Risk Assessment of Cyber Access to Physical Infrastructure in Cyber-Physical Systems (extended abstract of CPSS-16 keynote address)

David M. Nicol
Dept. of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign
Urbana, IL 61801
dmnicol@illinois.edu

Keywords

Risk assessment, computer networks, cyber-security, cyber-physical systems

1. OVERVIEW

Networks of computational devices are used increasingly to construct cyber-physical systems (CPS) that monitor and control significant physical infrastructures such as the electric grid, water supply systems, gas pipelines, maritime systems, etc. Reliance on these devices increases the means by which failure and/or malfeasance can adversely impact the integrity and operations of the infrastructure. A pressing question then asks how one can assess the risk to the infrastructure—or to the services it provides—through compromise of the cyber component of a CPS. The problem of answering this question is rife with issues. These include (but are not limited to)

- The traditional notion of risk is “probability of event” times “cost of event”. How do we estimate “probability” of a cyber-event?
- Cyber-physical systems are enormously complicated, and any risk model will have to abstract away some details. How can we do this with confidence?
- How can we combine models of the cyber system and infrastructure to determine the impact that different cyber-intrusions may have on the infrastructure?
- How can we identify the means of strengthening the cyber component of a CPS that has the greatest impact on risk?

This keynote address touches on several of these questions.

With respect to quantifying probability, the issue was once one of having enough events on record to build statistically significant hazard rate models. Unfortunately the number

of events is on the rise, but the problem now is one of data sharing, a side-effect of which might be development of such models. There are disincentives against infrastructure owners reporting cyber-incidents, including concerns for public perception and also liability, the new US Cyber-Information Sharing Act notwithstanding. “Probability of event” has two components however, one involving frequency of events, and the other having to do with successful penetration of a system *during* an attack. The work presented will focus on the former, in particular developing an “access function” that describes the ramification of access protection mechanisms on an attacker’s ability to compromise components of the network in order to reach and control devices in the physical infrastructure.

The question about model abstraction is technically and socially very real. Experience with infrastructure owners has shown us that information about risk upon which decisions can be made *must* be presented at a level that is accessible to the decision-makers. By needs this is significantly simpler than the level of detail required for an engineering analysis. Not only must this gap be bridged, but we must deal with the oft-encountered challenge in modeling of “making the model as simple as needed, but no simpler”; our risk model has to capture features to which risk is sensitive, and abstract away the others. Even if we have a handle on dealing with model resolution, the challenge remains (particularly for more abstract models) that knowledge of an infrastructure’s vulnerabilities may be held by its owner, but not shared with whoever is constructing the risk model. This is a social and legal problem, not a technical one.

The point is that a necessary but possibly insufficient piece of the puzzle is addressing the question of combining a model of a cyber system that controls an infrastructure with a model of the infrastructure itself, for the purposes of assessing the impact that cyber-disturbances may have. With a solution to that piece we can consider addressing the issue of strengthening the cyber-component in ways that have the greatest impact on resiliency of the infrastructure. The application of this solution is in many ways a motivator for the problem we do address.

2. APPROACH

The philosophy of the approach we’re developing takes to heart the need to provide tools for very large complex systems, but tools that are usable. We need to construct a system model in a reasonable amount of time, and at a

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CPSS’16, May 30-June 03 2016, Xi’an, China

© 2016 ACM. ISBN 978-1-4503-4288-9/16/05...\$15.00

DOI: <http://dx.doi.org/10.1145/2899015.2899025>

reasonable cost. This means the model must include details that are either extracted automatically, or are modeled at a fairly high level, perhaps assisted by automation. The measures computed by the model must be understandable to non-technical decision-makers.

We start with the assumption of some model of a computer network through which the infrastructure is accessed. That model includes physical and logical connections between computing entities (i.e., representation of networks, hosts, routers, switches, firewalls), and can include information about the software running on those devices. It is possible to use commercial (and open source) tools to discover elemental connections between hosts (connections which may pass through firewalls), and consider intrusions that use those connections as components of a stepping stone attack. Our model will account for the feasibility of a host being used as a stepping stone to essentially “switch” between one allowed inbound connection to another allowed outbound connection, and so be able to identify in principle all stepping stone pathways from attacker entry points to the computer-controlled actuators in the physical infrastructure. The risk analysis of interest analyzes those pathways and the impact an attacker may have on the infrastructure exercising them.

We have considered a number of different ways of quantifying the risk presented by the cyber-infrastructure to the physical infrastructure. We first considered modeling attacker behavior on the network as a Markov chain, and sought attack paths (sequences of compromised devices) that have the highest probability. The problem with this approach is that real networks give rise to Markov chains with enormously large state-spaces; the approach was not tractable for networks of any interesting size. We next considered creating a graph whose nodes are vulnerabilities on hosts, with weights derived from the vulnerability scores [3], and a directed edge defined between two nodes if the host on which the source vulnerability resides can reach the host on which the destination vulnerability resides through a port that enables an attacker on the source to exploit the vulnerability on the destination host [2]. We found problems with the basic formulation of seeking least cost paths, and on reflection, using individual paths to quantify overall exposure risk leaves room for improvement. We can efficiently enumerate the number of paths an attacker has available from network ingress points to physical infrastructure devices, but this too has problems. Are protection mechanisms that allow for 500,000,000 paths twice as secure as ones that allow for 1,000,000,000 paths? I would argue not. The approach developed in this talk focuses on *whether* an attacker can exploit a host, and if so, the difficulty of doing so, measured in terms of time. The quantitative description of an attacker’s access to a device (or to a set of devices) is a curve, flat-lined at 0 if there is no access, otherwise giving an monotone non-decreasing access-metric between 0 and 1 as a function of time. The access-metric captures both the difficulty of compromising hosts, and the breadth of ways an attacker has access to the physical system.

On the infrastructure side we require three things. First, we need to describe what a remotely connected user can cause an actuator to do. For example, a relay can open or close a circuit breaker and so take a transmission line out of service, or place it into service. The possible actions should include a temporal dimension, e.g., an attacker might be able

to open and close a breaker continuously in fast succession. Second, we need to identify sets of devices to consider as targets of a coordinated attack. There is potential here for combinatorial explosion; given N actuators there are $2^N - 1$ unique non-empty sets of devices. Decisions on which sets of devices are most meaningful are domain and system dependent. For example, if we wish to model an attack where a particular firmware version was compromised, we’d group together all the devices onto which that firmware was (or could be) installed. A third component is a means by which the state of the physical system can be assessed, as a function of all its boundary conditions, including the results of actions taken by the cyber-reachable devices. For example, in the electric power grid, the “performance index” is a well-accepted metric of the degree to which transmission lines are overloaded. A cyber-based attack that compromises one or more relays can cause them to open lines, inducing current overloads. The performance index is a measure of the impact on the physical system of the attack made through the cyber-system. We have built a toolset, Cypsa [1] for this very context, using that very metric.

Risk assessment needs now to combine the access-metric(s) of the cyber system with the assessment metric applied to the physical system, considering the impact of an attack. There are a variety of ways we might do this, but in some sense the particular method is not nearly as important as its behavior and use. We’d like the metric to identify the impact on risk of changing selected configurations, and we’d like to use the metric to identify the most cost-effective actions we might apply on the cyber-side to reduce the risk to the infrastructure.

3. SUMMARY

This talk¹ outlines an approach we are developing for the risk assessment of different types of physical infrastructures whose operations are exposed to cyber technology, and are threatened by the potential of that technology being compromised for the purposes of attacking the physical system. We are focused on making the approach usable, on making the metrics associated with the cyber-system understandable, and on making the overall approach useful for identifying the decision actions that will have the greatest (or most cost-effective) impact on reducing risk.

4. REFERENCES

- [1] K. Davis et al. Cyber-physical security assessment (cypsa) for electric power systems. *IEEE-HKN: THE BRIDGE*, 2016. to appear.
- [2] D. M. Nicol and V. Mallapura. Modeling and analysis of stepping stone attacks. In *Proceedings of the 2014 Winter Simulation Conference, WSC ’14*, pages 3036–3047, Piscataway, NJ, USA, 2014. IEEE Press.
- [3] M. Schiffman, G. Eschelbeck, D. Ahmad, A. Wright, and S. Romanosky. Cvss: A common vulnerability scoring system. *National Infrastructure Advisory Council (NIAC)*, 2004.

¹This work was supported in part by Department of Energy contracts DE-OE0000780 and DE-AR0000342, and by DHS contract 2015-ST-061-CIRC01. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.