Cyber/Physical Security Vulnerability Assessment Integration

Doug MacDonald, Samuel L Clements, Scott W Patrick,
Casey Perkins, George Muller, Mary J Lancaster, Will Hutton
Pacific Northwest National Laboratory
Richland, WA, USA

Abstract— **Securing high value and critical assets is one of the biggest challenges facing this nation and others around the world. In modern integrated systems, there are four potential modes of attack available to an adversary:**

- **physical only attack,**
- **cyber only attack,**
- **physical-enabled cyber attack,**
- **cyber-enabled physical attack.**

Blended attacks involve an adversary working in one domain to reduce system effectiveness in another domain. This enables the attacker to penetrate further into the overall layered defenses.

Existing vulnerability assessment (VA) processes and software tools which predict facility vulnerabilities typically evaluate the physical and cyber domains separately. Vulnerabilities which result from the integration of cyber-physical control systems are not well characterized and are often overlooked by existing assessment approaches.

In this paper, we modified modification of the timely detection methodology, used for decades in physical security VAs, to include cyber components. The Physical and Cyber Risk Analysis Tool (PACRAT) prototype illustrates an integrated vulnerability assessment that includes cyber-physical interdependencies. Information about facility layout, network topology, and emplaced safeguards is used to evaluate how well suited a facility is to detect, delay, and respond to attacks, to identify the pathways most vulnerable to attack, and to evaluate how often safeguards are compromised for a given threat or adversary type. We have tested the PACRAT prototype on critical infrastructure facilities and the results are promising. Future work includes extending the model to prescribe the recommended security improvements via an automated cost-benefit analysis.

Index Terms – **Risk analysis, Security, Modeling, Simulation, Power Industry**

I. Introduction

Both physical and cyber security domains offer solutions for the discovery of vulnerabilities through the use of various assessment methodologies and software tools. Each domain uses assessment tools that provide the ability to identify and categorize vulnerabilities and quantify risk within their own areas of expertise, but neither fully represents the true potential security risk to a facility, nor provides a comprehensive assessment of the overall security posture. We define three assessment levels that blend physical and cyber security vulnerability assessments where each level builds on the previous step.

- Level 1: Subject matter expert (SME) walk-down
- Level 2: High-level quantitative modeling
- Level 3: Detailed modeling

Level 1 assessments consist of using subject matter experts in both physical and cyber domains to perform a security review of an organization. The SMEs review security documentation and perform an on-site assessment. We found this to be most valuable to organizations that have a new or immature security program. It is relatively quick and the experts will be able to identify the major gaps in the security program.

Level 2 assessments help identify vulnerabilities that are not readily apparent without some analysis. The modeling begins to show the interplay between the cyber and physical security domains. We found this to be useful for organizations with a functioning security program in place and which want to ensure that their security investments are being allocated correctly.

Level 3 assessments model a system to a very granular level. This takes significant resources to ensure the model is accurate and properly represents the system. The candidates for performing this type of assessment would be organizations or facilities that protect information or materials that require a very high assurance of their safety and security. The value in modeling a system to this level allows for near real-time or continuous evaluation.

This paper focuses on the Level 2 methodology using a coarse approach to quantitative modeling and analysis of the integrated vulnerability assessment. Our effort produced a prototype software tool that provides insight into security systems and allows security managers to identify and quantify previously unaccounted for risk.

II. Related Work

The current state of the art for cyber attack modeling attempts to provide insight into how an attacker can navigate across layers of network infrastructure, exploiting vulnerabilities to reach an objective, such as disruption of a portion of business activities (e.g. distributed denial of service (DDoS) attack) or retrieval of valuable information (e.g. data theft).

Kuhl et al. (2007) used discrete event simulation to generate a network topology intrusion detection system based on a specified cyber-attack scenario.

Ingols et al. (2009) use attack graphs to model network architecture countermeasures and attacks on the network. This approach uses a series of rules that define the interactions between the attacker and the network. This approach offers the ability to model details of the network topology, including firewall rules, operating system configuration, and software vulnerabilities.

Sommestad et al. (2009) and Mo et al. (2009) use Bayesian approaches that extend attack and defense graphs to model risk in a cyber system. These approaches aim to establish a methodology for targeting improvements in the cyber security systems that are in place across infrastructure and business enterprises.

Jordan et al. (1998) developed a discrete event simulation of the evaluation of physical protection systems. This model extends previous approaches and provided an interface with an external model.

These approaches allow for an improved understanding of the vulnerabilities of the individual cyber and physical systems, but do not include the interactions and interdependencies that exist as a result of today's control systems that provide electronic management of physical systems. The goal of this research is to build an integrated vulnerability assessment of both cyber and physical domains to understand how attackers can exploit vulnerabilities in both domains to penetrate enterprise security systems more effectively than previously modeled.

III. Model Design

Our research process began by developing SMEs versed in both cyber and physical security methodologies so that each is aware of the vocabulary, terminology, risks and vulnerabilities of both the cyber and physical domains. These same subject matter experts were then used to perform a Level 1 joint vulnerability assessment of an organization to better understand each other's processes.

Following the initial assessment the team began creating a proof-of-concept software tool. The Adversarial Time-Line Analysis System (ATLAS) software suite was used as the basis for the formulation of this new security analysis tool. ATLAS was modified to reflect the new methodology and infused with a surrogate database to demonstrate the utility of including cyber elements within a principally physical assessment tool.

The limitations of the ATLAS software package as the basis for the new software tool is the linear fashion of the adversarial approach. In addition, ATLAS was designed to model physical security systems, though with significant modification it could model a cyber system. However, it did not have the capability to model the interaction between domains.

The Physical and Cyber Risk Analysis Tool (PACRAT) prototype developed by this project has the ability to explore the interactions between both the physical and the cyber domains. The tool allows the security professional to perform an overall vulnerability analysis on the entire system taking into account the previously unidentified and unaccounted for areas of physical/cyber inter-dependencies that are essential for both Level 2 and 3 type assessments.

In our initial work, the facility being assessed was a critical infrastructure electrical substation and the target was a transfer switch that needed to be opened, either physically or via computer, to disrupt the flow of electricity. No exfiltration by the adversary was required for the adversary to achieve their objective. In subsequent models, the target was a physical object which needed to be removed from a facility. In this case, the adversary had to enter the facility, acquire the object, and exit with the object to be successful. The adversary's goal in the first assessment was sabotage and in the second was theft. Differentiating between the two is important in calculating when the adversary has won as it affects the response force tactics as well as the time to interdict and neutralize the adversary.

The following elements of the system need to be documented to build the model:

- Physical security areas, physical safeguards and protection elements, and the connections to cyber areas
- Cyber security areas, cyber safeguards and protection elements, and the connections to physical areas

- Possible adversary actions in each area
- Target(s)

With this information we then populate the connections between areas and the safeguard sets protecting these areas which are placed the paths/connections between areas. The safeguard documentation includes the type/name of the safeguard, detection probability and delay value and the location of the safeguard (e.g., Locked Office Door; Detection Probability = 40%, Delay = 3 minutes; Location = between hallway and office #3).

Once all this information is in place we can begin to exercise our model.

IV. Modeling approach

The model is constructed as a Monte Carlo discrete event simulation, based on the original timely detection methodology. This methodology starts the "clock" at the point of the first detection and continues until the adversary either completes their task(s) or is neutralized by the response force. The timely detection methodology is to provide detection as early as possible and build in enough delay to allow responders the time needed to interrupt the chain of events before the adversaries can make it to the target and complete the final task.

The construction of the physical and cyber network structures focused on modeling unique paths to achieve the adversary's goal(s) and areas which provide a means to alter safeguard performances, or unique pathways to achieve goals. We removed redundant elements in the model to reduce compute time and eliminate identical combinations which would result in the same outcome. For example, we only model bypassing the access control to one enterprise desktop as they all use the same credentialing processes.

The baseline construction of the model is an adversary navigating a network in which areas are nodes and paths are the arcs between nodes. Residing on paths are safeguards, which are objects or devices that provide a means of detection and/or time delay. The adversary is provisioned with skill levels for cyber and physical attack modes, as well as a set of one or more objectives

Each of the various components of the model is described in more detail below:

Adversary – For modeling purposes we assumed a single intelligent adversary with moderate skill level in disabling cyber and physical safeguards. The adversary operates in one of two attack modes: stealth or speed. In stealth mode the adversary will attempt to defeat safeguards in whatever manner is least likely to be detected. For example, if needing to breach a door the adversary would take the time to try to pick the lock instead of using power tools or explosives. On the contrary, in speed mode the adversary chooses the fastest option available. The adversary is also set to operate in stealth mode until detected and then switch to speed.

Targets – The target is the adversary objective and may be in the physical or cyber domain. The target may be an abstract object, such as data or a computer-controlled switch, or a physical object. The adversary's objective may be to alter, damage, or steal the target.

Areas – In the context of a network problem, areas are network nodes. Areas are identified and modeled on the basis that they either offer the adversary an opportunity to alter the state of a system or provide access to additional areas.. In the physical layer, areas are most often rooms, collections of rooms, or buildings. For the cyber architecture, an area is more commonly an access-oriented definition, such as a network permission level or zone.

Paths – The connection between two areas are referred to as paths in the modeling framework. In network terminology, these are the arcs between nodes. An adjacency matrix or connection matrix is populated, indicated the ability of the adversary to access one area from another. Direction of adversary progress is an important factor in physical and cyber security, so paths are directional, resulting in an adjacency matrix that is not necessarily symmetrical (e.g. outbound network traffic does not have the same controls as inbound, or exiting a building is not subject to the same safeguards as entering). The adjacency matrix also includes the critical linkage between cyber designated areas and physical areas.

Safeguards - Safeguards provide the primary stochastic and time-based elements for the model. The time delay and probability of detection are the basic parameters for modeling each safeguard. Instances of safeguards in the system are modeled independently, allowing for isolated or system-wide changes in safeguard performance. By modeling safeguard performance based on adversary skill levels and allowing dynamic state changes, a great number of possible scenarios can be generated using the same system definitions. Multiple safeguards are often included on any given path and one or more of these safeguards must be defeated to traverse the path to the next area. This grouping of safeguards is defined as a safeguard set. The model then allows the adversary to choose which safeguard(s) from the set he will attempt to defeat.

Actions - To facilitate system wide state changes and non-safeguard related delays, actions can be assigned to area nodes. Upon entry into either cyber or physical areas, the adversary can perform pre-defined tasks. The tasks can be as simple as a time delay, or as complicated as degrading all instances of a particular safeguard type. One example would be the case where an adversary gains entry into an access control server, granting facility wide access to all password controlled locks. The action element provides a dynamic construct for modeling a variety of attack types.

Response – In our model the response is a time parameter. Once the adversary is detected a clock starts and once it reaches the defined response time that run ends. If the adversary has not been able to achieve his goal the system wins if the adversary has achieved his goal then the adversary wins.

Fig. 1 shows a potential attack path, including the safeguard values for detection and delay. The critical response point can be overlaid on this path to indicate the last point where the security team could launch a response and still interdict the adversary prior to the completion of the adversary task.

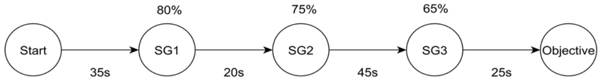

Figure 1. Attack path with notional safeguard detection likelihood and delay values.

Table I provides notional safeguard values for a modeled system. Each safeguard has a total of four values to capture the delay and detection information for the speed and stealth attack types.

Safeguard Model	**Notional Safeguard Values**			
	Delay		*Detection*	
Safeguard	*Speed*	*Stealth*	*Speed*	*Stealth*
Vibration Sensor	0.00	0.00	0.90	0.90
Barrier Fence	0.25	0.25	0.00	0.00
CCTV	0.00	0.00	0.90	0.90
ID Actuated Lock	0.75	0.75	0.20	0.2
Infrared Sensor	0.00	0.00	0.55	0.55
Agent Based NAC	6.00	3.00	0.25	0.25
Two-Factor Authentication	6.00	3.00	0.25	0.25
Host Firewall	6.00	3.00	0.25	0.25

Table I. Notional Safeguard Values

V. ANALYSIS AND RESULTS

In order to provide a meaningful analysis capability, a statistically significant number of runs are evaluated, providing a variety of analytical glimpses into the behavior of the overall system. Insight into area/path performance combinations commonly resulting in the adversary meeting their objective(s) will highlight potential vulnerabilities for each attack/iteration.

Fig. 2 illustrates the arcs or paths that are most often traversed in scenarios where the adversary wins, with red being the most frequently traversed paths, and yellow less frequently. This qualitative view of successful paths provides high-level information at a glance, which allows more refined analysis of detailed model output.

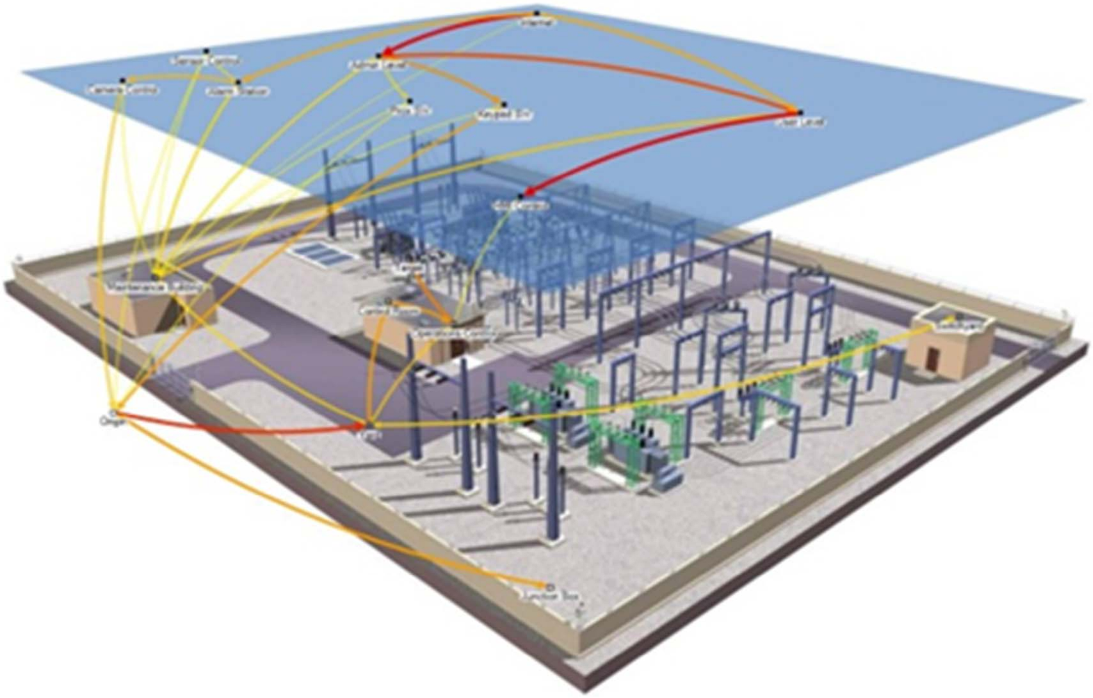

Figure 2. Graphical Representation of Modeled Network

The route of the adversary through the system is not predetermined in our model. By running the model a statistically significant number of times, paths which are successful most often can be identified without constraining adversary routes. This facilitates the discovery of novel vulnerabilities, and reduces the "failure of imagination" on the part of security planners by testing feasible pathways of a large connected network. This provides for the possibility of successful attacks, with low frequency of occurrence, to be included in the analysis and uncover how these attacks were executed. With the current analysis, these outliers can be identified through path length, duration of attack or other criteria.

Evaluation of vulnerabilities from insider threat is also possible using this methodology. By systematically starting the adversary at different points within the system and repeating the large number of iterations to evaluate the success rate from different starting points or scenarios, the model simulates the impact of an adversary gaining partial access through a security system before beginning an attack. This mimics the manner an insider could access a system and carry out an attack.

Current limitations of this approach occur in the technical and operational nature of the problem. The breadth of networked infrastructure and enterprise systems prohibit the full inclusion of all routes and connected systems. This is a result of systems that are connected electronically, but separated geographically or administratively. This can be improved by the definition of a repeatable method for drawing and configuring system boundaries with interconnected infrastructure.

Our research achieved the goal we had at the outset: to demonstrate that blending the cyber and physical VA processes into a single tool was possible and that it could provide valuable insight that might otherwise have been overlooked. For example, one system we evaluated was much more secure than the operators were taking credit for because of the separation of their alarm and access control systems. The methodology we developed is readily adaptable to multiple domains with cyber and physical security systems.

VI. Future Research

Future research will focus on testing cyber safeguards to determine accurate detection and delay values. This analysis has been done for physical safeguards but data are not widely available for cyber components.

Another area of research is integrating what-if analysis to show the effects of implementing different safeguards on different paths within the network to provide insight into the security benefits for specific safeguard investments or system designs.

Currently, the model only simulates one adversary attacking one facility. The model could be expanded to simulate multiple attackers working as a single team, or as multiple teams, which is a common tactic. Additionally, the model could be extended to analyze the overall security posture of an organization that spans multiple logical or physical facilities.

Similar to expanding the size and number of adversaries in the system, future efforts might include the integration of what is known in the physical security domain as a "force continuum" response for cyber attacks. The underlying principle here is to have flexible response options, varying on factors such as point of detection, type of detection, and cumulative detection statistics. Tailored response mechanisms, with varying response times, would further inject more accuracy into simulations.

Another promising research area could be the integration of near real-time threat and vulnerability information into the model. This information could be used to modify the detection and delay values of facility safeguards and their relationships to other safeguards to provide a current system effectiveness rating.

The current implementation of PACRAT focuses primarily on vulnerability and threat factors, through the inclusion of safeguard performance values and identification of system targets. The inclusion of the consequence factor, the impact of an attack, would allow PACRAT to operate in a more widely practiced and accepted risk analysis methods, conforming with DHS risk modeling approaches, where risk is measured as the product of threat, consequence and vulnerability. Additionally, the consequence factor would provide a common measure to fuse results from simulations of different targets, and the means to produce system wide risk metrics.

Acknowledgements

The research described in this paper was conducted under the Laboratory Directed Research and Development Program at Pacific Northwest National Laboratory, a multi-program national laboratory operated by Battelle for the U.S. Department of Energy.

A special thanks to Brad Key at QinetiQ North America and Dan McCorquodale at Rhino Corps for their keen insights into the timely detection methodology and modeling.

References

[1] M. E. Kuhl, J. Kistner, K. Costantini, and M. Sudit, "Cyber attack modeling and simulation for network security analysis," in *Proc 2007 IEEE Winter Simulation Conference*, pp 1180-1188.

[2] K. Ingols, M. Chu, R. Lippmann, S.Webster and S.Boyer, "Modeling modern network attacks and countermeasures using attack graphs," in *Proc 2009 IEEE Computer Security Applications Conference*, pp. 117-126.

[3] T. Sommestad, M. Ekstedt, and P. Johnson, "Cyber security risks assessment with Bayesian defense graphs and architectural models," in *2009 Proc 42nd Hawaii International Conference on System Sciences,* pp. 1-10.

[4] S.Y.K. Mo, P.A.Beling and, K.G. Crowther, "Quantitative assessment of cyber security risk using Bayesian network-based model," in *Proc. 2009 IEEE Systems and Information Engineering Design Symposium,* pp. 183-187.

[5] S.E. Jordan, M.K. Snell, M.M Madsen, J.S. Smith, and B.A. Peters, "Discrete-event simuation for the design and evaluation of physical protection systems" in *Proc 1998 IEEE Winter Simulation Conference*, pp. 899-905.