Mediating Cyber and Physical Threat Propagation in Secure Smart Grid Architectures

Clifford Neuman
Information Sciences Institute
University of Southern California
Marina del Rey, California, USA

Kymie Tan
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California USA

Abstract— The power grid is a federated system. Regions of the system are controlled by different organizations and security of the grid is imposed from above through regulation of the security techniques used by the federants. This approach will be less effective as we move to a smart grid, where control of some elements of the grid rests in the customer's home through technologies that enable remote access to appliances. These regions of the smart grid are less trusted, yet they interact in various ways with other parts of the grid. This paper demonstrates threat propagation in the smart grid from such regions, and discusses architectural approaches to mediating the impact of such flows.

I. INTRODUCTION

The power grid is a federated system. Although, regions of the system are controlled by different organizations, security of the grid as a whole has been (not very effectively) prescribed from above through regulations imposed on the security techniques used by each of the federants. This approach will be less effective as we move to a smart grid. Control of the smart grid extends into the home. Thus, control of some regions of the smart grid lie in the customer's home, through technologies that enable remote access to appliances. These regions are not as well trusted, and a secure smart grid architecture must address the impact of untrusted regions.

In this paper we discuss approaches for decomposing the smart grid into protection domains along organizational and cyber-physical boundaries. We show how threats in one domain have impact on other domains, and we suggest approaches to modeling and mitigating the impact of such threats.

Much of the existing work in security for cyber-physical systems focuses on providing security within the cyber-domain, and highlighting the impact on the physical domain of such breaches [1,2,3,4,5]. While deployment of secure smart grid architectures requires the deployment of security technologies within each cyber domain, these technologies are the subject of extensive study elsewhere, and are not unique to cyber-physical systems – they are just more important because of the potential consequences of a breach. If we want to understand what is different about security in a cyber-physical system we need to look not only at attacks that originate in a given cyber domain, but also those that originate in, or that are propagated through multiple domains, both cyber and physical [6]. Once those interactions are understood, both cyber and physical means must be used to mitigate the impact of cross-domain interactions.

II. TAXONOMY OF CYBER AND PYSICAL THREATS

To understand the new classes of threats in a cyber-physical system such as the smart grid, it is useful to characterize the interactions based on the domain that is the origin of the threat, and the domain where the impact is felt. We need to look also at intermediate domains that propagate the threat. In this section we present a couple of examples, and show impacts that are commonly expected, and some that might have received less attention.

While there will be far too many protection domains in a smart grid to effectively analyze, we will simplify the problem by grouping similar protection domains together. Thus we will utilize the following groups of domains in the rest of this discussion: 1) an untrusted domain which includes customer owned devices, and the open internet, 2) a utility distribution domain that includes devices related to the power distribution network, including AMI components, and 3) a utility business domain, which includes billing systems. Each of these three groups may have a cyber domain, and a physical domain associated with it, and the corresponding cyber and physical domains will be considered separately, although they are

This material is based upon work supported by the United States Department of Energy under Award Number DE-OE000012, the Los Angeles Department of Water and Power, and by the Department of Homeland Security and the Department of the Navy under Contract No. BAA-10-C-2018.

closely interrelated. For a larger system, one would have transmission domains, and operator domains (in fact, each separate utility or operator should be treated as having separate domains – both cyber and physical), but we are assuming a single utility smart grid deployment to simplify our explanation in this paper.

A. *Cyber – Cyber Threats (CC)*

Cyber threats include those to confidentiality or integrity of data in the system, including connecting to a device on a home area network and retrieving usage data, or modifying billing information, etc. They also include certain threats to the availability of services (including communication) within the computing components of the system. Cyber-Cyber threats can reside within a single protection domain, such as an intruder breaking into a home network controller, or they can cross multiple protection domains, such as exploiting firewall vulnerabilities to connect from an attackers site (on the open internet) to a utilities billing system to modify data or steal credit card numbers.

An important characteristic of cyber threats is that they are scalable, i.e. they are easily automated and replicated, and one should expect them to propagate freely across untrusted domains. For this reason, in the characterizations of attack chains that follow, we will reduce cases of repeated cyber domains, to a single cyber domain (e.g. CCP or CCCP will be treated the same as CP). We will see later how the ease with which these threats are replicated provides the means by which the threat propagates to other domains.

Defending against cyber-cyber threats is accomplished using techniques from the broader computer security community, often with added constraints resulting from the critical nature of the application domain [7]. To be effective in a smart grid environment these approaches must include not only communications security (encryption and integrity protection), data integrity (digital signatures), identity and policy management, network admission control (e.g. firewalls), but they must provide stronger mechanisms to enforce isolation at a very coarse level of granularity. From the perspective of utilities, the biggest problem in defending against cyber-cyber threats is that many of the impacted systems will be completely outside of their control – they will have neither the means to monitor what occurs on the customer's home network, nor the means to install security solutions. It is for this reason that smart grid architectures must consider customer owned devices and the home area network as untrusted.

B. *Cyber – Physical Threats (CP)*

Cyber-physical threats are those that originate within a cyber-domain, but which have an impact on the physical characteristics of a system. These are the threats that most people think of when they consider security in cyber – physical systems. In power systems, these are actions that are propagated through a programmable logic controller (PLC) or its equivalent (really almost any digital remote control device). For completeness, we need to also consider the "human PLC" and social engineering – whereby communications viewed by a user of the system through whatever channel is available, causes them to take a physical action. Given the percentage of energy consumption that goes to information technology today, one could also consider increased computation load on a sufficiently large number of computers as creating a transition from a cyber to a physical domain.

An important characteristic in considering the impact of CP threats is the identification of the physical domain affected. In the power community, more attention has been directed toward CP threats impacting the utility domain (whether distribution, transmission, or generation) because these are the obvious targets for an attack [8,9]. Less attention has been paid to CP threats on the customer's physical domain, and as we will see in section E, that is problematic.

Protection against cyber-physical threats is implemented at the PLC (or its equivalent) on both the cyber and physical sides. On the cyber side, the PLC must implement controls on access and strongly authenticate commands that it will process. But this is not enough – identities will be stolen, users will be impersonated. The PLC must include logic for governors that ensure that controlled devices will operate in a safe manner, even when sent commands to execute unsafe operations.

C. *Physical - Cyber Threats (PC)*

Physical – cyber threats (PC) are those that originate with physical actions and that impact the functioning of networking or information technology components of a system. The most obvious examples are the physical destruction of computers or network devices, and include cutting of communication lines, or shooting at pole-top repeaters. PC threats also include actions taken to cut power to computing or communications devices. If the power is turned off by remote control, it becomes an example of a CPC threat – which will be discussed in section F.

Physical measures are needed to prevent most physical attacks from impacting cyber-infrastructure. Devices can be hardened and battery backup provided. The impact of a physical attack can be mitigated through redundancy or over provisioning and reconfigurability. Multiple sources of power can be provided to critical devices. Extra servers can be deployed so that failed components can be easily replaced.

D. *Physical – Physical Threats (PP)*

Physical–physical threats (PP) in the power grid concern the way that physical actions (or conditions) in one part of the power grid impact other parts of the system. Most of the major blackouts in history resulted from such PP interactions (even though we don't know that any of them were intentional). A physical-physical threat exists when an adversary can cause a first physical action (or a set of physical actions in concert) in order to cause failure in *other* parts of a system. An example local to a customer's physical domain would be if an attacker turns on too many devices on a circuit with intent of tripping a circuit breaker. A cross-domain PP attack occurs when an attacker causes multiple neighbors to do the same with the intent of overloading part of the distribution network.

Defenses to physical–physical attacks reside in the physical domain. In power grids, these are usually load limiting devices, such as circuit breakers, that prevent the problem from

traveling further upstream, or that cut off the treat at its source (the breaker prevents a fire or electrocution at the source of the fault). It is important to keep in mind, that some protection devices may enable propagation of an attack. While a circuit breaker may prevent an attack from propagating upward in a domain hierarchy (discussed later), it explicitly propagates the effect downward to sibling domains fed from the same circuit. Thus, in the examples cited, an adversary could intentionally stimulate the countermeasure – which in turn causes the intended consequence. Such mechanisms might also generate their own, different, physical effect which will propagate upward; tripping the breaker stops an increase in demand from propagating upward, but the sudden reduction in load can propagate its own physical effects upward through the system.

E. Cyber – Physical – Physical Threats (CPP)

CPP threats cross multiple groups of domains and may be thought of as a PP threat, where the initial physical action is caused remotely. CPP threats are of particular concern in systems like the smart grid where the initial CP attack occurs on equipment owned by the customer and not likely controlled by the utility. The ease with which cyber-cyber attacks are replicated on untrusted networks enables synchronized CP attacks on many customers systems, which may be the vector through which the attack jumps to the distribution network.

To explain this more concretely with a smart grid vulnerability, consider the current generation of electric vehicles. One of the marketing "features" of the Nissan Leaf is the ability to control charging (and certain other functions) from your smartphone. As with all general purpose network connected computers, such smartphones are vulnerable to viruses, easily propagated by the lack of care shown by many users regarding the apps they install. An adversary could write a virus that would detect if a smartphone has the Leaf controlling app installed, and if so, at a designated time if could cause a large number of electric vehicles to simultaneously begin their charge cycle. This would in turn create a significant load spike on the utility's network – and even though the cyber-attack reaches no farther than the users home, the physical impact of the attack would be felt throughout the power grid as the grid responds (or is unable to respond) to the sudden load spike.

Until now, we have not considered the network of customer cell phones to be part of the control network for the utility's distribution system – but here we have shown how the characteristics of cyber-cyber threats, when coupled with CP and PP threats creates exactly that condition.

Mitigation of this kind of CPP threat is difficult at the CP boundary because the boundary exists on the customer owned and controlled equipment. Thus, mitigation of many CPP threats to the smart grid as whole needs to be focused on the PP boundary, with mechanisms put in place to prevent aggregated effects from multiple untrusted customer domains from causing failures in the distribution domains. Similar measures would be needed at higher levels in systems that include transmission and generation as well.

An area that has received little attention is methods for detecting such attacks. Since the cyber-attack is confined to the customer networks, it is not practical to expect the utilities data collection system for security and incident event management (e.g. intrusion detection) to have direct visibility of the attack. Instead, data must be collected at the boundaries of the utilities network that will allow them to draw inferences about such attacks. Because that boundary is a physical-physical boundary for attack propagation, the ability to detect such an attack requires correlation of data regarding changes in the load placed on the system at multiple meter locations. Such monitoring of physical properties becomes in part related to the topic of "situational awareness" in the grid, but also relates to specification based intrusion detection [10,11].

F. Cyber – Physical – Cyber Threats (CPC)

CPC threats are very similar to CPP threats except, in the case of the smart grid, the effect of the attack is felt primarily in cyber capabilities (communication and computation) of the utility's network. Because the attack is channeled through a physical component of a system, which lacks the detailed control the attacker would have through a purely cyber attack, the effect on the utilities network is similar to that for a CPP attack, i.e. denial of service.

An example would be for an adversary to remotely control a device on a customer's premises in such a way that the customer's meter transmits a status update on the AMI network. If the status of such devices across a large number of customers were to be toggled, the status updates could consume the available bandwidth on the AMI network, preventing other communications (e.g. billing or pricing information) from getting through. It is interesting to note that the method discussed to aid in detecting CPP attacks (transmitting event data for sudden load changes) might end up as a vector used to mount a CPC attack.

Defenses against CPC attacks are similar in philosophy to defenses against CPP attacks, in particular, mechanisms are needed to prevent aggregated effects from multiple untrusted customer domains from causing failures in the distribution domains. The particular defense in this case could be logical, and might include techniques for rate-limiting communication from certain parts of the distribution network.

G. Pysical – Cyber - Physical (PCP)

Physical – cyber – physical threats post an interesting class of failure. An example of such an attack leverages the response of a system to a collection of stimuli, such that if the adversary can cause the right set of stimuli, the system reacts in a way that is detrimental to its sustained operation. Consider a fire suppression system that is activated when the temperature at 3 sensors exceeds a certain threshold. The fire suppression system has a cyber component that receives readings, and filters out readings from up to two potentially failed sensors. With knowledge of such a system an adversary and two partners could each strike a match and hold it up to one sensor each. A similar attack could be mounted by manipulating the readings from the three sensors electronically; in that case it would be cyber-cyber–physical instead of physical–cyber–physical.

Defenses against physical–cyber–physical attacks can sometimes take place in the physical domain, protecting sensors and infrastructure where the attack can be initiated. Prevention of such attacks can also be affected with careful adversarial (red-reaming) analysis of the response logic employed within the grid itself.

H. *Transitive Threats and Grouping Domains*

To effectively model large systems, we need the ability to group domains for analysis. For cyber domains, similar domains must be grouped together when they have similar characteristics. For example, customer domains are similar and their characteristics modeled statistically, even though the management (or lack thereof) for those domains rests with different users. We do not yet know how to group these domains, but are looking at results from studies of malicious code propagation (viruses and worms) for potential avenues of investigation.

Similarly, physical domains may be grouped, but the grouping of physical domains should be hierarchical based on the physical topology of the system. This is important because many of the mitigation techniques to prevent propagation of physical threats are unidirectional, e.g. the breaker that prevents sudden load increases from propagating upward, cuts off power to physical domains lower in the hierarchy. Therefore we will need to create domains and sub-domains when modeling the physical part of the system.

The enumeration of cross-domain threats discussed in this paper is by no means complete. Longer chains can be created by appending the building blocks just described. Where two cyber domains come together in such chains, the analysis can follow that which would be done if the two Cyber domains were collapsed into one, but the mitigation strategies may be different because ownership and control of the domains differs.

III. ARCHITECTURAL GUIDANCE

The title of this paper mentions a secure smart grid architecture, rather than a security architecture for the smart grid. The biggest gains in security come not from the application of existing security technologies to a smart grid architecture that was designed without considering security, but instead from improvements to the basic information and control flow architecture created for the system itself.

This structure needs to be influenced by an understanding of the federated nature of the system. As these flows are designed, the architect needs to consider the fact that certain parts of the system will be untrusted. The smart grid requires the inclusion of devices that cannot be effectively secured. It is our belief that the greatest impact from the smart grid will be felt if it evolves in a manner similar to the internet, where there is core infrastructure, managed professionally, that provides services through a common open interface to "applications" that will be developed in a completive marketplace, and in a way that affords little control on those applications themselves (beyond the necessity of their using the open interfaces to obtain the core services). As we have seen on the internet, these "applications" will compete on the basis of functionality and cost, but probably not on security, and certainly not on their resistance to facilitating hypothetical attacks on the core infrastructure (security becomes "somebody else's problem").

Attempts to limit functionality of these "applications" by the utilities would be ineffective. The application market, or even enthusiasts would find ways to work around such limitations to do the neat things they want. Consider the earlier example of the Nissan Leaf smartphone interface; users are unlikely to accept a prohibition of remote management of their vehicles to prevent such a hypothetical attack.

With an understanding that regions of the smart grid will not be trustworthy, the rest of the architecture needs to be designed with strong boundaries between domains. In the long term, computers and networks need to be developed that provide stronger isolation between functional domains sharing common infrastructure. In the short term, such isolation across domains needs to be provided at the boundaries, with careful consideration given to each information or control flow that crosses such boundaries.

System architects will need to consider not only the information and control flows in the cyber domain, but those in the physical domain and those crossing domains – many of which will be dictated by physics. The resulting analysis can then be used to identify potential paths through the system that are unnecessary for the system's mission objectives, and measures can be implemented to eliminate or mitigate the impact of those flows.

IV. APPLICATION IN A SMART GRID DEMONSTRATION

We conducted our threat analysis in the context of the Los Angeles Department of Water and Power Smart Grid Regional Demonstration Project, which has partnered with NASA's Jet Propulsion Laboratory, the University of California at Los Angeles, and the University of Southern California to conduct smart grid research in the areas of customer behavior, cyber security, demand response, and electric vehicles.

Figure 1 shows several of the domains that need to be analyzed for such a system (though not all of the components are within the scope of the project). The shaded domains represent physical system components (which may also have a corresponding cyber-domain). Unshaded domains correspond to networking and information components in the system.

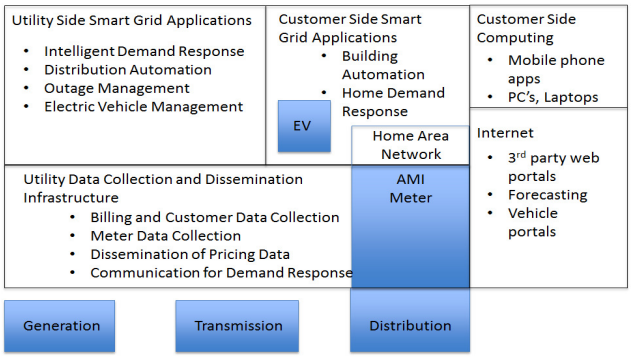

Figure 1. Domains in a Representative Smart Grid

In the figure, the boxes labeled AMI meter, Home area network, EV, and Customer Side Smart Grid Applications represent multiple use-specific domains, grouped according to the discussion in section H. The boxes labeled AMI meter (again) and the box labeled distribution corresponds to multiple domains, organized hierarchically this time, to model the fault tree of the physical system. Generation and transmission would be similarly modeled, though they are out of scope in the particular project. The customer side computing box represents a purely cyber domain over which the utility has no control and limited visibility. C-C threat propagation flows to the customer side smart grid application, as well as to the EV from the Internet domain, through the customer side computing domain. While not shown in this figure for topological reasons, threats can also propagate from the Internet domain to the Utility side domains through customer web portals, or poorly configured firewalls and VPNs used by the utility for other functions.

V. APPLICATION IN SYSTEM MODELING

An analysis of the cyber-physical security of a smart grid architecture must focus on the impact of faults and interactions that cross domains [12] rather than the localized response that might be seen in traditional penetration testing. This requires a capability to model large scale response to cyber-attack, as well as to perform modeling or simulation of the physical components of a system. To date, we have not seen both capabilities present in a single modeling environment. To improve the modeling capabilities of such systems, we recently formed a consortium called DEFT with initial members from the DETER lab at the University of Southern California, the US Department of Energy's Pacific Northwest National Laboratory, and the University of Illinois Urbana Champaign's research group on Trustworthy Cyber Infrastructure for the Power Grid (UIUC TCIPG).

When modeling the threats in this paper, it can be useful to decompose the problem into regions which are then modeled separately to learn local response. Some of those regions are better modeled in the cyber domain, accomplished on systems such as USC's DETER Testbed [13], a US Department of Homeland Security funded lab for evaluating the effectiveness of cyber-attacks and candidate defensive technologies. Through the DEFT consortium, the DETER Testbed will federate with other labs and testbeds that provide physical simulation and emulation tools for modeling the systemic response of the grid. Such experiments will span multiple sites and will enable the use of specialized resources to participate in large scale experiments. This capability may eventually provide the ability to expose only systemic response of subcomponents without fully disclosing sensitive details about parts of a system whose response will be studied.

The structure of the domains and cross domain threats presented in this paper can be useful when decomposing the problem into regions, and grouping those regions based on common characteristics or position in the physical system hierarchy. One can then draw on capabilities specific to each of the DEFT partners, with well-defined interfaces and boundaries at the domain transitions discussed earlier.

Domains	Characteristics and Impact	Prevention and Mitigation
C-C	Scalable, Automated, Easily Propagated Contributes to Transitive Threats	Computer Security Best Practices
C-P	Propagated by PLC, Social Engineering, Hard to defend in Customer Domain. Contributes to Transitive Threats	Cyber-defense at PLC Governers limit impact on physical side
P-C	Physical Destruction or Cutting Power from Cyber Components	Physical Security, Redundancy, Reconfigurability
P-P	Impact of physics of events in one region, affecting other regions. Contributes to Transitive Threats	Examples are Load Limiting devices Mitigation sometimes one-way
C-P-P	Physical action initiated in Cyber domain. Effect propagates to higher physical domain.	Difficult for utility to prevent or detect because first C-P bondary on customer premises
C-P-C	Less control by adversary. Usually results in denial of service to utility network.	Rate limitation of reporting events is one example of an effective mitigation stratey.
P-C-P	Physical stimuli exploiting programmed system response.	Physical protection of sensors Red-teaming response senarios

TABLE I. SUMMARY OF CROSS DOMAIN CHARACTERISTICS, IMPACT AND MITIGATION STRATEGIES

VI. SUMMARY

By necessity, the smart grid will include components that cannot be trusted. Smart grid architectures must take this into account, creating multiple protection domains, and managing the flow of information and commands across those boundaries. When defining these domains, each region of the system should be modeled with separate cyber and physical domains. Table I summarizes the cross-domain threat classes discussed in this paper. These classes can be used as an aid in understanding flows that might not be explicitly stated in the objectives of the system. Mitigation measures must then be deployed to limit the impact of such flows on the system's security and resiliency objectives.

ACKNOWLEDGMENTS

The authors thank David Alexander, Anas Almajali, Terry Benzel, Bob Braden, Natasha Brand, Marco Elizarraras, Ted Faber, Greg Finn, Jeff Gooding, Wiley Gustafson, Greg Horvath, Frank Kuykendal, Matt Lampe, Thom McVittie, Gaurav Mittal, Gracie Neuman, Don Paul, Eric Rice, Gordon Roesler, Mike Ryan, Tatyana Ryutov, Steve Schwab, Joe Touch, Zachary Tudor, Arun Viswanathan, John Wroclawski and Hain Zhou for discussions that helped us form the ideas presented in this paper and for feedback on earlier drafts.

REFERENCES

[1] F. Cleveland, *"Enhancing the reliability and security of the information infrastructure used to manage the power system,"* EEE Power Engineering Society General Meeting, 2007.

[2] A.R. Metke, R.L. Ekl, *"Smart Grid Security Technology,"* 2010 Innovative Smart Grid Technologies (ISGT), 2010.

[3] K. Moslehi, R. Kumar, *"Smart Grid - A Reliability Perspective,"* 2010 Innovative Smart Grid Technologies (ISGT), 2010.

[4] NIST 6728, "*Guidelines for Smart Grid Cyber Security:* Vol. 1, Smart Grid Cyber Security Strategy, Architecture, and High-Level Requirements," August 2010.

[5] D. Wei, L. Yan, M. Jafari, P. Skare, K. Rohde, *"An Integrated Security System of Protecting Smart Grid against Cyber Attacks,"* 2010 Innovative Smart Grid Technologies (ISGT), 2010.

[6] M. Amin (EPRI), Security Challenges for the electricity infrastructure. IEEE Computer (Security and Privacy Supplement) Volume 24, Number 4, pp 8-10. April 2002.

[7] H. Khurana, R. Bobba, T. Yardley, P. Agarwal, E. Heine, *"Design Principles for Power Grid Cyber-Infrastructure Authentication Protocols,"* 43rd Hawaii International Conference on System Sciences (HICSS-43), 2010.

[8] M.G. Lauby, *"Reliability Considerations for Application of Smart Grid Technologies,"* EEE Power & Energy Society General Meeting, p 4 pp., 2010.

[9] J. Stamp, A. McIntyre, B. Ricardson, *"Reliability impacts from cyber attack on electric power systems,"* IEEE/PES Power Systems Conference and Exposition, PSCE 2009, 2009.

[10] I. Balepin, S. Maltsec, J. Rowe, and K. Levitt, "*Using Specification Based Intrusion Detection for Automated Response*." Proceedings of the Recent Advance in Intrusion Detection (RAID), pp. 136-154. 2003.

[11] R. Berthier, W. Sanders, and H. Khurana. "*Intrusion Detection for Advanced Metering Infrastrctures: Requirements and Architectural Directions*" in 1st IEEE International Conference on Smart Grid Communication, pp 350-355, 2010.

[12] A. Faza, S. Sedigh, B. McMillin, *"Integrated Cyber-physical Fault Injection for Reliability Analysis of the Smart Grid,"* Computer Safety, Reliability, and Security. Proceedings 29th International Conference, SAFECOMP 2010, pp 277-90, 2010.

[13] T. Benzel, B. Braden, D. Kim, C. Neuman, A. Joseph, K. Sklower, R. Ostrenga, and S. Schwab, *"Experience with DETER: A Testbed for Security Research",* The 2nd IEEE Conference on Testbeds and Research Infrastructure for the Development of Networks and Communities. March 2006, Barcelona.