Challenges for Securing Cyber Physical Systems

Alvaro A. Cárdenas*, Saurabh Amin†, Bruno Sinopoli‡,
Annarita Giani* Adrian Perrig‡ Shankar Sastry*
*Department of Electrical Engineering and Computer Sciences. University of California, Berkeley
†Department of Civil and Environmental Engineering. University of California, Berkeley
‡Department of Electrical and Computer Engineering. Carnegie Mellon University

Abstract

We discuss three key challenges for securing cyber-physical systems: (1) understanding the threats, and possible consequences of attacks, (2) identifying the unique properties of cyber-physical systems and their differences from traditional IT security, and (3) discussing security mechanisms applicable to cyber-physical systems. In particular, we analyze security mechanisms for: prevention, detection and recovery, resilience and deterrence of attacks.

1. Introduction

Cyber-physical systems (CPS) have been at the core of critical infrastructures and industrial control systems for many decades, and yet, there have been few confirmed cases of computer-based attacks. CPS, however, are becoming more vulnerable to computer attacks for many reasons [5].

In this paper we analyze some of the growing concerns for the security of CPS. We first discuss the need to develop adversary models for CPS. Then, we identify some of the new and fundamentally different problems that we encounter in CPS as compared to traditional IT security. We end the paper by outlining some research directions for preventing, detecting, responding, surviving, and deterring computer attacks.

2. Adversary Model

A systematic study of the security of any system requires the description of the threats we expect to face. Developing an adversary model is a way to understand the scope of the problem and assess the risks.

We now describe some potential attackers, their motivations, and their resources.

Cybercriminals compromise computers anywhere they can find them (even in control systems). These attacks may not be targeted (i.e., they do not have the intention of harming control systems), but may cause negative side effects: control systems infected with malware may operate inappropriately. In 2006, for example, an attacker compromised a computer at a water filtering plant in Pennsylvania and used it as its own distribution system for spam and pirated software [9]. Another famous example of these type of attacks occurred in January 2003, when computers infected with the Slammer worm shut down safety display systems at the Davis-Besse power plant in Oak Harbor, Ohio. The Slammer worm was not designed to attack control systems, but the use of commodity IT software by control systems allowed this general-purpose worm to infect computers used in safety-critical systems.

Disgruntled employees are currently the major source of targeted computer attacks against control systems. The most well-known computer security incident in control systems is the attack on Maroochy Shire Council's sewage control system in Queensland, Australia, in 2000 [25]. The culprit of the attack on Maroochy Shire was a disgruntled ex-employee of the contractor company that had installed the control system, and who was trying to convince the water treatment company to hire him to solve the problem. Recently, there have been many more reported attacks caused by disgruntled employees [1], [2], [16], [22]. These attacks are important from a security point of view because they are caused by *insiders*: individuals with authorized access to computers and networks used by control systems; so even if control networks were completely isolated from public networks (and the Internet), attacks by insiders would still be possible. Because disgruntled employees generally act alone, the potential consequences of their attacks may not

be as damaging as the potential harm caused by larger organized groups.

Terrorists, activists, and organized criminal groups are another potential threat to control systems. While there is no concrete evidence that terrorists or activists have targeted control systems via computer attacks, there is some evidence on the possible involvement of criminal groups. In 2008, a senior analyst for the CIA mentioned that there was evidence of computer intrusions into some European power utilities followed by extortion demands [12]. Attacking control systems for extortion is not new. Physical attacks – for extortion and terrorism– are a reality in some countries [20]. Cyber-attacks are a natural progression to physical attacks: they are cheaper, less risky for the attacker, are not constrained by distance, and are easier to replicate and to coordinate.

Nation states may also be a possible threat to control systems. While some national-security officials claim that the electric grid in the U.S. has been penetrated by spies [11], others claim that the Soviet Union was the victim of an attack during the cold war (in 1982) when a logic bomb caused a gas pipeline explosion in Siberia [23]. In general, it should be no surprise that most military powers are looking into future attack technologies, including cyber-attacks against the physical infrastructure of other nations.

2.1. Attacks

Attackers may be able to launch unique attacks to control systems (i.e., attacks that are not possible in traditional IT systems). One possible example can be *resonance attacks*. In a resonance attack, an attacker that has compromised some sensors or controllers will force the physical system to oscillate at its resonant frequency.

We believe that a major research challenge is to identify and categorize the new type of attacks that are possible in control systems, and to examine their possible consequences.

2.2. Consequences of an Attack

To our knowledge there has not been a publicly-available analysis of the possible consequences to attacks against critical infrastructures. In our view, while some of the reports on SCADA security might appear overly alarmist (safety safeguards in most control systems might prevent major catastrophes), the fact that a user is able to obtain unauthorized privileges in a control system should be taken seriously.

The Maroochy Shire incident in 2000 showed some of the effects that attacks can have. We believe that an important direction for future research is on identifying the risks and consequences of a successful attack.

3. Differences between corporate IT security and CPS IT security

While it is clear that the security of control systems has become an active area in recent years, we believe that, so far, no one has been able to articulate what is new and fundamentally different in this field from a research point of view compared to traditional IT security.

In this paper we would like to start this discussion by summarizing some previously identified differences and by proposing some new problems.

The property of control systems that is most commonly brought up as a distinction with IT security is that software **patching and frequent updates, are not well suited for control systems**. For example, upgrading a system may require months of advance in planning of how to take the system offline; it is, therefore, economically difficult to justify suspending the operation of an industrial computer on a regular basis to install new security patches. Some security patches may even violate the certification of control systems.

In a recent anecdote, on March 7 of 2008, a nuclear power plant was accidentally shutdown because a computer used to monitor chemical and diagnostic data from the plant's business network rebooted after a software update. When the computer rebooted, it reset the data on the control system, causing safety systems to errantly interpret the lack of data as a drop in water reservoirs that cool the plant's radioactive nuclear fuel rods [17].

Another property of control systems that is commonly mentioned is the real-time requirements of control systems. Control systems are autonomous decision making agents which need to make decisions in real time. While availability is a well studied problem in information security, **real-time availability** provides a stricter operational environment than most traditional IT systems.

Large industrial control systems also have a large amount of **legacy systems**. Several research efforts have tried to provide lightweight cryptographic mechanisms to ensure data integrity and confidentiality [27], [30]. The recent IEEE P1711 standard is designed for providing security in legacy serial links [14]. Having some small level of security is better than having no

security at all; however, *we believe that most of the efforts done for legacy systems should be considered as short-term solutions*. For properly securing critical control systems the underlying technology must satisfy some minimum performance requirements to allow the implementation of well tested security mechanisms and standards.

Not all operational differences are more severe in control systems than in traditional IT systems. By comparison to enterprise systems, control systems exhibit comparatively **simpler network dynamics**: Servers change rarely, there is a fixed topology, a stable user population, regular communication patterns, and a limited number of protocols. Therefore, implementing network intrusion detection systems may be easier than in traditional enterprise systems [6].

3.1. New Security Problems in Control Systems

While all these differences are important, we believe that the major distinction of control systems with respect to other IT systems is the interaction of the control system with the physical world.

In general, information security has developed mature technologies and design principles (authentication, access control, message integrity, separation of privilege, etc.) that can help us prevent and react to attacks against control systems. However, research in computer security has focused traditionally on the protection of information. Researchers have not considered how attacks affect the *estimation* and *control* algorithms –and ultimately, how attacks affect the physical world.

We argue that while the current tools of information security can give *necessary* mechanisms for the security of control systems, these mechanisms alone are not *sufficient* for the defense-in-depth of control systems.

We believe that by understanding the interactions of the control system with the physical world, we should be able to

1) Better understand the consequences of an attack: so far there is no research on how an adversary would select an strategy once it has obtained unauthorized access to some control network devices.
2) Design novel attack-detection algorithms: by understanding how the physical process should behave based on our control commands and sensor measurements, we can identify if an attacker is tampering with the control or sensor data.
3) Design new attack-resilient algorithms and architectures: if we detect an attack we may be able to change the control commands to increase the resiliency of the system.

4. Countermeasures

Up to now, most of the effort for protecting control systems (and in particular SCADA) has focused on *reliability* (the protection of the system against random faults). There is, however, an urgent growing concern for protecting control systems against malicious cyber-attacks [4], [8], [10], [28], [29].

4.1. Prevention

We believe that the major research challenge for preventing the compromise of control systems is to identify ways in which asset owners and vendors of control systems will be motivated to follow best security practices. There are currently some efforts in this direction, in particular from the standards community.

Several sectors –including chemical, oil and gas, and water– are currently developing programs for securing their infrastructure. The electric sector is leading the way with the North American Electric Reliability Corporation (NERC) cybersecurity standards for control systems [19]. NERC is authorized to enforce compliance to these standards, and it is expected that all electric utilities are fully compliant with these standards by 2010.

NIST SP 800-53 Revision 3—the guideline for security best practices which federal agencies should meet—includes in its appendix a section on the security of industrial control systems. In addition, NIST has also published a Guide to Industrial Control System (ICS) Security [26]. Although these recommendations are not enforceable, they can provide guidance for analyzing the security of most utility companies.

ISA (a society of industrial automation and control systems) is developing ISA-SP 99: a security standard to be used in manufacturing and general industrial controls.

The use of wireless sensor networks in SCADA systems is becoming pervasive, and thus we also need to study their security. A number of companies have teamed up to bring sensor networks in the field of process control systems, and currently, there are two working groups to standardize their communications [13], [15]. Their wireless communication proposal has options to configure hop-by-hop and end-to-end confidentiality and integrity mechanisms. Similarly they provide the necessary protocols for access control and key management.

While all these efforts are important, we believe that we need to find other incentives (e.g., new policies, education and outreach efforts, or economic programs) to complement these self-regulated standards.

4.2. Detection and recovery

Because we can never rule out successful attacks, security engineering has recognized the importance of detection and response to attacks. While traditional intrusion detection systems look at network or computer system traces; control systems can provide a paradigm shift for intrusion detection. In particular, by monitoring the physical system for anomalies we may be able to detect attacks that are undetectable from the IT side.

In addition, intrusion *detection* systems have not considered algorithms for detecting *deception* attacks against *estimation* and *control* algorithms. In particular, previous detection of deception attacks launched by compromised sensor nodes assume a large number of redundant sensors [31]: they *have not considered the dynamics of the physical system* and how this model can be used to detect a compromised node. Furthermore, there has not been any detection algorithm to identify deception attacks launched by compromised controllers or sensors.

Another key aspect for detecting attacks is to provide enough information awareness to operators of control systems. Operators may need to be trained to detect possible attacks, and to have a properly defined protocol or guideline on how to respond and recover from them.

In the particular case of SCADA systems successful attacks aim to change the master station perception of the environment modifying the semantic of the information. At the same time the attackers want to minimize the manipulation of hardware and software. This is done deceiving some of the nodes. Traditional intrusion detection systems hardly work against these types of attacks. An approach to detection is to use linear classical estimation techniques based on models of the environment. Both physical attacks where the structural properties of the system are modified and integrity attacks on both sensor and control data must be detected. Nodes sense the environment and transmit the sensed information to the master station that performs aggregation and data analysis. Based on the sensor data the master station will compute an estimate of the state of the system. Such estimate can then be used to make decisions or control. Both sensing and control data are sent using communication networks and are therefore subject to denial of service and integrity attacks. In the latter, in particular, an attacker may alter the sensor reading approaching the node and modifying its neighborhood so that it senses artificial values. These are not technical attacks against the sensor network itself since hardware and software can remain genuine. So traditional intrusion detection techniques are not adequate.

Another approach is to implement a model-based detection scheme, cast as game between the detector and the attacker, where, given a desired probability of false alarm, the detector tries to maximize the detection probability, while the attacker will attempt at minimizing such probability via intelligent manipulation of the data sent by the compromised components. Initial results are available for dynamical systems that can be modeled as linear time-invariant with additive Gaussian noise.

To address the challenge of the interaction between the (cyber-) IT systems and the physical component methods from the area of hybrid systems has been proposed. Hybrid systems [18] have been a topic of intense research for the past decade, in the boundary between computer science and control engineering. They provide a unified framework for jointly modeling continuous systems (like the power transmission and distribution processes) and discrete systems (like the SCADA systems).

An essential task for facilitating the operator's response is its information awareness. Research on human-computer interaction for improving the awareness of the operator is a key research challenge.

Besides recovery with a human in the loop there is also a need for automatic recovery. Because CPS use *autonomous, real-time decision making algorithms* for controlling the physical world, attacks may introduce new challenges for the design and analysis of secure systems. We can bring ideas from control theory such as reconfiguration or fault-detection and isolation, to design autonomous and real-time detection and response algorithms for safety-critical applications that require real-time responses.

4.3. Resilience

There are several security design principles that can be useful for designing control systems that can survive attacks [3], [24]. **Redundancy**, for example, is a way to prevent a single-point of failure. **Diversity** is a way to prevent that a single attack vector can compromise all the replicas (the added redundancy). And the **principle of least-privilege**, and the **separation of privilege** (also known as separation of duty) principle

are design guidelines to limit the amount to privileges that a corrupted entity can have.

Physical and analytical redundancies should be combined with security principles (e.g., diversity and separation of duty) to **adapt** or **reschedule** its operation during attacks. For example, under sensor faults or when only intermittent sensory information is available, the system should be able to operate using open-loop control for a sufficient amount of time.

We also need to design novel *robust control and estimation algorithms* that consider more **realistic attack models** from a security point-of-view. These attack models should model deception and DoS attacks. Under the influence of such attacks, these algorithms should optimize the worst-case performance. **Game theoretic** techniques developed in economics for modeling rational adversaries might also be useful for this task.

Finally, the complex manner in which closed-loop dynamics of controlled systems (that are modeled by differential or difference equations) with the actions imposed by the behavior of networked components (the actions being random or malicious) mandates a re-thinking of the traditional "separation principle" in which security approaches to control systems are designed and deployed at enterprise or regulatory level and controllers are designed at physical or regulatory level. For example, it is not enough to design access control policies without taking into account the consequence analysis determining answers to "what if" type questions. Such a consequence analysis must factor the behavior of closed-loop dynamics under attacks on network components.

4.4. Deterrence

Deterrence usually depends on successful legislation, law enforcement, and international collaboration for tracking crimes committed outside our borders.

We believe that the identification of new deterrence mechanisms for the security of CPS is a promising area of research.

5. Building a Testbed for Cyber Physical Systems

In order to better understand how to protect cyber physical systems, it is imperative to perform vulnerability assessment and develop appropriate security mechanisms to protect them against attacks. To do so, developing a testbed is essential. In fact there are enormous limitation in testing attacks and countermeasures in a real system. Recently, a SCADA testbed for the power system has been developed at the University of Illinois at Urbana-Champaignin [21]. Sandia National Laboratories SCADA testbed is an example of a government sponsored testbed. The European community has also started working on creating a SCADA security testbed [7]. Within TRUST, an initiative between Vanderbilt University and the University of California at Berkeley focus on building a SCADA testbed. The testbed allows different variants for the various components of the system, see figure 1.

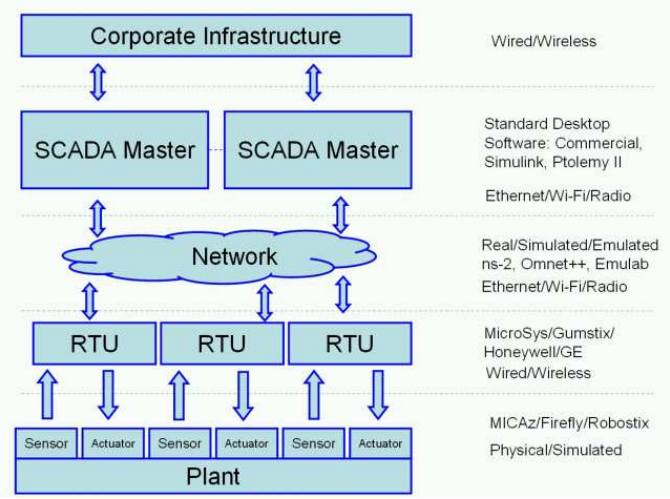

Figure 1. TRUST SCADA Tesbed

6. Conclusions

We have presented the current status of the field of secure control. We identified some unique properties that these systems have in comparison to traditional IT systems and proposed some new research challenges based on the physical models of the process being controlled. Our research challenges are mostly unsolved and we believe that future research in these areas can provide an additional level of security to control systems.

Acknowledgments

This work was supported in part by TRUST (Team for Research in Ubiquitous Secure Technology), which receives support from the National Science Foundation (NSF award number CCF-0424422) and the following

organizations: AFOSR (#FA9550-06-1-0244) Cisco, British Telecom, ESCHER, HP, IBM, iCAST, Intel, Microsoft, ORNL, Pirelli, Docomo Research USA, Qualcomm, Sun, Symantec, Telecom Italia, and United Technologies.

And by by the Center for Hybrid and Embedded Software Systems (CHESS) at UC Berkeley, which receives support from the National Science Foundation (NSF awards #CCR-0225610 (ITR),#0720882 (CSR-EHS: PRET), #0647591 (CSR-SGER), and #0720841 (CSR-CPS)), the U. S. Army Research Office (ARO #W911NF-07-2-0019), the U. S. Air Force Office of Scientific Research (MURI #FA9550-06-0312 and AF-TRUST #FA9550-06-1-0244), the Air Force Research Lab (AFRL), the State of California Micro Program, and the following companies: Agilent, Bosch, Lockheed Martin, National Instruments, Thales and Toyota.

We would also like to thank our reviewers for helpful comments, and for pointing out the section on control security in NIST SP 800-53 Revision 3.

References

[1] United States Attorney, Eastern District of California. Sacramento man pleads guilty to attempting ot shut down california's power grid. http://www.usdoj.gov/usao/cae/press_releases/docs/2007/12-14-07DenisonPlea.pdf, November 2007.

[2] United States Attorney, Eastern District of California. Willows man arrested for hacking into Tehama Colusa Canal Authority computer system. http://www.usdoj.gov/usao/cae/press_releases/docs/2007/11-28-07KeehnInd.pdf, November 2007.

[3] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing*, 1(1):11–32, January-March 2004.

[4] E. Byres and J. Lowe. The myths and facts behind cyber security risks for industrial control systems. In *Proceedings of the VDE Congress, VDE Association for Electrical Electronic & Information Technologies*, October 2004.

[5] A. A. Cárdenas, S. Amin, and S. Sastry. Research challenges for the security of control systems. In *Proceedings of 3rd USENIX workshop on Hot Topics in Security (HotSec)*, San Jose, CA, USA, July 2008.

[6] S. Cheung, B. Dutertre, M. Fong, U. Lindqvist, K. Skinner, and A. Valdes. Using model-based intrusion detection for SCADA networks. In *Proceedings of the SCADA Security Scientific Symposium*, Miami Beach, FL, USA, 2007 2007.

[7] H. Christiansson and E. Luiijf. Creating a european scada security testbed, 2007.

[8] J. Eisenhauer, P. Donnelly, M. Ellis, and M. O'Brien. *Roadmap to Secure Control Systems in the Energy Sector*. Energetics Incorporated. Sponsored by the U.S. Department of Energy and the U.S. Department of Homeland Security, January 2006.

[9] R. Esposito. Hackers penetrate water system computers. http://blogs.abcnews.com/theblotter/2006/10/hackers_penetra.html, October 2006.

[10] GAO. Critical infrastructure protection. Multiple efforts to secure control systems are under way, but challenges remain. Technical Report GAO-07-1036, Report to Congressional Requesters, September 2007.

[11] S. Gorman. Electricity grid in the U.S. penetrated by spies. http://online.wsj.com/article/SB123914805204099085.html, April 8 2009.

[12] A. Greenberg. Hackers cut cities' power. In *Forbes*, Jaunuary 2008.

[13] Hart. http://www.hartcomm2.org/frontpage/wirelesshart.html. *WirelessHart whitepaper*, 2007.

[14] S. Hurd, R. Smith, and G. Leischner. Tutorial: Security in electric utility control systems. In *61st Annual Conference for Protective Relay Engineers*, pages 304–309, April 2008.

[15] ISA. http://isa.org/isasp100. *Wireless Systems for Automation*, 2007.

[16] D. Kravets. Feds: Hacker disabled offshore oil platforms' leak-detection system. http://www.wired.com/threatlevel/2009/03/feds-hacker-dis/, March 18 2009.

[17] B. Krebs. *Cyber Incident Blamed for Nuclear Power Plant Shutdown*. Washington Post, http://www.washingtonpost.com/wp-dyn/content/article/2008/06/05/AR2008060501958.html, June 2008.

[18] J. Lygeros, K. H. Johansson, S. Simic, J. Zhang, , and S. Sastry. Dynamical properties of hybrid automata, 2003.

[19] NERC-CIP. *Critical Infrastructure Protection*. North American Electric Reliability Corporation, http://www.nerc.com/cip.html, 2008.

[20] B. News. *Colombia Rebels Blast Power Pylons*. BBC, http://news.bbc.co.uk/2/hi/americas/607782.stm, January 2000.

[21] H. Okhravi, C. Grier, M. Davis, Z. Tate, D. Nicol, , and T. Overbye. Cyber-security simulation testbed.

[22] K. Poulsen. Ex-employee fingered in texas power company hack. http://www.wired.com/threatlevel/2009/05/efh/, May 29 2009.

[23] T. Reed. *At the Abyss: An Insider's History of the Cold War*. Presidio Press, March 2004.

[24] J. H. Saltzer and M. D. Schroeder. The protection of information in computer systems. *Proceedings of the IEEE*, 63(9):1278–1308, September 1975.

[25] J. Slay and M. Miller. Lessons learned from the maroochy water breach. In *Critical Infrastructure Protection*, volume 253/2007, pages 73–82. Springer Boston, November 2007.

[26] K. Stouffer, J. Falco, and K. Kent. Guide to supervisory control and data acquisition (SCADA) and industrial control systems security. Sp800-82, NIST, September 2006.

[27] P. P. Tsang and S. W. Smith. YASIR: A low-latency high-integrity security retrofit for lecacy SCADA systems. In *23rd International Information Security Conference (IFIC SEC)*, pages 445–459, September 2008.

[28] R. J. Turk. Cyber incidents involving control systems. Technical Report INL/EXT-05-00671, Idao National Laboratory, October 2005.

[29] US-CERT. *Control Systems Security Program*. US Department of Homeland Security, http://www.us-cert.gov/control_systems/index.html, 2008.

[30] A. K. Wright, J. A. Kinast, and J. McCarty. Low-latency cryptographic protection for SCADA communications. In *Applied Cryptography and Network Security (ACNS)*, pages 263–277, 2004.

[31] Q. Zhang, T. Yu, and P. Ning. A framework for identifying compromised nodes in sensor networks. In *Proceedings of the 2nd International Conference on Security and Privacy in Communication Networks (SecureComm 2006)*, August 2006.