipulated in the malicious region.

The vulnerabilities in the BGP (and DNS) have been known for decades—the limitation in the routing protocol was first described in 1982.[1] So why do these vulnerabilities persist? The problem is not the lack of a technical solution.

There is no organization in charge of the global internet to dictate an answer, so a workable solution is elusive.

One problem is that there are competing solutions, with fierce advocates for each. Solutions differ, in part, with respect to exactly what security problem is being addressed, and it has been very hard to get agreement as to what the actual risk is. The security improvements are costly in performance (most of them involve encrypted messages among routers, which adds substantially to the overhead of processing the messages) so there is no enthusiasm for deploying them until it is clear that they are needed. And because BGP is a global system, many actors would have to agree on the solution and deploy it. Disagreements about how to improve BGP have persisted for years. There is no organization that is sufficiently in charge of the global internet that it can dictate an answer, so a workable solution is elusive.

DNS Vulnerabilities

The DNS maps a name (e.g., www.example.com) to an internet address. If the DNS is subverted by a malicious actor, it can be made to return the wrong address for a name so that traffic intended for a legitimate destination is instead sent to an attacker. With the DNS, there is perhaps less debate about what the threat is, but it has still been hard for the community to make progress on a solution.

Encryption can help mitigate DNS problems. If one user has an encryption key that is specific to another,

[1] In the design document describing the predecessor of BGP, the author stated: "If any gateway sends an NR [network reachability] message with false information, claiming to be an appropriate first hop to a network which it in fact cannot even reach, traffic destined to that network may never be delivered. Implementers must bear this in mind" (Eric C. Rosen, Exterior Gateway Protocol (EGP), Internet Request for Comment 827, October 1982, available at https://tools.ietf.org/html/rfc827, p. 32). The potential vulnerability was not flagged as a security risk, but more as a mistake to avoid.

intended user, then a rogue clone cannot decrypt what the user sends and the malicious interception will fail. (This is a successful attack on availability, but at least no other harms will occur.) So if attackers want to penetrate an encrypted communication, they must not only deflect the traffic to the rogue endpoint but also disrupt the key management system so the user under attack has the wrong key for the corresponding party (the key for the rogue clone and not the intended endpoint).

CA System Vulnerabilities

On the internet, encryption keys are most commonly managed by the CA system, so attackers are motivated to attack this system. It has not proven resistant to attack.

To oversimplify, the design of the CA system was based on a technically simplifying assumption that was flawed when deployed in the real world: that the servers that make up the CA system would be trustworthy. In a global system, that was an unrealistic assumption. Some CAs have proven corrupt, some have been penetrated, and some (specifically CAs operated by state actors) have been known to hand out false keys as part of a state-sponsored action.

Security is getting better. But attackers are getting better as well.

If attackers can both intercept traffic and hand out false keys, they can penetrate a connection that a user thought was protected by end-to-end encryption. This state of affairs holds today.

Denial of Service Attacks

A DoS attack disables a host, application, or region of the network by flooding it with so much extraneous traffic that legitimate traffic is squeezed out. The attack requires many simultaneous sources of traffic and so is usually called a distributed denial of service (DDoS) attack.

The first step in a DDoS attack is to penetrate and subvert many end-nodes on the network, installing malware that can later be commanded to undertake various tasks, from sending spam to participating in a DDoS attack. This collection of infected machines is colloquially called a botnet, and the person controlling it is a botmaster.

There are a number of ways to counter a DDoS attack (or to disrupt a botnet). Ideally, end-nodes on the internet would be secure enough to resist takeover by a botmaster. Indeed, security on traditional personal computers has greatly increased in the last decade. However, a new generation of inexpensive edge devices have come on the market—surveillance cameras, smart doorbells, and the like—many of which are designed with no thought to security. So botmasters have a new generation of easy targets to exploit.

The botmaster must control his infected machines, which requires some sort of command and control system. Defenders can attempt to disrupt that system, but the internet is a general-purpose data transport network, so botmasters invent new schemes to control their botnets as the old ones are disrupted.

Botnet disruption can also be prevented by replicating the potential target system until there are enough copies that the traffic from the botnet cannot overwhelm all of them at the same time. If there are 100 copies of a service scattered across the internet and the attacker targets all of them, then his botnet has to be 100 times as large as if there were only one copy. If the attacker concentrates on one copy, it can be disabled, but that leaves 99 still running.

Some of the most effective countermeasures to botnets have involved identifying the botmasters and taking them to court. But this approach is often thwarted by cross-jurisdictional issues and variation in law.

Conclusions

This picture may seem bleak. But at the level I have described, the problems are specific enough that there is a chance of action. And security is indeed getting better, although attackers are getting better as well.

Google has put in place a scheme to improve CA system security that does not require changes to the CA providers themselves. It is called certificate transparency, and while it has drawbacks, it is a path in the right direction. Other measures are emerging to improve email sender authenticity, and so on.

A number of lessons can be drawn for the future of cybersecurity:

- First, while encryption is a powerful tool, it is not magic. For encryption to be successful, it must be embedded in a system that can manage keys (such as

the CA hierarchy), and flaws are more likely to arise in that system than in the encryption algorithm itself.

- Second, some attacks are complex and multistage. A successful attack today does not come from simple exploitation of a vulnerability but may require several steps, with subversion of more than one system. This reality signals the sophistication of attackers—but also the possibility of thwarting an attack at multiple stages.
- Third, perfect security is not possible. Security is a multidimensional space of sometimes conflicting threats and responses. Is encryption a good idea or a bad idea in a specific context? Should an application allow a risky operation or block it?
- Fourth, the case of BGP illustrates two nontechnical barriers to better security: coordination problems and negative externalities. To undertake an improvement to BGP, the operators of the more than 70,000 regions that make up the internet have to agree to make the change. But why should they? They bear the cost of implementation and operation, but do not themselves benefit: it is users that are harmed when traffic is deflected by rogue routing. This is the negative externality—one actor bears the cost while another benefits. There are many examples in the cybersecurity space of coordination problems and negative externalities.
- Finally, there is nobody in charge. To a significant degree, this fact has been a strength, not a weakness. Since the internet evolves bottom-up by consensus, it is very hard for any single actor to dominate its future. But when a collective decision is needed, making that decision can be difficult and slow. Powerful actors (such as Google, with certificate transparency) are finding solutions to problems that they may hope to impose unilaterally. This outcome may be very beneficial for security, but reflects a change in internet governance toward a more centralized character. Increasing centralization (of many aspects of the internet) may be a byproduct of a push for better security.

Acknowledgments

I greatly appreciate the very thoughtful, detailed critique by Mary Ellen Zurko, as well as the meticulous editorial efforts of Cameron Fletcher.

The definition of insider threat must be expanded from the malicious human insider to include the unwitting human insider and the technological insider.

Cybersecurity:
Revisiting the Definition of Insider Threat

Nicole Lang Beebe and
Frederick R. Chang


Nicole Lang Beebe


Frederick R. Chang

The insider threat problem is older than the cybersecurity problem itself and has similarly proven to be exceedingly resilient to solution. Organizations work hard to establish adequate defenses to combat external cyber risk, but the insider threat may actually be a greater concern.

Redefining Insider Threat

As technological advances provide better tools to detect and prevent insider threat attacks, they also introduce new threats. They not only make it easier for adversaries to engage trusted human actors in a network but also introduce new, nonhuman trusted agents, such as mobile devices, internet-connected devices, and artificial intelligence (AI). Indeed, the definition and treatment of insider threat need to be expanded to include unwitting human actors *and technology* that act as trusted agents within networks.

Nicole Beebe is professor and chair, Department of Information Systems & Cyber Security, Melvin Lachman Distinguished Professor, and director of the Cyber Center for Security & Analytics at the University of Texas at San Antonio. Frederick Chang (NAE) is professor and chair, Department of Computer Science, Bobby B. Lyle Centennial Distinguished Chair in Cyber Security, and executive director of the Darwin Institute for Cyber Security at Southern Methodist University.

Types of Human Insiders

Human insiders are individuals with legitimate access to an organization's computers and networks and whose volitional actions put organizational data, processes, or resources at risk in an unwelcome or disruptive way, whether their intent is malicious or nonmalicious (e.g., policy violations motivated for organizational good without regard to unintended consequences that may introduce security risks) (Pfleeger et al. 2010). In recent years, the definition has expanded to include the unwitting insider (Costa et al. 2014; Guo et al. 2011; Maasberg et al. 2015; Willison and Warkentin 2013).

There are three broadly characterized types of human insider. Malicious insiders knowingly and willingly seek to harm their organization through espionage, theft of intellectual property (IP), fraud, or sabotage (Moore et al. 2013). In contrast, nonmalicious insiders may knowingly violate organizational policy but believe they are doing so for the greater good of their organization. They may, for example, circumvent security policies to get their job done more efficiently.

Insiders in the third class are often referred to as unwitting insiders. They neither intend their organization harm, nor even know they are doing anything wrong. Such insiders are dangerous to organizations (Greitzer et al. 2014), as they are susceptible to social engineering attacks (e.g., phishing), nefarious websites, or malware (Verizon 2019).

The Role of Technology

Malicious insider attacks appear to be on the rise, arguably due to, or at least enabled by, technological advancement. Exfiltrating information from an organization no longer requires surreptitiously photocopying large amounts of documents in small increments or secretly removing hard-copy data from an organization. Technological advances such as removable thumb drives, email, and cloud storage facilitate espionage and IP theft. The electronic ledger enables fraud and theft without stealing a physical dollar. The information technology (IT) infrastructure critical to organizations is a target for cybersabotage.

In criminal justice terms, technological advancement has increased the motive, opportunity, and means for the malicious insider. Insiders are (1) increasingly motivated by information-based targets, (2) equipped to commit such crimes through a robust market of point-and-click exploitation suites and ready-to-deploy malware, and (3) able to easily identify opportunities for information acquisition and financial gain in vulnerable internet-connected systems.

Technology has also resulted in an increase in external adversary use of unwitting insiders to gain a digital foothold in an organization: the adversary can then pivot, move laterally, and escalate privilege levels to obtain access and control over targeted digital asset(s). The unwitting insider has become a critical component in the system and process used by external hackers the world over. Until zero trust network model design principles (Kindervag 2010) become the norm, external hackers will continue to be enabled by the unwitting insider, by virtue of the operational trust given to the authorized user account. In fact, the most difficult insider threat to defend against is the unwitting insider (Verizon 2019).

Technology has resulted in an increase in external adversary use of unwitting insiders to gain a digital foothold in an organization.

Technology Itself as Insider Threat

As discussed, technology is both a target and an enabler (its role as a defender, detecting insider activity through data loss prevention appliances and security information event management devices, is assumed). Looking ahead, however, it is poised to become a perpetrator—the next insider. This prediction is based not on machine-initiated malfeasance in the evolution of a far-flung view of sentient AI, but rather on technology as a trusted insider in a complex system: a machine is given access and its outputs are trusted inputs to other machines (and humans) in a larger system.

Trusted Machines and Systems

The prevailing paradigm to trust the machine remains. In industrial manufacturing environments, for example, system and manufacturing process engineers design cyberphysical systems whose mechanical devices and machinery are automated by human-designed computer-based algorithms (commonly referred to as Industry 4.0). Unfortunately, these systems are designed

for reliability without sufficient regard to cybersecurity (Thames and Schaefer 2017). In effect, the computer-based algorithm becomes a trusted agent within the system, as do the robotic manufacturing devices carrying out the computer-based directives. Both the software and the hardware are treated as trusted agents within the system.

In critical infrastructure contexts, research has illuminated an array of security vulnerabilities introduced by interconnecting operational technology (OT) and IT networks (Lun et al. 2016; Murray et al. 2017). These challenges are becoming evident as 50 billion Internet of Things (IoT) devices (Afshar et al. 2017) are increasingly interconnected—to each other, to manufacturing systems, to traditional networks, and beyond. The problem is only going to worsen when fifth generation (5G) networks come online and machine-to-machine interconnectivity promised by Industry 4.0 becomes a reality.

The insider is a trusted actor on a network, whether that actor is human, an embedded device, the software, the network, or the AI.

Last, the added layer of evolving AI in next-generation manufacturing environments (Industry 5.0), with its envisioned human-machine symbiosis, means that AI will become the next trusted agent in the system. But research on adversarial AI techniques (Carlini and Wagner 2017; Dalvi et al. 2004; Lowd and Meek 2005) and on efforts to detect and counter them (Tramèr et al. 2017; Yuan et al. 2019) makes it clear that AI cannot necessarily be trusted. Yet AI *is* still widely trusted.

In sum, a system-based view of the insider threat necessitates an evolved perspective of "who" the insider is. A continued human-centric approach that focuses solely on malicious actors is myopic and dangerous. The insider is the trusted actor on a network, whether that actor is human, an embedded device, the software, the network, or the AI, and its risk should be considered regardless of whether the action is volitional or nonvolitional and whether the motive is malicious or nonmalicious.

The Need for a Systems-Level Approach

The complexity and interdependency of the modern organization necessitates a systems-level view of security that is integrated into the full systems design from the beginning. Security practitioners have long argued that security must be "baked in," not "tacked on" at the end, but unfortunately this philosophy has largely focused on device-level design without taking into account a systems engineering perspective. The security of systems of systems and across interconnected systems remains underaddressed.

Furthermore, the conceptualization of the insider has remained focused on the vulnerable but trusted human. Not only is it inadvisable to trust all insiders to act without malicious intent, it is essential to increase vigilance against the unwitting insider.

In advocating for an expanded view of the insider threat to include the technological insider, we are informed by important earlier work that has described similar concepts—including some of the pioneering computer security work from the 1970s that proposed the notion of subverted software or hardware as a serious security threat to computing systems of the time (Anderson 1972; Lampson 1973; Schell 1979). The problem was critically important then—and it still is. Indeed, the case can be made that in today's advanced technological landscape, greater awareness of and attention to the problem are long overdue.

Mechanisms of Technological Insider Threats

The US National Insider Threat Policy (DNI 2012), written in response to Executive Order 13587, "Structural Reforms to Improve the Security of Classified Networks and the Responsible Sharing and Safeguarding of Classified Information" (Obama 2011), sets expectations and identifies best practices for deterring, detecting, and mitigating insider threats. However, it focuses entirely on the human insider and turns to technology only to deter and detect the human insider threat. It calls for program establishment, training of program personnel, monitoring of user and network activity, and employee training and awareness. The more recent Insider Threat Guide (NITTF 2017) and maturity framework (Belk and Hix 2018) from the National Insider Threat Task Force (NITTF) continue this trend. Yet the risk posed by the technological insider is clear and present.

Personal Devices Used for Work

According to a 2018 study of 500 IT executives, CEOs, and other senior managers in the United States, commissioned by and conducted in partnership with Samsung, nearly 80 percent of employees cannot accomplish their required work effectively without a mobile device (Oxford Economics 2018). The vast majority of these devices are personally owned and managed, and 61 percent of the respondents expected their employees to be available remotely. According to the study, companies save approximately 15 percent with a bring your own device (BYOD) policy over an enterprise-owned device policy, so many companies opt for BYOD.

Poor Security Practices

With so many organizations dependent on personally owned mobile devices as part of their workflow, user smartphone security awareness and implementation are critical. Unfortunately, however, many users continue to demonstrate poor security awareness and hygiene on their mobile devices (Parker et al. 2015), putting their organizations at risk as smartphones become part of an organizational network infrastructure through email access, file activity, workflow processing, and much more. Mobile device users incur risks by

- using poor password policies;
- failing to use device and screen locks (and when they do, often employing weak authentication mechanisms), virus protection, encryption, remote device locator services, and/or remote wiping services; and
- visiting unsafe websites and/or installing risky software (Mylonas et al. 2013; Parker et al. 2015; Sebescen and Vitak 2017).

Defenses

To combat poor user security on BYOD devices, companies are deploying mobile device management (MDM) software that enables software network access control, identification of outdated or nonexistent virus detection software, enforcement of passcode requirements, detection of "jailbroken" phones, remote location and wiping, application-level security through VPN tunnels, white-/blacklisting, and dynamic policy enforcement. Other organizations are opting for a containerization approach, wherein company data, communications, and work-related applications are stored in an encrypted partition/area of the device.

This all sounds promising, except that 70 percent of small businesses fail to implement MDM or containerization (Parizo 2018). Without such protections, millions—and potentially billions—of smartphones are connected to organizational networks as trusted insiders, gaining access to sensitive data, propagating malware, and giving adversaries footholds in their networks.

Embedded Devices

Embedded devices affect enterprises through supply chain security risks and through interconnectivity between IoT devices and enterprise networks.

Supply Chain Vulnerabilities

Supply chain vulnerabilities jeopardize enterprise networks when individual computing components in systems are compromised (e.g., through integrated circuit chips from untrusted sources). The resulting challenges to the security of the computational system may affect the authenticity of materials and components, physical tracking and antitampering efforts, secure communications, and the installation of backdoors, among others.

How trustworthy are the computing components vital to the vast array of devices used daily in work and personal life?

How trustworthy are the computing components vital to the vast array of computational devices used daily in work and personal life? This is a significant and growing concern, as evidenced by several major research funding programs by US federal agencies. In 2007 the Defense Advanced Research Projects Agency committed nearly $25 million to the Trust in Integrated Circuits Program (Adee 2007). In 2011 it invested another $49 million in the Integrity and Reliability of Integrated Circuits Program (Rawnsley 2011). The Intelligence Advanced Research Project Agency followed suit with a Trusted Integrated Chips Program.[1] The National Institute of Standards and Technology has issued comprehensive guidance on

[1] https://www.iarpa.gov/index.php/research-programs/tic

cyber supply chain risk management (Boyens et al. 2015), and supply chain security is a major thrust of the Department of Energy's latest $70 million funding announcement for a cybersecurity institute for energy efficient manufacturing (DOE 2019).

Interconnectivity Risks

Embedded devices are increasingly connected in every area of life, especially with the rapid adoption of IoT devices and their interconnectivity with enterprise networks. This interconnectivity may be direct (enterprise IoT devices as part of the network) or indirect (personal IoT devices connected to personal smartphones, which are then connected to enterprise networks). Gartner (2016) predicts that by 2020 over 25 percent of identified security attacks in enterprises will involve IoT devices, a prediction that seems to be borne out by the exponential increase—nearly 3700 percent—in IoT malware: Kaspersky Lab (2018) saw 3,219 pieces of IoT malware in 2016, 32,614 in 2017, and 121,588 in the first half of 2018. It appears clarion calls about the problem were well founded (e.g., Schneier 2014).

Research into detecting and defending against adversarial AI attacks remains limited, in both quantity and success.

The Mirai botnet, for example, leveraged IoT devices such as webcams, DVRs, and routers to launch a distributed denial of service attack on internet service provider Dyn, taking down a significant portion of the internet on October 21, 2016, and preventing millions of people from accessing over 1,200 websites, including Twitter and Netflix, for nearly an entire day (Kolias et al. 2017). Home and building automation systems have been shown to have significant vulnerabilities in critical systems such as those for HVAC; fire detection, suppression, and alerts; security and access control, such as cameras/CCTVs and door locks; and lighting (Peacock and Johnstone 2014).

Network Devices

The problem of the personal device insider and the embedded device insider is about to get worse with the rollout of 5G, which will enable greater device-to-device connectivity, interconnectivity between physical and virtual devices, and greater/faster bandwidth. 5G's enhanced encoding, air interface, channel frequencies, and antenna technologies mean that cellular devices will be able to do even more and IoT devices will be able to interconnect on a much grander scale. An individual's smartphone will be connected to his home automation system, vehicle, peers when gaming—and BYOD work environment (and everything in between).

Software (AI/Autonomy)

The last technological insider is associated with AI and autonomous systems. The scenario has been fictionally illustrated on the big screen and in novels for quite some time, but the threat is real.

Most existing machine learning (ML) classifiers are not particularly robust or immune to adversarial attacks (also known as adversarial learning or adversarial AI) (Kurakin et al. 2018). This was well known and of sufficient concern for Google Brain to organize a competition at the 2017 Conference on Neural Information Processing Systems to generate new adversarial attack samples and develop defenses to counter them.[2] Research into detecting and defending against adversarial AI attacks, however, remains limited, in both quantity and success (Athalye et al. 2018; Kurakin et al. 2018).

In addition to adversarial input attacks, data poisoning (Biggio et al. 2012) and model stealing attacks (Tramèr et al. 2017) against ML systems must be addressed. Organizations that use AI, even with human-in-the-loop designs, must carefully consider the trust models associated with AI and autonomous systems and vigilantly monitor for nefarious concept drift (systems slowly moving outside of reasonable parameters).

The risk of AI or autonomous systems becoming a *witting* insider may be many years off, but the very near reality is that they can be manipulated nefariously, much like the nonmalicious/nonvolitional (i.e., unwitting) human insider. The quest for reliable adversarial AI detectors and defense systems must accelerate, because the widespread deployment of AI, ML-based, and autonomous systems is already underway.

[2] https://kaggle.com/c/nips-2017-non-targeted-adversarial-attack

Addressing the Challenge through Systems Engineering

The complexity of the insider threat problem—three classes of human insiders as well as the technological insiders discussed here—necessitates a true systems-engineered solution.

When systems are designed with a systems engineering perspective, their constituent components and interconnectivity are fully and intentionally integrated synergistically to optimally perform a collective function. Security may be viewed as both a subsystem and a system design characteristic. However, all too often it is treated only as a design characteristic at the subsystem level rather than being considered in the design across the entire system of systems.

The relatively new field of security engineering has attempted to remedy this, but still falls short. Security engineering typically takes a multilevel approach that considers software security, information security, physical security, and technical surveillance separately rather than in an integrated system design. Integrated security engineering from a systems engineering perspective should consider information/data flows, subsystem interconnectivity, human-to-system and system-to-system access controls, and privilege design and protection.

Information and data flow security has largely been researched within, not between, systems. Examples include system call monitoring (Hofmeyr et al. 1998) and anomaly detection (Bhatkar et al. 2006), as well as the more recent concept of data flow assertions (Yip et al. 2009) for application layer security, wherein an explicitly identified data flow plan is checked for compliance at runtime. There has been little attention to the need to integrate security and systems engineering throughout the system development process (Mouratidis et al. 2003).

Most integrated system development security has been geared toward identifying and mitigating security risks associated with OT-IT interconnectivity or between legacy and modern systems. But the problem goes much deeper. As attackers gain a foothold in an inconsequential system and then move to more consequential targets through pivoting and lateral movement, escalating privileges along the way, they are able to obtain access disturbingly easily wherever the target may be within a network.

Consider the design of an advanced manufacturing system, for example. Manufacturing subsystems are interconnected through secure versions of machine-to-machine transport protocols (e.g., MTConnect and MQTT) or secure 5G wireless interconnectivity, secured with software, trusted hardware components, and intrusion detection systems. In addition:

- Interdevice trust and dataflow are secured through trusted protocols, access management, and security analytics.
- Predictive analytics based on the application of machine learning and artificial intelligence are routinely calibrated and verified, not blindly trusted.
- Digital twin and digital thread technologies are employed for cyberphysical vulnerability detection.
- Proactive mitigations, response, and resiliency are performance requirements that drive functional requirements such as isolation, containerization, moving target defense, advanced sensing, anomaly detection, self-healing, forensic data retention, rapid information sharing, and full system visualization.

Enterprises are connecting more and more weak links to their network through personal, embedded, network, and autonomous devices.

Conclusion

A system is only as strong as its weakest link, and enterprises are connecting more and more weak links to their network through personal, embedded, network, and autonomous devices. Progress in addressing the insider threat has been made since the establishment of the NITTF in 2011, but unfortunately, the focus has been on the malicious human insider. The scope must be broadened to include the unwitting insider and the technological insider as well.

The wave of 5G interconnected IoT devices, interconnectivity between home and work through the notoriously undersecured smartphone in an increasingly BYOD enterprise, and the development and spread of AI-enabled software are all making it increasingly dif-

ficult to defend digital networks. These trends, along with the growing interconnectedness of critical infrastructures, are cause for great concern.

Acknowledgments

Special thanks to Steve Lipner for his constructive comments that materially improved the perspectives and contributions of this article. Additional gratitude to Cameron Fletcher and Jenni Simonsen for their careful editing to improve the readability of the ideas and thoughts conveyed herein.

References

Adee S. 2007. Contracts awarded for DARPA's Trust in Integrated Circuits Program. IEEE Spectrum: Technology, Engineering, and Science News, Dec 6.

Afshar V. 2017. Cisco: Enterprises are leading the Internet of Things innovation. HuffPost, Aug 28.

Anderson JP. 1972. Computer Security Technology Planning Study, Vol I. Report ESD-TR-73-51. AFSC Electronic Systems Division, Hanscom AFB, Bedford MA.

Athalye A, Carlini N, Wagner D. 2018. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. Proceedings, 35th International Conf on Machine Learning, Jul 10–15, Stockholm.

Belk RW, Hix TD. 2018. Insider Threat Program: Maturity Framework. McLean VA: National Insider Threat Task Force.

Bhatkar S, Chaturvedi A, Sekar R. 2006. Dataflow anomaly detection. 2006 IEEE Symposium on Security and Privacy, May 21–24, Berkeley/Oakland.

Biggio B, Nelson B, Laskov P. 2012. Poisoning attacks against support vector machines. ArXiv:1206.6389.

Boyens J, Paulsen C, Moorthy R, Bartol N. 2015. Supply Chain Risk Management Practices for Federal Information Systems and Organizations. NIST Special Publication 800-161. Gaithersburg: National Institute of Standards and Technology.

Carlini N, Wagner D. 2017. Adversarial examples are not easily detected: Bypassing ten detection methods. Proceedings, 10th ACM Workshop on Artificial Intelligence and Security, Nov 3, Dallas.

Costa DL, Collins ML, Perl SJ, Albrethsen MJ, Silowash GJ, Spooner DL. 2014. An ontology for insider threat indicators: Development and applications. Online at https://resources.sei.cmu.edu/asset_files/ConferencePaper/2014_021_001_426817.pdf.

Dalvi N, Domingos P, Sanghai S, Verma D. 2004. Adversarial classification. Proceedings, 10th ACM SIGKDD International Conf on Knowledge Discovery and Data Mining, Aug 22–25, Seattle.

DNI [Director of National Intelligence]. 2012. National Insider Threat Policy. Online at https://www.dni.gov/files/NCSC/documents/nittf/National_Insider_Threat_Policy.pdf.

DOE [US Department of Energy]. 2019. DOE announces $70 million for Cybersecurity Institute for Energy Efficient Manufacturing. Energy.Gov, Mar 26.

Gartner. 2016. Gartner says worldwide IoT security spending to reach $348 million in 2016. Apr 25. Stamford CT.

Greitzer FL, Strozer JR, Cohen S, Moore AP, Mundie D, Cowley J. 2014. Analysis of unintentional insider threats deriving from social engineering exploits. 2014 IEEE Security and Privacy Workshops, May 17–18, San Jose.

Guo KH, Yuan Y, Archer NP, Connelly CE. 2011. Understanding nonmalicious security violations in the workplace: A composite behavior model. Journal of Management Information Systems 28(2):203–236.

Hofmeyr SA, Forrest S, Somayaji A. 1998. Intrusion detection using sequences of system calls. Journal of Computer Security 6(3):151–180.

Kaspersky Lab. 2018. New IoT-malware grew three-fold in H1 2018. Press release, Sep 18.

Kindervag J, with Balaouras S, Coit L. 2010. No more chewy centers: Introducing the zero trust model of information security. Cambridge MA: Forrester Research.

Kolias C, Kambourakis G, Stavrou A, Voas J. 2017. DDoS in the IoT: Mirai and other botnets. Computer 50(7):80–84.

Kurakin A, Goodfellow I, Bengio S, Dong Y, Liao F, Liang M, Pang T, Zhu J, Hu X, Xie C, and 13 others. 2018. Adversarial attacks and defences competition. The NIPS '17 Competition: Building Intelligent Systems, eds Escalera S, Weimer M. Cham, Switzerland: Springer International Publishing.

Lampson B. 1973. A note on the confinement problem. Communications of the ACM 16(10):613–615.

Lowd D, Meek C. 2005. Adversarial learning. Proceedings, 11th ACM SIGKDD International Conf on Knowledge Discovery in Data Mining, Aug 21–24, Chicago.

Lun YZ, D'Innocenzo A, Malavolta I, Di Benedetto MD. 2016. Cyber-physical systems security: A systematic mapping study. ArXiv:1605.09641.

Maasberg M, Warren J, Beebe NL. 2015. The dark side of the insider: Detecting the insider threat through examination of dark triad personality traits. 48th Hawaii International Conf on System Sciences, Jan 5–8, Kauai.

Moore AP, McIntire D, Mundie D, Zubrow D. 2013. Justification of a pattern for detecting intellectual property theft by departing insiders. Technical note CMU/SEI-2013-

TN-013. Pittsburgh: Carnegie Mellon University Software Engineering Institute.

Mouratidis H, Giorgini P, Manson G. 2003. Integrating security and systems engineering: Towards the modelling of secure information systems. In: Active Flow and Combustion Control 2018 (Vol 141), ed King R. Cham: Springer International Publishing.

Murray G, Johnstone MN, Valli C. 2017. The convergence of IT and OT in critical infrastructure. Proceedings, Australian Information Security Management Conf, Dec 5–6, Perth.

Mylonas A, Kastania A, Gritzalis D. 2013. Delegate the smartphone user? Security awareness in smartphone platforms. Computers and Security 34:47–66.

NITTF [National Insider Threat Task Force]. 2017. Insider Threat Guide: A Compendium of Best Practices to Accompany the National Insider Threat Minimum Standards. McLean VA.

Obama B. 2011. Structural Reforms to Improve the Security of Classified Networks and the Responsible Sharing and Safeguarding of Classified Information (Executive Order 13587). Federal Register 76(198). Online at https://www.dni.gov/files/NCSC/documents/nittf/EO_13587.pdf.

Oxford Economics. 2018. Maximizing Mobile Value: Is BYOD Holding You Back? Oxford.

Parizo C. 2018. What is MDM? Does your small business need it? Samsung Business Insights, Oct 23.

Parker F, Ophoff J, Belle JV, Karia R. 2015. Security awareness and adoption of security controls by smartphone users. 2nd International Conf on Information Security and Cyber Forensics, Nov 15–17, Cape Town.

Peacock M, Johnstone MN. 2014. An analysis of security issues in building automation systems. Proceedings, 12th Australian Information Security Management Conf, Dec 1–3, Perth.

Pfleeger SL, Predd JB, Hunker J, Bulford C. 2010. Insiders behaving badly: Addressing bad actors and their actions. IEEE Transactions on Information Forensics and Security 5(1):169–179.

Rawnsley A. 2011. Can DARPA fix the cybersecurity "problem from hell"? Wired, Aug 5.

Schell RR. 1979. Computer security: The Achilles' heel of the electronic Air Force? Naval Postgraduate School, Monterey CA.

Schneier B. 2014. The Internet of Things is wildly insecure—and often unpatchable. Wired, Jan 6.

Sebescen N, Vitak J. 2017. Securing the human: Employee security vulnerability risk in organizational settings. Journal of the Association for Information Science and Technology 68(9):2237–2247.

Thames L, Schaefer D, eds. 2017. Cybersecurity for Industry 4.0: Analysis for Design and Manufacturing. Cham: Springer.

Tramèr F, Kurakin A, Papernot N, Goodfellow I, Boneh D, McDaniel P. 2017. Ensemble adversarial training: Attacks and defenses. ArXiv:1705.07204.

Verizon. 2019. Insider Threat Report. Online at https://enterprise.verizon.com/resources/reports/insider-threat-report/.

Willison R, Warkentin M. 2013. Beyond deterrence: An expanded view of employee computer abuse. MIS Quarterly 37(1):1–20.

Yip A, Wang X, Zeldovich N, Kaashoek MF. 2009. Improving application security with data flow assertions. Proceedings, ACM SIGOPS 22nd Symposium on Operating Systems Principles, Oct 11–14, Big Sky MN.

Yuan X, He P, Zhu Q, Li X. 2019. Adversarial examples: Attacks and defenses for deep learning. IEEE Transactions on Neural Networks and Learning Systems.

Previous incidents highlight the need for clear, specific guidance on cybersecurity best practices and incentives for their use.

Learning from Cybersecurity Breaches:
The Trouble with Recommended Best Practices

Josephine Wolff is an assistant professor of cybersecurity policy at the Fletcher School of Law and Diplomacy at Tufts University.

Josephine Wolff

One of the recurring themes in discussions of cybersecurity is how rapidly the landscape of threats is evolving and how difficult it is for defenders to keep pace with ever-changing attack vectors and vulnerabilities. While it is true that security threats and controls change over time as technologies continue to develop, that idea sometimes leads to the dismissal of older security breaches as irrelevant in the context of trying to defend against tomorrow's threats. But there is considerable value in reexamining past attacks for clues about the ways they remain relatively static over time and how those patterns can be used to do a better job of defending against future attacks.

Introduction

Past incidents are particularly helpful in efforts to understand the strengths and weaknesses of the many recommended cybersecurity best practices. There are recommendations issued by government agencies, industry consortiums, international standard-setting organizations, private companies, nonprofits, and individual security experts. With the wealth of recommended best practices, many organizations are unsure which policies and security controls to implement.

This article looks at selected cybersecurity incidents through the lens of the perpetrators' motivations—financial theft, espionage, and public humiliation of the victims—to show how they shape the trajectory of the incidents,

creating repeated attack patterns over time. These patterns can help guide defensive interventions aimed at preventing similar incidents. They also shed light on why existing cybersecurity best practices have often failed to clarify effectively the essential security responsibilities and controls for organizations.

Although recommended practices provide a useful starting point for organizations trying to implement stronger security protections, they are designed to target very specific stages of breaches and draw on the capabilities of only a single stakeholder involved in the incident—the organization directly targeted. This narrow focus limits the effectiveness of the best practices because they take advantage of only one small part of both the larger defensive ecosystem and the range of involved stakeholders online.

The recommended practices can also be extremely challenging for organizations to adopt and implement. For instance, the most recent, fourth revision of the NIST 800-53 catalogue[1] includes 115 low-impact security controls, 159 moderate-impact controls, and 170 high-impact controls—more than 440 options. For a small or medium-sized enterprise looking for guidance, this can be an oppressive and impractical quantity of controls.

Furthermore, the lack of empirical evidence that these best practices actually reduce the risk of cybersecurity incidents makes it difficult for organizations to decide which ones to use.

What Drives Cyberattacks?

There are many ways to categorize cybersecurity incidents—according to their targets, their perpetrators, or the technical exploits the perpetrators use, for instance. Each classification scheme is useful in different contexts, such as trying to understand which types of organizations are most likely to be targeted, or which criminal groups require the most attention from law enforcement, or which technical vulnerabilities are most frequently exploited. For identifying attack patterns, or repeated sequences of behavior that persist over years and years in, it is especially useful to consider perpetrators' motives because those goals shape the final, and most essential, stages of their attacks.

[1] Available from the National Institute of Standards and Technology National Vulnerability Database, online at https://nvd.nist.gov/800-53.

Types of Motivation

The three classes of attacker motivation discussed here—financial gain, espionage, and public humiliation—are not comprehensive or exclusive. For instance, the Stuxnet worm was used for none of these purposes but instead to cause physical damage to an Iranian uranium enrichment plant (Zetter 2014). This goal of physical sabotage is yet another potential attacker motivation—one that may become increasingly common as more physical devices are connected to the Internet of Things and it becomes possible to control them through computer networks. However, at present, there are too few examples of such incidents to allow for extensive analysis.

It is useful to consider an attacker's motives because those goals shape the final, most essential stages of an attack.

Profit, espionage, and public shaming are not mutually exclusive motives. For instance, in the cases of both the 2014 Sony Pictures data breach and the 2015 Ashley Madison breach, the theft of data was motivated by the perpetrators' desire to publicly shame the victims, but the resulting dumps of sensitive information were then used by others to commit financially motivated crimes involving identity theft and financial fraud. Similarly, the GameoverZeus botnet operated by Evgeniy Bogachev's organization in Russia was used by Bogachev and his associates to steal millions of dollars through the Cryptolocker ransomware, but some of the information was provided to the Russian government to aid espionage efforts (Schwirtz and Goldstein 2017).

Motivation and Methods

Despite the potential for overlap, goals still provide a useful framework for considering how best to defend against future attacks because the final attack stages are often not replaceable for the perpetrators and are therefore repeated year after year, breach after breach, even as computing technology evolves.

Many of the earlier, technical stages of the breaches that enable an intruder's initial access to target systems

(e.g., phishing, exploiting software vulnerabilities, or stealing credentials) are easily replaced: If one avenue of access is cut off by effective security controls, the perpetrators simply switch to another. But the final stages of their attacks—when they carry out their ultimate aims—are not so easily substituted to achieve the same goal. That makes these stages the most crucial for defenders to cut off and also the most static over time.

Financially Motivated Incidents

Most reported cybersecurity incidents are motivated by money (Verizon 2019). The specific techniques whereby cybercriminals steal data they can use to make money have certainly evolved over time, but the desire to make money and the mechanisms for converting stolen data into financial gain have changed much more slowly. One of the earliest and most common forms of these crimes is the theft of large volumes of payment card information that can be sold on the black market and used for large-scale financial fraud.

The TJX Breach (2005)

The 2005 breach of TJX, Inc. by Albert Gonzalez and a few coconspirators resulted in the theft of 45.6 million payment card numbers, making it one of the largest such breaches at the time of its discovery.

Recommended best practices often do not address the stages of an incident that are most susceptible to effective defensive interventions.

Gonzalez and his team had first identified the firm as a potential target while "wardriving" on a Miami highway using a long radio antenna to detect insecure wireless networks. They found a Marshalls store, owned by TJX, that had a wireless network encrypted using older, less secure WEP encryption. They parked in the store's lot and proceeded to collect packets off the wireless network until they had sufficient information to reverse-engineer the encryption key and retrieve in plaintext the necessary credentials to connect to the company's main servers in Framingham, MA (Verini 2010).

Several payment card networks and issuing banks that bore the brunt of covering the costs of the resulting payment card fraud sued TJX for its negligence in failing to secure the card data more effectively. The lawsuits focused on the company's failure to encrypt its stores' wireless networks using WPA encryption.

But TJX's failure to use up-to-date WiFi encryption was only one missed opportunity in a long chain of events that led up to the execution of the breach—from packet sniffing, credential theft, and remote logins on TJX servers to exfiltration of stored payment card data and the manufacture and sale of fraudulent payment cards. The insecure wireless network was no more responsible for the entirety of the breach than any of the other decisions by TJX and other stakeholders that enabled the later stages of the breach. And yet, because WPA encryption was specified in the Payment Card Industry Data Security Standard (PCI DSS) at the time, both the media and the courts viewed that particular decision as, in many ways, the crucial mistake that demonstrated the company's negligence (Wolff 2018).

The recommended best practices that organizations are routinely faulted (in the aftermath of a breach) for not implementing often do not address the stages of an incident that are most susceptible to effective defensive interventions. In this case, the emphasis on wireless network encryption in the legal proceedings that followed the breach shifted focus away from the incident's monetization stages in which Gonzalez and his coconspirators sold the stolen data and repatriated their profits to the United States—the process by which law enforcement officials were ultimately able to identify and arrest them (Verini 2010).

The SCDOR Breach (2011)

As explained above, often it is not clear which practices an organization should adopt. In the wake of the 2011 breach of millions of tax records from the South Carolina Department of Revenue (SCDOR) critics questioned whether the SCDOR had adhered to required standards and recommended best practices. Then–South Carolina governor Nikki Haley blamed the IRS for failing to instruct the state to encrypt its tax records. The IRS, in turn, invoked NIST as the responsible entity for setting technical security standards for government agencies. Meanwhile, the South Carolina state legislature expressed outrage that the SCDOR had not required multifactor authentication, which might

have defended against the phishing attack that initiated the breach (Mandiant 2012).

The uncertainty around which stakeholders were responsible for determining security best practices and which best practices were essential for the protection of sensitive information contributed to the inability to clarify responsibility and liability for the incident.

The GameoverZeus Botnet and Cryptolocker Ransomware (2013–2014)

The GameoverZeus botnet that distributed the Cryptolocker ransomware program in 2013 and 2014 was able to bypass many forms of multifactor authentication and monetize previously worthless data by selling it back to its original owners for a cryptocurrency ransom. The cryptocurrency payments allowed the perpetrators to evade the centralized financial intermediaries, such as banks and payment card networks, that had powerful defensive forces to combat payment card fraud and identity theft. And because Cryptolocker targeted individuals' computers it distributed the costs to thousands of disparate users rather than concentrating them in the fraud ledgers of credit card companies and issuing banks. This cost distribution diminished the incentives for any intermediaries to intervene and eliminated the class action lawsuits that had driven liability regimes and the associated need for clear recommended best practices to distinguish between negligent and unlucky breach victims.

Liability and Cyberinsurance Challenges

The challenge with relying on catalogues or lists of best practices to determine even partial liability for security incidents is that it narrows the scope of organizations' security responsibilities to relatively confined, nonexhaustive recommendations that are often based more on generally accepted consensus than on empirical evidence. In fact, not all recommended security controls provide measurable improvements in system security, and lists of best practices can sometimes be used to propagate controls (e.g., requirements to regularly change passwords) that do more harm than good in the long term (Wolff 2016).

The lack of clear correlation between security best practices and security outcomes is a significant challenge for the rapidly growing number of insurance firms offering cyberinsurance policies. These firms typically do not have the in-house security expertise that would allow them to audit potential customers' security postures or identify the safeguards those customers should take as part of their coverage, and find they cannot rely on existing best practices as guidance. This has led many insurers to partner with security firms to provide customer assessments and security services (Wolff and Lehr 2018).

Not all recommended security controls provide measurable improvements in system security, and some do more harm than good in the long term.

The monetization stages of financially motivated cybercrimes have traditionally been the most vulnerable to defensive intervention because they relied on a small set of intermediaries, such as well-known online black markets for stolen cards and fraudulent card manufacturers. For instance, law enforcement authorities were able to take down Gonzalez's operation through his fence for stolen card data, a man named Maksym Yastremskiy. Later models of financial cybercrime shifted to enable other types of monetization that relied less on fencing operations like Yastremskiy's.

Cyberespionage

Defending against political or economic cyberespionage typically requires restricting the exfiltration of data and segmenting sensitive portions of networks to contain intrusion attempts and subsequent access to high-value information. But adhering to these best practices is not always sufficient, especially when the espionage efforts are state-sponsored and carried out by actors with considerable expertise and resources.

The following cases highlight both the vulnerability of a company that rigorously followed recommended security standards and the absence of serious consequences for organizations that failed to do so. These examples do little to encourage others to pay more attention to these recommendations.

The Case of DigiNotar (2011)

The Dutch certificate authority DigiNotar was compromised in 2011 as part of what was later hypothesized to

be an espionage operation by the Iranian government (Hoogstraaten et al. 2012). An intruder penetrated DigiNotar's multiple lines of defense to generate rogue certificates for Google's domain, and these certificates were then likely used to capture credentials for thousands of Iranian Google accounts (Hoogstraaten et al. 2012).

The DigiNotar compromise is an example of how many stages can be involved in completing espionage attacks. Not only did the perpetrators have to compromise a certificate authority, they then had to redirect users to fraudulent web pages, probably through Domain Name System (DNS) cache poisoning, to use the rogue certificates they had generated (Hoogstraaten et al. 2012).

Officials blamed other agencies for hindering their ability to implement security upgrades or establish clear expectations for information security.

Besides highlighting the numerous opportunities for defensive intervention, the DigiNotar attack illustrates the limitations of adhering to industry best practices. DigiNotar had a rigorously segmented network structure to separate the high-security certificate-issuing portion of the network from the company's outward-facing web presence. Every request for a new certificate had to be approved by at least two company employees and the servers used to generate certificates were stored in a secure room that could be accessed only by using biometric hand recognition, a key card, and a PIN (Hoogstraaten et al. 2012).

DigiNotar's security setup reads like an excerpt from a manual on how to design a secure system, and yet the intruders were able to find a way to tunnel into the most secure portion of the network and generate rogue certificates.

The OPM Data Breach (2015)

In other cases of espionage, organizations demonstrated much less rigorous adherence to basic tenets of security best practices. A series of espionage efforts led by the Chinese People's Liberation Army (PLA) Unit 61398 and directed at more than 100 private companies and government agencies, primarily based in the United States, were described in a 2013 report by security firm Mandiant. Some of the incidents were confirmed in an indictment of several PLA officers the following year; like the Mandiant report, the indictment pointed to phishing emails and other social engineering campaigns as the primary means of access for the state-sponsored economic espionage efforts (Mandiant 2013).

In 2015, in another espionage operation attributed to the Chinese government, the US Office of Personnel Management was breached and information about 21.5 million people who had worked for the federal government or received security clearances was stolen (Chaffetz et al. 2016). Congressional hearings highlighted OPM's lack of encryption, multifactor authentication, and intrusion-monitoring technology, but in a move reminiscent of the SCDOR breach aftermath, OPM officials deflected blame to the Department of Homeland Security, NIST, the Office of Management and Budget, and other agencies that they felt had hindered their ability to implement security upgrades or establish clear expectations for information security (Chaffetz et al. 2016).

Public Shaming Incidents

Security incidents aimed at publicly shaming the victims often involve publicly denouncing or embarrassing the target before as large an audience as possible. This broadcasting stage can be an especially tricky element of cyberattacks to regulate because such regulations may resemble speech or press restrictions that run counter to many countries' fundamental principles.

The Sony Pictures Breach (2014)

When North Korea breached Sony Pictures in 2014 and released large volumes of stolen emails and internal records, several people, including lawyers hired by Sony, suggested that it was, or should be, illegal for reporters to write about the stolen information because it supported the mission of the attackers and infringed on Sony's intellectual property (Boies 2014). While it was true that the widespread media attention to the breach aided North Korea's supposed mission of humiliating Sony Pictures and undermining the company's business, regulating the media and online intermediaries that helped distribute the stolen information would have been a problematic solution to a cyberattack that aimed to spread discord and wreak havoc on its target.

Spamhaus DDoS (2013)

By contrast, the massive distributed denial of service (DDoS) attacks against Spamhaus in 2013, by a group frustrated by the organization's widely used spam blocklists, did not rely on media coverage or popular attention. Rather, they depended on DNS operators that had failed to restrict their servers to only resolve queries from computers in their administrative domains. These open resolvers enabled the Spamhaus attackers to send DNS queries that pretended to be from Spamhaus servers, causing the open DNS servers to respond to Spamhaus with large DNS files that bombarded the antispam organization's servers and forced them offline (Prince 2013a).

Not operating open resolvers was a known security best practice for DNS operators but there were inadequate incentives for many of them to adhere to the practice. At the time of the Spamhaus DDoS attacks, the Open Resolver Project estimated there were 21.7 million open resolvers online (Prince 2013b). But it was Spamhaus, not the open resolver operators, that bore the brunt of the DDoS attacks, pointing again to the challenge of incentivizing organizations to implement security standards when that investment does not directly benefit them.

Lessons and Takeaways

Several recurring themes and lessons emerge from the high-profile cybersecurity incidents of the past decade and the failure of best practice recommendations to prevent them.

First, it is often extremely difficult for organizations to navigate the many sets of security best practices both because there are so many and because so few data exist to indicate which are actually effective at preventing or mitigating bad outcomes.

Second, the uncertainty about which best practices to follow creates a loophole of sorts for breached organizations, such as the SCDOR and OPM, to blame other agencies for failing to tell them exactly which security controls they should have been using.

Third, the TJX and SCDOR breaches show that regulators and courts are often reluctant to close the loophole by clarifying the security expectations for firms because they do not want to be responsible for dictating those expectations. Although this approach makes clear that there is no set combination of controls that obviates an organization's liability, it makes it more difficult for organizations to determine which recommendations they should follow.

Fourth, examining the controls that courts, legislative hearings, and media reports most often emphasize in cases such as those discussed in this article reveals a tendency to highlight the absence of controls that might have blocked early, technical stages of the intrusions rather than more systemic interventions that would have involved third-party intermediaries, law enforcement, and/or regulators. This is in part because those controls can be most easily implemented by the breached party rather than requiring cooperation from other stakeholders. But that is exactly what makes them less effective and more likely to address stages of attacks that can easily be substituted for new attack vectors by perpetrators.

Cooperation among stakeholders is needed to make controls more effective.

Conclusions

Improving cybersecurity best practices requires defining clearer guidelines with less onerous implementation, collecting better data on their efficacy, and fighting against the narrative that any given security incident is the result of one particular missing control. More than that, though, it requires a more comprehensive understanding of the entire security ecosystem that these best practices aim to strengthen, so that recommendations for merchants, nonprofits, and government agencies are not developed separately from those for DNS operators, payment processors, or browser manufacturers.

So long as best practices are limited in scope to individual organizations and do not include mechanisms to enable cooperation and support from all stakeholders, they will continue to serve only a narrow function and defend against only a small subset of the stages of cyberattacks.

References

Boies D. 2014. RE: Your possession of privileged and/or confidential information stolen from Sony Pictures Entertainment. Deadline, Dec 15.

Chaffetz J, Meadows M, Hurd W. 2016. The OPM Data Breach: How the Government Jeopardized Our National Security for More than a Generation. House Committee on Oversight and Government Reform, 114th Congress. Washington.

Hoogstraaten H, Prins R, Niggebrugge D, Heppener D, Groenewegen F, Wettinck J, Strooy K, Arends P, Pols P, Koupprie R, and 3 others. 2012. Black Tulip report of the investigation into the DigiNotar Certificate Authority breach. Delft: Fox-IT BV.

Mandiant. 2012. South Carolina Department of Revenue: Public Incident Response Report. Alexandria VA.

Mandiant. 2013. APT1: Exposing One of China's Cyber Espionage Units. Alexandria VA.

Prince M. 2013a. The DDoS that knocked Spamhaus offline (and how we mitigated it). Cloudflare Blog, March 20.

Prince M. 2013b. The DDoS that almost broke the internet. Cloudflare Blog, March 27.

Schwirtz M, Goldstein J. 2017. Russian espionage piggybacks on a cybercriminal's hacking. New York Times, March 12.

Verini J. 2010. The great cyberheist. New York Times, Nov 10.

Verizon. 2019. Data Breach Investigations Report. New York.

Wolff J. 2016. Perverse effects in defense of computer systems: When more is less. Management Information Systems 33(2):597–620.

Wolff J. 2018. You'll See This Message When It Is Too Late: The Legal and Economic Aftermath of Cybersecurity Breaches. Cambridge: MIT Press.

Wolff J, Lehr W. 2018. Roles for policy-makers in emerging cyber insurance industry partnerships. TPRC46: Research Conference on Communications, Information and Internet Policy, Sep 21, Washington.

Zetter K. 2014. Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon. New York: Crown Publishing Group.

We believe it is critical to involve and work with our community as much as possible to support the mission of the university.

Cybersecurity in Higher Education:
One University's Approach

Christian Hamer is chief information security officer at Harvard University.

Christian Hamer

Universities tend to be complicated organizations, diverse and decentralized, with a number of business units. The analogy of a "miniature city" is appropriate: Universities have housing, dining, retail operations, hospitals, power and other utility plants, and even police forces with real jail cells. These operations come with their own information technology systems, risks, and compliance regimes, all of which are in addition to the student, research, and administrative data that are core to supporting the institution's mission of teaching and research.

Universities thus have a vast array of systems and data that must be protected. In addition, faculty or administrators may be targets for cyberattackers because of the research they do, the organizations with which they are affiliated, or unpopular opinions that they share publicly. The challenge is to apply appropriate protection to the data, systems, and people without obstructing the school's mission.

Types of Cyberattacks and Cyberattackers

Cyberattacks can be categorized into two types: unauthorized access and business disruption. Unauthorized access may result in the theft or inappropriate exposure of confidential information. Business disruption usually involves rendering systems or data inaccessible or altering them in some way.

Cyberattackers can be categorized into four types, based on their motivations. Nation-state attackers are interested in espionage and/or disruption. Cybercriminals are interested in financial gain. "Hacktivists" and terrorists are motivated by their cause. And hobbyists are in it for their own entertainment. While there are some groups that blur the lines a bit, especially in the nation-state and cybercriminal areas (e.g., some nation-states have pursued financial gain via cyberattacks), most attackers fit into one of these categories.

Historically, nation-states have had access to the most sophisticated tools and techniques, followed by cybercriminals. However, advanced tools are becoming more available to the general public, making it easier for novices to carry out attacks that once required a high level of sophistication.

Elements of a University Cybersecurity Program

Universities may be subject to each type of attack and targeted by all of the groups. Furthermore, there are constantly new threats to cybersecurity: every day new vulnerabilities are discovered, as are attacks that exploit them (often with confusing or scary names, such as "dragonblood") and new groups thought to be behind them. It can be very difficult to know what to worry about, let alone keep up with the latest attacks.

Advanced tools are becoming more available to the public, making it easier to carry out attacks that once required a high level of sophistication.

In cybersecurity efforts to defend a university's data, systems, and people, the idea that "one size can fit all" is unrealistic. But the problem can be made manageable by breaking the complexity into simpler subcomponents (e.g., based on type of attack or attacker), using consistent, repeatable processes to understand risk, using frameworks to guide the selection of controls, focusing on the basics, and involving the whole community in the process.

First, a risk-based approach is appropriate. What does such an approach mean or look like in practice? At Harvard, we have created an annual process to gather information about risks from numerous sources across the university: the Institutional Risk Management program, a self-assessment process, our faculty and research communities, and a team that monitors and responds to cybersecurity incidents. We analyze the information and use it to evaluate and update, as appropriate, our service catalogue and our project roadmap each year.

We use an industry-recognized framework to guide our program. There are many available; the NIST Cybersecurity Framework, ISO 27000 series, and COBIT are among the most popular. Harvard uses the NIST Cybersecurity Framework (CSF),[1] which covers five functions: identify, protect, detect, respond, recover. To help prioritize which program elements to address first, we use the Center for Internet Security's Critical Security Controls (CIS Controls),[2] which map to a subset of the CSF functions.

Whichever framework one chooses, it is critical to do the basics first. While that statement may sound obvious, it can be easy to get lost amid the challenges of the latest attack or product and ignore the importance of cybersecurity "blocking and tackling": understanding the computer and software assets that must be protected, keeping them protected and patched regularly, and ensuring strong identity and access management processes.

Finally, we believe it is critical to involve and work with our community as much as possible. We view our role as enabling the mission of the university. Therefore, we need to devote time to ensure that we understand what our partners in the community are trying to accomplish, what their challenges are, and how we can help them meet their goals securely. This approach is much more effective than being perceived as a barrier, and is critical to successfully partnering with our community.

Solutions That Work

Among hundreds of products, technologies, projects, and initiatives, three stand out as the most effective in our environment: multifactor authentication, endpoint detection and response, and end user awareness. I will describe each, explain why we implemented it at Harvard, and report what results we observed. Every

[1] https://www.nist.gov/cyberframework

[2] https://www.cisecurity.org

environment is different, of course, and these may not be appropriate in other environments, but they have proven very effective in ours.

Multifactor Authentication

What Is It? Why Use It?

Three "factors" can be used to authenticate the user of a computer or application: "something you know" (e.g., a password), "something you have" (e.g., a phone or token), and "something you are" (e.g., a fingerprint). Multifactor authentication (MFA) requires two or more of these factors for a user to log in to a computer or application. Multiple versions of the same factor (e.g., security questions in addition to a password—both "something you know") are not, by definition, MFA.

Over the past couple of decades, time-based one-time password tokens that generate a different number every minute were the most well-known implementation of "something you have," used with a password ("something you know") to provide a second authentication factor. Today, mobile devices are often used instead of a token, through a smartphone application or even a text message.

At Harvard, we implemented MFA because we had a problem with accounts compromised by phishing, password reuse, and brute force (a trial-and-error method of password guessing). While there are means to address each of these issues, we didn't think any of them would be as effective or efficient as adding a second factor to passwords for authentication.

How to Get User Buy-In

Implementation of MFA was a big change for our community, so we strongly emphasized the user experience. This focus started with the name: We called the new approach *two-step verification* because that's the term Google had used and it was familiar to at least some in our community. There was no reason to introduce another technical term when we could use a familiar one.

The next challenge was with onboarding. During user testing with the native experience offered by our vendor, it became clear that it was not simple enough and would not allow us to deploy at scale without a negative impact on the community. So our Identity and Access Management team developed a simpler and easier interface that they integrated with the vendor's product. Once we had simplified the experience as much as possible, we invested a lot of time and effort in clear and complete documentation in a variety of formats, from checklists to videos.

Our last challenge was to make sure that our community understood both why the new approach was important and that it would be easy for them to use. Most of the community wasn't aware of the attacks we were experiencing and, if they were aware, didn't think they could or would be a target. We changed that by relating anonymized stories of real incidents that had happened to others on campus, and made it clear what might be at stake for them personally. Then we talked about how easy and unobtrusive the solution was, addressing some of the perceived pain points we had heard about.

Finally, we brought the solution to the community in the form of clinics across the university to ensure that those who wanted help with onboarding had very easy and convenient access to it. In addition to the clinics, we provided documentation and support. We onboarded some 65,000 members of our community in about seven weeks.

People generally want to do the secure thing and the best way to leverage that intent is to make it easy for them to do so.

Our success was largely due to our focus on the user experience. This project demonstrates that people generally want to do the secure thing and the best way to leverage that intent is to make it as easy as possible for them to do so.

The problem we set out to solve was that of compromised accounts. The results have been exactly as we hoped: In the nearly three years since we rolled out MFA, the number of protected accounts that have been compromised is almost zero.

Endpoint Detection and Response

Detection Challenges

Like most organizations of a few years ago, we relied on various configurations of antivirus and antimalware to protect our systems and as a means to detect intrusions. While these technologies were never perfect, they seemed adequate for detecting simple attacks. However, most of them are based on the concept of "enumerating

badness": they use cryptographic signatures of known malicious software as a means of detection. Such detection requires that someone has already identified the software as malicious. It also means that the databases of known malicious software will continue to grow at an accelerating rate.

This approach provides attackers with at least two opportunities. The first is to slightly modify their malicious software so that its signatures go from "known malicious" to "unknown." In fact, there is a category of software (called "packers") that obfuscates malicious software so that it will not be detected by signature-based tools.

EDR provides critical early warning of malicious activity and allows us to take steps to stop it.

The second opportunity for attackers is to "live off the land": to use the tools and accounts of legitimate system administrators (used to operate and maintain systems) to attack the systems. This is a practice favored by more sophisticated attackers and against which traditional antivirus and antimalware are completely ineffective.

Our Experience with EDR

We confronted a series of attacks that used both of these techniques and quickly realized that we lacked visibility into when or where the attacks were happening. At the same time, new endpoint detection and response (EDR) products were evolving. These products analyze the behavior of software, not just its cryptographic signature, and look for anomalies and patterns that are indicative of malicious activity. That makes them equally adept at catching both malicious software that has been slightly modified and attackers using built-in tools for system administration in malicious ways.

When we deployed EDR on our computer systems, it did two things. First, and most importantly, it gave us the visibility we needed into the kinds of malicious activity described above. Second, it highlighted how poor a job the more traditional protective tools were doing. While EDR doesn't always stop attacks on its own, it provides critical early warning of malicious activity. This allows us to take steps to stop the activity before the attackers can move to their ultimate targets. An important feature is that EDR has no impact on our users unless there is a real security issue.

While EDR may be a point-in-time technology that is replaced in the future, it has earned its place on this list by helping us detect and quickly respond to sophisticated attacks significantly better than any other technology we have tried and without any negative impact on our community.

End User Awareness

End user awareness has long been a controversial topic in cybersecurity circles. People often say that "users are the weakest link," and it is true that many attacks focus on end users ("phishing" may be the most well known of them). But we have found it valuable to treat our community as a resource and a partner rather than consider them "the weakest link."

From User Awareness to Action

We know that our community members want to do the right thing and just need the proper resources to know what the right thing is and how to do it. Our goal is to identify a small set of behaviors that we can influence to significantly increase security outcomes. We work on influencing behavior because awareness without action does no good. We called our campaign "Small Actions, Big Difference" to emphasize the simple things people can do to make themselves a lot more secure.

We adopted the Fogg Behavior Model[3] as a framework for the campaign. According to this model, behavior change requires motivation, ability, and prompts. While the prompts are important—users have to recognize the situation—we have found that the motivation and ability are the components on which we need to focus the most.

To ensure that we provide users appropriate motivation to take a desired action, such as applying updates when they are available, we highlight the benefits to the user. Emphasis on the personal benefits of a particular action has been the most effective approach for us.

Realistic Requests for Action

The "ability" component of influencing behavior addresses two questions, an obvious one and a more subtle one that can be overlooked. The obvious question is "Does the user have the skills and abilities to take

[3] https://www.behaviormodel.org/

the desired action and have we made the process easy enough that s/he can follow it?" This again speaks to the criticality of good user experience in fostering security.

The more subtle question is "Is what we are asking realistic?" For example, the advice "don't click on links in emails or open attachments" may be realistic for some, depending on their job function, but it is not realistic for most people. We always want to make sure we do not provide advice that sounds simple but is effectively impossible to follow.

Asking the community to "click wisely" is an important part of our awareness campaign. Although it will never be 100 percent effective, it is possible to enlist the people who do recognize phishing emails to protect those who do not. When even one person reports these emails to us, we can take technical steps to protect the entire community. Now thousands of messages are forwarded to us each month and fed into an automated processing system to prioritize the most important ones for immediate action.

Talent and Collaboration

Recruitment and retention of talented people are a challenge in cybersecurity. One of the most effective ways to meet that challenge is to look for people who have diverse backgrounds and experiences and can help find novel solutions to problems.

Because universities, given their scope and scale, offer some of the most interesting problems in cybersecurity and have a huge spectrum of systems connected to their networks, they offer a variety of challenges for those interested in cybersecurity. And it's in keeping with the mission of higher education to invest in training and development. The diversity of challenges and focus on professional development are among the benefits that universities can offer information security professionals.

Universities also value collaboration and information sharing. In my experience, those in higher education are willing to take time to help others learn, within the team, across departments, or across universities. Higher education has a formal Information Sharing and Analysis Center (ISAC) that is used to exchange information and best practices. The spirit of collaboration and sharing is one of the things that I and my colleagues most treasure about higher education.

What's on the Horizon?

I mentioned above that multifactor authentication was an effective means to prevent certain unauthorized access attacks. However, there are no "magic silver bullets" in information security, so it is not a surprise to find that this technology has flaws that are starting to be exploited.

MFA implementations that use a phone call or text message are vulnerable to attacks in which criminals trick a mobile phone service provider into moving a number to a different phone. Attacks like these have been responsible for the losses of large sums of cryptocurrencies (Chavez-Dreyfuss 2019). Though less common, attackers have also set up fake portals to capture and "relay" MFA responses, and it is likely these will occur more often. As these attacks evolve, technologies will change to phase out phone calls and emails in favor of more secure means of authentication.

Business disruption attacks typically destroy data and systems, as happened in the well-publicized Sony Pictures Entertainment attack of 2014 (Peterson 2014). Business disruption is also the effect of ransomware, wherein files are locked by an attacker and promised to be returned in exchange for payment, often in cryptocurrency that is difficult to trace.

We work on influencing behavior because awareness without action does no good.

The next iteration of these attacks is likely to be those that alter data in subtle but important ways. Recently, researchers used artificial intelligence (AI) to add the appearance of cancerous regions to 3D CT scans that fooled medical professionals (Mirsky et al. 2019). Such attacks will probably become more common, so the industry needs to increase efforts to protect and prove data and system integrity.

Vendors in the cybersecurity space have been promoting AI and machine learning as the solutions to today's and tomorrow's problems. There are some practical applications of machine learning to analyze large datasets and produce actionable information, and these technologies promise to make cybersecurity analysts much more productive. It is natural, then, to expect that attackers may use the same techniques to make their tools and attempts more effective.

Conclusion

The challenge of information security in higher education is to protect the critical data, systems, and people in support of—and without obstructing—the mission of teaching and research. To confront that challenge, the information security team at Harvard tries to simplify the complexity, use a repeatable risk-based approach to identify and mitigate risks, and use a framework to organize and guide controls, focusing on the basics and involving the community in the process. One of the most important contributing factors to our most successful projects has been ensuring a good user experience.

With the pace of change in cybersecurity continually increasing, and new threats always on the horizon, the path to success is in simplifying the landscape, making good security easier, and identifying and executing on the most important priorities.

References

Chavez-Dreyfuss G. 2019. US investor awarded $75 million in cryptocurrency crime case. Reuters, May 10.

Mirsky Y, Mahler T, Shelef I, Elovici Y. 2019. CT-GAN: Malicious tampering of 3D medical imagery using deep learning. 28th USENIX Security Symposium, Aug 14–16, Santa Clara. Online at https://arxiv.org/pdf/1901.03597.pdf.

Peterson A. 2014. The Sony Pictures hack, explained. Washington Post, Dec 18.

Efforts to address cybersecurity challenges require theoretical perspectives, practical and organizational approaches, and policy understanding.

Policy Dimensions of Cybersecurity Engineering Challenges

Fred B. Schneider and
Lynette I. Millett

Fred B. Schneider

Lynette I. Millett

Computing systems are increasingly subject to attack in support of criminal activities as well as nation-state espionage, sabotage, and now information-influence campaigns. New technical developments are providing ingredients for better defenses, but there is still much work to be done. However, a focus only on the technical aspects of cybersecurity will produce solutions that are unlikely to be effective in practice. Organizational, social, economic, and policy issues must be considered in concert with any technical developments.

Case studies are one way to understand the critical role of nontechnical issues in what, at first, appear to be technical matters. Three are discussed in this paper: quantum-resistant encryption (NASEM 2017), managing data breaches (NASEM 2016), and microarchitectural vulnerabilities in processors (NASEM 2019a). Each was the subject of a workshop organized by the National Academies' Forum on Cyber Resilience (www.cyber-forum.org).

A Goal: Trustworthy Systems

For a system to be considered trustworthy, stakeholders require a basis to believe that a system will do what is expected and not do the unexpected.

Fred Schneider (NAE) is the Samuel B. Eckert Professor of Computer Science in the Department of Computer Science at Cornell University and chair of the Forum on Cyber Resilience, Computer Science and Telecommunications Board, National Academies of Sciences, Engineering, and Medicine. Lynette Millett is director of the forum.

Cybersecurity is the dimension of trustworthiness concerned with resisting attacks. Expected and unexpected behaviors here are often characterized broadly in terms of three classes of properties:

- *Confidentiality*. Which principals are allowed to learn what information or, equivalently, what must be kept secret?
- *Integrity*. What changes to information and uses of resources are allowed?
- *Availability*. When must inputs be read and outputs produced?

Whether a system satisfies these properties depends on the soundness of assumptions made by designers, developers, and users. Moreover, by making additional assumptions about the deployment environment, it generally becomes easier to build a system, although assumptions can have cybersecurity consequences. For example, if a component behaves as expected only while an assumption holds, then an attacker can cause disruption by falsifying the assumption.

For a system to be considered trustworthy, it must do what is expected and not do the unexpected.

Other cybersecurity consequences of assumptions are more subtle. For instance, to express expectations about system behaviors in terms of operations and interfaces (as is common) is to make an implicit assumption that an attacker's actions are limited to certain restricted modes of access for controlling the system or for learning its state. In summary, assumptions create room for vulnerabilities.

The design of a secure system can be viewed as an exercise in identifying assumptions and assessing whether they reflect reality and whether they can be violated by motivated adversaries with certain capabilities. Security mechanisms can create environments where stronger assumptions hold (e.g., because a broader range of behaviors are precluded or handled appropriately) and, therefore, some component runs in a more benign setting than it would have absent those mechanisms. The improved environment created by such security mechanisms does come at some cost, though—the mechanisms themselves need to be secure and additional assumptions might have to hold to ensure their correct operation.

With this approach, design choices can be justified, as befitting an engineering discipline. As in other fields of engineering, design will be an iterative process, deemed finished only when the remaining assumptions are (believed) difficult or unlikely to be violated.

Quantum-Proof Encryption and Cryptoagility

Nearly all computing systems in use today rely on cryptosystems to protect confidentiality and integrity of information, whether that information is "at rest" (i.e., in some form of storage) or "in motion" (i.e., in transit on a network). A few public key cryptosystems are prevalent, and the security they provide depends on an assumption that specific mathematical computations, such as factoring or computing discrete logarithms on sufficiently large numbers, are computationally infeasible given the processing power that may be available to an adversary. Based on current mathematical understanding, this is a safe assumption for digital computers that can be built today and are expected in the future. Even though digital computers are likely to get ever more powerful, they will not have the capacity to factor the large numbers used in public key cryptosystems.

The Threat of Quantum Computers

A sufficiently large general-purpose quantum computer, if it could be built, would invalidate assumptions about the difficulty of performing the mathematical operations that underpin today's widely used public key algorithms. For example, if a powerful quantum computer could be built (NASEM 2019b), any actor could use it to execute Shor's algorithm (Shor 1994) to factor large integers. This would be detrimental to both network communications and the protection of long-term data encrypted and stored using algorithms that depend on the difficulty of factoring large integers. The possibility of breakthroughs in quantum computing is thus cause for concern. Cryptography researchers are developing new public key algorithms that are resistant to cryptoanalysis by adversaries using any combination of classical or quantum computers.

The good news is that it appears that such encryption algorithms can be built and with acceptable performance on modern computers. The bad news is that simply having quantum-resistant encryption algorithms available is not enough. Installed software must use them,

requiring upgrades to an enormous amount of software. And those upgrades will not be straightforward. For one thing, cryptographic keys and encrypted messages will undoubtedly be different sizes with these new cryptosystems, potentially requiring changes to any software that stores keys or manipulates encrypted messages. Protocols to set up encrypted connections or do other provisioning also have to be upgraded.

In short, transitioning the internet ecosystem to quantum-resistant cryptosystems will not simply be a matter of replacing the few routines that perform encryption and decryption; the upgrade likely will require making and distributing changes to all software that uses these routines, too.

Most software systems now in use were designed to allow upgrades after deployment. But few were designed to readily support *cryptoagility*, wherein encryption and decryption routines (and perhaps a small amount of other code) are all that must be updated to replace a cryptosystem.

Previous Cryptosystem Challenges

History offers insight into the costs and challenges of replacing a cryptosystem widely used by the internet ecosystem. The TLS (Transport Layer Security) protocol was standardized in 1999 (and version 1.3 was recently finalized) as a revision of the early SSL (Secure Sockets Layer) protocol; TLS encrypts all types of internet traffic. Yet there are still systems in place that support SSL 2.0. So, even after two decades, some systems have not been fully transitioned. The general consensus among experts is that it takes about a decade to substantially roll out a new cryptosuite and retire the old one.

The experience with TLS suggests that, if usable quantum computers may be achieved within a decade, now is the time to start upgrading to quantum-resistant cryptosystems. The most sensible course would be to upgrade systems to improve cryptoagility, because even if quantum computers never become available, cryptosystems might need to be changed to defend against attacks developed based on new insights in the underlying mathematics or on bugs in the software implementation.

National Security Concerns

Efforts to update the protocols and software underpinning the internet ecosystem bring many technical and operational challenges to software producers and their customers. And if the upgrades are connected to a competitive or national security advantage, then important nontechnical factors come into play, too. Global platform companies—even those based in the United States—respond to the needs and requirements of their customers around the world. Some of these customers may be governments, which may not want to use the same cryptographic tools as others. Other customers may be subject to foreign government regulations aimed at either improving security (e.g., by prohibiting risky practices) or undermining it (e.g., by mandating backdoors or other sorts of access).

The general consensus is that it takes about a decade to substantially roll out a new cryptosuite and retire the old one.

In the interests of its national security, a country may wish to regulate surveillance-resistant cryptosystems or insist on its own locally created cryptosuites—for example, to facilitate law enforcement access to encrypted content or to enable government surveillance of network communications. A platform provider could be caught in one of the following binds if the cryptosystem it uses to protect its software upgrades is not compliant with a nationally mandated cryptosuite:

- For consistency, scale, and to avoid having to trust a customer's government, a company would want to distribute software upgrades using its own cryptography and keep those keys secret. The company would find it unacceptable to use a cryptosystem it did not trust or to provide a government with keys used by its customers' systems, because that would risk losing control of both the integrity of its installed systems and the confidentiality of its intellectual property.
- A software upgrade capability would have to be able to change locally installed cryptosuites. Governments, fearing supply chain attacks, would find that unacceptable, since a (potentially foreign) company would then have control over what cryptography the platform uses and would have the capability to either install backdoors for exfiltrating content or remove government-provided backdoors.

There are thus at least two parties—a company and a government—neither comfortable with the other having needed control. Yet one of them must have that control. In addition, end users and companies that provide networking capabilities may have interests that conflict with each other, with software companies, and/or with governments. A technical solution might satisfy all these constraints, but none is currently known. Most likely, addressing these and other policy and international complications will require not just new forms of cryptoagility but also new thinking about what entities should have jurisdiction over encryption policies and what should be the goals of those policies. New policies and norms formulated by international standards bodies as well as other regulatory and political entities are likely to be part of a solution.

Managing Data Breaches

Two Approaches: Prevention, Recovery

Prevention, which takes the conservative stance of prohibiting actions that could violate some security policy, is one approach to managing the risk of cyberattacks. With the other, recovery, the risk is made tolerable by having the means to mitigate the effects of an attack. Recovery typically involves both technical mechanisms to detect attacks and to return a computing system to a preattack state, and nontechnical means to reverse effects of the attack on the system's environment.

What entities should have jurisdiction over encryption policies and what should be the goals of those policies?

Prevention receives the lion's share of attention in the cybersecurity research community, but the inevitable failure of any defense implies that even prevention-based defenses benefit from implementations of recovery measures.[1] Data breaches and identity theft commonly inspire discussions about recovery, with incident reports like the following now a regular front page feature:

- In early 2019, Marriott International revealed that several million of its customers' unencrypted passport numbers were stolen in a data breach that exposed personal information of nearly 500 million individuals (Gressin 2018).
- Breaches of credit card information at major retailers are now so common that the process for issuing new cards has become routine, often completed even before a consumer is aware of a compromise.

In these examples, recovery would superficially appear to be straightforward to implement. Marriott said that it will cover the cost of new passports when fraud can be shown. Banks incorporate the cost of frequent card replacement into their business model. However, passport replacement does not remedy the harm if the theft was undertaken to track movements of certain passport holders. And there remain unaddressed costs to consumers for credit card disclosures, given the time it takes to correct and update accounts.

The Need for Other Approaches

Identity theft sometimes goes well beyond the disclosure of a single identifying number, and with not just financial but all manner of other harms to individuals and institutions, as illustrated in the following examples:

- The breach of sensitive security clearance information at the Office of Personnel Management (OPM) in 2015[2] caused harm to national security, not least in that it revealed the names and personal networks of individuals with clearance to the attacker.
- Exposure of user identities and profiles in the Ashley Madison data breach in 2015 (Krebs 2015) caused people to suffer humiliation and reputational harm (but not identity theft).
- An Equifax data breach in 2017 revealed the social security numbers and birthdates of nearly 150 million users of the credit bureau, undermining confidence in the company itself and in the entire credit management infrastructure.[3]

[1] The challenges associated with supporting recovery were surveyed in a recent workshop (NASEM 2018).

[2] OPM Cybersecurity Resource Center, Cybersecurity Incidents: What Happened (https://www.opm.gov/cybersecurity/cybersecurity-incidents/).

[3] The US Federal Trade Commission provided guidance for consumers and business owners in a series of blogs, September 2017–June 2018 (https://www.ftc.gov/equifax-data-breach).

- A national database in Singapore containing the names and addresses of more than 14,000 people who had tested positive for HIV was breached in early 2019 and the data were made public (Ives 2019).

In each example, approaches such as replacing credit cards, freezing credit accounts, or even remedying significant identity theft are not sufficient to resolve the harms caused. The OPM example is especially problematic because security clearance applications hold vast amounts of personal and sensitive details, including information about family and social networks. People who provided this information were assured that the government would keep it confidential. Some individual or entity now has a vast trove of information on people in positions of trust with the US government. To remedy that national security harm in any meaningful way may not be possible.

The straightforward, and now reflexive, responses after a data breach—replacing identifier numbers, hardening systems, freezing certain credit accounts, notifying users—are based on an assumption about possible harms. But as should now be clear, that assumption is unsound: The toolbox of easy responses is insufficient to cope with the scale, scope, nature, and duration of potential harms from many major data breaches.

Furthermore, new kinds of data manipulation and information distortion are now being deployed on social platforms, the effects of which may often be individually intangible but globally intractable and socially toxic. The credibility and/or targeting of these postings is enhanced by attacker access to information gleaned from data breaches—data stolen in one breach may be used in another.

Policy, legal, and regulatory approaches are needed to address the myriad ways in which data breaches can cause harm well beyond an individual's credit rating.

Vulnerabilities in Computer Architectures

The recent discovery of the security vulnerability Spectre (Kocher et al. 2018) and others (Canella et al. 2018) showed that software engineers and compiler designers have, for the past two decades, been making unsound assumptions about what computers do—and don't do—when executing programs.

Spectre exploits speculative execution, a microarchitectural processor performance optimization through which instructions are executed for predicted execution paths before it is known with certainty whether the corresponding instructions will be executed. The speculation allows processors to accomplish useful work rather than idling, and this allows for parallel work. Results that are not needed because the program follows a path other than the predicted one are simply discarded.

Unfortunately, instructions that are erroneously executed during speculative execution exhibit side effects by affecting resource availability or changing the state of the processor's cache. Spectre and a related vulnerability named Meltdown (Lipp et al. 2018), both discovered in 2017 and revealed publicly in early 2018, exploit these side effects to deduce private memory contents. Since then, numerous other attacks have been described that take advantage of differences between the architectural properties that programmers assume, the architectural specifications, and the actual microarchitectures of modern processors.

The toolbox of responses is insufficient to cope with the scale, scope, nature, and duration of potential harms from many major data breaches.

Although some limited mitigations have been deployed, satisfactory comprehensive technical solutions for these kinds of vulnerabilities do not exist; research is ongoing. Moreover, there is some urgency in finding mitigations, because speculative execution is an important performance feature of all modern processors, and solutions are likely to require hardware changes that take a long time to deploy. Nearly all modern computing systems are vulnerable to these kinds of hardware-focused attacks, although software bugs are widespread and often easier to exploit, so for the time being, easier-to-exploit avenues for attack are plentiful.

Policy and Practical Challenges

Independent of technical mitigations for existing processors, vulnerability and the orchestration of any associated upgrades pose some formidable policy dilemmas.

- Disclosure of a vulnerability is itself a vexing issue. Who should be notified during the initial embargo

(confidentiality) period, and, more controversially, who should not be informed? The processor manufacturer? Users? Wholesale customers of the manufacturer? Operating system vendors? Software developers? Cloud computing vendors whose products depend on effective isolation? Disclosing to everyone who may have a reasonable need to know can be dangerous because notifying too many people during an embargo period can result in attackers finding out about a vulnerability while defenders are in the dark.

- Who is responsible for the repairs? For vulnerabilities discovered by those working on offense (e.g., intelligence agencies and companies that supply them with zero-day exploits[4]), the "embargo" objectives are inverted. The goal is to maximize the power of the exploit, which includes keeping the bug from being patched.

Notifying too many people about a vulnerability during an embargo period can mean attackers find out about it while defenders are in the dark.

- Because military and government computers use modern processors, vulnerabilities constitute national security risks from a defense perspective as well as opportunities to attack vulnerable systems used by rivals. When should governments be notified? Which governments? There are additional policy challenges when vendors are expected or required to inform a government about vulnerabilities.

A recent workshop on this topic (NASEM 2019a) covered the technical dimensions and explored how researchers, manufacturers, and vendors managed the disclosure for Spectre. The workshop included discussion of the challenge of creating effective vulnerability disclosure processes, especially for vulnerabilities whose mitigations affect multiple layers of the computing stack (e.g., hardware, operating system software, applications, and cloud service implementations) and could therefore require assistance from stakeholders ranging from small software companies to multinational corporations to national governments.

Spectre and its ilk result from the success of computer engineering itself—processor innovations driven by performance needs have quietly and inadvertently undermined basic software engineering assumptions about how data are handled. Those now incorrect assumptions can be exploited, leading to the possibility of an attack that strikes at the heart of the interface between software and hardware.

What Next?

Engineering any artifact involves decisions about trade-offs, and these decisions are made relative to expectations about the environment in which the artifact will be deployed. Changes to the environment make it crucially important to revisit and change a design.

Computing systems and the internet ecosystem are a case in point. In the early days of computing, the environment in which computer systems were deployed was assumed to be relatively benign. Performance, time to market, and features were paramount; security was a concern in only a very small number of cases. But as dependence on computing systems—by nations, communities, and individuals—has grown, so have the threats. Assumptions about the environment that once were acceptable are no longer sound, and can have dangerous consequences. Insofar as assumptions beget vulnerabilities, old designs need to be rethought.

An added challenge is the increasing sophistication and capabilities of adversaries. New assumptions need to be robust against not only today's but also tomorrow's adversaries, because components built today will remain fielded for a long time and will become increasingly difficult to replace.

Security trade-offs are not just technical choices; they have important policy implications, too. Efforts to address underrecognized and multidisciplinary aspects of evolving cybersecurity challenges will require theoretical perspectives (e.g., taxonomies of harm and remedies for data breaches, analyses of what system parameters affect what sorts of recoverability and update capabilities) as well as practical and organizational approaches (e.g.,

[4] A zero-day exploit is a cyberattack that occurs before a software weakness becomes known to the software provider who would be responsible for providing a mitigation.

how does the anticipated lifecycle affect the organizational structure that maintains the system; what mechanisms are in place to ensure an orderly path to update cryptography suites when needed). Moreover, absent a continuing dialogue across these disciplines, attempts at solving cybersecurity problems are likely to fail.

What's needed are people who are not only expert in specific technical or policy matters but also comfortable in discussions concerning the entire spectrum of technical and policy issues and solutions. The National Academies' Forum on Cyber Resilience is one such effort, but there is much ground to cover, so more efforts are needed.

Acknowledgments

Much of this paper is based on discussions in the National Academies' Forum on Cyber Resilience; the authors are grateful to forum members for their thoughtful contributions at those meetings. Additional thanks are due to Paul Kocher, Jon Eisenberg, and Emily Grumbling for their comments on a draft version of this paper. The forum is supported by the National Science Foundation (award number CNS-14194917), the Office of the Director of National Intelligence (award number 10004154), and the National Institute of Standards and Technology (award number 60NANB16D311). We are grateful for their participation and vote of confidence in the work and activities of the forum.

References

Canella C, Van Bulck J, Schwarz M, Lipp M, von Berg B, Ortner P, Piessens F, Evtyushkin D, Gruss D. 2018. A systematic evaluation of transient execution attacks and defenses. arXiv:1811.05441v1.

Gressin S. 2018. The Marriott data breach. Consumer Information blog, Dec 4. Washington: Federal Trade Commission.

Ives M. 2019. Singapore says records for 14,200 HIV patients, held by an American, were leaked. New York Times, Jan 28.

Kocher P, Genkin D, Gruss D, Haas W, Hamburg M, Lipp M, Mangard S, Prescher T, Schwarz M, Yarom Y. 2018. Spectre attacks: Exploiting speculative execution. arXiv:1801.01203.

Krebs B. 2015. Extortionists target Ashley Madison users. Krebs on Security, Aug 21.

Lipp M, Schwarz M, Gruss D, Prescher T, Haas W, Fogh A, Horn J, Mangard S, Kocher P, Genkin D, and 2 others. 2018. Meltdown: Reading kernel memory from user space. Proceedings, 27th USENIX Security Symposium, Aug 15–17, Baltimore.

NASEM [National Academies of Sciences, Engineering, and Medicine]. 2016. Data Breach Aftermath and Recovery for Individuals and Institutions: Proceedings of a Workshop. Washington: National Academies Press.

NASEM. 2017. Cryptographic Agility and Interoperability: Proceedings of a Workshop. Washington: National Academies Press.

NASEM. 2018. Recoverability as a First-Class Security Objective: Proceedings of a Workshop. Washington: National Academies Press.

NASEM. 2019a. Beyond Spectre – Confronting New Technical and Policy Challenges: Proceedings of a Workshop. Washington: National Academies Press.

NASEM. 2019b. Quantum Computing: Progress and Prospects. Washington: National Academies Press.

Shor PW. 1994. Algorithms for quantum computation: Discrete logarithms and factoring. Proceedings, 35th Annual Symposium on Foundations of Computer Science, Nov 20–22, Santa Fe.

The massive scale and decentralized nature of the IoTT provide attackers with a large attack surface for exploitation.

Raising Awareness of Security Challenges for the Internet of Trillions of Things

John A. Stankovic and Jack Davidson

John A. Stankovic

Jack Davidson

Computer security problems have evolved over the last 50 years from a minor concern to major operational risks. Every day new devices are added to the Internet of Things (IoT). Conservative projections have 50 billion devices on the internet by 2020, but—with autonomous vehicles, smart phones, smart wearables, smart cities, numerous other smart applications, and nanotechnology—we foresee an "Internet of Trillions of Things (IoTT)" before long. If computer security problems are formidable now, consider when there is an IoTT!

The proliferation of devices and applications will give rise to many new complications and research challenges, especially in cyberphysical system (CPS) security. Because the smart devices of IoTT systems will be so numerous and easily accessible, and will interact directly with the physical world (including humans), they will exhibit tremendously large attack surfaces with increasing types and numbers of vulnerabilities. Attacks on these systems may cause the inoperability of major infrastructures such as transportation or energy, great financial losses, and many other negative impacts, even death.

CPS technology is required to build IoTT systems. It is therefore necessary to increase awareness of CPS security problems and to address them

John Stankovic is the BP America Professor and Jack Davidson is a professor, both in the Department of Computer Science at the University of Virginia.

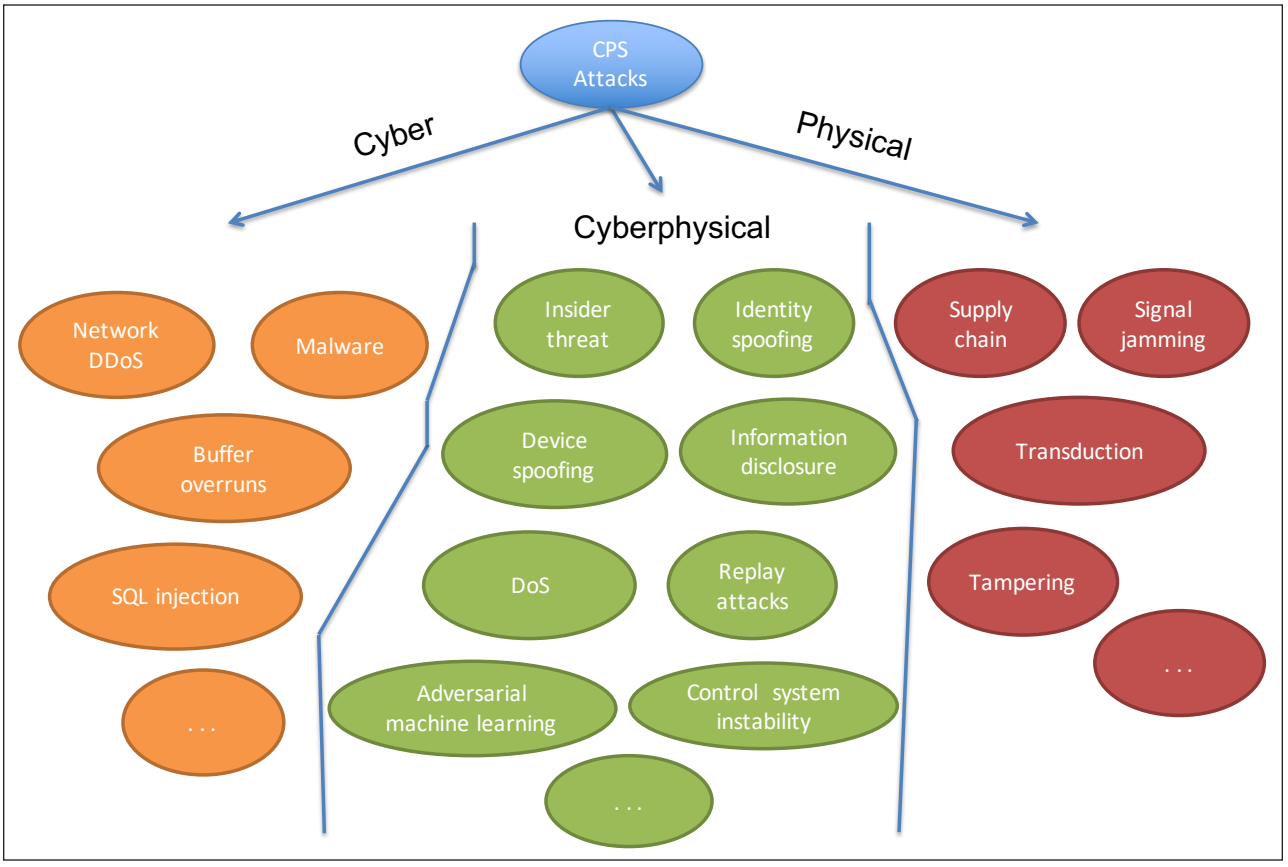

FIGURE 1 Taxonomy of sample attacks on cyberphysical systems (CPS). DDoS = distributed denial of service; DoS = denial of service; SQL = structured query language.

before trillions of devices are deployed with insufficient security protections.

In this paper we define a CPS security attack taxonomy based on new problems for the cyber, physical, and cyberphysical aspects of the IoTT, and discuss examples of new security problems for the physical and cyberphysical areas. For each example, we explain potential consequences and present approaches to solutions. We do not discuss cyberattacks, such as network distributed denial of service and malware, as they are well covered by existing literature.

CPS Attack Taxonomy

Cyberphysical systems are engineered systems that are built from, and depend on, the seamless integration of computational and physical components. They are often connected to the internet, with applications in many domains, such as smart cities, health care, transportation, energy, emergency response, agriculture, and defense. CPS attacks can thus have serious and harmful physical results (e.g., an air bag being activated during normal driving). This characteristic sets them apart from traditional IT systems, where confidentiality is usually the most important property.

Figure 1 illustrates a taxonomy of CPS attacks. It does not show a complete range of attacks but rather provides examples in three categories: cyber, physical, and cyberphysical.

At the cyber layer, the IoT is susceptible to many common types of attacks, such as malware and distributed denial of service. A particularly serious concern with the IoTT will be the timely application of critical patches and system updates, which often require temporarily disabling system security protections. Given the massive scale and decentralized nature of the IoTT, this update process provides an attacker with a window of vulnerability and a large attack surface for exploitation.

Cyberphysical attacks combine software intrusion/alteration with effects on the physical aspects of a system. Attacks by insiders, information disclosure, replay, and denial of service (DoS) have been common and can now also be applied to the physical aspects of

the IoT. Identity and device spoofing and control system instability are newer attacks, created—or significantly increased in frequency—with the emergence of the IoT.

Many novel types of attacks have appeared at the physical layer of systems. These include ways to disrupt systems by attacking different steps in the supply chain. Attacks based on transduction (discussed below) are particularly debilitating. Physical tampering is a risk because smart devices often operate in open environments. And because most smart devices communicate wirelessly, they are vulnerable to jamming.

Physical Attacks on the IoT

Supply Chain and Other Tampering

Many security solutions assume that hardware is trustworthy, but with the common practice of outsourcing the construction of hardware platforms, the supply chain can be a source of security attacks (Ray et al. 2018). Trojan horse circuits and embedded software might be included in delivered products to cause harm or surreptitiously transmit data to an adversary. For example, an internet router might not only transmit packets to the intended destination but also send copies to the adversary.

Many companies automatically collect data for maintenance and performance control. Such data may be a source of security attacks.

Compounding the problem is that many companies, such as those involved in communications and the manufacture of printers, automobiles, and aircraft, want to (or already do) automatically collect data for maintenance and performance control. Such data may be a source of security attacks.

Supply chain attacks can affect all parts of society and almost all applications: financial records and company secrets can be stolen, automobiles and planes can be made to crash. New tools are needed to validate that delivered hardware/software platforms do not include hidden circuits or embedded software that can cause attacks.

With billions (or more) of IoT devices and easy access to them, other types of tampering are also possible. For example, an attacker can physically move sensors to an unwanted location, point a fixed-direction camera in the wrong direction, impede an actuator from its full range of motion, or jam wireless communications. These changes may result in fires not detected, missed detection of a serious crime at a previously monitored location, safety doors that don't close properly, or complete inaccessibility of an IoT application system.

Solutions for such attacks must be developed or improved. Unwanted movement of devices could be detected with additional motion sensors or accelerometers. Correlation among data from a set of sensors could reveal that one sensor has been moved away from the others. Related sensing modalities (e.g., temperature, pressure, and volume sensors relate to each other by physics) could be used to detect attacks on one modality. And frequency hopping, spread spectrum, and other techniques can provide some resilience to the jamming of wireless communications.

Transduction Attacks

One especially complex class of attacks for smart devices is transduction attacks, which exploit the physics and unintended functions of circuits and sensors to alter a sensor's output (Fu and Xu 2018).

Manipulation through Voice Recognition

The Dolphin Attack (Zhang et al. 2017) uses an inaudible sound wave to trick a speech recognition system such as Siri, Google Now, or Alexa into taking action that was not requested. It takes advantage of the fact that, while microphones are built to primarily hear the human voice, they can also detect (unintended) inaudible sounds, making them vulnerable to transduction attacks (Roy et al. 2017).

Because more and more IoT systems have or are developing voice interfaces, the consequences of transduction attacks are unbounded. The attacks may permit illegal entry to a location, open the door locks of a home or business, or provide harmful advice via a medical cognitive assistant.

Solutions must include better frequency filters and signal processing to avoid the appearance of ultrasound resonances in the voice frequency range.

Backdoor Coupling

Another type of transduction attack is called backdoor coupling: signals enter a system indirectly via coupling between wires or components, so a sensor designed to detect one modality may be activated by another. For example, it was shown that playing sounds embedded in a YouTube video allowed an adversary to control a smartphone's accelerometer, thanks to a mechanical coupling between the speaker and the resonant frequency of the sensor (Fu and Xu 2018). This can have negative consequences for any application that relies on the accelerometer, such as a step counter, dead reckoning location estimators, or a sensor that monitors an elderly person's level of activity.

Most security solutions are ineffective for transduction attacks because they are designed for digital risks, not analog. Circuits must be manufactured to reduce the effects of resonance, increase the frequency of checking sensor output by software, and enhance the layout of components on system boards to minimize unwanted coupling.

Cyberphysical Attacks

Types of security attacks that are exacerbated in the IoT are denial of service, spoofing, adversarial machine learning, and control system attacks.

Denial of Service

A DoS attack diminishes or eliminates a system's capacity to perform its expected function. Hardware failures, software bugs, resource exhaustion, environmental conditions, or any complicated interaction between these factors can cause a DoS. When a bug is detected attackers often use the internet to initiate DoS attacks.

As an example of a DoS in communications, attackers may induce a collision in only one byte of a transmission to disrupt an entire packet. A naïve link-layer implementation may attempt retransmission repeatedly, culminating in the exhaustion of batteries in one or more connected smart devices and disrupting system function, such as transportation monitoring in a smart city.

Solutions must detect such attacks and create power-aware smart devices that avoid using all the power when under attack.

Spoofing

In a spoofing attack, a person or program successfully masquerades as another to gain an illegitimate advantage. With the IoTT it could also be a smart device.

The attacking device can act as a data source, presenting fake data streams to the system, or even pretend to be a particular person and issue commands as if it were that person. Spoofing can cause an IoT system to produce wrong, often safety-critical, results. For example, an IoT-based process control plant may be erroneously informed that chemical vats are overheating and shut them down, causing loss of revenue, or indicate that they are operating well when they are not, resulting in overheating and even explosion.

Adversarial machine learning can cause misidentification of a stop sign, with severe consequences for a self-driving car.

Solutions require standard techniques such as authentication, encryption, and anomaly detection, as well as IoT-specific measures that coordinate among properly acting smart devices. The use of redundant sensors that are not accessible through the internet may also prove useful.

Adversarial Machine Learning

Machine learning (ML) is essential to the functionality of many cyberphysical systems. For example, autonomous vehicles use deep neural networks (DNNs) to detect, identify, and locate objects in the environment and navigate the vehicle. Deep learning algorithms are also used to analyze and recognize speech for voice-activated cyberphysical systems and to recognize people and objects in security systems.

The extensive use of ML algorithms has enabled a new type of CPS attack: adversarial machine learning. Attackers can develop adversarial inputs with small (even imperceptible) perturbations that cause a trained ML model to misclassify an object, such as a road sign (Eykholt et al. 2018); misidentification of a stop sign, for example, could have very serious consequences for a self-driving car.

Adversarial ML attacks can also target speech recognition systems (Carlini and Wagner 2018) and image recognition systems (Kurakin et al. 2017). A targeted

attack causes the ML system to assign a specific (i.e., targeted) label to the object; for example, an attack on a voice-activated command and control system would enable a command of the attacker's choosing. Similarly, a targeted attack on a facial recognition system that matches faces against a whitelist of approved people (e.g., to control access for large events) could admit an unknown, potentially malicious person.

Many cyberphysical systems (e.g., self-driving cars, voice-activated command and control systems) can easily be acquired legally or illegally by attackers to carry out black box attacks. A black box attack is defined as one that does not have direct access to the underlying ML model (as in a white box attack) to develop sophisticated attacks; the attacker can only deduce the CPS operation by providing specific inputs and observing the output. Black box attacks on complex systems are more difficult than white box attacks, but they have been demonstrated in a number of domains.

A chemical plant control system that becomes unstable might rapidly open and close a critical valve, causing it to fail.

Adversarial ML attacks are a relatively new CPS threat. Research is needed to understand them and to build robust ML models that are not susceptible to adversarial manipulation of the physical artifacts (e.g., signs, images, audio) that are inputs to the system.

Control System Attacks

Control systems are a critical component of many types of cyberphysical systems. Examples include industrial control systems, supervisory control and data acquisition systems, autonomous vehicles, and medical devices.

Instability and Physical Damage

In a control system attack, the adversary seeks to move the system from a region of stability to one of instability where control outputs may fluctuate arbitrarily and exceed normal operating parameters. For example, gain scheduling is often used to control nonlinear systems. Essentially, the system is controlled by a family of linear controllers or gains, each of which is designed for a particular operating region. Gain scheduling attacks can be effected using techniques such as sensor spoofing and denial of service.

Consider an autonomous aerial vehicle (UAV). It uses carefully constructed gains for operating modes such as takeoff, landing, cruising, or hovering. By tampering with the inputs, an attacker can cause a transition from a gain that is appropriate for the UAV's operating mode to one that is inappropriate. For example, when the UAV is hovering, a change in the gain computed for cruising could result in loss of the UAV.

The ability to spoof a sensor reading or to delay receipt of a signal opens the possibility of a control system instability attack. Many control systems are designed assuming that sensor readings are within certain operating thresholds and that the communication channel to send both data and control signals operates as intended. By sending carefully constructed inputs using a replay attack, the adversary may make the control system unstable—and the system could enter a state where returning to a stable state is not possible.

Similarly, using DoS techniques, an attacker could cause a control system to become unstable by delaying packets that contain control information needed to stabilize a system. The instability could result in severe oscillations that could cause physical damage. For example, a chemical plant control system that becomes unstable might begin rapidly opening and closing a critical valve, causing it to fail.

Unlike purely cyber systems, cyberphysical systems open the door to attacks that cause severe physical damage—on par with the damage caused by kinetic weapons. However, unlike kinetic attacks, CPS attacks can be stealthy and precise identification of the attacker is often difficult. These attacks require a deep understanding of the physics of both the process being controlled and the logic that controls the equipment. A well-publicized example of a CPS control system attack is the Stuxnet attack on the Iranian Natanz enrichment facility (Langner 2013).

Sophisticated stealthy attacks on physical infrastructure can also exploit the physics of the process being controlled. A compelling example is an attack on an industrial pump in which the attack payload is a stream of cavitation bubbles created via malicious control of an upstream valve. Over time the stream of

bubbles will pit the pump's impellers and eventually cause the pump to fail (Krotofil 2017).

Protective Approaches

The attacks described in this section rely on malicious inputs that disrupt or hijack the control system. Standard software engineering techniques, such as rigorous testing, can and do reduce the threat surface, but they often do not provide the necessary coverage—especially for a complex system with a variety of sensors and actuators.

An approach called *fuzzing* shows promise in uncovering CPS vulnerabilities not found by traditional techniques. In fuzzing, inputs are automatically generated to force coverage of unexplored code (Miller et al. 1990). The technique is particularly applicable to cyberphysical systems where the range of possible inputs is difficult to enumerate or bound.

For stealthy physical attacks, solutions include redundant sensors and consistency checks to detect a deteriorating system.

Summary

The CPS-based IoTT presents an enormous increase in potential attack surfaces. Many of these systems will interact with humans, further expanding the attack surface and resulting in significantly more vulnerabilities and potential negative impacts on society. If the past is any harbinger of the future, the security attack-solution competition will continue. Highly inventive attackers will exploit the physical, cyberphysical, and cyber layers to their advantage. Developers of smart devices and smart applications must be aware of new classes of potential attacks and build solutions for them as first principles, not only after problems are uncovered.

Some new solutions are promising, but of course they will be useful only if actually implemented. Diverse techniques have proven helpful in computer security and should also help in the IoTT. However, too often speed to market, cost in dollars, or the benefits of homogeneity keep the development of devices and systems from incorporating known security solutions. If these conditions persist and there are trillions of smart devices, there may be widespread chaos and increased risks of physical danger.

References

Carlini N, Wagner D. 2018. Audio adversarial examples: Targeted attacks on speech-to-text. Proceedings, 2018 IEEE Security and Privacy Workshops, May 21–24, San Francisco.

Eykholt K, Evtimov I, Fernandes E, Li B, Rahmati A, Xao C, Prakash A, Kohno T, Song D. 2018. Robust physical-world attacks on deep learning visual classification. Proceedings, 2018 IEEE/CVF Conf on Computer Vision and Pattern Recognition, Jun 18–22, Salt Lake City.

Fu K, Xu W. 2018. Inside risks: Risks of trusting the physics of sensors. Communications of the ACM 61(2):20–23.

Krotofil M. 2017. Evil bubbles or how to deliver attack payload via the physics of the process. Black Hat USA, Jul 22–27, Las Vegas.

Kurakin A, Goodfellow IJ, Bengio S. 2017. Adversarial examples in the physical world. 5th International Conf on Learning Representations, Apr 24–26, Toulon.

Langner R. 2013. To Kill a Centrifuge: A Technical Analysis of What Stuxnet's Creators Tried to Achieve. Dover DE: Langner Group.

Miller B, Fredriksen L, So B. 1990. An empirical study of the reliability of UNIX utilities. Communications of the ACM 33(12):32–44.

Ray S, Peeters E, Tehranipoor MM, Bhunia S. 2018. System-on-chip platform security assurance: Architecture and validation. IEEE Proceedings 106(1):21–37.

Roy N, Hassanieh H, Choudhury RR. 2017. BackDoor: Making microphones hear inaudible sounds. MobiSys, Jun 19–23, Niagara Falls.

Zhang G, Chen Y, Ji X, Zhang T, Zhang T, Xu W. 2017. DolphinAttack: Inaudible voice commands. Proceedings, ACM Conf on Computer and Communications Security, Oct 30–Nov 3, Dallas.

We describe cyberattack surfaces and potential solutions for securing connected and automated vehicles and transportation infrastructure.

Security of Connected and Automated Vehicles

Mashrur Chowdhury, Mhafuzul Islam, and Zadid Khan

Mashrur Chowdhury

Mhafuzul Islam

Zadid Khan

The transportation system is rapidly evolving with new connected and automated vehicle (CAV) technologies that integrate CAVs with other vehicles and roadside infrastructure in a cyberphysical system (CPS). Through connectivity, CAVs affect their environments and vice versa, increasing the size of the cyberattack surface and the risk of exploitation of security vulnerabilities by malicious actors. Thus, greater understanding of potential CAV-CPS cyberattacks and of ways to prevent them is a high priority.

In this article we describe CAV-CPS cyberattack surfaces and security vulnerabilities, and outline potential cyberattack detection and mitigation strategies. We examine emerging technologies—artificial intelligence,

Mashrur Chowdhury is the Eugene Douglas Mays Professor of Transportation, director of the Center for Connected Multimodal Mobility (C^2M^2), and codirector of the Complex Systems, Analytics and Visualization Institute, and Mhafuzul Islam and Zadid Khan are PhD students, all in the Glenn Department of Civil Engineering at Clemson University.

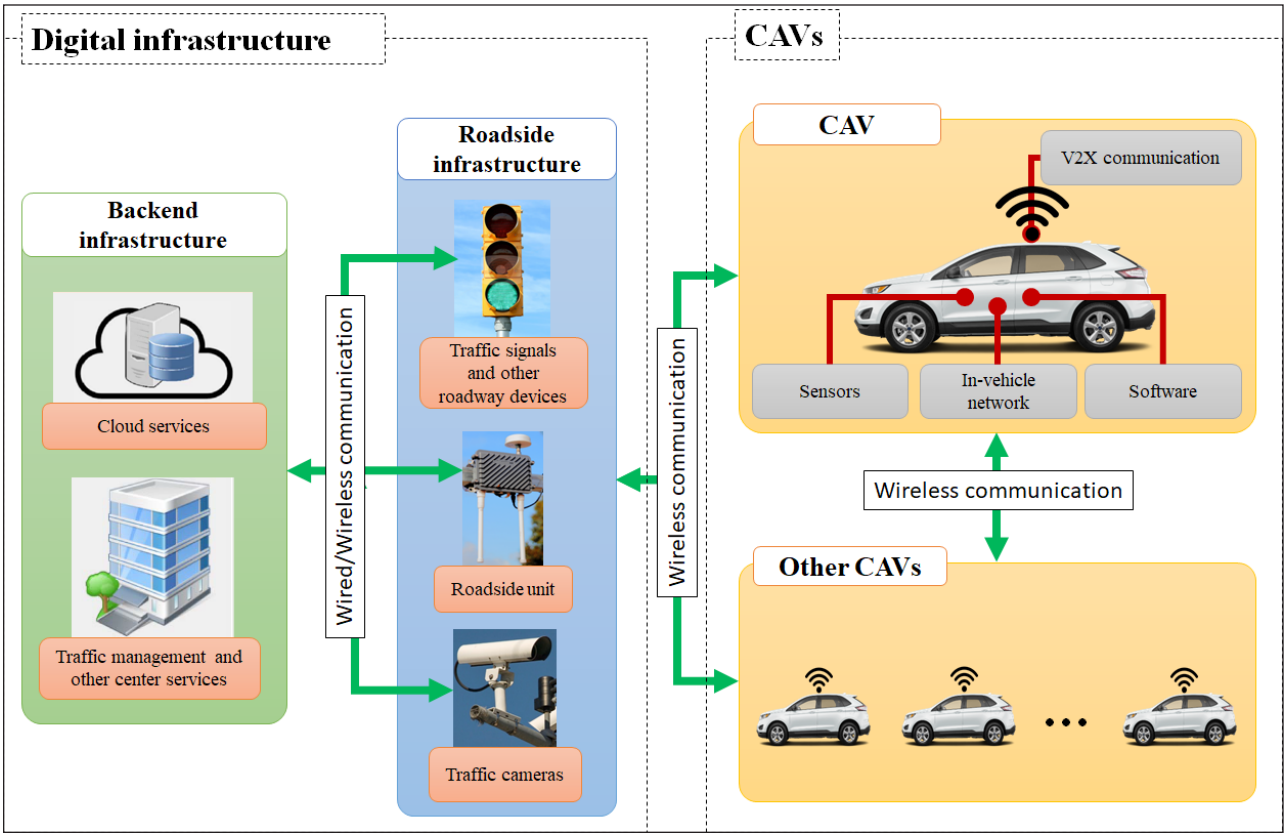

FIGURE 1 Overview of a cyberphysical system for connected and automated vehicles (CAVs). V2X = vehicle-to-everything.

software-defined networks, network function virtualization, edge computing, information-centric and virtual dispersive networking, fifth generation (5G) cellular networks, blockchain technology, and quantum and postquantum cryptography—as potential solutions aiding in securing CAVs and transportation infrastructure against existing and future cyberattacks.

Introduction

CAVs are the subject of research and development by both academia and industry because of their potential to improve traffic safety and operations.

The limitations of human perception prevent motorists from discerning what is beyond the range of human sight, such as roadway incidents and work zones around a corner or in the far distance. Even autonomous vehicles' sensors fail to discern many such impediments. Vehicle-to-everything (V2X) communication provides a 360-degree view that transcends the limited capabilities of both human-driven and automated vehicles.

CAVs are planned to be part of a broader connected city initiative, communicating with other CAVs and with smart infrastructure and services (figure 1). CAV-CPS will use connectivity to improve roadway operational efficiency using real-time roadway traffic information (e.g., about traffic signal phasing and timing, traffic incidents and queues) while improving safety in numerous ways, such as redundancy in case an automated vehicle's onboard sensors fail.

However, the highly interconnected CAV-CPS will introduce challenging security issues and vulnerabilities (Deka et al. 2018), with far-reaching consequences for a poorly secured system. For example, attackers could gain access to CAV control systems and cause catastrophic multivehicle crashes. Thus, it is critical to develop security solutions to protect CAVs, their occupants, other road users, and the associated infrastructure.

The Society of Automotive Engineers has created a cybersecurity guidebook of recommended practice, SAE J3061, that establishes a common terminology for security threats, vulnerabilities, and risks across the vehicular CPS (SAE 2016). SAE J3061 recommended practices are based on ISO 26262, an established standard for automotive functional safety (ISO 2018). The

TABLE 1 Examples of potential attacks and countermeasures on CAV sensors

Sensor type	Attack type	Attack description	Potential countermeasures
GPS	Jamming	Blocking or interfering with the satellite signals received by the GPS	• Multiple GPS receivers and GPS activity monitoring (Parkinson et al. 2017) • Antijamming devices as GPS frequency filters (Chien 2015)
	Spoofing	Modifying the position, velocity, and time values of the GPS receiver	• Navigation message authentication and cross-correlation of multiple GPS receivers (Psiaki et al. 2013) • Integrated navigation systems combining GPS and inertial navigation system (Jwo et al. 2013) • GPS data validation using V2X communication (Anouar et al. 2017)
Camera	Illusion and blinding	Adversarial physical objects compromising the vision-based systems	• Redundant sensors to verify camera data (Petit et al. 2015) • Machine learning models trained with adversarial data (Islam et al. 2019)
LiDAR	Jamming	Turning off or degrading its performance	• Protective glasses around a LiDAR that act as light filters (Tuchinda et al. 2006)
	Spoofing	Tricking it into detecting false objects	• Misbehavior detection system incorporating other sensor data, such as from the camera (Petit et al. 2015)

rapid evolution of CAV technologies requires adapting these standards and developing new ones to address CAV-CPS security challenges.

CAV Security Vulnerabilities and Solutions

We divide CAV-CPS cyberattack surfaces into three main groups—

- in-vehicle systems, which include sensors, software, and in-vehicle network;
- V2X communication networks; and
- supporting digital infrastructure—

and outline countermeasures to cyberattacks in each category.

In-Vehicle Systems

Given the fatal consequences that may result if a CAV's in-vehicle systems are compromised, ensuring their security is of paramount importance to the automotive industry.

Sensors

In addition to V2X systems, a CAV is generally equipped with light detection and ranging (LiDAR), camera, radio detection and ranging (RADAR), global positioning system (GPS) and inertial measurement unit sensors. They collaborate to improve vehicle safety and operational efficiency and are responsible for localization, obstacle avoidance, and trajectory and path planning (Schwarting et al. 2018). Table 1 presents examples of potential cyberattacks on GPS, camera, and LiDAR, as automated vehicle safety and operation largely depend on these sensors.

Software

Modern vehicles are equipped with multiple electronic control units (ECUs) to manage vehicle functionality through signal acquisition, processing, and control. ECUs may have vulnerabilities that can be exploited by an attacker, and CAVs, with their additional sensors and functionalities, have more ECUs than non-CAVs do (Wyglinski et al. 2013). One such additional ECU is the navigation control module; if compromised, it may misdirect the vehicle toward an unintended destination or even off the road (Chattopadhyay and Lam 2018).

Because ECUs are interdependent, if one is compromised, others are affected. For example, an attack on the engine control module may result in the transmission of false data about wheel speed to the electronic brake control module, which may inappropriately activate the brakes (Parkinson et al. 2017).

One of the biggest security concerns with CAV software is the over-the-air update, which may allow malware injections and thus enable an attacker to gain remote access and control of the vehicle (Nie et al. 2017). Another security threat is software bugs in the source code of the ECU software/firmware.

A simple way to ensure software/firmware security is to use strong cryptographic solutions, which include

message authentication, to prevent malware injections (Pike et al. 2017). Signature- and behavior-based models can secure the communication between ECUs (Parkinson et al. 2017). However, using secured over-the-air updates, these mitigation strategies can be kept up to date to maintain and strengthen the security of the CAV software/firmware (Mayilsamy et al. 2018).

In-Vehicle Network

ECUs in a CAV communicate with each other using communication message protocols, such as FlexRay, or controller area network (CAN). CAN is a universal real-time messaging protocol widely used by the automotive industry given its lower implementation cost (Liu et al. 2017).

Although the CAN bus has some basic security features, such as firewalls, it is vulnerable to cyberattacks. The absence of authentication and encryption allows unauthorized devices to join the CAN through an onboard device port. As a broadcast-based network, CAN messages have neither source nor destination address, which means that every in-vehicle device can listen to any unencrypted CAN message. The CAN bus is also vulnerable to denial of service (DoS) attacks. For example, an attacker generating high-priority CAN messages can prevent ECUs from communicating with each other because of the prioritization rule of the CAN protocol (Koscher et al. 2017).

Another alternative in-vehicle communication system is the automotive Ethernet, which provides a higher data transmission rate, better reliability and adaptability (Huo et al. 2015), and more security than the CAN (because the Ethernet protocol contains source and destination addresses) (Jadhav and Kshirsagar 2018). In addition, the highly secure hardware security module, trusted platform module (Corbett et al. 2018), and hardware authentication (Intel 2019) are viable solutions for in-vehicle network security.

Vehicle-to-Everything Network

In a V2X network, to communicate with external agents such as pedestrians, other vehicles, roadside transportation infrastructure, and servers (cloud-based or in-house), a CAV uses vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication. Cyberattacks on the V2X network can compromise its availability, integrity, confidentiality, and authenticity (Alnasser et al. 2019). Following are the most likely attacks on a V2X network:

- Black- and greyhole attacks: The compromised vehicle stops forwarding all packets (blackhole) or some packets (greyhole) to other CAVs, so that other CAVs cannot receive safety-critical information, such as forward collision warnings.
- DoS and distributed DoS (DDoS): Attackers disrupt the V2X network by data flooding, causing a delay in the transmission of safety-critical information and making the network services unavailable.
- Jamming: An attacker broadcasts signals to corrupt data or disable communication channels.
- False message injection: The compromised vehicle creates fake messages or alters received messages and broadcasts them.
- Eavesdropping: The compromised vehicle uses false identities to capture data packet information, thus acquiring sensitive and confidential data.
- Certificate replication: An attacker conceals itself from certification authorities by replicating the certificates of legitimate vehicles.
- Sybil attack: An attacker creates multiple identities to gain the trust of legitimate CAVs.
- Impersonation: An attacker establishes itself as trustworthy in the network by impersonating a trusted entity to gain access to sensitive information.

Solutions for V2X network security can be based on cryptography, behavior, or identity.

- Cryptography-based solutions include encryption, secure key management and authentication. For example, a security credential management system can prevent impersonation, false identity, and eavesdropping attacks (Ravi and Kulkarni 2013; Whyte et al. 2018).
- Among behavior-based models, one of the most popular is the weighted-sum method, in which a level of trust of a connected vehicle is based on the weighted sum of relevant criteria, such as the transmission range, vehicle speed, and vehicle direction (Sugumar et al. 2018). For example, a blackhole attack can be detected by a reputation-based global trust model, where a CAV's past behavior is considered in determining its current trust value (Li et al. 2013). Behavior-based cooperative awareness between vehicles can prevent greyhole attacks (Ali

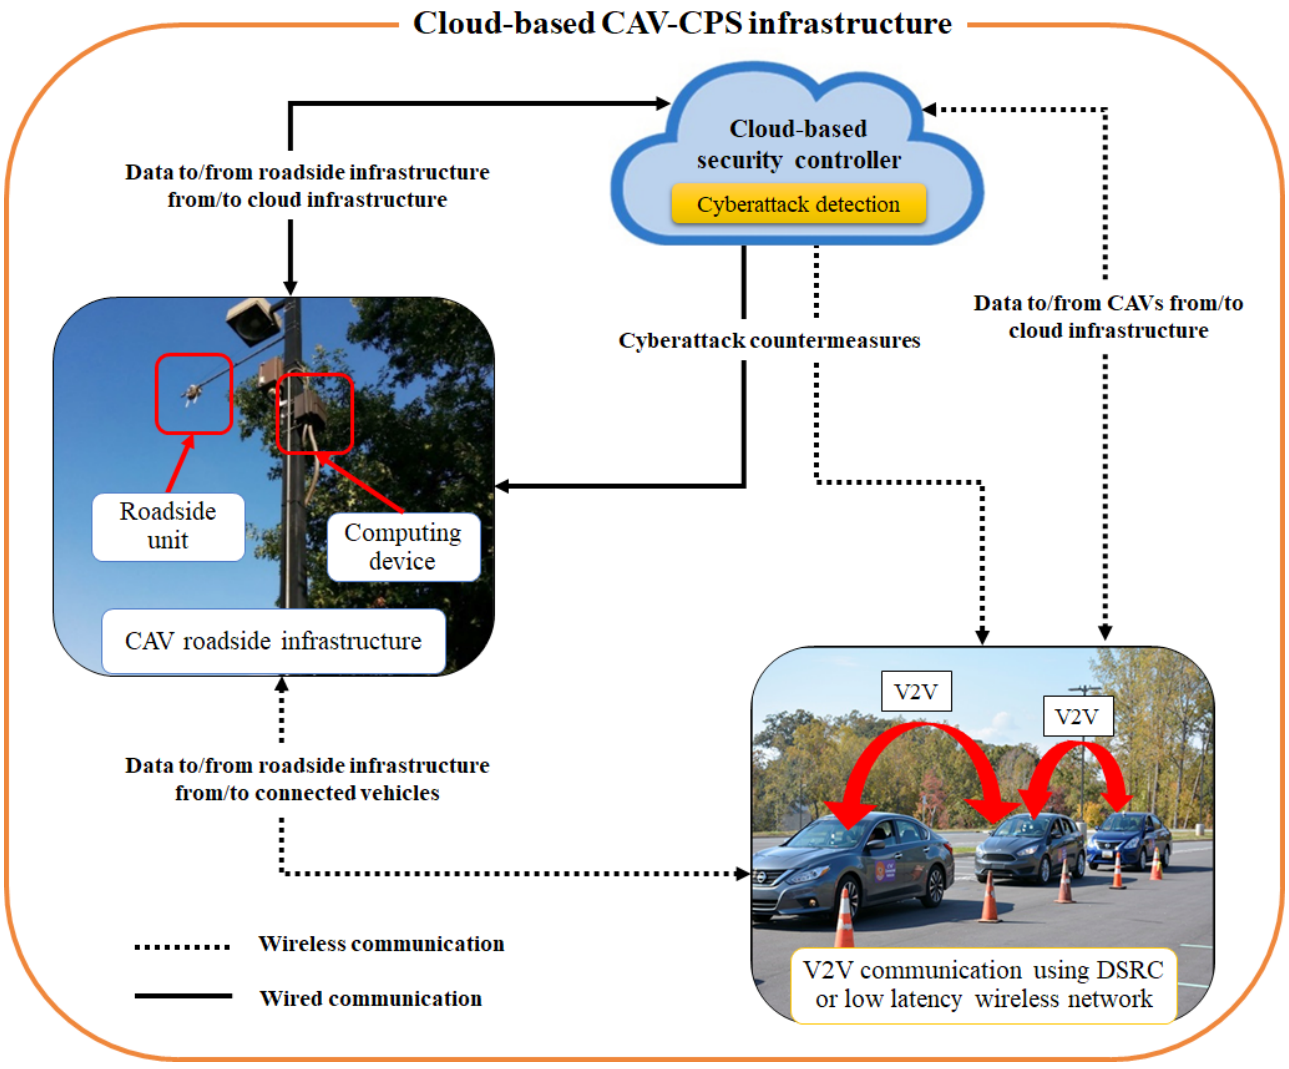

FIGURE 2 Cloud-based security for a connected and automated vehicle (CAV)–cyberphysical system (CPS) at the Center for Connected Multimodal Mobility (C^2M^2) testbed, Clemson, SC. DSRC = dedicated short-range communications; V2V = vehicle-to-vehicle.

Alheeti et al. 2016). DoS and DDoS attacks can be detected and mitigated using policy- and rule-based techniques (Islam et al. 2018).

- Identity-based security solutions can be used to detect eavesdropping and secure vehicle privacy (Kang et al. 2016).

Digital Infrastructure

In addition to V2X security, the security of the CAV-CPS digital infrastructure is of major concern. Digital infrastructure includes roadside units, traffic signals and cameras, traffic management centers, and cloud infrastructure (figure 1). Digital infrastructure enables cloud-based security solutions for different components of the V2X network and CAVs (figure 2), such as

- connected vehicles communicating with each other using dedicated short-range communications or any other low-latency wireless network;
- roadside infrastructure communicating with connected vehicles via a low-latency wireless network and other roadside infrastructure using wireless/wired networks;
- computing devices running connected vehicle applications;
- cloud servers hosting a security controller for detecting and mitigating cyberattacks to the V2X network.

We used a cloud-based attack detection and mitigation system to detect cyberattacks on the V2X network and roadside infrastructures in a testbed at the Center for Connected Multimodal Mobility (C^2M^2),

a US Department of Transportation (USDOT) Tier 1 University Transportation Center headquartered in Clemson, South Carolina (Islam et al. 2018). Although the cloud would add some vulnerabilities to the system, cloud security is well studied, and protective techniques are available (Basu et al. 2018). However, an attacker can take advantage of the vulnerabilities of the protocols associated with message transfers between the CAVs/roadside infrastructure and cloud servers to cut off the communication link between them (Dinculeană and Cheng 2019). Distributed security solutions using the software-defined network can solve this issue by giving autonomy to the CAVs or roadside infrastructure in terms of operating their corresponding security solutions independently (Darabseh et al. 2015).

Future Landscape for CAV Cybersecurity

The rapid evolution of information technology has led to the development and use of emerging technologies to enhance CAV security. Although all emerging technologies are not fully tested and have their own challenges, their potential for securing CAVs is beyond dispute. When creating CAV security solutions using these emerging technologies, a security-by-design process that integrates cyberattack detection and countermeasures needs to be adopted at the outset (Chattopadhyay and Lam 2018).

Artificial Intelligence

Advances in computer hardware technologies and distributed computing facilities (e.g., cloud computing) have led to numerous innovations in AI-based cybersecurity that can be used to protect CAVs.

Deep neural networks (DNNs) are used in the development of intrusion detection systems (IDS) for connected vehicles (Aloqaily et al. 2019). Recurrent neural networks, such as a long short-term memory, are perhaps the most relevant and widely used deep learning model for IDS (Levi et al. 2018). Another use of AI for security is context-aware user-behavior analytics, in which CAV behavior data are collected and used to detect security threats (Wasicek et al. 2017).

However, AI itself can be subject to cyberattacks. For example, by placing a small amount of graffiti or a sticker on a road traffic sign, an attacker can cause the DNN-based traffic sign recognition system of an automated vehicle to misclassify the sign, which can lead to traffic crashes (Eykholt et al. 2018). The attacker specifically targets the DNN to compromise the performance of the automated vehicle even as the changes to the sign are imperceptible to humans (Papernot et al. 2017).

When designing a CAV security solution using AI, the AI system itself must be secured by improving its robustness.

When designing a CAV security solution using AI, the AI system itself must be secured by improving its robustness.

Software-Defined Network

Traditional network management is nonprogrammable, complex, and error-prone (Modieginyane et al. 2018), but a programmable software-defined network (SDN) can be organized for more secure network management and enhanced CAV security (Nobre et al. 2019).

SDN involves a centralized controller for the data flow between vehicles and the roadside infrastructure (Jaballah et al. 2019), making it an important tool for CAV security. During a cyberattack, the roadside infrastructure routes the data to the SDN controller, and a detection and mitigation application in cyberspace detects the attack and selects an appropriate mitigation strategy (figure 3). An SDN controller can push the mitigation strategy to SDN-enabled switches (e.g., OpenFlow switches) in the physical space to deploy the updated mitigation strategy. Furthermore, SDN combined with AI can be used to detect and mitigate cyberattacks on the in-vehicle system (Khan et al. 2019).

Network Function Virtualization

Unlike legacy systems, in which network functions, such as firewalls and IDS, are deployed in proprietary hardware, network function virtualization (NFV) can provide a cost-effective approach to CAV security (Alwakeel et al. 2018). Virtualized network functions can run on top of commodity/off-the-shelf hardware, such as industry-standard servers, storage, and switches. Virtualization of network security functions can be used to design security solutions that provide the same

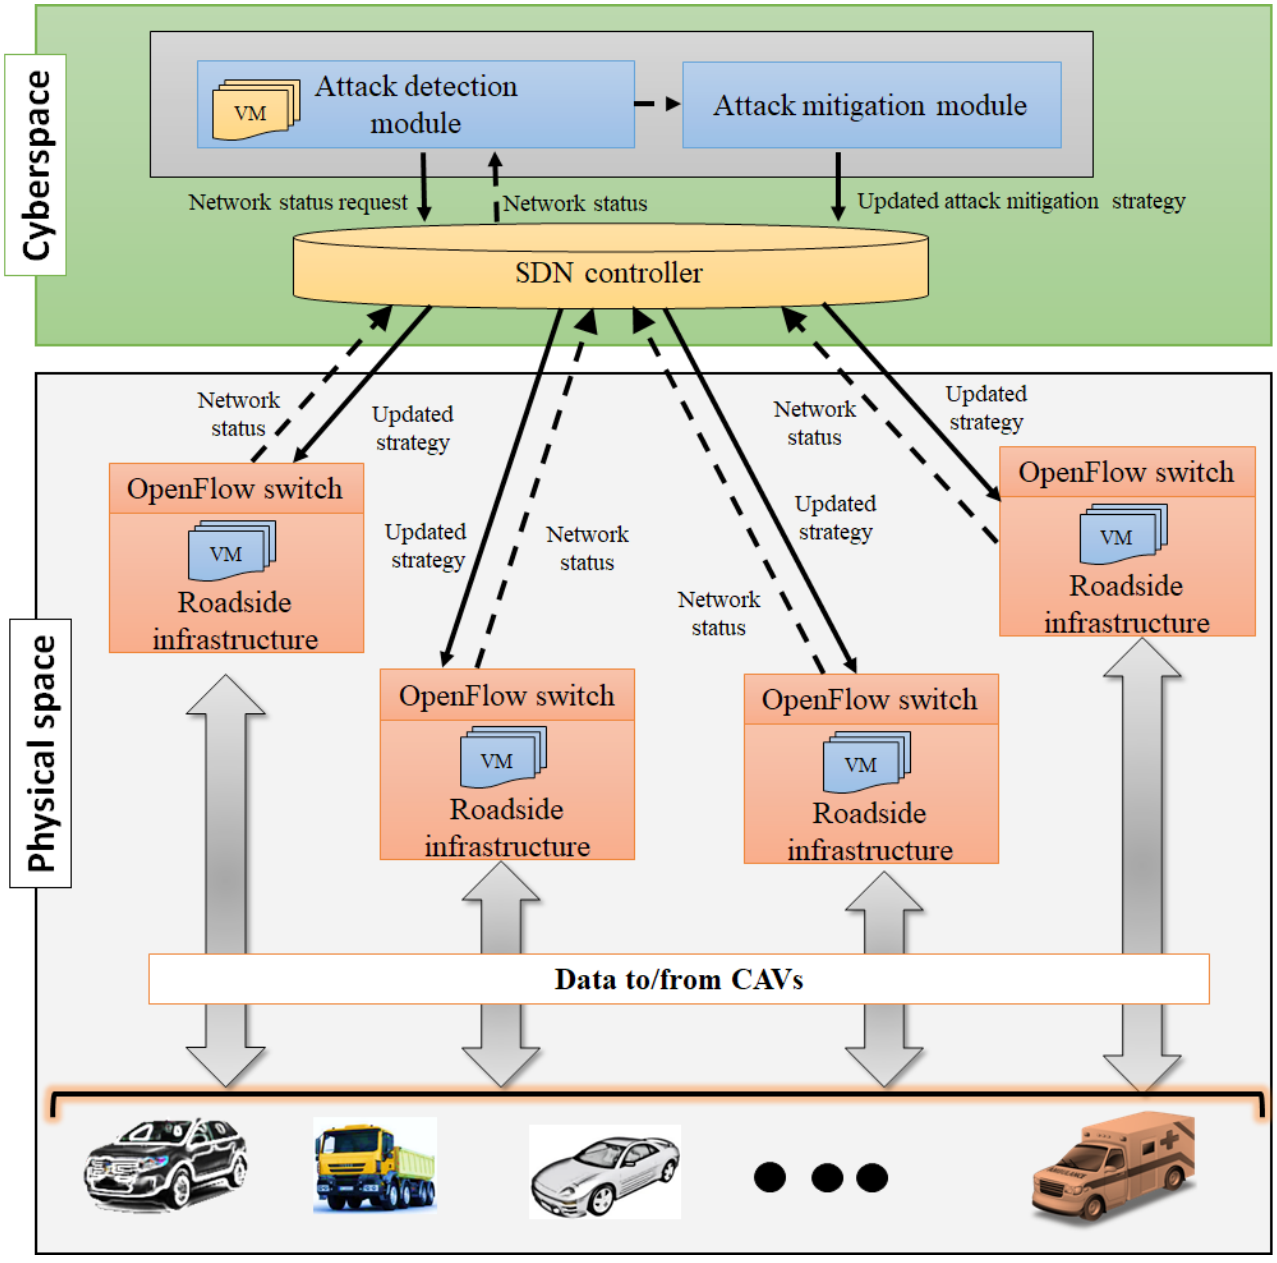

FIGURE 3 Connected and automated vehicle (CAV)–cyberphysical system (CPS) cybersecurity using software-defined network (SDN) and network function virtualization (NFV). VM = virtual machine.

or better performance compared to security modules in proprietary hardware (Han et al. 2015).

In addition, SDN can provide a platform to virtualize network security functions using virtual machines (VMs). With NFV and SDN, security software in a VM can be launched in an SDN-enabled roadside infrastructure within a CAV network (figure 3). SDN and NFV thus act together to improve CAV security in terms of granularity, flexibility, scalability, and resiliency.

Edge Computing

Edge computing, in which data are processed in any computing device (e.g., edge device, aggregation centers) close to the data source, ensures high bandwidth usage and distribution of the computational tasks among edge devices (Shi and Dustdar 2016). For example, CAV security modules are deployed near the data source (e.g., in CAVs or roadside devices) for faster processing and effective protection. Edge computing– and

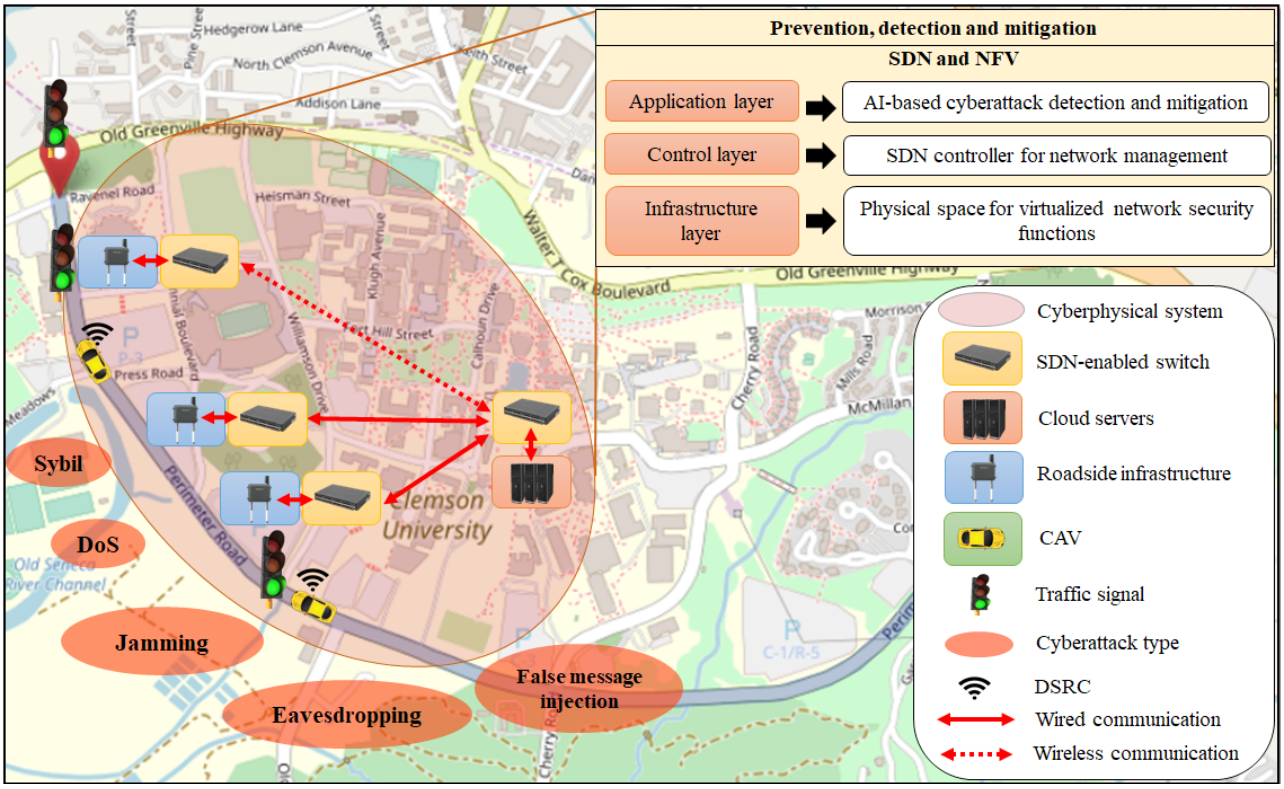

FIGURE 4 Cybersecurity research using edge computing, software-defined network (SDN), and network function virtualization (NFV) at the Center for Connected Multimodal Mobility testbed at Clemson University. CAV = connected and automated vehicle; DoS = denial of service; DSRC = dedicated short-range communications.

AI-based security solutions can be updated/modified in real time to improve cyberattack detection and mitigation. Although the addition of edge devices adds new attack surfaces, edge computing can provide faster security solutions for CAVs by having the solutions closer to CAVs and thus reducing the data transmission delay.

At C^2M^2 we use an edge computing–based testbed for CAV cybersecurity research (figure 4) (Chowdhury et al. 2018). Different cyberattacks are created against connected vehicles and transportation infrastructure (e.g., traffic signal controllers and roadside devices) to explore effective detection strategies using AI and mitigations using edge computing, SDN, and NFV.

Information-Centric and Virtual Dispersive Networking

Information-centric networking (ICN) is a new paradigm of communication networking. Unlike traditional host-centric networking, ICN centers on information or content and enables in-network caching and replication; expected benefits include improved network efficiency, scalability, and robustness (Kalogeiton and Braun 2018). ICN in combination with edge computing has the potential to add new dimensions to the network management of connected vehicles (Grewe et al. 2017).

Virtual dispersive networking is another new technology that differs from traditional networks as messages are divided into numerous parts, encrypted, and sent through multiple paths so that attackers can't intercept the whole message, preventing eavesdropping (Twitchell 2013).

5G Network

Fifth generation cellular communications are expanding through the Internet of Things (Li et al. 2018). Researchers are adopting 5G in CAVs to pioneer a more secure, reliable, and resilient CAV-CPS (Dey et al. 2018). 5G features such as ultralow latency, higher throughput, and trustworthiness will increase CAV robustness and security compared to the 4G network (Jover and Marojevic 2019).

However, a few 5G protocol specifications are unchanged from 4G, which may transfer some security vulnerabilities to 5G. Moreover, new vulnerabilities are bound to emerge, leading to new security requirements

in 5G. The Next Generation Mobile Networks Alliance has identified additional security requirements of 5G wireless networks (NGMN 2015).

Among advantages of the 5G network, edge devices can leverage its higher bandwidth to offload data into the cloud server in real time to facilitate cyberattack detection and mitigation. In addition, 5G features will make it possible to support SDN and NFV in improving network scalability and management. Thus, 5G can make the integration of security solutions easier in a CAV-CPS.

Researchers are exploring the use of quantum random number generators to develop ECUs for automotive security.

Blockchain

The blockchain is a distributed and trust-based security solution. Using blockchain, CAVs can establish trust with each other, the roadside infrastructure, and cloud servers. The distributed nature of the blockchain makes it difficult to launch an attack as the blocks of data are protected by consensus protocols (Cachin and Vukolić 2017). The blockchain technology may enhance CAV security (Singh and Kim 2017) and privacy (Dorri et al. 2017) by providing a decentralized trusted V2X communication network. However, more research is needed to determine the potential of blockchain technology to improve CAV cybersecurity.

Quantum and Postquantum Cryptography

Although the development of quantum computers is still in its early stage, quantum computing has the potential to create breakthroughs in AI, optimization, and cryptography (Smedley 2018). One potential use of quantum cryptography for CAVs is quantum key distribution, which can ensure secure key exchange between ECUs in the in-vehicle network (Nguyen et al. 2019). However, this may require an optical fiber connection in the network, which may not be feasible because of the higher production cost (Dennison 2019).

One of the biggest threats to in-vehicle security is the encryption system based on nonrandom numbers. Researchers are exploring the use of quantum random number generators to develop ECUs for automotive security (Nguyen et al. 2019).

As with other technological advances, quantum computing may be used to create CAV cyberattacks. CAVs should leverage postquantum cryptography algorithms to protect against such attacks.

Conclusions

Researchers in academia and industry are actively working on developing new methods for the detection and mitigation of CAV-CPS cyberattacks. However, challenges persist in this rapidly evolving area.

Given security concerns regarding the acquisition of sensitive information by an attacker, more research should focus on V2X security. CAV data privacy is a concern because CAVs will connect with in-vehicle personal devices and the outside world. Therefore, ensuring V2X security is extremely important to prevent the spread of cyberattacks from CAVs to connected infrastructures and vice versa.

Meeting these challenges requires a concerted effort in academia, industry, and government. Such collaboration can lead to improved standards for a secure CAV-CPS and wide adoption of best security practices across the CAV industry. It will also pave the way for better privacy, safety, and security countermeasures that complement and strengthen each other.

Acknowledgments

This study is supported by the Center for Connected Multimodal Mobility (C^2M^2), a USDOT Tier 1 University Transportation Center, headquartered at Clemson University, South Carolina. Any opinions, findings, conclusions, or recommendations expressed in this article are those of the authors and do not necessarily reflect the views of the C^2M^2, and the US government assumes no liability for the contents or use thereof. The authors thank Cameron Fletcher for her time and effort in editing the article thoroughly and improving its quality.

References

Ali Alheeti KM, Gruebler A, McDonald-Maier K. 2016. Intelligent intrusion detection of grey hole and rushing attacks in self-driving vehicular networks. Computers 5(16).

Alnasser A, Sun H, Jiang J. 2019. Cyber security challenges and solutions for V2X communications: A survey. Computer Networks 151:52–67.

Aloqaily M, Otoum S, Al Ridhawi I, Jararweh Y. 2019. An intrusion detection system for connected vehicles in smart cities. Ad Hoc Networks 90:101842.

Alwakeel AM, Alnaim AK, Fernandez EB. 2018. A survey of network function virtualization security. Conf Proceedings, IEEE Southeastcon, Apr 19–22, St. Petersburg FL.

Anouar B, Mohammed B, Abderrahim G, Mohammed B. 2017. Vehicular navigation spoofing detection based on V2I calibration. IEEE Colloquium on Information Science and Technology, Oct 24–26, Tangier.

Basu S, Bardhan A, Gupta K, Saha P, Pa M, Bose M, Basu K, Chaudhury S, Sarkar P. 2018. Cloud computing security challenges & solutions: A survey. IEEE 8th Annual Computing and Communication Workshop and Conf, Jan 8–10, Las Vegas.

Cachin C, Vukolić M. 2017. Blockchain consensus protocols in the wild. arXiv:1707.01873.

Chattopadhyay A, Lam K-Y. 2018. Autonomous vehicle: Security by design. arXiv:1810.00545.

Chien YR. 2015. Design of GPS anti-jamming systems using adaptive notch filters. IEEE Systems Journal 9(2):451–460.

Chowdhury M, Rahman M, Rayamajhi A, Khan SM, Islam M, Khan Z, Martin J. 2018. Lessons learned from the real-world deployment of a connected vehicle testbed. Transportation Research Record 2672(22):10–23.

Corbett C, Brunner M, Schmidt K, Schneider R, Dannebaum U. 2018. Leveraging hardware security to secure connected vehicles. WCX World Congress Experience, Apr 10–12, Detroit.

Deka L, Khan SM, Chowdhury M, Ayres N. 2018. Transportation cyber-physical system and its importance for future mobility. In: Transportation Cyber-Physical Systems, eds Deka L, Chowdhury M. Cambridge MA: Elsevier.

Dennison C. 2019. How fiber protection is enabling next-generation automotive systems. PPC blog.

Darabseh A, Al-Ayyoub M, Jararweh Y, Benkhelifa E, Vouk M, Rindos A. 2015. SDSecurity: A software defined security experimental framework. 2015 IEEE International Conf on Communication, London.

Dey K, Fries R, Ahmed S. 2018. Future of transportation cyber-physical systems – Smart cities/regions. In: Transportation Cyber-Physical Systems, eds Deka L, Chowdhury M. Cambridge MA: Elsevier.

Dinculeană D, Cheng X. 2019. Vulnerabilities and limitations of MQTT protocol used between IoT devices. Applied Sciences 9(5):848.

Dorri A, Steger M, Kanhere SS, Jurdak R. 2017. BlockChain: A distributed solution to automotive security and privacy. IEEE Communications 55(12):119–125.

Eykholt K, Evtimov I, Fernandes E, Li B, Rahmati A, Xiao C, Prakash A, Kohno T, Song D. 2018. Robust physical-world attacks on deep learning visual classification. 2018 IEEE/CVF Conf on Computer Vision and Pattern Recognition, Jun 18–22, Salt Lake City.

Grewe D, Wagner M, Arumaithurai M, Psaras I, Kutscher D. 2017. Information-centric mobile edge computing for connected vehicle environments: Challenges and research directions. Proceedings, Workshop on Mobile Edge Communications, Aug 21, Los Angeles.

Han B, Gopalakrishnan V, Ji L, Lee S. 2015. Network function virtualization: Challenges and opportunities for innovations. IEEE Communications 53(2):90–97.

Huo Y, Tu W, Sheng Z, Leung VCM. 2015. A survey of in-vehicle communications: Requirements, solutions and opportunities in IoT. IEEE 2nd World Forum on Internet of Things, Dec 14–16, Milan.

Intel. 2019. Intel Authenticate Technology: Hardware-enhanced security. Online at https://www.intel.com/content/www/us/en/security/authenticate/authenticate-is-hardware-enhanced-security.html.

Islam M, Chowdhury M, Li H, Hu H. 2018. Cybersecurity attacks in vehicle-to-infrastructure applications and their prevention. Transportation Research Record 2672(19):66–78.

Islam M, Chowdhury M, Li H, Hu H. 2019. Vision-based navigation of autonomous vehicle in roadway environments with unexpected hazards. arXiv:1810.03967.

ISO. 2018. ISO 26262-1:2018 - Road Vehicles—Functional Safety. Geneva: International Organization for Standardization.

Jaballah WB, Conti M, Lal C. 2019. A survey on software-defined VANETs: Benefits, challenges, and future directions. arXiv:1904.04577.

Jadhav S, Kshirsagar D. 2018. A survey on security in automotive networks. 4th International Conf on Computing Communication Control and Automation, Aug 16–18, Puna, India.

Jover RP, Marojevic V. 2019. Security and protocol exploit analysis of the 5G specifications. IEEE Access 7:24956–24963.

Jwo DJ, Chung FC, Yu KL. 2013. GPS/INS integration accuracy enhancement using the interacting multiple model nonlinear filters. Applied Research and Technology 11(4):496–509.

Kalogeiton E, Braun T. 2018. Infrastructure-assisted communication for NDN-VANETs. 19th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks, Jun 12–15, Chania.

Kang J, Yu R, Huang X, Jonsson M, Bogucka H, Gjessing S, Zhang Y. 2016. Location privacy attacks and defenses in

cloud-enabled internet of vehicles. IEEE Wireless Communications 23:52–59.

Khan Z, Chowdhury M, Islam M, Huang C-Y, Rahman M. 2019. In-vehicle false information attack detection and mitigation framework using machine learning and software defined networking. arXiv:1906.10203.

Koscher K, Czeskis A, Roesner F, Patel S, Kohno T. 2017. Experimental security analysis of a modern automobile. High Energy Physics 2017(11):1–16.

Levi M, Allouche Y, Kontorovich A. 2018. Advanced analytics for connected car cybersecurity. IEEE 87th Vehicular Technology Conf, Jun 3–6, Porto.

Li X, Liu J, Li X, Sun W. 2013. RGTE: A reputation-based global trust establishment in VANETs. Proceedings, 5th International Conf on Intelligent Networking and Collaborative Systems, Sep 9–11, Xi'an.

Li S, Xu LD, Zhao S. 2018. 5G Internet of Things: A survey. Industrial Information Integration 10:1–9.

Liu J, Zhang S, Sun W, Shi Y. 2017. In-vehicle network attacks and countermeasures: Challenges and future directions. IEEE Network 31(5):50–58.

Mayilsamy K, Ramachandran N, Raj VS. 2018. An integrated approach for data security in vehicle diagnostics over internet protocol and software update over the air. Computers & Electrical Engineering 71:578–593.

Modieginyane KM, Letswamotse BB, Malekian R, Abu-Mahfouz AM. 2018. Software defined wireless sensor networks application opportunities for efficient network management: A survey. Computers and Electrical Engineering 66:274–287.

NGMN [Next Generation Mobile Networks Alliance]. 2015. 5G White Paper. Frankfurt am Main.

Nguyen HN, Tavakoli S, Shaikh SA, Maynard O. 2019. Developing a QRNG ECU for automotive security: Experience of testing in the real-world. 2019 IEEE International Conf on Software Testing, Verification and Validation Workshops, Apr 22–23, Xi'an, China

Nie S, Liu L, Du Y. 2017. Free-fall: Hacking Tesla from wireless to CAN Bus. Black Hat USA, Jul 27.

Nobre JC, de Souza AM, Rosário D, Both C, Villas LA, Cerqueira E, Braun T, Gerla M. 2019. Vehicular software-defined networking and fog computing: Integration and design principles. Ad Hoc Networks 82:172–181.

Papernot N, McDaniel P, Goodfellow I, Jha S, Celik ZB, Swami A. 2017. Practical black-box attacks against machine learning. Proceedings, 2017 ACM Asia Conf on Computer and Communications Security, Apr 2–6, Abu Dhabi.

Parkinson S, Ward P, Wilson K, Miller J. 2017. Cyber threats facing autonomous and connected vehicles: Future challenges. IEEE Transactions on Intelligent Transportation Systems 18(11):2898–2915.

Petit J, Stottelaar B, Feiri M, Kargl F. 2015. Remote attacks on automated vehicles sensors: Experiments on camera and LiDAR. Black Hat Europe, Nov 10–13, Amsterdam.

Pike L, Sharp J, Tullsen M, Hickey PC, Bielman J. 2017. Secure automotive software: The next steps. IEEE Software 34(3):49–55.

Psiaki ML, O'Hanlon BW, Bhatti JA, Shepard DP, Humphreys TE. 2013. GPS spoofing detection via dual-receiver correlation of military signals. IEEE Transactions on Aerospace and Electronic Systems 49(4):2250–2267.

Ravi K, Kulkarni SA. 2013. A secure message authentication scheme for VANET using ECDSA. 4th International Conf on Computing, Communications and Networking Technologies, Jul 4–6, Tiruchengode, India.

SAE. 2016. Cybersecurity Guidebook for Cyber-Physical Vehicle Systems. SAE International.

Schwarting W, Alonso-Mora J, Rus D. 2018. Planning and decision-making for autonomous vehicles. Annual Review of Control, Robotics, and Autonomous Systems 1(1):187–210.

Shi W, Dustdar S. 2016. The promise of edge computing. Computer 49(5):78–81.

Singh M, Kim S. 2017. Intelligent vehicle-trust point: Reward based intelligent vehicle communication using Blockchain. arXiv:1707.07442.

Smedley P. 2018. Autonomous vehicles in the quantum age. Connected World, Apr 3.

Sugumar R, Rengarajan A, Jayakumar C. 2018. Trust based authentication technique for cluster based vehicular ad hoc networks (VANET). Wireless Networks 24(2):373–382.

Tuchinda C, Srivannaboon S, Lim HW. 2006. Photoprotection by window glass, automobile glass, and sunglasses. Journal of the American Academy of Dermatology 54(5):845–854.

Twitchell RW. 2013. Virtual dispersive networking systems and methods. US Patent 9,071,607.

Wasicek A, Pesé MD, Weimerskirch A, Burakova Y, Singh K. 2017. Context-aware intrusion detection in automotive control systems. 5th ESCAR USA Conf, Jun 21–22.

Whyte W, Weimerskirch A, Kumar V, Hehn T. 2018. A security credential management system for V2V communications. IEEE Transactions on Intelligent Transportation Systems 19(12):3850–3871.

Wyglinski AM, Huang X, Padir T, Lai L, Eisenbarth TR, Venkatasubramanian K. 2013. Security of autonomous systems employing embedded computing and sensors. IEEE Micro 33(1):80–86.

Cybersecurity is not an end in itself but an ongoing set of practices throughout the system lifecycle to achieve system goals and requirements.

What Every Engineer Should Know about Cybersecurity

Thomas A. Longstaff and
Noelle K. Allon

Thomas A. Longstaff

Noelle K. Allon

Computer science and engineering—which includes computer, computational, communication, and information science and engineering—is the branch of engineering that concerns itself with cybersecurity. However, for the safe and secure development and deployment of engineering systems, attention to and knowledge of cybersecurity should extend beyond the domain of computer science and engineering to other branches of engineering. The participation of a cybersecurity engineer on a systems design team can ensure mitigations and modifications that will increase system resilience and longevity.

The Need for Cybersecurity Is Everywhere

Engineers develop many capabilities by embedding software in the systems they produce, from kitchen appliances and baby monitors to home security systems and cars, to name just a few. A large subset of these systems offer internet connectivity and comprise devices that constitute the Internet of Things. There are already more than 25 billion "things" connected to the internet, and the number is constantly growing.

Thomas Longstaff is chief technology officer of the Software Engineering Institute at Carnegie Mellon University. Noelle Allon is an analyst in the CERT Division of the Software Engineering Institute.

Even if a system is produced with no apparent need for security, it can be vulnerable to misuse. For example, a smart lightbulb may not seem to require traditional security, but if its computational resources (e.g., for its settings and remote control) are misused, the light will not function as users expect it to. Engineers must ensure that the embedded software of all systems meets safety and security requirements in the environments in which the systems are deployed.

As system complexity increases, the number of vulnerabilities does as well.

In addition to growing in number, devices are becoming more complex, especially as engineers incorporate machine learning and other artificial intelligence capabilities. As the complexity increases, the number of vulnerabilities does as well. As a result, each and every device is vulnerable to misuse, creating the potential for harm unless engineers can secure and continually update devices. As of April 2019, there were over 120,000 Common Vulnerabilities and Exposure entries in the National Vulnerability Database.[1] These vulnerabilities can exist in the code, the environment, or the system, and many can result in harm to the system itself or to people or other systems. Organizations can incur multiple costs—such as the breach of sensitive information, loss of money, reputation, and time—when attackers successfully exploit a vulnerability.

While it is not reasonable to expect that every engineer will become an expert in cybersecurity, some awareness of cybersecurity threats, risks, and trade-offs can help any engineer understand cybersecurity requirements and how to work with cybersecurity engineers throughout the system lifecycle. The participation of cybersecurity engineers in systems engineering can help organizations reap benefits such as enhanced value of the system being engineered, improved system flexibility in a changing environment, protection of legitimate system users from harm, and greater likelihood that engineers will meet system requirements.

We review cybersecurity engineering goals and tools and offer five questions to guide implementation of cybersecurity in engineering endeavors. We discuss how organizations can address each question.

[1] This database, hosted by the National Institute of Standards and Technology, is available at https://nvd.nist.gov/.

What Is Cybersecurity Engineering?

Cybersecurity engineering involves efforts to secure systems from both intentional and unintentional harm.

Researchers learned early lessons in cybersecurity engineering from incident responses to detected attacks and breaches, and these lessons have motivated better integration of cybersecurity in traditional systems and software engineering. From a financial and safety perspective, investing in cybersecurity up front is always cheaper than incurring the costs and risks associated with insufficient security.

Cybersecurity is not an end in itself but an ongoing set of practices in every stage of the system lifecycle to achieve all system goals and requirements.

Five Questions to Determine Whether Cybersecurity Engineering Is Needed

The following five questions are a starting point for incorporating cybersecurity goals and requirements in any engineered system:

1. What is the input to the system and who controls it?
2. What value does the system hold or need to protect?
3. What harm can adversaries do if they take control of the system?
4. Is there a fail-safe when a cybersecurity event is detected?
5. Can the system adapt to an evolving environment of use, attacks, and interoperability?

These questions do not represent a comprehensive set of cybersecurity concerns. Rather, they focus on (1) attack surface, (2) magnitude of consequence, (3) hazard, (4) resilience, and (5) system evolution, reflecting the most common ways that vulnerabilities and unintended consequences are introduced into systems. Their answers can help an organization determine whether to include a cybersecurity engineering specialist on the engineering team.

What Is the Input to the System and Who Controls It?

Keeping systems secure involves more than automating security rules to reduce human labor costs and error rates. It is important to know what the required inputs to the system are and who is controlling the system. Inputs to the system constitute the attack surface, which

presents opportunities for both an adversary's control of a system's behavior and misuse by an authorized user.

Attackers can exploit vulnerabilities to force a system to violate safety constraints, divulge information, change or destroy valuable information, or use unauthorized resources. A cybersecurity engineer can use knowledge of the source of the data and which users control the system to develop an appropriate threat model and determine the controls necessary to restrict, detect, mitigate, and recover from adversarial use of the inputs.

When a system contains confidential, private, or valuable information, the data and actions that control system behavior must yield predictable outcomes. To achieve such outcomes, cybersecurity engineers will model the system behavior under all foreseeable conditions to ensure that the model closes in a safe and secure state.

Complex systems with many interacting components can create conditions that are difficult to model; focusing on the portions of the system that are reachable through the attack surface can help to manage the complexity. In addition, engineers can leverage advances in model-based system engineering (McDermott et al. 2019) and formal methods such as static analysis tools to explore mitigations for vulnerabilities in the attack surface.

The incorporation of machine learning capabilities complicates this question because both the training data and operational data will determine the system behavior, and both sets of data may be susceptible to attacks. It is often difficult to discover attacks on these types of data through traditional monitoring.

A common example that illustrates this problem involves self-driving cars. They must be able to distinguish between different traffic signs. But researchers have discovered that when attackers place just a few stickers on a stop sign in a certain configuration, self-driving cars incorrectly classify it as a yield sign (Silver 2017), potentially creating an unsafe condition.

What Value Does the System Hold or Need to Protect?

A system can house various types of data, from public material to medical histories of patients with the same health insurance to highly classified intelligence on government adversaries. Knowledge of the value of the information that a system needs to protect will be used to frame the goals and requirements for protecting that information, both while it is at rest and while it is being processed or transmitted. The answer to this question also calls for identifying both the type of adversary who would try to extract or manipulate the information held by the system and the attacker's motivation for doing so.

The impacts of a successful attack are tied to the value of the information in a system. Evaluation of the extent of the impacts can inform engineering trade-offs in the development and deployment of mitigation technologies and the incorporation of data protections such as encrypted storage, secure protocols, and protected processing.

Many systems have failed by incorporating or implementing cryptography with exploitable flaws in the cryptographic components.

When a system holds information of significant value, a cybersecurity engineer familiar with the use of cryptography should participate in its design and development. Many systems have failed by incorporating or implementing cryptography with exploitable flaws. The participation of a cybersecurity engineer is also important when the system coordinates between distributed components, as detection of data in transit is a well-known attack strategy.

Other elements of a cybersecurity strategy include appropriate encryption key management, access control through identity management, and authentication and authorization of appropriate assets per user. In addition, a monitoring and resilience strategy should be developed to look for and mitigate any attempts to subvert the proper behavior of the system.

What Harm Can Adversaries Do If They Take Control of the System?

It is relevant to know whether the system controls a safety-critical function or physical action that could benefit an adversary. For example, software that controls an alarm system and electronic locks may depend on an external time source that an adversary could manipulate to allow unauthorized individuals to bypass the system and cause harm. For cyberphysical systems, such harm may affect physical materials or services managed by the system (e.g., electric power, water, alarm systems).

With knowledge of the context in which the system operates and the harm that could result if an adversary controls the system behavior, it is possible to create a threat model to drive a realistic risk analysis to reduce risks and harms to acceptable levels. This analysis could motivate the incorporation of cybersecurity capabilities that limit the system's physical behavior to ensure that it always adheres to safety and security constraints regardless of software vulnerabilities or an adversary's control of the input.

Is There a Fail-Safe When a Cybersecurity Event Is Detected?

Recognizing that a system has entered an undesired state is a key area of software and systems engineering, and developing a detection and response strategy is an area for a cybersecurity engineer.

Subversion or corruption is an unintended consequence of a system modification because many cybersecurity capabilities remain silent until the system is attacked.

Common cybersecurity strategies include developing detection and response capabilities when designing software-intensive systems to improve resilience. For a system to be resilient, engineers must design it to operate reliabily (with a specified minimum set of capabilities) in foreseeable adverse conditions. Analysis of these conditions may indicate the incorporation of a resilient fail-safe mechanism to manage potential cybersecurity events.

A resilient design may involve a physical mechanism (e.g., a circuit breaker), but in software-intensive systems, cybersecurity engineers more often establish resilience through software exception handling. The dependency analysis necessary to safely and securely implement resilience in a complex system is nontrivial.

A cybersecurity engineer can ensure inclusion of appropriate system monitoring and automation to activate resilient mechanisms regardless of the input conditions that lead to an undesired state. Implementation of resilience tends to run more smoothly the earlier a cybersecurity engineer is involved in the project.

Can the System Adapt to an Evolving Environment of Use, Attacks, and Interoperability?

Finally, it is simply not possible to anticipate all future uses and enhancements of the system under design. To achieve resilience, it is standard practice to design and build a system that can adapt over time to an evolving environment. Traditional attributes such as modifiability, detailed models, architecture, and documentation all contribute to the ability of a system to easily evolve based on changing conditions.

However, a system's evolution must not be accessible to an adversary who could introduce vulnerabilities and undesired behavior to a system. Furthermore, the authorized modification of the system must not subvert or corrupt any previous cybersecurity capability. The modification could introduce undesired consequences such as exposing encryption keys, removing monitoring capabilities, introducing previously mitigated vulnerabilities, and creating a new attack surface for an adversary.

Subversion or corruption is frequently an unintended consequence of a system modification because many cybersecurity capabilities remain silent until the system is attacked. Engineers cannot therefore easily test these capabilities through unit testing of the new system modification.

Adaptability is key for system resilience, but to the extent possible engineers must maintain cybersecurity goals and requirements for any modifications they make.

Conclusion

Investing in cybersecurity does not mean that all engineers will or should become cybersecurity experts. However, it does mean that engineers in different fields will be able to create a resilient system with safety and security assurances. Asking some simple questions can help to determine when and how a cybersecurity engineer should be added to the team.

The questions above are designed to be easily answered by any trained engineer working with a system. Often, however, the answers can reveal an unanticipated level of complexity. Although this complexity may introduce up-front cost in the design, the resulting system will be safer, more adaptable, and more secure than it would be without those answers—i.e., if engineers rely only on patching the system after having designed and built

it. Mitigations at every layer—from code to software to network—can help ensure that a system functions correctly even under adverse conditions.

The addition of a cybersecurity expert to a system design team will prompt needed thinking about security. As engineers incorporate more machine learning and other artificial intelligence capabilities in the systems they build, these questions are crucial to ensure the safe and secure operation of the systems that people encounter and use every day.

References

McDermott TA, Canedo A, Clifford MM, Quirós G, Sitterle VB. 2019. System assurance in the design of resilient cyber-physical systems. In: Design Automation of Cyber-Physical Systems, eds Al Faruque M, Canedo A. Cham, Switzerland: Springer.

Silver D. 2017. Adversarial traffic signs. Medium.com, August 15.

EES Perspective

When White Hats Wear Black Hats: The Ethics of Cybersecurity

Dianne Martin is professor emeritus of computer science, George Washington University.

C. Dianne Martin

Codes of ethics for computer professionals have been evolving over the past four decades. The profession of computer science has matured to the extent that well-developed ethical principles have emerged to guide the general practice of the discipline.

Background

Since the first code of computer ethics was adopted by the Association for Computing Machinery in 1973, there has been a realization across the profession that the hardware, software, and networking systems developed by information technology (IT) experts are embedded with serious social and ethical implications, and those who develop and maintain them must adhere to a high moral standard if the interest of a dependent public is to be protected. This is particularly true in the area of cybersecurity, where the tools used by the guardians of cyberspace are often the same as those used by perpetrators of evil intent.

The technical capabilities of computers in general and the internet in particular continue to develop rapidly beyond the human ability to guarantee

---

This column is produced in collaboration with the NAE's Center for Engineering Ethics and Society to bring attention to and prompt thinking about ethical and social dimensions of engineering practice.

a safe and ethical environment in cyberspace. While legal, medical, accounting, and other established professions have legally binding codes of conduct overseen by longstanding regulatory bodies, IT security professionals have yet to establish formal guidance or universal checks and balances. In addition, the industry lacks an independent register maintained by an oversight entity to determine who can practice ethical hacking or security research.

Cyberspace is plagued with rogue individuals, groups, and even state-sponsored actors intent on fraud, crime, espionage, damage to infrastructure, even terrorism. This causes the cybersecurity landscape to shift quickly as organizations seek to fill a growing gap for security experts amid a shortfall of skilled graduates (Knowles 2016b).

To combat this threat, it is necessary to train computer security experts along two fronts. The first is to provide the technical training to use the same tools and strategies as the bad actors in order to neutralize such threats. The second, often neglected, is to provide rigorous ethical training for cybersecurity professionals.

In this charged climate, the tendency is to focus on fast-tracking the development of technical knowledge in order to deploy new talent to the front line quickly, not considering the lack of maturity of the new recruits, who may end up abusing their abilities. Lacking awareness of the context of cybersecurity ethics, many have to defer to their personal moral compass, which can lead to mistakes as often as it leads to good decisions.

"Rainbow of Hackers"

In addressing this issue, Aidan Knowles (2016a), an Ethical Hacking Engineer for IBM in Ireland, defines what he calls the "rainbow of hackers": black hats (the bad guys), white hats (the good guys), and grey hats (somewhere in between):

> Generally, white-hat and black-hat hackers do similar tasks. Both target applications, networks, computer systems, infrastructure and occasionally even people; often, both camps use the same tools and resources. But their work is not completely homogeneous, differentiating on some major points—including motivation, permission, legality and time.
>
> A white hat is commonly employed or contracted to carry out an attack under explicit permission and clear-cut boundaries. The goal of white hats' work is to research, find and test vulnerabilities, exploits and viruses in their defined targets. The findings of these professional engagements are reported directly to the target to enable them to fix any holes and strengthen their overall security posture. White hats are also sometimes involved in developing security products and tools….
>
> In contrast, black hats cause great intentional damage and profit at the expense of their targets.… [They include] cybercriminals, cyber spies, cyber terrorists and hacktivists…. Malicious actors may not always be operating externally from their victim. Research suggested that the insider threat within an organization's networks and premises, including from current or former employees and contractors, is responsible for a large portion of successful hacks. To carry out attacks, black hats may develop their own malicious tools but will frequently employ or repurpose existing white-hat software.
>
> Grey hats, as the name suggests, are more ambiguous in their definition. Their work may be classified as leaning toward good or bad on the spectrum depending on your perspective. The term grey hat is sometimes used to describe those who break the law but without criminal intent. This definition may include cyber vandals who deface websites and so-called rogue security researchers who publicly share discovered vulnerabilities without notifying or receiving prior permission from their targets.

Individuals' reliance on their personal moral compass leads to mistakes as often as it leads to good decisions.

Without clear ethical standards and rules, cybersecurity professionals may be almost indistinguishable from the black-hat criminals against whom they seek to protect systems and data.

Standards and Codes of Ethics

At present cybersecurity managers have to rely on reputation and background checks to determine the trustworthiness of potential hires. If new hires betray this trust by behaving unethically, there is no third-party committee or board to evaluate the consequences of these actions and to rule in the context of the profession as a whole. In most cases, rogue security professionals cannot be struck from a register or removed from a database because an industrywide database does not exist.

What are appropriate ethical standards for cybersecurity experts and how can those standards be integrated in their training and job experiences? Several

associations—the Information Systems Security Association (ISSA), International Information Systems Security Certification Consortium [(ISC)²], and SysAdmin, Audit, Network and Security (SANS)—have voluntarily developed ethics codes for cybersecurity (those of ISSA and (ISC)² are reproduced below). However, in general industry professionals are not required to subscribe to these bodies or adhere to their codes of conduct. The exception is the (ISC)² Certified Information Systems Security Professional (CISSP), who could lose this certification if evidence reveals violation of the (ISC)² Code of Ethics.

What are appropriate ethical standards for cybersecurity experts and how can they be integrated in training and job experiences?

ISSA Code of Ethics

To fill the compelling need for an ethics code for cybersecurity professionals, the Information Systems Security Association was the first to establish the following Code of Ethics[1] for its members in 2006:

1. Perform all professional activities and duties in accordance with all applicable laws and the highest ethical principles;
2. Promote generally accepted information security current best practices and standards;
3. Maintain appropriate confidentiality of proprietary or otherwise sensitive information encountered in the course of professional activities;
4. Discharge professional responsibilities with diligence and honesty;
5. Refrain from any activities which might constitute a conflict of interest or otherwise damage the reputation of or is detrimental to employers, the information security profession, or the Association; and
6. Not intentionally injure or impugn the professional reputation or practice of colleagues, clients, or employers.

Although the code is quite general, it does provide a moral framework for ethical cybersecurity practice based on the principles of integrity, respect for confidentiality and privacy, and avoidance of conflicts of interest.

(ISC)² Code of Ethics

Similarly, (ISC)² recognized the need for an ethics code to cover its certification of expertise in cybersecurity. Information security professionals certified by (ISC)² are informed that such certification is a privilege that must be earned and maintained. In support of this principle, all (ISC)² members are required to commit to fully support the (ISC)² Code of Ethics. Members who intentionally or knowingly violate any provision of the code are subject to action by a peer review panel, including possible revocation of certification. (ISC)² members are obligated to follow an ethics complaint procedure to report any action by another member that breaches the code. There are only four mandatory canons in the code, which are high level and general. The (ISC)² Code of Ethics Preamble and Canons[2] are shown below:

(ISC)² Code of Ethics Preamble:

- The safety and welfare of society and the common good, duty to our principals, and to each other, requires that we adhere, and be seen to adhere, to the highest ethical standards of behavior.
- Therefore, strict adherence to this Code is a condition of certification.

(ISC)² Code of Ethics Canons:

1. Protect society, the common good, necessary public trust and confidence, and the infrastructure.
2. Act honorably, honestly, justly, responsibly, and legally.
3. Provide diligent and competent service to principals.
4. Advance and protect the profession.

The basic elements of the canons are similar to all codes of ethics for IT professionals: the need to serve and protect a dependent public, the requirement of competence to perform the work, and a focus on overall integrity. As stated on the (ISC)² website, such high-level guidance is not intended to be a substitute for the ethical judgment of the professional.

Examples of Actions Requiring an Ethical Response

Using the general ethical principles delineated by ISSA and (ISC)² as guidelines, following are examples

[1] https://www.issa.org/issa-code-of-ethics/

[2] https://www.isc2.org/Ethics

of inappropriate ("low road") and appropriate ("high road") actions that fall within typical cybersecurity duties (Martin 2017; Tull 2016).

Denial of Service Attack Recovery

In the course of doing firewall security scans, the security team may discover a port call that results in a denial of service (DoS) attack. The low road response to such an attack is to hack and attack back at the host. However, that could result in other sites being caught in the DoS crossfire. The high road response—in keeping with ISSA standards 1 and 2 as well as (ISC)² canons 1 and 2—is to block the attack and gather forensic evidence to respond to it legally and ethically.

Penetration Testing and Response

Cybersecurity professionals often do penetration testing to determine the robustness of firewalls and security features in a system. If they detect a weakness or vulnerability that could allow a remote hacker to take control of the system and cause significant harm, there are two possible responses.

The low road response is immediate full disclosure—publishing full details of the vulnerability as soon as possible and making the information available to everyone without restriction. This could enable black-hat hackers to exploit the weakness before it is fixed.

The high road response, "responsible disclosure," is more nuanced. Responsible disclosure requires the security expert to confidentially report the weakness to the company, work with the company to develop a fix within a given timeframe, and then publicly disclose the vulnerability and the fix at the same time (Tull 2016). This response would be in keeping with ISSA standards 3, 5, and 6 and (ISC)² canons 1 and 3.

Fighting Malignant Worms with Benign Worms

A cybersecurity expert believes that a benign worm might be able to patch a known vulnerability, inoculate systems to protect them from a malignant worm, and keep it from spreading. Should she release it "in the wild"? The low road decision would be to release it and hope for the best. The high road approach, consistent with ISSA standards 2, 4, and 5 and (ISC)² canons 1, 2, and 4, would be to make the benign worm code publicly available with caveats that knowledgeable professionals should use it with care.

Conclusion

Cybersecurity experts—the white hats—work with sensitive data, have access to company and national secrets and generally wield much power over networks, systems, and data. How individuals handle this responsibility comes down to their ethical yardstick, and reinforcing that ethical yardstick is a fundamental responsibility of the programs that train these experts.

One way to move in this direction would be to require certification for all cybersecurity experts, with mandated ethics training and periodic updates to this training required to maintain certification. Although this would not be a guarantee to prevent malicious hacking, it would at least ensure that all trained professionals were made aware of ethical issues related to the exercise of their technical skills.

The high road response to a DoS attack is to block it and gather forensic evidence to respond to it legally and ethically.

It is very likely that rogue actors will continue to proliferate on the internet regardless of the presence of legally binding codes of conduct and other measures, just as criminal acts persist even in the presence of well-established legal systems. This makes an even stronger case for rigorous ethics training to be required for all cybersecurity professionals as one line of defense against such attacks.

In a recent comprehensive white paper on cybersecurity and ethics (Yaghmaei et al. 2017, p. 3), the authors observe that

> the ethics of cybersecurity is not an established subject. In all domains, cybersecurity is recognized as being an *instrumental value*, not an end in itself, which opens up the possibility of trade-offs with different values in different spheres. The most prominent common theme is the existence of trade-offs and even conflicts between reasonable goals, for example between usability and security, accessibility and security, privacy and convenience.

They go on to state that one of the most important features of cybersecurity is to sustain trust in institutions and maintain the integrity of data. Given the inherent

tension between possible end goals and the importance of preserving overall trust in the internet, it can be argued that developing ethically aware cybersecurity experts is as important as developing technically competent cybersecurity experts.

Acknowledgments

The author gratefully acknowledges the thoughtful suggestions and edits made to improve this article by Keith Miller, University of Missouri–St. Louis, and *The Bridge*'s managing editor, Cameron Fletcher.

References

Knowles A. 2016a. How black hats and white hats collaborate to be successful: The hacker rainbow. Security Intelligence, May 4. Online at https://securityintelligence.com/how-black-hats-and-white-hats-collaborate-to-be-successful/.

Knowles A. 2016b. Tough challenges in cybersecurity ethics. Security Intelligence, Oct 12. Online at https://securityintelligence.com/tough-challenges-cybersecurity-ethics/.

Martin CD. 2017. Black hat, white hat: The ethics of cybersecurity. ACM Inroads 8(1).

Tull J. 2016. A Snapshot in Cybersecurity Ethics. Formerly available at-security-ethics.

Yaghmaei E, van de Poel I, Christen M, Gordijn B, Kleine N, Loi M, Morgan G, Weber K. 2017. CANVAS White Paper 1: Cybersecurity and Ethics. Vrije Universiteit Brussel. Online at https://ssrn.com/abstract=3091909.

There are opportunities for engineering to make transformative contributions to the curtailment of human trafficking.

A Call to the Engineering Community to Address Human Trafficking

Jonathan P. Caulkins, Matt Kammer-Kerwick, Renata Konrad, Kayse Lee Maass, Lauren Martin, and Thomas Sharkey


Jonathan Caulkins


Matt Kammer-Kerwick


Renata Konrad


Kayse Lee Maass


Lauren Martin


Thomas Sharkey

Human trafficking (HT) is a horrific and seemingly intractable problem that is typically construed as falling beyond the purview of engineers.

This paper argues that engineering systems analysis can produce important insights concerning HT operations and ways to reduce its frequency. Three cases of such systems analysis illustrate (a) the limitations of individual-level interventions against sex trafficking, (b) the benefits of applying network analysis and interdiction models to HT supply chains, and (c) options to reduce the use of trafficked labor in the preparation and distribution of fish products.

Jonathan Caulkins (NAE) is Stever University Professor of Operations Research and Public Policy at Carnegie Mellon University's Heinz College. Matt Kammer-Kerwick is a research scientist at the Bureau of Business Research at the University of Texas at Austin. Renata Konrad is associate professor of operations and industrial engineering at Worcester Polytechnic Institute. Kayse Lee Maass is an assistant professor of mechanical and industrial engineering at Northeastern University. Lauren Martin is an associate professor in the University of Minnesota's School of Nursing. Thomas Sharkey is associate professor of industrial and systems engineering at Rensselaer Polytechnic Institute.

The International Labour Office (ILO 2017) has estimated that there are 25 million victims in forced labor around the world, including 4.8 million in forced sexual exploitation. There are opportunities for engineering to make transformative contributions to the curtailment of human trafficking.

Introduction

Human trafficking is a fundamental human rights abuse and far more common than the average person imagines. It has not traditionally been viewed as a problem for engineers to address, but it can be thought about as a complex system and deserves attention alongside other significant challenges for engineering.

Terminology

Despite the word "trafficking" in the term, HT is defined not by the movement of victims but rather by their inability to escape exploitation. Victims are restrained by debt bondage, threats to family, fear of deportation, and shame more often than actual shackles.[1]

Human trafficking can be thought about as a complex system.

HT is not the same as *human smuggling*, the facilitation of illegal border crossing by willing travelers. While some HT victims are smuggled across borders, HT victims can be exploited in their own country, and undocumented immigrants may not work at all or work under conditions that they are free to leave.

Sex trafficking is commercial sexual exploitation through force, fraud, or coercion. Because only adults can give consent, any commercial sexual exploitation of minors is considered sex trafficking; the adult sex industry is composed of both trafficked and nontrafficked individuals (Martin et al. 2017). Sex trafficking networks in the United States differ depending on the citizenship of the victim (Busch-Armendariz et al. 2009). Organizations that exploit US citizens often both recruit and exploit their victims, and are often fairly autonomous; networks that traffic foreign citizens often include multiple interdependent organizations, including some that specialize in recruitment and transport and others that "employ" the victims in sex work.

Labor trafficking is most common in agriculture, health and beauty services, landscaping, domestic work, construction, and manufacturing (Polaris 2017). Victims in the United States include citizens and foreign nationals, for whom recency of immigration is a significant risk factor. Imported products may have been produced by trafficked labor abroad.

Trafficking Goods vs. Services

It is easy to spot similarities between HT and trafficking in drugs, weapons, and wildlife. They are all transactional crimes and, while HT clearly has victims, they usually do not report the crime. Yet there are important differences. The other forms of trafficking involve criminals delivering goods to consumers, whereas, from the traffickers' perspective, the "object" in human trafficking is more akin to a capital asset: it is exploited to produce services (sex acts, labor) and the services are sold to consumers.

This distinction has implications that some simple calculations can elucidate. The parameter values used here and elsewhere in this paper are not laboratory-measured with high precision; it is difficult to gather reliable data on HT because it is hidden, illegal, and dangerous. Yet even approximate values can produce insight through the sorts of calculations engineers do.

Consider how often a gang with $500,000 in annual revenues purchases a trafficked "product" (human or drug) from its "suppliers." A gang that generates such revenue by forcing HT victims to provide commercial sex services might at any time be exploiting five HT victims, assuming a typical HT victim in the United States can produce about $100,000 in revenue per year, e.g., by averaging ten $30 sex acts per day (Kara 2009). If victims were exploited for an average of 2½ years before being released, sold, or killed (Kara 2009), the gang would need to acquire only two HT victims per year, on average. In contrast, a drug gang with $500,000 in annual revenues might purchase drugs from its supplier almost daily if, as is typical, drugs were acquired in lots of 50 retail units that sold for $30 each (Caulkins et al. 2016), since $500,000/($30 * 50) = 333 days of purchases. That difference in the number of "product acquisitions" (two per year vs. almost daily) has implications for the supply chains' structure, scale, and vulnerability to law enforcement.

This introduction underscores three points:

- Most HT is business activity that involves supply chains in ways that street crimes such as assault or robbery do not.

[1] See UNODC (2000) for a formal definition of human trafficking.

- HT differs from other market crimes and must be analyzed from first principles, rather than presuming that conclusions concerning drug trafficking, for example, apply to HT.
- HT encompasses diverse industries, and the particulars matter. Polaris (2017, p. 2) identifies 25 types of HT in the United States, "Each [with] its own business model, trafficker profiles, recruitment strategies, victim profiles, and methods of control that facilitate human trafficking."

This article offers three case studies illustrating how engineering systems analysis can shed light on different segments of HT and a brief guide to how engineers can help address it.

Victim-Level Interventions and the Problem of Replacement

Human trafficking is a complex criminal activity, and interventions that address one component (e.g., root causes, services for victims, or interdiction of trafficking operations) may affect other components. We show how victim-level interventions may paradoxically create a net increase in the number of victims. For simplicity's sake, we focus on sex trafficking victims. A more thorough exploration would include independent sex workers, who may experience exploitation but not trafficking per se.

Conceptual Model of Operational Costs and Market Equilibria

Sex trafficking operations profit by supplying sex to buyers through the control and exploitation of victims who provide sex. Their operations are shaped by market forces such as demand elasticity and the cost of replacing victims who escape their trafficking situation (Martin and Lotspeich 2014). It is relevant to consider how victim-level interventions affect operational costs, the market clearing price and volume, and the number of HT victims.

From the perspective of an HT operation, interventions that remove existing HT victims prompt the acquisition of replacement "workers." This replacement imposes some costs that are passed on to sex buyers, potentially deterring some of them.

If replacement costs are high, victim-level interventions could drive up prices enough to shrink the market and decrease the number of trafficking victims. However, if replacement costs are low, the trafficker could obtain a replacement without much increase in the price to purchase sex. That could leave the number of victims being exploited at any given time nearly unchanged, while reducing the average duration of exploitation, leading to a greater flow of new victims to meet demand. Thus while comprehensive services and support for victims are important and needed, they may not be sufficient to reduce impacts on victims and society.

Bringing Numbers to Bear

There is no definitive research on pricing structures for sex acts, lengths of time that victims are exploited, or victim replacement costs, all of which vary by context, locale, country, type of sex act, market segment, and more (Dank et al. 2014).

The number of human "acquisitions" has implications for HT supply chains' structure, scale, and vulnerability to law enforcement.

For this exercise, we consider estimates for a brothel-type operation in New York City, where a victim costs $3,000 to recruit, is trafficked for 2½ years, and performs ten $30 sex acts per day (Kara 2009). In this case, the traffickers' revenues per victim recruited are $273,750 and the $3,000 recruitment cost thus consumes approximately 1 percent of the traffickers' gross revenue.

When replacement costs are so low relative to revenues per victim, victim-level interventions may actually increase the number of HT victims: if the interventions reduce by half the average time an individual is trafficked, then traffickers must obtain two people to provide the same volume of activity. Suppose the demand elasticity is such that the ~1 percent increase in suppliers' cost structure reduces market volume by 4 percent. Then, where there had been 100 organizations recruiting 2 HT victims per year, there are now 96 organizations recruiting 4 per year. The number of victims at any point in time has been reduced by 4 percent from 500 to 480, but the number drawn into HT each year *increases* by 92 percent from 200 to 384.

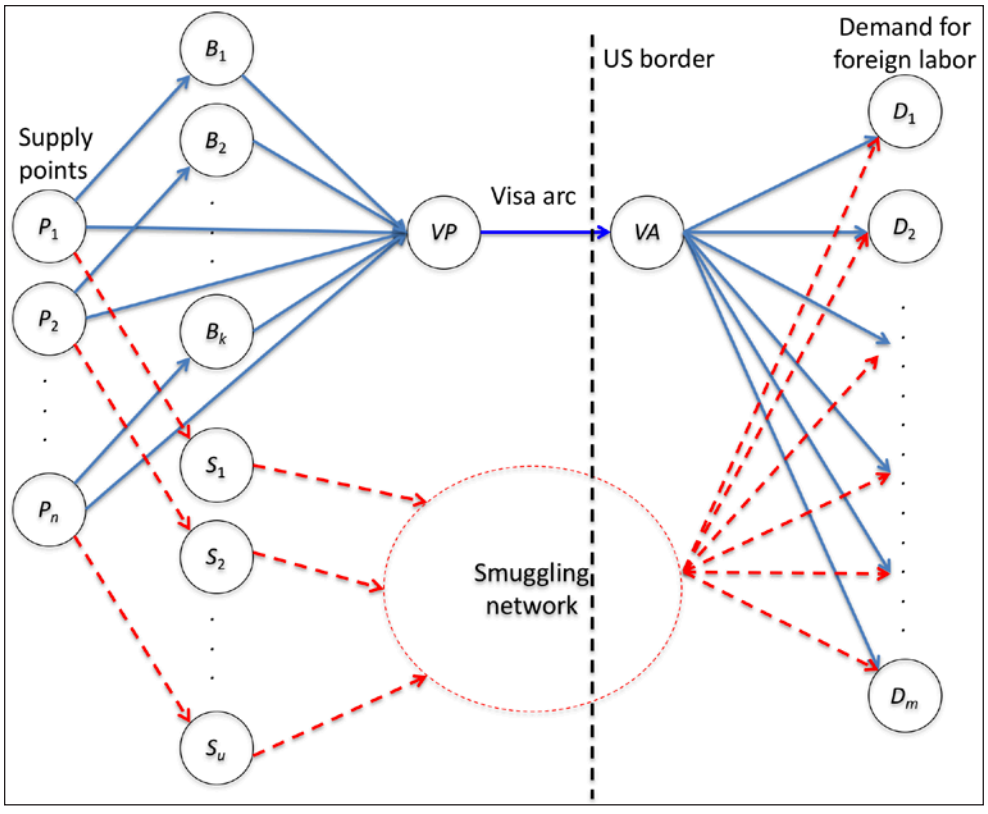

FIGURE 1 Network model of labor supply chain of foreign nationals. Blue lines denote a legal network and red dashed lines an illegal network for foreign laborers entering the United States. *B* = labor broker; *D* = employer (demand); *P* = source of potential laborers (town, region, country); *S* = smuggler (trafficker); *VA* = visa accepted; *VP* = visa process.

Of course, not all commercial sex is supplied by HT victims, and not all HT organizations are identical (Martin and Lotspeich 2014). HT is a dynamic system and improvement in one area may worsen conditions in another. While it is imperative to support victims in exiting an HT situation, such efforts may not deter traffickers from continuing their operations by recruiting replacements, particularly in market segments where the cost of replacing a victim is low, as in the New York example.

Network Analysis for Disrupting Human Trafficking

Network analysis can shed light on ways to disrupt human trafficking. A network is a collection of nodes and arcs where nodes model potential HT states, locations, or processes and arcs model the relationships between the nodes.

Figure 1 provides a conceptual network of the supply chain of foreign nationals entering the United States for work. Nodes P_1 through P_n are populations of potential laborers and countries, regions, or towns. *D* nodes represent US employers that are willing to hire foreign laborers, who move through either a legal or illegal network to work in the United States.

The network of solid blue lines represents the process of legally obtaining a visa, whether directly (the arc from the supply point to the "visa process," *VP*, node) or through a labor broker (represented by the *B* nodes). Note that the network bottlenecks at the visa arc, where the "visa accepted" (*VA*) node represents entry into the country. The network of red dashed lines represents persons illegally entering the country to work.

Laborers may be vulnerable to HT whether they enter the United States through the legal or illegal network. Brokers may demand large sums to begin the visa process, creating a form of debt bondage that makes it difficult for the laborer to leave the US employer no matter the working conditions. After Hurricane Katrina, for example, brokers required Indian welders to pay a recruitment fee of $10,000–$20,000 to help obtain an H-2B visa to work for Signal International, and these workers became indentured to the brokers (Desai 2015).

Laborers who enter the country illegally are more likely to be victims of HT. For example, undocumented farm laborers in North Carolina are 2.8 to 5.2 times more likely to be victimized than are documented laborers, depending on the type of victimization (Barrick et al. 2013, 2014). Another study found that 31 percent of surveyed unauthorized migrant laborers in San Diego County had been HT victims (Zhang 2012).

Therefore, one potential intervention for reducing labor trafficking would be to decrease the flow of

laborers in the illegal network. Interventions in the illegal network might increase its cost, making the flow less economical than legal employment practices. Alternatively, policies could increase the capacity of the legal network by reducing the visa bottleneck. Expanding the supply of visas should reduce HT among foreign nationals in the illegal network and will increase flow in the legal network, where HT occurs at a much smaller rate. Further, it may shift the balance between foreign and domestic laborers in the workforce.

The effectiveness of policy efforts to decrease HT will depend on the nature and extent of changes and on whether the legal network is properly designed, resourced, and incentivized. Network analysis provides a framework for accounting for these effects in order to identify the most cost-effective way to have the desired impact on disrupting HT.

Seafood Supply Chains and Labor Trafficking

Fishing is big business: aquaculture and a global fishing fleet of roughly 4.6 million vessels annually produce over 160 million tons of fish for global consumption (FAO 2014). Demand for inexpensive seafood and poor regulation have created conditions ripe for labor exploitation (Sutton and Siciliano 2016), leading to pervasive and well-documented human trafficking in the seafood supply chain (ILO 2012; IOM 2011; Pearson et al. 2006; Stringer et al. 2014), although, because illegal labor exploitation is intentionally hidden, it is difficult to quantify the extent.

US retail seafood sales are approximately $100 billion per year, of which only about 10 percent is domestically produced (White 2016). The US government is not blind to problems in the seafood supply chain. Its Seafood Import Monitoring Program (SIMP) imposes reporting and recordkeeping requirements to prevent illegal, unreported, and unregulated-caught or misrepresented seafood. Effective January 1, 2018, it requires the tracing of 13 priority species from the point of harvest to entry in US commerce. But this new program applies only to selected species and is motivated primarily by environmental concerns.

We offer the following questions to guide engineering design and to determine what it would take to effectively inspect the entire US fish supply chain to eliminate HT labor exploitation.

What is the least costly way to deploy a portfolio of monitoring tools to combat HT?

The diverse participants and the scale and complexity of the international fishing supply chain create immense challenges for monitoring HT; maritime enforcement institutions cannot simply post observers on all 4.6 million vessels operating internationally. What is the least costly way to deploy inspections, audits, automatic identification systems (AIS), voluntary compliance, satellite technology, or other methods to supervise fish harvesting and processing to guarantee that the supply chain is void of HT? How can such tools be combined in a cost-efficient manner to identify both illegal fishing and human trafficking?

How can data analytics identify vessels that use HT victims?

Algorithms could be developed to identify characteristics associated with HT in deep seas and then (using tools such as Skytruth) detect vessels exhibiting suspicious behavior. For example, gaps in satellite data may indicate that a fishing vessel turned off its AIS signal to avoid disclosing its location (Sutton and Siciliano 2016). Likewise, a meeting between two ships at sea may signal offshore transshipment to avoid landing at port, or a fishing vessel that remains at sea for extended periods of time may be more likely to be associated with human rights abuses (e.g., trapping its crew aboard).

How can engineering design enhance inspection of the US fish supply chain to eliminate HT labor exploitation?

How can a systems approach be used to determine whether monitoring efforts could be paid for with an excise tax on fish products?

Monitoring the seafood supply chain, even with an efficient design, would not be free. Would consumers be willing to pay for the cost to implement a supply chain monitoring system to ensure that fish products sold in the United States are fair trade?

Superficially, consumer financing of supply chain monitoring seems feasible. If the estimated $100 billion seafood consumer market were taxed at 0.5 percent,

$0.5 million would be available to support monitoring and enforcement—far more than SIMP costs. However, such a tax would affect sales. Consumer demand, environmental sustainability, labor trafficking, aquaculture, and socioeconomic development interact interdependently and form a large system with complex, dynamic, diverse, and nonlinear characteristics. A simulation model could provide insights related to the effects of an excise tax on US fish consumption, global fish stocks, and trafficked labor, thereby addressing a range of policy and behavior questions.

Guiding Thoughts for Engineers

We offer the following thoughts for those whose interest has been piqued to consider working in this domain.

First, it is important to take account of the complexity and ambiguity inherent in these systems. Related crimes and other misconduct against HT victims are common (Aronowitz et al. 2010)—identity theft, forgery, assault, theft of documents, and forced participation in other crimes—and HT overlaps other problem domains, such as drugs, warfare, other forms of violence, and labor exploitation.

There are risks of unintended consequences to well-intentioned interventions.

Further complicating matters, substantial challenges exist in identifying and measuring victimization. For example, the criteria needed for prosecutorial proof of victimization differ from the criteria needed by those who provide services to victims. For engineers working in this field, it is important to dig deep and understand the origins of numbers, available data, and their proper interpretation.

Much current scholarship focuses analysis solely on HT victims and perpetrators. Engineers can contribute by studying the entire system within which HT occurs. There are multiple relevant decision makers—employers, agents and recruiters, public officials, retailers, and consumers—with different goals and levels of information. As illustrated by the seafood supply example, an ecosystem view may offer solutions, particularly if it incorporates the dynamics between ecosystem participants and recognizes the alternatives that are available when these actors make decisions. These dynamics and alternatives include the limited choices available to victims and the various coercive strategies used by perpetrators and their agents to lure victims.

Interventional targets for such complex systems vary in scale and effort to deploy. Technology-driven initiatives, like the monitoring tools discussed above for illegal fishing, are complicated solutions in absolute terms, but the benefits can be enhanced by parallel investment in changes to social norms through community education and training about labor and basic human rights (Battista et al. 2018). Such fundamental changes to social norms likely require more time to diffuse through communities compared to the already substantial time to design and deploy monitoring tools.

There are risks of unintended consequences to well-intentioned interventions. Anyone working in this domain needs to listen and look for potential harms as they bring important new insights from the application of powerful systems thinking. For example, a fishing tax as an intervention to fund antitrafficking programs might also reduce demand and suppress legal fishing operations with fair employment policies. Caution is warranted particularly given the lack of empirical data on so many aspects of HT.

Conclusions

Human trafficking is a morally wicked practice, and stopping it is a wicked problem in the formal sense of the term (Rittel and Webber 1973). It is a domain in which sincere efforts may absorb enormous energy without significantly shrinking the problem.

Efforts to reduce the scale of HT must be grounded in a systems-level understanding of its many interconnected and dynamically interacting parts. The primary actors—the organizations that recruit, transport, and exploit HT victims—are fundamentally business operations with supply chains. Engineers understand how to describe and manipulate supply chains. We usually work to make them more efficient, but the same systems analysis applied in reverse can shed light on ways to make those supply chains less efficient and, ideally, uncompetitive with supply chains that provide the same services without exploiting forced labor. HT is both a business operation and a violation of basic human rights.

To date, engineers have not been much involved in combating HT. But the discipline now recognizes

engineers' responsibility to address society's grand challenges (Vest 2008), and defeating HT can certainly be seen as a similarly worthy challenge for engineers.

Acknowledgments

We are grateful to Kelle Barrick for her thoughts and help in the preparation of this paper. The National Science Foundation provided funding for a four-day workshop to explore new applications of Operations Research and Data Analytics to End HT (CMMI-1726895), which was the impetus for this paper (Kammer-Kerwick et al. 2018).

References

Aronowitz A, Theuermann G, Tyurykanova E. 2010. Analysing the business model of trafficking in human beings to better prevent the crime. Vienna: Organization for Security and Co-operation in Europe Office of the Special Representative and Coordinator for Combating Trafficking in Human Beings.

Barrick K, Lattimore PK, Pitts W, Zhang SX. 2013. Indicators of Labor Trafficking among North Carolina Migrant Farmworkers. Research report prepared for the National Institute of Justice, Washington.

Barrick K, Lattimore PK, Pitts WJ, Zhang SX. 2014. Labor trafficking victimization among farmworkers in North Carolina: Role of demographic characteristics and acculturation. International Journal of Rural Criminology 2(2):225–243.

Battista W, Romero-Canyas R, Smith SL, Fraire J, Effron M, Larson-Konar D, Fujita R. 2018. Behavior change interventions to reduce illegal fishing. Frontiers in Marine Science 5:403.

Busch-Armendariz N, Nsonwu M, Cook Heffron L. 2009. Understanding human trafficking: Development of typologies of traffickers, Phase II. 1st Annual Interdisciplinary Conf on Human Trafficking, Oct 29–31, Lincoln NE.

Caulkins JP, Disley E, Tzvetkova M, Pardal M, Shan H, Zhang X. 2016. Modeling the structure and operation of drug supply chains: The case of cocaine and heroin in Italy and Slovenia. International Journal of Drug Policy 31:64–73.

Dank M, Khan B, Jay J, Downey CPM, Mayer D, Owens C, Yu L. 2014. Estimating the Size and Structure of the Underground Commercial Sex Economy in Eight Major US Cities. Washington: Urban Institute.

Desai R. 2015. Landmark human trafficking case ends with bankruptcy for Signal International, Inc. Human Rights First, Jul 24.

FAO [UN Food and Agriculture Organization]. 2014. The State of World Fisheries and Aquaculture: Opportunities and Challenges. Rome.

ILO [International Labour Office]. 2017. Global Estimates of Modern Slavery: Forced Labour and Forced Marriage—Executive Summary. Geneva.

IOM [International Organization for Migration]. 2011. Trafficking of Fishermen in Thailand. Bangkok.

Kammer-Kerwick M, Busch-Armendariz N, Talley M. 2018. Disrupting Illicit Supply Networks: New Applications of Operations Research and Data Analytics to End Modern Slavery. Workshop report and proposed research agenda. University of Texas at Austin.

Kara S. 2009. Sex Trafficking: Inside the Business of Modern Slavery. New York: Columbia University Press.

Martin L, Lotspeich R. 2014. A benefit-cost framework for early intervention to prevent sex trading. Journal of Benefit-Cost Analysis 5(1):43–87.

Martin L, Melander C, Karnik H, Nakamura C. 2017. Mapping the Demand: Sex Buyers in the State of Minnesota. Minneapolis: University of Minnesota.

Pearson E, Punpuing S, Jampaklay A, Kittisuksathit S, Prohmmo S. 2006. The Mekong Challenge: Underpaid, Overworked and Overlooked—The Realities of Young Migrant Workers in Thailand, vol 1. Bangkok: International Labour Organization.

Polaris. 2017. The Typology of Modern Slavery: Defining Sex and Labor Trafficking in the United States. Washington.

Rittel HW, Webber MM. 1973. Dilemmas in a general theory of planning. Policy Sciences 4(2):155–169.

Stringer C, Simmons G, Coulston D, Whittaker DH. 2014. Not in New Zealand's waters, surely? Linking labour issues to GPNs. Journal of Economic Geography 14(4):739–758.

Sutton T, Siciliano A. 2016. Seafood Slavery. Washington: Center for American Progress.

UNODC [United Nations Office on Drugs and Crime]. 2000. United Nations Convention against Transnational Organized Crime. Vienna.

Vest CM. 2008. Context and challenge for twenty-first century engineering education. Journal of Engineering Education 97(3):235–236.

White C. 2016. American seafood consumption up in 2015, landing volumes even. Seafood Source, Oct 27.

Zhang SX. 2012. Looking for a hidden population: Trafficking of migrant laborers in San Diego County. San Diego State University.

An Interview with . . .

Deanne Bell, TV Host and Founder-CEO of Future Engineers

RON LATANISION (RML): Deanne, thank you for joining us today. We're very happy to have this opportunity to talk with you. Let's begin right at the beginning. You got your undergraduate degree in mechanical engineering from Washington University in 2002. How did you become interested in mechanical engineering?

DEANNE BELL: As a little girl I loved to tinker and invent and build. I participated in an afterschool problem-solving program and was on a team of mainly girls. The team would come over to my house after school and my parents let us build and invent in my backyard. We were very resourceful—we would go to thrift stores, ask them for junk they weren't using, and take it apart. For example, for a "wash world" theme that we dreamt up, we made a massive musical washing machine with eyelashes that could be played remotely; it spun and shot out golf balls painted to look like bubbles that played drums. We usually presented our inventions in the context of a skit, so there was a stylistic element to our creations.

Each of us on the team had our different specialties. The engineering side of things resonated with me and another girl and we both ended up becoming engineers. I loved using power tools when I was little.

RML: I can sense that, given your involvement with the DIY Network. Were your mom and dad engineers?

MS. BELL: My father is an electrical engineer and my mother was a math teacher before becoming a full-time homemaker. I was raised in a family that embraced the sciences and engineering, but I was also raised to be very well-rounded. I did lots of sports and music lessons and all kinds of things. I wasn't necessarily encouraged to pursue engineering, but it was something I really wanted to pursue.

It's funny, though, because when I was little I didn't know what an engineer did for a living. I remember being in elementary school and my dad offering to help me with math homework. I scoffed at him and said, "No, Mom's the math teacher." For a long time I didn't know what an engineer was, and I think that's part of the reason I'm committed now to communicating it to the younger generation, so we can build that identity and give students an idea of what engineering is even at the youngest ages.

RML: That's certainly something that the National Academy of Engineering shares in terms of our interest in getting young people, particularly young women, involved in engineering. Tell us about Future Engineers.

MS. BELL: I'm an engineer and worked in mechanical design before I became a television host. Between TV shows there was some down time so I found myself collaborating with a local nonprofit to create engineering workshops and I thought, 'How can I extend this? On TV there is an engaged audience and a lot of students watching. How can I take this further and integrate my engineering work on TV into curriculum?' That simmered for a long time. And when I was speaking at corporate

---

This conversation took place May 14. It has been edited for concision and clarity.

events I kept hearing "We want to get more involved in STEM but we don't know the best way to plug in." Then I thought, 'How can I combine my experience communicating engineering in the media, and my experience developing workshops for students, and my experience with industry? How can we come up with challenges based on research and emerging technologies to create something fun, new, and relevant for kids out there?'

The idea simmered until one day I decided, 'Let's make this happen.' I saw that NASA had a 3D printer going to the International Space Station and I cold-called NASA's In-Space Manufacturing Manager and said, "I think we should do a national challenge and get students all across the country to design objects for astronauts to 3D print on the Space Station."

Most people would probably think I'm crazy, but the person I spoke with loved the idea. She had two daughters herself. I ended up bringing in the ASME Foundation. It took a number of months to get a Space Act Agreement[1] in place, which was the critical link to get the first challenge off the ground. And Future Engineers was born.

We hosted the first 3D Printing in Space challenge and it was well received. It happened at the same time federal agencies were coming up with commitments for the White House Maker Initiative—it was the perfect synergy and timing.

Initially, Future Engineers was synonymous with 3D printing challenges. Based on that experience, Future Engineers then applied for, and was awarded, SBIR funding from the US Department of Education to develop an online technology platform that can host engineering student challenges of all kinds.

RML: How does it work? For example, how did the students interface with NASA and how did the challenge roll out?

MS. BELL: We brainstorm to come up with a theme. Sometimes the theme might seem simplistic, but the intent is to develop a challenge where students can participate not only at a very elevated level but also at a very basic level. Everyone can complete the challenge, regardless of experience level.

For example, our second challenge was to design a container to 3D print in space. We had high school students designing hydroponic plant growth chambers, and younger students designing little boxes to collect rocks on Mars. We left it open so that students from kindergarten to 12th grade could participate and feel a sense of accomplishment in completing the challenge.

I cold-called NASA and proposed a national challenge for students to design objects for 3D printing on the Space Station.

Students enter all challenges online. Everything is online—the challenge, submission portal, curriculum resources, videos, links to free design tools that teachers and students can use in classrooms. Our goal is always to have challenges be free or low cost. Students can participate either at home or in class—we're now at about 80 percent participation from classes. Students submit their entries to the website and get a gallery page dedicated to their entry. They can see all of the innovative submissions received from across the country.

For the first challenge, students designed objects that they wanted to be 3D printed in space. The winning design, a multipurpose precision maintenance tool, became the first student-designed 3D print in space. The winner, Robert Hillan, was invited to NASA's Marshall Space Flight Center, where he went into payload operations and spoke to astronauts live on the ISS and saw his 3D print floating in space.

RML: What a great experience. So NASA actually printed in space the design that your challenger developed?

MS. BELL: Yes.

RML: This is a great experience for kids. They learn about not only computers and software and design but also interaction with a huge organization like NASA.

MS. BELL: There are also finalist interviews online with judges. Students in the final interview present

[1] "The National Aeronautics and Space Act (the Space Act) provides NASA with the unique authority to enter into a wide range of 'other transactions,' commonly referred to as Space Act Agreements (SAAs). The Agency enters into SAAs with various partners to advance NASA mission and program objectives, including international cooperative space activities." (https://www.nasa.gov/partnerships/about.html)

their designs to experts—and they don't just get a pat on the head: the experts really dig in and ask about how the students created the design, how it could be useful, or how it could be improved.

The 3D printing challenges were actually part of NASA's research in terms of thinking about how they're going to leverage 3D printing technologies on future missions. K–12 students have extremely creative ideas, so it's great that NASA's public engagement efforts could inspire these young innovators in such a meaningful way.

RML: What's the frequency of the challenge?

MS. BELL: Well, we're at a transition point. Recently Future Engineers launched our 2.0 platform, which can host as many challenges as we want simultaneously and aligned to all kinds of topics, themes, and education standards, not necessarily space or 3D printing. We plan to work with all kinds of industry partners to come up with challenges that showcase the diversity of STEM. For example, we just hosted a Bright Art Challenge focused on properties of light, a Weather Balloon challenge focused on data visualization and payload design, a Future Creatures challenge focused on adaptations, and a Name That Molecule challenge focused on 3D molecular visualizations. Students' entries go into their portfolio on our site, and teachers can build classroom portfolios as well.

Currently, we're hosting a challenge with NASA where K–12 students all across the country are asked to name the Mars 2020 Rover. The contest is open in fall 2019 so make sure to tell your kids or grandkids to go to FutureEngineers.org and submit a name.

CAMERON FLETCHER (CHF): You've been working so far with NASA as your primary institutional partner, but now you're branching out to other institutions or companies?

MS. BELL: We started by working with the ASME Foundation, with technical assistance from NASA, and now we have a partnership with NASA on the Name the Rover contest via a Space Act Agreement. Also, the support of the US Department of Education's SBIR program has been pivotal. But we're very excited to be branching out to develop partnerships with new organizations to create other kinds of challenges.

CHF: Can you tell us who some of those other entities are?

MS. BELL: Not yet. But I can mention others that we've previously worked with on themes or educational prizes. One theme was aligned with Star Trek's 50th anniversary. Makerbot has graciously donated many 3D printers to schools of winning students. We've had students visit with SpaceX engineers at their headquarters. We collaborated with Digital Domain, a visual effects company in Los Angeles, where the winning students got to suit up in motion capture suits to see how digital models are used to make movies. They made their own movie of themselves on Mars, which was really fun.

For each of our challenges there's a different theme and we try to come up with educational experiences and educational prizes to illustrate for students what they might do for a living as an engineer and how they can apply the skills they've learned in the challenge.

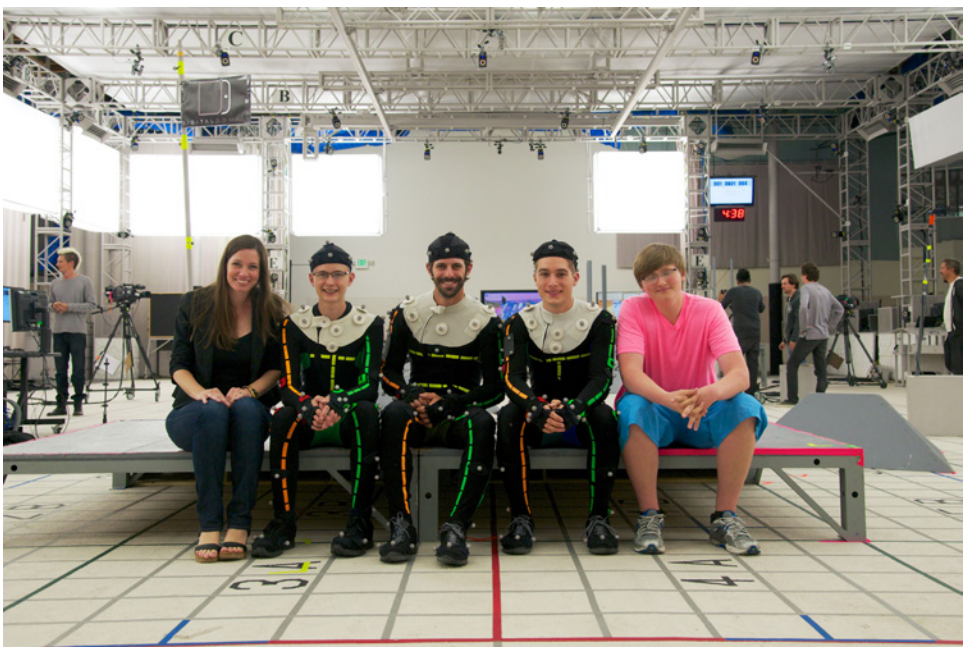

Deanne (left) with winning students in motion capture suits at Digital Domain, where they made a simulated movie on Mars.

CHF: Do you have any sense of how many students you've reached and engaged over the years? And is there any kind of follow-up to create a community of these participants, like an alumni group?

Girl Scouts at the Future Engineers pop-up makerspace, 2017 G.I.R.L. National Convention, Columbus, OH.

MS. BELL: We definitely have a growing community of students and educators, but once the students graduate from high school we don't currently have a way to engage them, except as judges. We're looking at how we can leverage college students as mentors and expand to college challenges as well. We're excited to bridge that gap, especially with women in engineering. A lot of research shows that girls need to be engaged before they're 12, so we try to do challenges that engage students at the youngest ages to get them excited and inspired, to help create an engineering identity and then keep building on the momentum to strengthen the pipeline as students get older. We recently completed a research study using six prototype challenges with middle school classrooms, which showed significant gains in enjoyment of engineering. The gains were primarily driven by girls, so we're on a good path.

CHF: You acknowledged earlier that there are separate challenges for, say, kindergartners and other youngsters as opposed to high school participants. Are there different criteria or are different submissions selected from different age groups?

MS. BELL: Sometimes we use the same criteria for the different competitive age groups, but we select a junior and a teen winner. With our new platform, we can do a kindergarten-only challenge or a high school challenge, or only DC public schools or only Florida students, for example. We can design challenges for different ages, grades, demographics, or regions.

Education standards are aligned to different grades so we make sure our challenge themes align to grade-appropriate standards. Our focus is on creating challenges that can be done as a classroom activity. Teachers are amazing so we strive to support them as best we can.

CHF: You've mentioned engaging girls and you just used the word demographics. Do you have outreach efforts to other demographic groups, like underrepresented minorities or socioeconomically disadvantaged kids?

MS. BELL: We do public engagements focusing on different demographics. We've focused a lot on girls and women to date; for example, in 2017 we were at G.I.R.L., the Girl Scouts National Convention where we collaborated with the ASME Foundation to teach over 2,000 girls to use CAD. And last year we helped coordinate an event in Washington at the Smithsonian National Air & Space Museum, where local students asked an astronaut questions live as she was on the International Space Station.

Also, most of our challenges involve multimedia content. I'm a TV host and we make launch videos for our challenges as well as educational media for the website. We make sure that the media we create is representative of engineers of all kinds. We're committed to diversity and to making sure our challenges are accessible and

free so that there's no barrier to entry to participate in engineering.

RML: We've spent a lot of our time talking about a topic that is of great interest to us and to the NAE: encouraging young people to become involved as engineers. What you're doing sounds very commendable.

I'd like to turn the conversation to your TV hosting. How did you make the transition from being a practicing engineer with Raytheon and then with a startup in Boston to TV host?

MS. BELL: It was a leap of faith. I was working in Los Angeles. I'm from Florida originally and went to college in St. Louis. I got my first job offer for a job in LA and came here. I was working as an engineer at Raytheon. It was wonderful—I had fantastic mentors, I was given a lot of responsibility, I was thrown in the fire.

I just needed that one person to say yes, to take that leap of faith. It changed my life.

Of course, there's another industry in Los Angeles: the media industry. I didn't really have much involvement with it until I went on vacation for a week and saw an ad posted for an engineer to host a television show. I was kind of joking with a friend and I sent in some stuff and gave my landline phone number. A week later I got back from vacation and I had all these messages from the casting director saying "We want to bring you in."

I put on a suit, brought my resume, went in—and was put in front of a green screen and given a cue card and told to read lines. I thought, 'Whoa, I have no idea what I'm doing.' Then they sat me down on a stool and asked me to talk about technology that I had worked on. And I just talked.

I remember the casting agent being really surprised that I existed. I was disheartened by the fact that this person hadn't encountered many female scientists or engineers when casting for these kinds of shows. But when I asked myself, 'What female engineer have I seen on TV?,' I couldn't think of one. I thought, 'We need to change this. Women need to have more representation in the media.'

I went through a series of auditions for that show, but the show never got greenlit, it just disappeared.

A year later I decided to take a leave of absence to travel. I'm a big traveler. I had saved my pennies and wanted to hike in Tibet and in New Zealand and just travel. While in the Philippines, I was thinking 'I really need to get my act together and figure out what I'm doing next in life.' I was at a hostel, looking at online job boards—and again, an ad popped up: "Looking for an engineer to host a TV show on PBS." What are the chances of that?!

So I sent in my application from the Philippines, a picture and a few sentences, and I heard from the casting agent. "You look great. Our last audition is in two days in Boston." I wrote back, "I'm in the Philippines, can you wait?" And they responded, "Let us know if you can make it."

That night I had a glass of wine with a German pig farmer and he said, "You know what, I think you should go." And I agreed. I decided, 'I need to do this.' I called the travel agency and got everything rearranged. I flew back and arrived in Boston at about 9:00 p.m., auditioned the next morning at 9:00 a.m., and ended up getting the job. That was my very first job as a television host.

CHF: Did you send a thank you note to the German pig farmer?

MS. BELL: I don't even know his name. But I think about that pig farmer often. There I was on the cusp of taking this big risk, this leap, and I just needed someone to say 'yes.' There're so many reasons to say 'no,' to think 'This is probably a waste of my time. It's going to cost money. It isn't going to be worth it.' I just needed that one person to say *yes*, to speak to the inner voice in my head to take that leap of faith. It changed my life.

RML: Which show was it?

MS. BELL: My very first hosting job was on a show called *Design Squad* on PBS. It still exists online, but I only cohosted the first season. It's basically a reality show where they would give students a design challenge. For example, "You have a red wagon or a tricycle and two drills and you have to make a dragster. You have two days to design and build." Then we took them to the speedway and tested their cars. The students competed and at the end of all the challenges the student with the most points won a college scholarship. It was produced in Cambridge on the MIT campus, in the solar car garage, I think.

CHF: I see from your website that you've been involved in quite an array of activities—*Smash Lab* and *Money*

Hunters, *The Egyptian Job* and *Rise Up*. How on earth did you get involved in each of these?

MS. BELL: I'm fortunate that the momentum just kept going. I got called to host a Discovery Channel show not long after hosting the first season of *Design Squad*. Then at some point you get an agent and go from there.

I am a very "niche" talent on TV, but the beauty of science and engineering is that they're in everything. I hosted a show about trying to figure out how an Egyptian pyramid that was built in the 19th century BC was broken into—I ran a bunch of physics calculations about how certain rocks were moved back when. And then a show on ESPN—who would think you could be an engineer on ESPN? We renovated athletic facilities for schools and my cohost was a professional football player; he worked on the athletics program side of it and I worked with the contractors to renovate the physical facilities. That was *Rise Up*. It was kind of like a home makeover show but for athletic facilities.

Deanne (right) with nurse-inventor Bobbie Sue McCollum and her Goldilocks Valves prototype, on *Make Me a Millionaire Inventor*.

Smash Lab is probably the most sensational show I worked on. I like to say that I specialized in the science of smashes, crashes, and destruction. It was a difficult show to make—and definitely an adrenaline rush every day. We had to stage a massive crash or disaster, use instrumentation to figure out how bad it was, and then invent something that could potentially make it safer. It could be the most ridiculous invention ever. We had a couple of weeks or less to do these big inventions, so it was just proof-of-concept stuff.

We went out and simulated these things in the desert—I spent a lot of time out in the desert in Southern California. We did one show involving a high-rise escape out of a building that was on fire. There were also big rigs running into things, train crashes, we simulated a logging truck rolling over on a sharp turn, and we simulated the equivalent of a shoe bomb on an airplane. For that one we rented airplanes in the Mojave Desert to conduct the final test. One plane had a blast-proof material applied and the other didn't. The planes were at ground level, so we had to inflate them to create a pressure differential as if they were at altitude, but we found that the planes were riddled with holes, so first we had to patch the planes, then inflate them, and then our explosives expert initiated the explosions. We used pressure sensors, temperature sensors, and a bunch of high-speed photography to analyze the events, but our attempted solution was a flop against that very difficult scenario.

Most recently I cohosted a show on CNBC called *Make Me a Millionaire Inventor*. It's a business show. We worked with hardware inventors who had an idea for a product and maybe had a great prototype but didn't know how to take it to the next level. We had an in-house product development team who would bring their idea to life with a fully fleshed-out prototype like you would see on a showroom floor. I mentored the inventors on the product development process as well as the process of developing a business plan. At the end of each episode the inventors would pitch their idea to an investor for a chance at getting money to launch their product into a business, and many of them did.

Some of the products didn't receive an investment on the show but the inventors persevered after being rejected. For example, a nurse named Bobbie Sue McCollum works in the ER where they use bag valve masks to manually put air into patients' lungs. A nurse is supposed to squeeze it and then wait 6 seconds, then squeeze again and wait 6 seconds, then squeeze again…. But in the adrenaline of the moment the bag might get squeezed too soon and overinflate the patient's lungs. She invented a device that attaches to the bag valve

mask and regulates the reinflation so the nurse knows when it's safe to squeeze the bag again.

She didn't get an investment on the show, but she got great advice about the healthcare market and how to pursue licensing opportunities with larger medical manufacturers, and she's now adapting her business plan to go after these opportunities.

RML: How do inventors find their way to the show? For example, how did Bobbie Sue discover the show and get into your broadcast?

MS. BELL: There was a lot of work by the production company to find inventors that were a good fit. Once the show was on the air, the show itself became a recruitment tool for inventors to reach out.

RML: Many people have very good ideas, they just don't know how to implement them in order to solve a problem. We all see problems every day and sometimes you have a sparkling idea that might pan out but people don't know how to follow through. I just wonder how a person who sees a problem, like the inflation valve in the ER, and then is clever enough to decide that there must be a way to solve this and regulate the device, finds her way to your show. There must be a screening process to determine whether there's an invention that has some merit or if it's blue sky and totally unworkable.

MS. BELL: There's definitely a screening process. There's screening in terms of intellectual property and what exists on the market, and in terms of viability for the product in the marketplace. And then screening for the TV show: Is this going to make a good TV show about this journey?

Engineering can evoke emotion. It can empower people.

I think what's interesting about *Make Me a Millionaire Inventor* is that it illustrates how engineering can evoke emotion. It can empower people. We had these inventors, some of them had been sitting on an idea for 15 years, paralyzed because they didn't know how to take the next step.

It's about having the courage to do it and of course the means or the resources. But it's also about mindset. I'm so grateful every day for my engineering mindset. I'm trained as an engineer to ask the right questions, to find the answers and to figure things out. Great ideas are common to all the inventors we've worked with. But engineering is what takes great ideas to the next level.

A lot of people talk to me about inventions and how do you come up with the next big idea. I think sometimes people think too far in advance—they think about how to create the next Facebook—when really almost every prototype is something that's close to you. With Facebook, it was kids in a dorm room trying to connect with other students in their college. That's what I tell people: Think of a problem around you that inspires you, solve that, and then figure out how to scale it.

What's great on the show is that these inventors come in and we show them their prototype—and the tears roll. It's amazing that engineering can do that for someone. Engineering can take something that's been an idea that they've been working on for so long, it's so personal to them—and there it is: they can see it, touch it and hold it. It's awesome to show that engineering can have such an emotional element. It's one of my favorite shows that I've worked on.

CHF: What are some inventions that you've pursued as an entrepreneur?

MS. BELL: I stumbled into becoming an entrepreneur with Future Engineers and it took me a while to realize that I'm the founder and CEO of an education technology company. I'm trained as a mechanical engineer but my day-in and day-out work to build Future Engineers is as chief architect of a sophisticated online platform that navigates the complexities of running innovation challenges for students of all ages, at home and in class, on a national scale. We have to deal with privacy, working with minors, student IP, and contest/promotion regulations. I dove into driving the software development and had to wear a million different hats to get Future Engineers up and running. I just filed my first patent application, so I guess I'm officially an inventor myself now too.

I tell people I became an entrepreneur by accident. I was following my passion. I really wanted to inspire more engineering outreach and I thought, 'How can I best facilitate this?' Doing in-person workshops was hard to scale, and not cost effective. I decided that a platform was the best way to connect students all across the country on a common goal, a common challenge. And I think students love that they can see how a kid in Boston and a kid in Florida solved the same prob-

lem, and they can learn by looking at the innovations of other students and get inspired. A lot of students who have won our challenges say they participated as a beginner and saw these amazing inventions that other students were coming up with and that encouraged them to hone their skills.

CHF: So you have repeat participants?

MS. BELL: Yes, very much so.

RML: I get the sense, Deanne, that you have a lot of things going on all at the same time—and you're handling them all beautifully. That's in itself quite a story. I understand that you also speak to audiences of young people?

MS. BELL: Yes, I speak to audiences of all ages.

RML: About engineering as a career, or what are you asked to address?

MS. BELL: Often with students I talk about my journey—my passions, but also my frustrations and how I persevered, and how I now have this crazy job that I love. You can look at my career and see successes and atypical engineering paths that I've taken, but there are a lot of times I failed or second-guessed myself or thought maybe I didn't fit.

In college, there was a time in my junior year when I thought about dropping engineering entirely. I could do the classes just fine, but I didn't know if engineering was the right path for me. I felt like my career ambitions were very different than my peers.

Thank goodness for a female professor who mentored me in college. She showed me that what I was learning were tools that I could apply in industry and that I would thrive with my personal skills and passions. She was 100 percent right! I try to communicate that to students.

RML: That's very important. I think most engineers are not typically in the public eye. They're great at solving problems, they can bring things to the market, they can build all sorts of interesting devices—but they don't have the same inclinations or skills to interact with the public that I think can be helpful. You're providing public access to the engineering enterprise in a way that can be inspirational to young people. You bring an important dimension that isn't typical of the engineering community. It just isn't typical of engineers to be outright publicly visible. And yet, all the things engineers do affect the public—we build engineering systems that serve society—but we don't necessarily interface very well with society. I think what you're doing is terrific.

MS. BELL: Thank you.

RML: Where do you see yourself in 5 or 10 years?

MS. BELL: I often talk about imagining the impossible—if you can dream it, you can build it, you can make it happen. You just have to dream it up first. I'm fortunate in that I imagined the impossible of being a TV host and it happened.

I became an entrepreneur by accident. I was following my passion.

In 5 to 10 years I see myself, in addition to being a television host, having a role in producing a television show, being on the creative side of the production as well as in front of the camera—or maybe just behind the camera—and inspiring girls to pursue science and engineering.

I also see an enormous future for Future Engineers. I see us hosting hundreds of challenges, simultaneously engaging students all across the globe, building engineering portfolios, developing an engineering identity with students as young as 5 years old, highlighting the awesome science and technology work that's being done in our world, and showing students that there is engineering in *everything*. It's everywhere around us!

I see Future Engineers having engineering challenges paired with movies, paired with emerging technologies, paired with global challenges or needs, and engineering challenges that inspire students to learn technology tools from 3D design to data visualization to designing augmented reality software—a whole world of engaging students in challenges.

CHF: Thank goodness your mentor talked you out of dropping out of engineering. I have one more question for you, Deanne. Would you like to offer any particular thoughts to the readers of *The Bridge*?

MS. BELL: I guess to summarize I'll say that we need to engage students with engineering early and often. This absolutely requires investments in the future, but it also applies to our everyday life. Think about your

vocabulary. Just as I can paint something without being a professional painter, or I can nurse someone back to health without being a registered nurse, we all have the capacity to engineer. Next time you see a young girl or boy iterate, improve, or innovate something new, tell them they were doing engineering. We learn from context and kids need to see and hear more about engineering to build an understanding of what engineers do and who we are. And just because someone becomes a TV host or a politician—or anything, really—it doesn't mean they left engineering, it means they chose to apply their engineering mindset in a new way, to solve a new problem. The more we celebrate the zillions of things engineers do in this world, the more inclusive and diverse our profession will be, and the greater problems we can solve—together.

The more we celebrate the zillions of things engineers do, the more inclusive and diverse our profession will be and the greater the problems we can solve—together.

RML: Deanne, we want to thank you for this wonderful conversation. I think what you're doing is truly inspirational. I hope a lot more young men and women will get to know what you're doing. I have five grandchildren and I'm going to make sure they all tune in to your show because I think this is really important stuff.

MS. BELL: That's wonderful to hear.

CHF: Deanne, it sounds like you are having fun on company time.

MS. BELL: I am! And I should add that when I auditioned for CNBC I was 8 months pregnant and when I launched Future Engineers I was 8 months pregnant. It has been a journey navigating that as well, but I'm a huge proponent of giving women the flexibility they need. When you do, they will thrive and excel. It's been great to be able to create my own situation that fits me perfectly, with TV hosting and doing education work, and at the same time building my family. I have a 4-year-old and another who just turned 1.

And if you ever want to talk with another atypical engineer, my husband is a software engineer but works as a visual effects supervisor at Marvel—he recently worked on *Avengers: End Game* and *Infinity War*—so we're an interesting duo. That's kind of a funny story, actually. When we met, I told him I was an engineer—and he said, "I'm an engineer." So I added, "But I don't have a traditional engineering job." And he said, "Neither do I." Then I asked him, "What do you do?" He said, "I blow up stuff in movies." And I said, "I blow up stuff on TV." It was a match made in heaven. We went on our first date the next day and every day since.

CHF: That's great.

MS. BELL: Wonderful to speak with you both.

CHF: Thank you, Deanne.

RML: Thank you.

NAE News and Notes

NAE Newsmakers

Gilda A. Barabino, dean of the Grove School of Engineering and Daniel and Frances Berg Professor, City College of the City University of New York, is the recipient of the 2019 AIChE Award for Service to Society. Dr. Barabino is being recognized for her approach in using engineering principles to solve medical issues that include disease therapies and tackling health disparities; for her public policy leadership to advance the engineering profession; and for her career-long efforts and transformative impact to broaden participation in the engineering fields through advocacy, mentorship, and professional development of underrepresented minority students and faculty. The award will be presented November 10 at the annual meeting of the American Institute of Chemical Engineers.

Marsha J. Berger, Silver Professor of Computer Science and Mathematics, Courant Institute, New York University, and **Arkadi S. Nemirovski**, John Hunter Chair and professor, School of Industrial and Systems Engineering, Georgia Institute of Technology, received **2019 Norbert Wiener Prizes** from the American Mathematical Society (AMS) and the Society for Industrial and Applied Mathematics. The prizes were presented at the 125th annual meeting of the AMS in Baltimore in January. Dr. Berger was honored "for her fundamental contributions to adaptive mesh refinement and to Cartesian mesh techniques for automating the simulation of compressible flows in complex geometry." Dr. Nemirovski received the prize "for his fundamental contributions to high-dimensional optimization and for his discovery of key phenomena in the theory of signal estimation and recovery."

Barry W. Boehm, TRW Distinguished Professor of Software Engineering, University of Southern California, was recently chosen to receive the **INCOSE 2019 Pioneer Award**, the International Association of Top Professionals **Top Professor of the Year 2019 Award**, and the **Marquis Who's Who Lifetime Achievement Award 2019**.

Pablo G. Debenedetti, dean of research and Class of 1950 Professor in Engineering and Applied Science, Princeton University, has been selected for the AIChE **Alpha Chi Sigma Award for Chemical Engineering Research**. He was cited for "seminal research contributions on the metastable liquid-liquid phase transition and properties of water at supercooled conditions." He will receive the award November 10 at the AIChE annual meeting in Orlando.

The Semiconductor Industry Association has announced **Robert H. Dennard**, IBM Fellow Emeritus, IBM Thomas J. Watson Research Center, as winner of the **2019 Robert N. Noyce Award**. Dr. Dennard is recognized for inventing the memory technology underpinning every computer, smartphone, tablet, and other ubiquitous electronic devices and for pioneering the scaling physics that have allowed the semiconductor industry to produce exponentially advanced products at lower cost. He will accept the award on November 7 in San Jose.

Ann P. Dowling, president of the Royal Academy of Engineering and professor of mechanical engineering at the University of Cambridge, is to receive the **Royal Medal** from the Royal Society. Three Royal Medals, also known as the Queen's Medals, are awarded annually by the sovereign on the recommendation of the Council of the Royal Society. Dame Ann is recognized for her leading research on the reduction of combustion emissions, aerodynamic noise, and the design of aircraft, as well as her distinguished service to engineering.

Rensselaer Polytechnic Institute has honored **Nancy D. Fitzroy**, retired, General Electric Corporate Research and Development. On May 16, the Admissions Building was dedicated and renamed the **Nancy Deloye Fitzroy '49 and Roland V. Fitzroy Jr. Admissions Building**. Dr. Fitzroy was the first woman to graduate from the institute with a degree in chemical engineering.

Naomi Halas, Stanley C. Moore Professor of Electrical and Computer Engineering and founding director of the Laboratory for Nanophotonics at Rice University, has been elected a **fellow of the Royal Society of Chemistry**. To be eligible for RSC fellowship, one must have made "a substantial contribution to the improvement of natural knowledge, including mathematics,

engineering science, and medical science."

Tobin J. Marks, Vladimir N. Ipatieff Professor of Catalytic Chemistry and professor of materials science and engineering, Northwestern University, has been elected a **foreign fellow of the European Academy of Sciences** in recognition of his scientific contributions. Dr. Marks is a world leader in the fields of organometallic chemistry, chemical catalysis, materials science, organic electronics, photovoltaics, and nanotechnology. The induction ceremony will take place in Madrid in October.

Chad A. Mirkin, director, International Institute for Nanotechnology, and George B. Rathmann Professor of Chemistry, Northwestern University, has been selected to receive one of two 2019 international Kabiller Awards. The awards biennially recognize two top scholars, one pioneer and one rising star in the field of nanoscience and nanomedicine. Professor Mirkin was chosen to receive the **Kabiller Prize in Nanoscience and Nanomedicine** for outstanding achievement in the field of nanotechnology and its application to medicine and biology. The award will be presented November 13 in Chicago.

Washington University in St. Louis will present the **Chancellor's Award for Innovation and Entrepreneurship** to **Yoram Rudy**, Fred Saigh Distinguished Professor of Engineering and professor of biomedical engineering, at a ceremony on November 8. Professor Rudy is honored for his inventions that have changed the way cardiologists measure deadly irregular heartbeats.

Emanuel M. Sachs, chair, Scientific Advisory Board, 1366 Technologies Inc., and Professor Post Tenure, Massachusetts Institute of Technology, has been awarded the **2019 SME Industry Achievement Award**. Dr. Sachs is recognized for being instrumental in transforming global manufacturing; he and colleagues at MIT first developed the concept of 3D printing in the late 1980s while inventing and establishing binder jet technology.

Bruce Watson, Institute Professor of Science, Department of Earth and Environmental Sciences, Rensselaer Polytechnic Institute, has been awarded the **Roebling Medal** by the Mineralogical Society of America. The MSA's highest honor is bestowed for scientific eminence in the broad field of mineralogical science. The award was presented during MSA's joint annual meeting with the Geological Society of America in Indianapolis November 4–7, 2018.

Eric Horvitz, Distinguished Scientist and director, Microsoft Research, and **Eric E. Schmidt**, chair of the Defense Innovation Board and technical advisor to the board, Alphabet Inc., have been appointed to serve on the **National Security Commission on Artificial Intelligence**. The commission will review advances in artificial intelligence, machine learning, and associated technologies with national security implications, including the competitiveness of US efforts, international trends and cooperation, workforce and education incentives, data standards, and ethical considerations for future application.

2019 China-America Frontiers of Engineering Hosted by Qualcomm in San Diego

The 2019 China-America Frontiers of Engineering Symposium was held June 20–22 at Qualcomm in San Diego. NAE member **Zhenan Bao**, K.K. Lee Professor and chair of chemical engineering, and Jing Cheng, Cheung Kong Professor of Biomedical Engineering at Tsinghua University, cochaired the symposium. The NAE carries out this activity with the Chinese Academy of Engineering.

Consistent with the design of the bilateral FOEs, this meeting brought together approximately 60 early-career engineers from US and Chinese universities, companies, and government labs for a 2½-day meeting to discuss leading-edge developments in four engineering fields. The session topics were Frontiers of Neuroengineering for Restoring Human Sensory and Motor Functions, Smart Cities, 5G Wireless Communication Technology, and New Materials.

Advances in neuroengineering have transformed neural interfacing technology for recording and stimulating neurons, thereby enhancing both computational capacity in neural decoding/encoding and the translation of neurotechnology to clinical applications. This session focused on system-level research in neuroengineering that aims to restore

2019 China-America Frontiers of Engineering Symposium attendees get to know each other during a "speed dating"–type session on the first morning.

human motor or sensory function. The four talks covered bidirectional brain-computer interfaces; passive brain-computer interfaces, which provide information from user mental activity to a computerized application without the need for the user to control his brain activity; invasive and noninvasive stimulation technologies to treat hearing loss; and neuromodulation, an emerging field that offers a cross-disciplinary approach for investigating neural circuits and treating neurological and psychiatric disorders.

The world's urban population is expected to grow by 2.5 billion people between now and 2050, with much of the growth concentrated in urban areas. Technology-focused smart cities approaches are being explored to improve the efficiency, sustainability, and resilience of urban infrastructure systems needed to support this expanding and urbanizing population. The approaches taken by China and the United States in the development of smart cities technologies are often distinct, with top-down approaches needed in China to achieve scale quickly, while organic approaches are more common in the United States. The first speaker described how to exploit smart cities technologies to make urban infrastructure systems more resilient to natural hazards and malicious attack. This was followed by a presentation on mobile crowdsensing for smart cities. The third presenter described the challenges of meshing cyberphysical and sociotechnical systems. The session concluded with a talk on a method for city planning that promotes the physical and mental health of residents.

In both the United States and China, 2019 marks the year for the commercial rollout of 5G wireless communication systems based on the global 3GPP 5G New Radio (5G-NR) standard. The expansive range of 5G wireless communication technology and applications has resulted in very active R&D efforts in academia and industry. Talks in this session covered baseband design and implementation for 5G, the state of the art for 5G with a focus on low-latency and high-reliability applications, next-generation applications enabled by ultra-low-power integrated circuits, and the impact of 5G on high-speed railway communications.

Materials are the building blocks for societal advances, from com-

puting miniaturization to smart, energy-efficient homes. Underlying these technological advances is the development of new materials, inspired by natural systems, motivated by synthetic and manufacturing innovations, and driven by sustainability concerns. The first presentation covered nanomaterials that can be used to solve the pressing problem of water scarcity. Next was a talk on biomimetic adhesives that retain performance under water. The third speaker discussed additive manufacturing of complex materials, for example in bioprinting and multimaterial printing. The session concluded with a talk on mass production of high-quality flexible metamaterials through a roll-to-roll process for large-area assembly of photonic crystals, which has applications in sensors, 5G, and medical treatments.

In addition to the formal presentations, an icebreaker on the first morning in a "speed-dating" format helped attendees from the two countries become acquainted. In the afternoon, a poster session preceded by flash poster talks provided an opportunity for all participants to share information about their research and technical work. The posters were left up throughout the meeting, facilitating further discussion and exchange during the coffee breaks. On the second afternoon attendees toured the Qualcomm Museum and areas devoted to the Internet of Things and Qualcomm's Snapdragon™ platform. The symposium program, list of attendees, and presentation slides are available at the 2019 CAFOE link at www.naefrontiers.org.

Funding for this activity was provided by Qualcomm, The Grainger Foundation, and the National Science Foundation. The next CAFOE meeting will be held in China in 2021.

The NAE has been hosting an annual US Frontiers of Engineering meeting since 1995 and, in addition to CAFOE, has bilateral FOE programs with Germany, Japan, and the European Union. The meetings bring together highly accomplished early-career engineers from industry, academia, and government and provide an opportunity to learn about developments, techniques, and approaches at the forefront of fields other than their own. The program also facilitates the establishment of contacts and collaboration among the next generation of engineering leaders.

For more information about the activity, or to nominate an outstanding engineer to participate in future FOE meetings, go to www.naefrontiers.org or contact Janet Hunziker at JHunziker@nae.edu.

Summer Interns in the NAE Program Office

SOUMYA CHAPPIDI is a University of Virginia undergraduate student in the School of Engineering and Applied Sciences, studying for a BS degree in systems engineering and a BA in mathematics, with a minor in computer science. She has a side business in photography, is passionate about STEM education, and hopes to serve in the Peace Corps after finishing her undergraduate studies. Soumya is bilingual in English and Telugu. She worked with the Grand Challenges Scholars Program.

SUDHIR (SID) SHENOY earned his bachelor of engineering in electronics and communication from Jain University School of Engineering and Technology in Bangalore.

Interns Sid Shenoy and Soumya Chappidi with President John Anderson.

He is finishing his master's degree in computer engineering at the University of Virginia, where he has studied machine learning, design and analysis of algorithms, computer architecture, autonomous mobile robots, robots and humans (HRI), cooperative autonomous systems, sensors and perception, cloud computing, and cyberphysical systems, and supervised project research on cryptography in cloud computing. His particular interest is in driverless cars/trucks, robotics, and man-machine interfaces. For extracurricular activities Sid engages in swing dancing, ballroom dancing, filmmaking, and debate. He worked with the NAE's Center for Engineering Ethics and Society in the area of ethics of AI and machine learning.

Soumya and Sid were a great addition to the NAE Program Office and contributed meaningfully to the GCSP and CEES efforts. We wish them well as they continue their studies!

Calendar of Meetings and Events

October 4–5	NAE Council Meeting
October 5	NAE Peer Committee Meeting
October 6–7	**NAE Annual Meeting**
October 22–23	Working Ethics into the Conversation: Introducing STEM Faculty to Teaching Ethics
October 24–25	OEC (Online Ethics Center) Advisory Group and Editorial Board Joint Meeting
November 7–9	EngineerGirl Ambassador Training Meeting Anaheim, California
November 10	EngineerGirl Steering Committee Meeting Anaheim, California
November 11–12	NAE Grand Challenges Scholars Program Annual Meeting
November 18–20	EU-US Frontiers of Engineering Symposium Stockholm

All meetings are held in National Academies facilities in Washington, DC, unless otherwise noted.

In Memoriam

KEN AUSTIN, 87, founder, A-dec Inc., died May 1, 2016. Mr. Austin was elected in 1999 for inventing, designing, manufacturing, and marketing innovative dental equipment systems and facilities.

VLADIMIR V. BOLOTIN, 82, head, Laboratory of Reliability, Russian Academy of Sciences, died May 28, 2008. Dr. Bolotin was elected a foreign member in 1996 for contributions to technical leadership in the dynamics, stability, strength, and reliability of structure and machines.

JAMES E. BROADWELL, 97, consultant, Redwood City, CA, died June 22, 2018. Dr. Broadwell was elected in 1987 for contributions to the understanding and management of turbulent mixing with application to chemical laser design.

LLOYD S. CLUFF, 85, retired director, Earthquake Risk Management, Pacific Gas and Electric Company, died June 4, 2019. Mr. Cluff was elected in 1978 for contributions in identifying active seismic faults and their potential motions with applications to earthquake engineering.

BRUCE G. COLLIPP, 88, marine engineering consultant, died November 29, 2017. Mr. Collipp was elected in 1991 for pioneering contributions to the semisubmersible offshore floating drilling platform and for continued leadership in the development of innovative ocean engineering.

FERNANDO J. CORBATO, 93, professor emeritus of computer science and engineering and senior lecturer, Massachusetts Institute of Technology, died July 12, 2019. Dr. Corbato was elected in 1976 for contributions to the development of multiple-access computer systems.

HARVEY G. CRAGON, 89, Ernest Cockrell Jr. Centennial Chair Emeritus in Engineering, University of Texas at Austin, died September 7, 2018. Mr. Cragon was elected in 1978 for contributions to the development of large-scale digital computer systems.

CHARLES B. DUKE, 81, retired vice president and senior fellow, Xerox Corporation, and research professor of physics, University of Rochester, died June 28, 2019. Dr. Duke was elected in 1993 for providing the theoretical foundations for developments in xerography.

ELMER G. GILBERT, 89, professor emeritus of aerospace engineering, University of Michigan, died June 16, 2019. Dr. Gilbert was elected in 1994 for contributions to the theory and practice of multivariable, optimal, nonlinear, and computer control systems and to control engineering education.

ROBERT C. HANSEN, 91, president, R.C. Hansen Inc., died February 9, 2018. Dr. Hansen was elected in 1992 for contributions to the understanding of electromagnetic near-fields, electrically small antennas, and phased arrays.

LEE A. IACOCCA, 94, Lee Iacocca and Associates, died July 2, 2019. Mr. Iacocca was elected in 1986 for outstanding engineering management and coordination of engineering, manufacturing, marketing, and corporate planning activities in the automotive industry.

SRINIVASA H. IYENGAR, 85, retired partner, Skidmore, Owings & Merrill, died July 4, 2019. Mr. Iyengar was elected in 2000 for leadership and contributions in structural design of tall buildings and long-span structures and advances in fire-resistive construction.

GEORGE W. JEFFS, 94, University Professor and Thomas Lord Professor of Operations Research, Carnegie Mellon University, died March 18, 2019. Dr. Jeffs was elected in 1978 for contributions to integer programming and its applications to the scheduling and planning of industrial facilities.

JACK L. KERREBROCK, professor emeritus of aeronautics and astronautics, Massachusetts Institute of Technology, died July 19, 2019. Dr. Kerrebrock was elected in 1978 for contributions in the development of propulsion and energy conversion systems design and research, education, and national service.

CHRISTOPHER C. KRAFT JR., 95, retired director, NASA ohnson Space Center, died July 22, 2019. Dr. Kraft was elected in 2003 for contributions to engineering through the planning and direction of the nation's manned space flight missions.

JOHN L. MASON, 95, independent consultant, died June 3, 2019. Dr. Mason was elected in 1988 for pioneering developments in space power generation, vehicle power plant heat exchangers, and environmental control systems.

KARL H. NORRIS, 98, retired research leader, Instrumentation Research Laboratory, US Department of Agriculture, died July 17, 2019. Mr. Norris was elected in 1980 for research and development of systems for simple and rapid analysis of quality factors in food products and similar materials.

NORMAN F. PARKER, 87, retired president and CEO, Varian Associates Inc., died March 29, 2011. Dr. Parker was elected in 1976 for leadership in electronics and contributions in inertial navigation, including the Navaho, Nautilus, Minuteman, and submarine fleet system.

DONALD W. PEACEMAN, 91, retired senior research advisor, Exxon Production Research Company, died June 19, 2017. Dr. Peaceman was elected in 1999 for contributions to the development and usage of transient three-dimensional multiphase simulators for predicting performance of petroleum reservoirs.

JACQUES PETERS, 95, professor emeritus, Department of Mechanical Engineering, Catholic University of Louvain, died December 12, 2018. Dr. Peters was elected a foreign member in 1977 for contributions to the engineering of machine tool and manufacturing processes.

ROBERT PLUNKETT, 100, retired professor of mechanics, University of Minnesota, Minneapolis, died March 19, 2019. Dr. Plunkett was elected in 1974 for contributions to experimental and analytical mechanics to solve noise, vibration, and fatigue problems.

CALVIN F. QUATE, 95, Leland T. Edwards Professor of Engineering Emeritus, Stanford University, died July 6, 2019. Dr. Quate was elected in 1970 for research, teaching, and management in microwave and solid-state electronics.

RICHARD J. ROBBINS, 85, president, The Robbins Group LLC, died May 30, 2019. Dr. Robbins was elected in 1991 for pioneering development of tunnel boring machines, and for enhancing the worldwide US techno-

logical leadership position in rock tunnel construction.

JOSEPH C. SALAMONE, 79, chief science officer, Rochal Industries, died July 9, 2019. Dr. Salamone was elected in 2011 for advances in ophthalmological devices and wound healing therapies and for distinguished academic and professional service.

WALTER J. SCHRENK, 85, retired senior research scientist, Dow Chemical Company, died October 31, 2018. Mr. Schrenk was elected in 1994 for invention, research, and development of polymer fabrication processes that have enabled broad use of plastics in new applications.

RICHARD N. WRIGHT, 87, retired director, Building and Fire Research, National Institute of Standards and Technology, died May 31, 2019. Dr. Wright was elected in 2003 for sustained leadership in building research, for the development of standards, and for representing the US building industry and research community worldwide.

Publications of Interest

The following reports whose authoring committees included NAE members were recently published by the National Academy of Engineering or the National Research Council. Unless otherwise noted, all publications are for sale (prepaid) from the National Academies Press (NAP), 500 Fifth Street NW–Keck 360, Washington DC 20055. For more information or to place an order, contact NAP online at <*www.nap.edu*> or by phone at (888) 624-6242. *(Note: Prices quoted are subject to change without notice. There is a 10 percent discount for online orders when you sign up for a MyNAP account. Add $6.50 for shipping and handling for the first book and $1.50 for each additional book. Add applicable sales tax or GST if you live in CA, CT, DC, FL, MD, NC, NY, VA, WI, or Canada.)*

Beyond Spectre: Confronting New Technical and Policy Challenges – Proceedings of a Workshop. In 2017 researchers discovered a vulnerability in microprocessors used in computers and devices all over the world. The vulnerability, named Spectre, combines side effects from caching and speculative execution, techniques that have been used for many years to increase the speed at which computers operate. The discovery upends a number of common assumptions about cybersecurity and draws attention to the complexities of the global supply chain and global customer base for the vast range of devices and cloud capabilities that all computer users rely on. In October 2018 the Forum on Cyber Resilience hosted a workshop to explore the implications of this development. This publication summarizes the workshop presentations and discussions.

Fred B. Schneider (chair), Samuel B. Eckert Professor of Computer Science, Cornell University, and **Butler W. Lampson**, technical fellow, Microsoft Research, served on the workshop committee. Free PDF.

Framing the Challenge of Urban Flooding in the United States. Flooding is the natural hazard with the greatest economic and social impact in the United States, and these impacts are becoming more severe. Catastrophic flooding from recent hurricanes, including Superstorm Sandy in New York (2012) and Hurricane Harvey in Houston (2017), caused billions of dollars in property damage, adversely affected millions of people, and damaged the economic well-being of major metropolitan areas. Flooding takes a heavy toll even in years without a named storm or event. Major freshwater flood events from 2004 to 2014 cost an average of $9 billion in direct damage and 71 lives annually, and the cumulative costs of frequent, small floods can be similar to those of infrequent extreme floods. This report examines examples in specific metropolitan areas and identifies commonalities and variances among the areas in terms of causes, adverse impacts, unexpected problems in recovery, or effective mitigation strategies. It also relates, as appropriate, causes and actions of urban flooding to existing federal resources or policies.

David R. Maidment (chair), Hussein M. Alharthy Centennial Chair in Civil Engineering, Center for Water and Environment, University of Texas at Austin, and **Dara Entekhabi**, professor, Parson Laboratory, Massachusetts Institute of Technology, served on the study committee. Paper, $58.00.

Quantum Computing: Progress and Prospects. Quantum mechanics, the subfield of physics that describes the behavior of very small (quantum) particles, provides the basis for a new paradigm of computing. First proposed in the 1980s as a way to improve computational modeling of quantum systems, quantum computing has recently garnered significant attention thanks to progress in building small-scale devices. However, significant technical advances will be required before a large-scale, practical quantum computer can be achieved. This report provides an introduction to the field, including the unique characteristics and constraints of the technology, and assesses the feasibility and implications of creating a functional quantum computer capable of addressing real-world problems. The report considers hardware and software requirements, quantum algorithms, drivers of advances in quantum computing and quantum devices, benchmarks associated with relevant use cases, the time and resources required, and how to assess the probability of success.

Mark A. Horowitz (chair), Yahoo! Founders Chair, EE and CS, Stanford University, and **Dan Boneh**, professor, computer science and electrical engineering, Stanford University, served on the study committee. Paper, $55.00.

Astrobiology Strategy for the Search for Life in the Universe. Astrobiology is the study of the origin, evolution, distribution, and future of life in the universe. The inherently interdisciplinary field encompasses astronomy, biology, geology, heliophysics, and planetary science, with complementary laboratory activities and field studies conducted in a wide range of terrestrial environments. Combining scientific interest and public appeal, the search for life in the solar system and beyond provides a scientific rationale for many activities of NASA and other national and international agencies and organizations. This study, requested by NASA, offers a science strategy for astrobiology that outlines key scientific questions, identifies the most promising research in the field, and indicates the extent to which mission priorities in decadal surveys address the search for life's origin, evolution, distribution, and future in the universe. The consensus report makes recommendations for advancing research, obtaining measurements, and realizing NASA's goal to search for signs of life in the universe.

Philip M. Neches, founder, Teradata Corporation, served on the study committee. Paper, $80.00.

Environmental Engineering for the 21st Century: Addressing the Grand Challenges. Environmental engineers support the well-being of people and the planet in areas where the two intersect. Work in this field has improved countless lives through innovative systems for delivering water, treating waste, and preventing and remediating pollution in air, water, and soil. These achievements are a testament to the multidisciplinary, pragmatic, systems-oriented approach that characterizes environmental engineering, This report outlines the crucial role for environmental engineers and identifies five pressing challenges that environmental engineers are uniquely poised to help address: sustainably supply food, water, and energy; curb climate change and adapt to its impacts; design a future without pollution and waste; create efficient, healthy, resilient cities; and foster informed decisions and actions.

Craig H. Benson, dean, School of Engineering, Hamilton Endowed Chair in Civil and Environmental Engineering, University of Virginia; **G. Wayne Clough**, secretary emeritus, Smithsonian Institution, and president emeritus, Georgia Institute of Technology; **John C. Crittenden**, director, Brook Byers Institute for Sustainable Systems, Georgia Institute of Technology; **Maxine L. Savitz**, retired general manager, Technology/Partnerships, Honeywell Inc.; **Norman R. Scott**, professor emeritus, Department of Biological and Environmental Engineering, Cornell University; and **R. Rhodes Trussell**, chair and founder, Trussell Technologies Inc., served on the study committee. Paper, $45.00.

Strategic Investments in Instrumentation and Facilities for Extraterrestrial Sample Curation and Analysis. The United States has a treasure trove of extraterrestrial samples that were returned to Earth via space missions over the past four decades. Analyses of these samples have led to major breakthroughs in understanding of the age, composition, and origin of the solar system. Having the instrumentation, facilities, and qualified personnel to analyze such samples, especially from missions that take up to a decade or longer from launch to return, is thus of paramount importance if NASA is to capitalize on its investment in these missions and to achieve the full scientific impact afforded by these extraordinary samples. This consensus report assesses current planetary science capabilities for

sample return analyses and curation, and determines what capabilities are missing and will be needed for future sample return missions. It evaluates whether current laboratory support infrastructure and NASA's investment strategy are adequate to meet these analytical challenges and advises how the community can keep abreast of evolving and new techniques in order to stay at the forefront of extraterrestrial sample analysis.

James H. Crocker, retired vice president, international, Lockheed Martin Space Systems Company, served on the study committee. Paper, $65.00.

Negative Emissions Technologies and Reliable Sequestration: A Research Agenda. To achieve goals for climate and economic growth, negative emissions technologies (NETs) that remove and sequester carbon dioxide (CO_2) from the air will need to play a significant role in mitigating climate change. Unlike carbon capture and storage technologies that remove CO_2 emissions directly from large point sources such as coal power plants, NETs remove CO_2 directly from the atmosphere or enhance natural carbon sinks. Storing the CO_2 from NETs has the same impact on the atmosphere and climate as simultaneously preventing emission of an equal amount of CO_2. Analyses show that deploying NETs may be less expensive and less disruptive than reducing some emissions, such as a substantial portion of agricultural and land-use emissions and some transportation emissions. This report assesses the benefits, risks, and "sustainable scale potential" for NETs and sequestration. It also defines the essential components of a research and development program, including estimated costs and potential impact.

Mark A. Barteau, vice president for research, Texas A&M University, College Station, served on the study committee. Paper, $115.00.

Minority Serving Institutions: America's Underutilized Resource for Strengthening the STEM Workforce. There are over 20 million young people of color in the United States whose representation in STEM education pathways and in the STEM workforce remains far below their numbers in the general population. Their participation could help reestablish US preeminence in STEM innovation and productivity while also increasing the number of well-educated STEM workers. Nearly 700 minority-serving institutions (MSIs) provide pathways to STEM educational success and workforce readiness for millions of students of color. They vary in their origins, missions, student demographics, and levels of institutional selectivity, but overall provide a gateway to higher education and the workforce for underrepresented students of color and those from low-income and first-generation-to-college backgrounds. The challenge is to capitalize on the unique attributes of these institutions and to equip them with the resources, faculty talent, and vital infrastructure needed to educate and train current and future generations of scientists, engineers, and health professionals. This consensus report examines the nation's MSIs and identifies promising programs and effective strategies that have the highest potential return on investment for the nation by increasing the quantity and quality of MSI STEM graduates. It also provides information about the importance of MSIs to other stakeholders in the nation's higher education system and the organizations that support them.

Wesley L. Harris, Charles Stark Draper Professor of Aeronautics and Astronautics, Massachusetts Institute of Technology, served on the study committee. Paper, $65.00.

Offshore Well Completion and Stimulation: Using Hydraulic Fracturing and Other Technologies – Proceedings of a Workshop. While the public is generally aware of the use of hydraulic fracturing for unconventional resource development onshore, it is less familiar with the well completion and stimulation technologies used in offshore operations, including hydraulic fracturing, gravel packs, "fracpacks," and acid stimulation. Just as onshore technologies have improved, well completion and stimulation technologies for offshore hydrocarbon resource development have progressed over many decades. To increase public understanding of these technologies, a NASEM workshop in October 2017 examined the unique features of operating in the US offshore environment, with discussions of well completion and stimulation technologies, environmental considerations and concerns, and health and safety management. Participants from government, industry, academia, and nonprofit sectors shared their perspectives on operational and regulatory approaches to mitigating risks to the environment and to humans in the development of offshore resources. This publication summarizes the workshop presentations and discussions.

David A. Dzombak, Hamerschlag University Professor and head, Department of Civil and Envi-

ronmental Engineering, Carnegie Mellon University, was a member of the workshop planning committee. Paper, $60.00

Continuous Manufacturing for the Modernization of Pharmaceutical Production: Proceedings of a Workshop. In July 2018 a NASEM workshop explored business and regulatory concerns associated with the use of continuous manufacturing techniques to produce biologics such as enzymes, monoclonal antibodies, and vaccines. Participants also discussed challenges to integration across the manufacturing system, including upstream and downstream processes, analytical techniques, and drug product development. The workshop addressed these challenges broadly across the biologics domain but focused particularly on drug categories of greatest FDA and industrial interest such as monoclonal antibodies and vaccines. This publication summarizes the workshop presentations and discussions.

Gintaras V. Reklaitis, Burton and Kathryn Gedge Distinguished Professor of Chemical Engineering, Purdue University, chaired the workshop planning committee. Paper, $65.00.

Research to Improve Estimates of Impacts of Changes in Truck Size and Weight Regulations. This TRB report defines a program of 27 coordinated research projects in six areas to reduce major sources of uncertainty in projections of the consequences of proposed changes in truck size and weight limits. The authoring committee acknowledges that improved models for projecting impacts of changes in truck size and weight limits, while necessary, will not guarantee the success of future truck size and weight policy studies. Such studies will be useful as guides for decisions only if policy objectives and practical policy options are clearly defined, the analysis is logically structured to reveal the most promising policies, and uncertainties are properly characterized. The committee's first report, in 2018, summarized the research recommendations of past truck size and weight limit studies and identified criteria for deciding the priority of topics for inclusion in the research plan.

Dennis F. Wilkie, retired corporate vice president, Ford Motor Company and Motorola Inc., and **Sharon L. Wood**, dean, Cockrell School of Engineering, University of Texas at Austin, served on the study committee. Free PDF.

Metrics for Successful Supercritical Water Oxidation System Operation at the Blue Grass Chemical Agent Destruction Pilot Plant: Letter Report. The supercritical water oxidation system at the Blue Grass Chemical Agent Destruction Pilot Plant near Richmond, Kentucky, is a secondary waste processing reactor that is an important unit in the function of the Blue Grass Chemical Agent Pilot Program. The system is designed to reactively destroy the primary products of agent hydrolysis, thus preventing chemical reformation of the original agents. This letter report develops metrics that can be used to determine the success or risks of failure of the system, focusing on safety, corrosion, performance, and reliability, availability, and maintainability.

Ronald M. Latanision, senior fellow, Exponent Inc., served on the study committee. Free PDF.

Final Report of the Committee on a Strategic Plan for US Burning Plasma Research. Fusion offers the prospect of virtually unlimited energy, and the United States and many nations around the world have made enormous progress toward achieving fusion energy. With ITER scheduled to go online within a decade and demonstrate controlled fusion ten years later, now is the time for the United States to develop plans to benefit from its investment in burning plasma research and take steps to develop fusion electricity for the nation's future energy needs. At the request of the Department of Energy, a NASEM committee was organized to develop a strategic plan for US fusion research. This report's two main recommendations are that the United States should (1) remain an ITER partner as the most cost-effective way to gain experience with burning plasma at the scale of a power plant and (2) start a national program of research and technology to support the construction of a compact pilot plant that produces electricity from fusion at the lowest possible capital cost.

C. Paul Robinson, president emeritus, Sandia National Laboratories, served on the study committee. Paper, $95.00.

PLANNED GIVING

Your **LEGACY...** it's **Personal**

Did you know that charitable gift annuity (CGA) rates increased last year for the first time since 2012?

Rates increased by 0.30% to 0.50% for those ages where most annuity contracts are done.

What is a CGA?

A CGA is a simple contract between you and the National Academy of Engineering (NAE) that will provide you with an income tax deduction and a fixed stream of payments to you and/or a family member for the rest of your/their life.

How will your gift inspire others?

Your CGA gift to the NAE supports our future efforts to advance engineering education and the public's awareness of the crucial role of engineering in our lives and world. You will also demonstrate your commitment and leadership by participating in the Campaign for the NAE.

With a charitable gift annuity, you can:

- Receive payments for the rest of your life, a portion of which will be tax-free;
- Deduct the income tax deduction calculated based on your age;
- Get the flexibility to make your gift with cash or marketable securities;
- Save capital gains taxes for gifts of appreciated securities;
- Support the future of the NAE; and
- Become a member of the NAE's Heritage Society.

Contact Radka Nebesky, Director of Development, at RNebesky@nae.edu or 202.334.3417 for more information about the Campaign for the NAE and how you can establish a CGA today. You can also visit www.NationalAcademiesGiving.org for more information.

The BRIDGE

(USPS 551-240)

Periodicals
Postage
Paid

National Academy of Engineering
2101 Constitution Avenue NW
Washington, DC 20418