Cyber-physical Systems Security- Challenges and Research Ideas

Partha Pal, Rick Schantz, Kurt Rohloff, Joseph Loyall
BBN Technologies, Cambridge, MA 02138
{ppal,schantz,krohloff,jloyall}@bbn.com

1. Introduction

Cyber-physical systems (CPS) are at a crossroads. Typical CPSs with multi-loop control consisting of device controllers, plant-level distributed control, and system-wide SCADA components are pushed (see Figure 1 below) by market forces to:

- Connect many diverse (e.g., control and business) systems leading to an internetworked system of systems spanning wide geographic areas,
- Use off-the-shelf networks (sometimes involving public infrastructure like the Internet, and often shared) for communication, and general purpose computing hardware and software for information processing, and
- Empower end-users and customers with more information and control.

As a result of this push, restrictive and purpose-built interfaces typical of CPSs are replaced by more open interfaces, and new dependencies among interconnected subsystems are established (sometimes unintentionally). Access to more information and control surfaces for users and supervisors in an open and interconnected situation amplifies the impact of erroneous or malicious actions. Because today's cyber-physical system of systems is often part of critical national infrastructure, the stakes are even higher. Additionally, unlike many distributed systems (e.g., e-commerce or logistics planning systems) where a delay or unavailability of a few seconds is usually tolerable, CPSs must often meet strict timing requirements during normal operation as well as during recovery. We argue that securing the emerging internetworked and information-rich CPSs is harder than securing typical distributed information systems because a number of cyber-security issues that have been addressed individually or in isolation within a subsystem come together in a new context with many additional challenging requirements. In this position paper we present some of them.

2. Challenges and Solution Approaches

In many ways, the direction taken by CPSs is reminiscent of the early days of the Internet. The Internet is arguably the most successful internetworked environment and example of user empowerment. Our view of CPS security is informed by our Internet experience [1], and based on our expertise in multiple generations of cyber-security research [2, 3, 5], especially in survivability [4], and on our control systems experience [14, 15]. We claim that consideration of CPS security must accept that a) it is

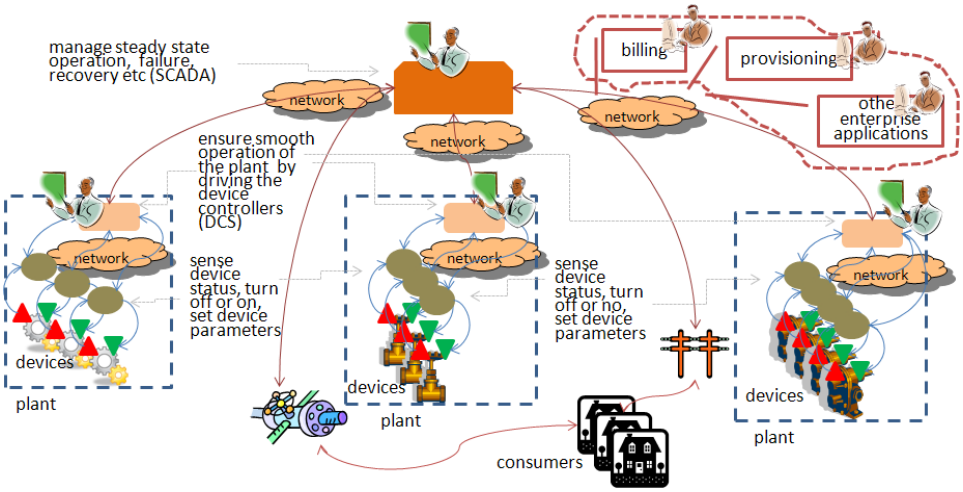

Figure 1: Modern CPS as internetworked system of systems

impossible to prevent or eliminate the threat of all attacks, and b) it may not even be possible to detect attacks accurately and early enough to be stopped. Attempts to build survivable versions of distributed information systems based on these assumptions have resulted in considerable progress in a number of technical problem areas. However, transferring these successes to build survivable cyber-physical systems are faced with additional challenges. Examples of security issues that are further exacerbated in the context of CPS security include:

1. **Data Interpretation**: Making sense of large volumes of low-level data covering *both* the *cyber* and *physical* aspects of the CPS, some of which may be incomplete and imperfect, to determine the security state of the CPS.
2. **Information and control sharing**: Effectively sharing information and control authority in the context of internetworked system of systems where individual systems, users and operators belong to different organizations, and perform different tasks or have disparate roles and responsibilities.
3. **Containing compromises**: Minimizing the system-wide *physical* impact of *cyber*-level compromises or privilege abuses by malicious actors within individual subsystems
4. **Maintaining timeliness**: Ensuring that the timeliness properties of the base system are maintained during normal no-attack condition as well as in the presence of malicious activities, in spite of the overhead and interference introduced by defense mechanisms.
5. **Validation**: Validating the defenses applied to a CPS, especially the dynamic defensive behavior, so that it can be accepted for use in safety critical applications.

We now discuss these challenges in more detail, offering preliminary ideas to move forward. The ordering reflects ease of exposition, not importance or difficulty.

2.1. Data Interpretation

The flexibility and ease accorded by modern computing hardware, general purpose software, and ubiquitous connectivity encourage collection, storage and dissemination of large volumes of data. This has been observed in cyber domains (e.g., the number of log messages and alerts generated [6]) as well as non-cyber domains (e.g., information collected by the intelligence community [7]). Empowering operators and users, not only to utilize the full potential of the CPS, but also to enable them to effectively handle malicious attacks, provides additional motivation to collect even more data. Noise and uncertainty in such data resulting from sensor and communication failures is unavoidable. Because the data is collected over wide area distributed system, it is likely that at any point an observer (i.e., the sensors) will only have partial information about the system state or event of interest. The involvement of a malicious adversary that can choose to attack the CPS to cause real failures in the physical level and suppress reporting of such failure in the cyber-level, or inject misinformation at the cyber-level about non-existent failures in the physical level adds a level of complexity that does not arise in the defense of typical information systems. Therefore, the large volumes of available data covering both the cyber and physical aspects of the CPS that are supposedly collected to assist the operators and users in their efforts to use and manage the CPS more efficiently presents both a challenge and an opportunity. Making sense of partial and imprecise information that may reflect various combination of the compromises in the physical as well as the cyber aspects of the CPS is a significant technical challenge. On the other hand, our ability to collect, disseminate and store such data presents the opportunity to develop and employ sophisticated algorithms to process and interpret such data.

We argue that CPSs of the future will need to incorporate automated mechanisms to handle the large volume of collected data. Such mechanisms will need to employ far more sophisticated algorithms than those that are currently used for data visualization in the typical HMI applications today. These mechanisms will act more like an expert assistant to users and operators, as opposed to a display tool, considering among others, the possibilities that what is reported and observed cannot be trusted. Reasoning about the evidentiary (e.g., task completion and occurrence of scheduled events) and accusatory (e.g., alerts and fault reports) aspects of collected data, such mechanisms will evaluate and present the most likely explanation in terms of the operational state of the system (i.e., for each component in the CPS whether the component is up, down or compromised) that fits the current set of observations. We have done initial work in the cyber realm [8]. In the context of CPSs, additional research is needed to construct uniform knowledge representation (KR) and reasoning mechanisms covering both cyber (e.g., hosts and networks) and physical (e.g., devices and plants) aspects of the CPS,

and supporting different types of concurrent workflow and operational goals (e.g., control and business).

2.2. Information and Control Sharing

Next generation CPSs will span multiple organizations, connect different types of systems ranging from individual plant control systems to business systems like billing and customer management. Consequently, they need to support various types of end users and operators, each with their specific need to access information and control interfaces. The need to control information flow across organizational boundaries, and to prevent unauthorized sharing and use of sensitive information and control surfaces will act as a counterbalance to the urge to collect and share large volumes of data. This leads to our second challenge: how to effectively share information and control responsibility in the context of internetworked system of systems where individual systems, users and operators belong to different organizations, perform different tasks or have disparate roles and responsibilities, without violating applicable organizational policies and regulatory requirements? For a specific example consider a CPS that connects medical equipments in patient rooms in a hospital. The CPS generates and updates electronic patient records, but also controls drugs and therapeutic interventions into patients' bodies. While patient records are shared among hospitals, pharmacies, doctors' offices, payers, researchers, etc. not everybody should have the same level of access. A whole different set of rules apply to coordinate and regulate access to the control surfaces. Clearly, an open internetwork of various subsystems with the CPS is unacceptable. On the other hand, it is impractical and a software engineering nightmare to create a system that offers a customized stovepipe for each and every case of sharing and interoperation. How does one design and implement an internetworked system where various cooperating external subsystems can have only the access that is stipulated for them?

In terms of way forward, we note that federated [9] and cross-domain solutions [10] are being developed to address analogous issues in the cyber realm. It is envisioned that these solutions will be applicable in the CPS context. However, new research is needed to facilitate federating control—the federated information management solutions typically are concerned about information management. The existing cross-domain solutions also assume a fairly restrictive environment in terms of available interconnections among domains and about "need to know", which may not be true for CPSs.

2.3. Containing Compromises

Empowering end-users by composing individual subsystems that may offer adequate control and security within themselves incurs the risk of unintended consequences. In the context of the electric grid, variable pricing may lead to a situation where all consumers attempt to run their power hungry devices at a time when the price is low, leading to a new peak instead of peak shaving, or even causing a brownout. The mechanism to push pricing data can be hijacked by a malicious actor to trigger such an onslaught; unsuspecting users can be tricked to interact with and divulge private information to phishing portals. Similarly, having the utility company (operator) control smart devices in consumers' homes is also a risk: compromised utility systems can leak consumers' private information, or control the smart home devices in an abusive manner. Consequently, our third challenge is to minimize the system-wide impact of cyber-level compromises or privilege abuses by malicious actors within individual subsystems.

As a potential way forward, let us note that use of a survivability architecture incorporating multiple overlapping detection capabilities and containment regions with potentially overlapping control hierarchies has showed promising results in the context of information systems [4, 11, 17]. New research is needed to understand the transitive reach into the physical realm from the cyber realm, especially when a cyber component is compromised and abused by a malicious adversary, and to develop technology and design techniques to minimize and dynamically restrict the cyber-based interdependent control surfaces that cut across organizational and subsystem boundaries. The detection mechanisms used in survivability architectures need to be extended to cover physical aspects of the system. The semantics of containment may also need to be extended. For example, each activity in the cyber realm (such as blocking network access of a device because it is corrupt) may need to closely synchronize with activities in the physical realm (such as load shedding).

2.4. Maintaining Timeliness

The timing requirements associated with various CPS functions offer an interesting attack surface. Without any deep penetration into the system or its semantics, a malicious attacker can subvert its operation by simply manipulating the time taken to transport or process information. Use of off-the-shelf platforms and their published vulnerabilities provide

the attacker an easy starting point. Interconnection over potentially shared network infrastructure further increases the risk. Therefore, our fourth challenge is to ensure timeliness of the base system (e.g., collecting of data, analysis, etc.) as well as timeliness of the defense (e.g., ability to detect malfunction and mounting an effective response) in the presence of malicious actors. A specific example arises in synchro-phasor-based wide area measurement systems. Synchro-phasors produce ~30 measurement samples (e.g., voltage amplitude and phase angle, frequency) per second, and for some situation assessment and state estimation applications these measurements need to be delivered with a latency of ~1 second. Not only an attacker can disrupt the flow of measurement updates and increase the latency beyond useful range, he can also insert fake and incorrect updates. How do we ascertain that the delivered samples are genuine (i.e., from an authenticated source) and not modified in transit all the while remaining within the specified latency requirement?

Regarding the way forward, we cite emerging technologies such as fast set up (and recovery) of circuit style connections over optical networks [12], and complex event processing [13] that minimizes memory access and other bottlenecks as candidates for further investigation. TDM style scheduling of communication and CPU, separation of service delivery and control paths, and anytime algorithms with quality measures are also potential candidates. Lightweight and optimistic mechanisms that ascertain source authenticity and message integrity of a stream by oversampling as opposed to a per message basis present another interesting avenue of investigation.

2.5. Validation

Most of the potential solution approaches described above will require additional software. As it is, the reliability of software components is an issue. Then, the state of the art in cyber-security validation is still maturing. Infusion of the cyber components and the existence of intelligent adversaries make the approaches to validate physical systems insufficient for next generation CPSs. While it is possible to develop new defenses and devise survivability architectures for CPSs, ascertaining that they work as expected will remain an overarching problem. This leads to our fifth and final challenge: How to validate the defenses applied to a CPS? We note that infusion of more cyber-based control introduces opportunity for additional unassisted (e.g., software controlled) operation in future CPSs. To defend against attacks on the cyber components, additional defense mechanisms will also be added, and some of which will also introduce unassisted dynamic behavior. A specific issue in this context is validating dynamic behavior that cannot be enumerated and tested exhaustively a-priori, especially for CPSs that have safety critical applications.

In terms of way forward, we argue that until a new methodology emerges, validation of next generation of CPS security and survivability must combine existing best practices in both the cyber and physical realms such as adversarial testing using red teams and model based studies [14, 16]. Identifying various dependencies in the internetworked system of systems, and studying the multi-layered dependencies using stochastic models to detect and study cascades, surges and common mode failures is a potential new avenue of research.

3. Conclusion

Securing CPSs is a hard problem, even with the recognition that the goal is not absolute security. Advances made in survivability and security research in the cyber realm, emerging software and networking technologies, and the lessons learnt from the Internet experience gives us hope that a) we can build a high-water mark, and b) keep raising the bar by means of constant innovation. There are sufficient technical challenges to energize the research community; to move forward we need support from the funding agencies and commitment from the stakeholders of the CPSs of critical importance.

10. References

[1] R. Schantz, “BBN’s Network Computing Software Infrastructure and Distributed Applications (1970-1990),” *IEEE Annals of the History of Computing*, Vol. 28, No. 1, IEEE, 2006, pp. 72-88.

[2] S. Kent, “Security and the Internet, circa 1980-1990,” in A Technical History of the Internet, Tutorial at the annual conference of the ACM Special Interest Group in data Communication (SIGCOMM) 1999, Cambridge, MA (http://www.cs.utexas.edu/users/dragon/sigcomm/t1/kent.slides.ppt).

[3] D. F. Vukelich, D. Levin, J. Lowry, "Architecture for Cyber Command and Control: Experiences and Future Directions," DARPA Information Survivability Conference and Exposition,, vol. 1, pp. 0155, DARPA Information Survivability Conference and Exposition (DISCEX II'01) Volume I-Volume 1, 2001.

[4] P. Pal, F. Webber, R. Schantz, “Survival by Defense-Enabling”, Foundations of Intrusion Tolerant Systems

(Organically Assured and Survivable Information Systems), IEEE, 2003, pp 261-269.

[5] K. Rohloff, T. Başar. “Deterministic and Stochastic Models for the Detection of Random Constant Scanning Worms,” *ACM TOMACS (Transactions on Modeling and Computer Science)* Special Issue on Simulation, Modeling and Security, Vol. 18, No. 2, April 2008, pp. 1-24.

[6] G. Conti, K. Abdullah, J. Grizzard, J. Stasko, J. Copeland, M. Ahamad, H. Owen, C. Lee, “Countering Security Analyst and Network Administrator Overload Through Alert and Packet Visualization," *IEEE Computer Graphics and Applications (CG&A),* Vol. 26, No. 2, IEEE, 2006, pp. 60-70.

[7] M. MacDonald, A. Oettinger, “Information Overload: Managing Intelligence Technologies,” Harvard International Review, Fall 2002.

[8] P. Benjamin, P. Pal, F. Webber, P. Rubel, M. Atighetchi. “Using A Cognitive Architecture to Automate Cyberdefense Reasoning,” in Proceedings of the 2008 ECSIS Symposium on Bio-inspired, Learning, and Intelligent Systems for Security (BLISS 2008), IEEE Computer Society, August 4-6, 2008, Edinburgh, Scotland.

[9] G. Mitchell, J. Loyall, J. Webb, M. Gillen, A. Gronosky, M. Atighetchi, A. Sinclair, “A Software Architecture for Federating Information Spaces for Coalition Operations,” in Proceedings MILCOM 2008, San Diego, CA, November 17-19, 2008.

[10] DoD cross-domain solutions access controlled web page: http://iase.disa.mil/cds/index.html)

[11] P. Pal, F. Webber, R. Schantz, "The DPASA Survivable JBI- A High-Water Mark in Intrusion-Tolerant Systems," in Proceedings of the EuroSys Workshop on Recent Advances in Intrusion-Tolerant Systems, Lisbon, March 23, 2007.

[12] The DARPA CORONET program http://www.darpa.mil/sto/strategic/coronet.html

[13] StreamBase page on Complex Event Processing: http://www.streambase.com/complex-event-processing.htm

[14] Kurt Rohloff, Joseph Loyall, Richard Schantz. Quality Measures for Embedded Systems and Their Application to Control and Certification. *ACM SIGBED Review*, Special Issues on Workshop on Innovative Techniques for Certification of Embedded Systems. Volume 3, Number 4, October 2006.

[15] K. Rohloff, S. Lafortune. On the Synthesis of Safe Control Policies in Decentralized Control of Discrete Event Systems. *IEEE Transactions of Automatic Control*. 48:6, pg.1064-1068. June, 2003

[16] A. M. Sanchez and F. J. Montoya. Safe supervisory control uder observability failure. *Journal of Discrete Event Dynamical Systems: Theory and Applications*, 16:493{525, 2006.

[17] Kurt Rohloff, Richard Schantz and Yarom Gabay. "High-Level Dynamic Resource Management for Distributed, Real-Time Embedded Systems." 5th Symposium on Design, Analysis and Simulation of Distributed Systems (DASD), San Diego, CA, 2007.