

Chapter 14

MISSION-BASED ANALYSIS FOR ASSESSING CYBER RISK IN CRITICAL INFRASTRUCTURE SYSTEMS

Thomas Llanso, Gregg Tally, Michael Silberglitt and Tara Anderson

Abstract Adversaries with the appropriate expertise and access can potentially exploit the large attack surface provided by the cyber component of critical infrastructure assets to target operations across the various sectors and significantly impact society. This paper describes a family of cyber risk methodologies known as “mission-based analysis” (MBA) that assist system designers in identifying the threats that pose the highest risk to mission execution and in prioritizing mitigation actions against the threats. This paper describes our experiences applying MBA and discusses its benefits and limitations. Also, it describes future enhancements of MBA and compares the approach with other assurance methodologies.

Keywords: Mission-based analysis, cyber security, risk assessment

1. Introduction

The Patriot Act of 2001 defines critical infrastructures as “systems and assets, whether physical or virtual, so vital to the United States that the incapacitation or destruction of such systems and assets would have a debilitating impact on security, national economic security, national public health or safety, or any combination of those matters.” The dependence of critical infrastructure systems on the cyber component brings both benefits and risks. “Risk” in this context is the product of the impact of a cyber attack on a critical infrastructure mission and the likelihood that the attack will occur. Cyber attacks on critical infrastructures have occurred for decades, but are steadily increasing in scope and frequency. Examples include the Siberian pipeline sabotage in 1982, the Stuxnet attack in 2010, and the incidents at Aramco [13] and at U.S. financial

institutions [4] in 2012. To minimize risks from cyber attacks, it is necessary to identify the vulnerabilities, the likelihood that they will be exploited and their potential impact on the critical infrastructure mission, so that it is possible to identify mitigation actions that increase survivability and resiliency. The 2013 U.S. Presidential Executive Order, Improving Critical Infrastructure Cybersecurity, highlighted the importance of critical infrastructure security and emphasized the need to implement risk-based standards.

At the Johns Hopkins University Applied Physics Laboratory (JHU/APL), we have developed three mission-related risk methodologies that fall under the umbrella of mission-based analysis (MBA). This paper briefly describes our experiences applying MBA and discusses its benefits and limitations in critical infrastructure contexts. MBA is designed to be general enough to analyze cyber threats as well as other threats, such as electronic jamming and physical attacks. However, the primary focus of this paper is cyber threats.

The goal of MBA is to analyze an operational mission, cyber threats to the mission, and information technology systems that support the mission in order to answer four questions: (i) If a threat were to be carried out, what would be the impact to the mission? (ii) What is the estimated level of effort for an adversary to realize a given threat? (iii) What mitigation actions are possible for the so called “hot spots” – threats that have a high mission impact and are relatively easy for an adversary to conduct? (iv) What are the operational costs of the mitigation actions?

The MBA processes share many similarities with NIST risk assessment methodologies [11]. However, one significant difference is the use of likelihood in the NIST methodologies versus the attacker level of effort in MBA. While the NIST methodologies are more generic and abstract, MBA is specifically focused on evaluating the risk of cyber attacks.

To date, MBA has been applied to a number of real-world contexts, primarily within the U.S. Department of Defense. Examples include analyses of Navy ship-board and submarine systems, satellite systems and a homeland security application. These applications have given us experience in optimizing and calibrating MBA, as well as identifying areas for improvement.

At a high level, all MBA variants follow a similar sequence of steps (Figure 1). A set of analytical models are first populated with data from the target problem domain. These models include adversary, mission, system and network models. The models are then scored in order to estimate risk. The scoring typically involves assigning numeric values to the mission impact and adversary level of effort (LOE) corresponding to a threat.

The scoring results are combined to provide an estimate of the risk to the mission due to cyber attack. The results are structured to show a prioritization of threats. Mitigation of threats is then considered, followed by an evaluation of the efficacy of the mitigation actions. The results enable an analyst to develop an action plan that mitigates the highest risk threats first. As threats, missions and cyber systems evolve, it is important to re-evaluate the risk periodically by repeating the steps in the methodology.

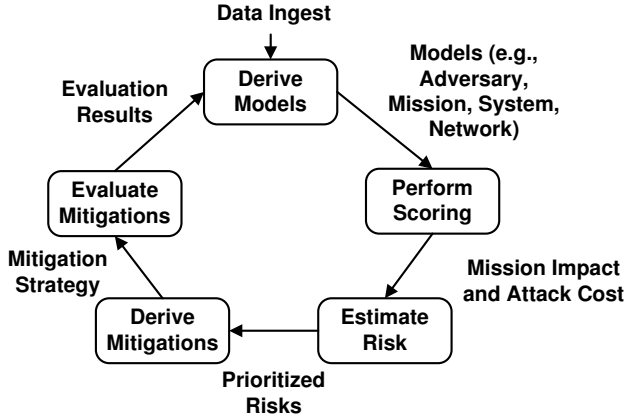


Figure 1. MBA activities.

2. Related Work

Several methodologies have been developed for evaluating the risk to critical infrastructure systems. The practical risk assessment methodology (PRAM) [9] analyzes the risk to safety-critical systems using historical data on accident probability and consequences. The goal is to ensure that new systems are at least as safe as existing systems. However, cyber attacks are fundamentally different from accidents in that the likelihood of occurrence is based on the attacker's skill, resources and motivation, not on random events. It is inappropriate to assume that historical attack data is indicative of future attack methods, frequency or success.

Ralston, *et al.* [12] have proposed risk assessment approaches for SCADA and distributed control systems. McQueen, *et al.* [7] have employed compromise graphs in which the nodes represent attack stages and the edges represent the time to compromise. They discovered that, when this approach was applied to a SCADA system, an 86% reduction in the number of vulnerabilities resulted in only a 3% to 30% increase in the time to compromise, depending on the attacker skill level and target. This methodology could be a useful future enhancement to the MBA process in order to assess the attacker level of effort.

Various mission-focused cyber methodologies have been proposed (see, e.g., [6]). The cyber security risk management (CSRM) methodology [3] uses a risk management approach to qualitatively assess and prioritize cyber security risks. The approach, which is based on NIST 800-30 [11], encompasses four processes: risk management planning, risk assessment, risk mitigation, and risk monitoring and control. It addresses all the phases of the system development lifecycle, using a cost-benefit approach to assess countermeasures that may reduce security risks. The process includes threat assessment and assignment of likelihood values to potential attacks. Threats are identified using a threat database to identify threat categories relevant to the system environment. For a given threat, the analysis estimates the consequence (impact)

Table 1. MBA variants.

Variant	Focus
Cyber Investment Analysis Methodology (CIAM)	Enterprise
Network Mission Assurance (NMA)	Network
Mission Information Risk Analysis (MIRA)	System

score (on a scale of one through five) for the attributes: mission objectives, system functions, harm, operational cost and programmatics. The assessment uses the highest consequence score for the overall impact. The likelihood and impact scores are plotted in a matrix, similar to a “heat map” used in the MBA process. This process allows alternative countermeasures to be compared for their overall consequence scores that include all five attributes. The CSRM approach appears to be comparable in some aspects to the cyber investment variant of MBA, which is described below.

3. MBA Variants

Table 1 presents the three variants of MBA. The first variant is the cyber investment analysis methodology (CIAM) [5], which considers risk and mitigation actions at the enterprise level by analyzing forensic data on attacks, vulnerabilities, CVSS scores, protection strategies and protection costs to estimate an optimal investment level by protection type. The second is network mission assurance (NMA) [2], which focuses on the availability of network bandwidth and how cyber attacks that impact network capacity can affect mission. The third is mission information risk analysis (MIRA) [6].

This paper focuses on MIRA because it is highly relevant to critical infrastructure environments. MIRA constructs three main models for the mission, system architecture and adversary. The abstraction level of each model is chosen based on the desired fidelity of the results and the amount of time and resources allowed for the analysis.

The mission model MM is a five-tuple (M, A, B, F, D) . M is a set of distinct mission types supported by the target information technology system. A is a set of quantitative mission measures of effectiveness (MOEs) that describe critical performance requirements that must be met to achieve missions in M . B is a set of quantitative system-level MOEs that are required to realize the mission-level MOEs in A . F is a set of mission-essential functions whose invocation directly impacts MOEs in B and transitively in A . D is a set of required information elements that are acted on by the functions in F . In general, there is a many-to-many mapping of MM elements.

The system architecture model SM is a four-tuple (T, N, L, C) . T is a set of node types, instances of which are found in the system. A node is defined as an active entity capable of carrying out computation and/or communication operations (e.g., router, switch, desktop, laptop, server or wireless device). N

is a set of node instances in the architecture. L is a set of links between node instances. C is a mapping of data types from D in the mission model to N . A given link represents connectivity, typically in a network or communications context, between two node instances.

The adversary model AM consists of an estimate of the maximum LOE that the anticipated worst-case cyber adversary might muster against the target system. The value ranges from one to ten, where one indicates minimal exertion on the part of the attacker and ten indicates the maximal level of exertion corresponding to the capabilities possessed by a top nation-state attacker.

We now describe the MIRA activities that map to Figure 1. In the Derive Models activity, analysts populate MM , SM and AM . They can do so using a number of techniques, such as reviewing relevant documentation, interviewing mission and system experts, and, if permitted for an existing system, by running automated discovery analytics against operational mission system environments to identify the nodes (N), links between nodes (L), and data types (D) on nodes (C).

In detailing the models, analysts also capture how mission data flows over the system nodes in the context of different mission-essential functions. To characterize the adversary, the analyst considers the mission in the context of different kinds of potential adversaries who might wish to harm the mission.

In the Perform Scoring activity, analysts derive two distinct types of scores: mission impact scores and attack LOE scores. In the most detailed case, a mission impact score is assigned for each viable tuple from (M, A, B, F, D, N, CT) where M , A , B , F , D and N are defined above and CT denotes the type of compromise, e.g., confidentiality (adversary accesses data), integrity (adversary modifies data or service function) or availability (adversary prevents the use of data or a service). A mission impact score is an ordinal value in the range one to five, where one denotes “fully mission capable” (i.e., no mission impact) and five denotes “not mission capable” (i.e., mission fails). An LOE score is assigned to each viable tuple (N, D, CT, AV) where N , D and CT are defined above and AV is the attack vector. Typical attack vectors are network, insider and supply chain implant.

For LOE scoring, analysts have employed the global information grid information assurance portfolio (GIAP) scoring approach that estimates the required cyber attacker capability on an absolute scale ranging from one (script kiddie) to ten (nation-state). Research into platform vulnerability history and threat reports help analysts narrow the estimated capability requirements. Several approaches are available for determining the types of scoring, each with its own strengths and weaknesses. MIRA allows for the use of alternative scoring approaches. One approach we are currently studying involves constructing models of the mission and the related cyber system. A simulated Monte Carlo attacker repeatedly constructs and directs cyber attacks at the system model; the impact of these attacks are automatically assessed by functions that compute MOEs from the system state and mission state. Scoring is a rich area for




Figure 2. Heat map.

future research. A more detailed discussion of scoring procedures is omitted in this paper for reasons of space.

In the Estimate Risk activity, analysts create a risk matrix called a heat map (Figure 2). The x-axis of the heat map depicts mission impact scores and the y-axis represents LOE scores, with the highest LOE scores located closest to the origin. In this arrangement, cyber attacks, that are both highly mission impacting and that require a relatively low LOE, appear in the upper right-hand quadrant of the heat map (outlined in Figure 2). These attack contexts represent the highest risk to the mission and become key inputs to the remaining activities.

MIRA currently considers different attack steps/contexts in heat maps individually. An interesting topic for future research is the analysis of combinations of attack steps from an impact/LOE perspective.

In the Derive Mitigations activity, analysts consider different combinations of security controls, such as those described in NIST 800.53 [10], and the effectiveness of these controls at countering the prioritized attack contexts identified in the heat map. Such controls include people, processes and technologies. MIRA enables analysts to prioritize mitigation actions based on the risk reduction achieved. After the mitigation actions have been derived, it is necessary to evaluate the actions.

In the Evaluate Mitigations activity, analysts re-factor the revised system architecture model with the derived mitigation actions and rescore the LOE values of the affected attack contexts. A key goal in this activity is to increase the LOE values via well-placed mitigation actions. Increased LOE values can help reduce the overall mission risk, especially if the recomputed LOE values exceed the estimated capabilities of the cyber adversary in the adversary model *AM*. Ultimately, analysts produce a set of recommendations based on the overall analysis, including the mitigation actions. For existing systems, penetration testing can provide additional evaluation and validation of mitigation actions after they have been implemented in the cyber architecture.

As stated earlier, of the three variants that make up MBA, MIRA and NMA are most applicable to critical infrastructure systems. CIAM is the least relevant given its general information technology enterprise focus. NMA is useful for critical infrastructure system analysis, especially in cases where network bandwidth and availability are key performance considerations. However, MIRA is the best approach for conducting risk analyses of critical infrastructure systems because of its strong mission/system/data ties. The remainder of this paper focuses on MIRA in the context of critical infrastructure systems.

4. Applications of the MIRA Variant of MBA

This section briefly describes two examples involving the application of MIRA to critical infrastructure systems, one is a real-world satellite system and the other is a hypothetical scenario involving a railroad control system.

4.1 Satellite Control System

The MIRA analysis of the satellite control system integrates both cyber and electronic warfare attacks, enabling an end-to-end mission risk analysis of the system. The satellite control system consists of purpose-built hardware devices, including ground antennas and special communications processors that are easily represented in the MIRA system model. For reasons of sensitivity, the satellite control system is described at a high level. However, the lessons learned with regard to the application of MIRA to critical infrastructure systems are clear.

The following components are involved in the analysis:

- **Mission Model:** The mission model defines the overall mission objectives of the satellite control system, key mission-level and system-level MOEs, actors or role players engaged with the satellite control system, mission-essential functions and mission data processed by the satellite control system. The satellite control system has a worldwide footprint capable of controlling designated satellites for various purposes on a global basis. The control system provides support for satellite operation, control and maintenance, including functions such as telemetry, tracking and commanding operations, pre-launch checkout and simulation, launch support and early-orbit support when satellites are in their initial or transfer orbits and require maneuvering to their final orbits. Time on the satellite control system is strictly scheduled based on reconciling conflicting, prioritized demands and the availability of ground stations located around the globe.

The satellite control system mission model captures three primary mission threads with a defined set of quantitative MOEs. An example of a satellite control system mission is: “Help diagnose faults and restore designated satellites to full operation.” The MOEs highlight the critical mission requirements needed to obtain mission success. An example of a

satellite control system MOE is: “The probability of success in completing a scheduled contact.”

Mission-critical functions are then defined for the satellite control system mission set. An example of a mission-essential function is the “Submit updated schedule function.” As noted above, the mission areas are analyzed and described as a series of mission threads invoking mission-essential functions. Examples of mission-essential functions defined for the satellite control system analysis are: “Submit satellite communications session request,” “Reserve overhead resources for communications fabric,” “Submit urgent needs schedule request,” “Configure and track link,” “Establish communications session” and “Internal and external space operations center communications with satellite.”

- **System-Level Model:** The satellite control system is decomposed into node types, node instances and node interconnectivity. The centralized command and control, scheduling and remote uplink/downlink nodes are captured. The resulting system model comprises more than two dozen data types and four dozen computational nodes, such as the schedule data type and the data type indicating the identity and location of the target satellite system.
- **Adversary Model:** The adversary model assumes nation-state level attackers of the advanced persistent threat variety. The threat is applied to the criticality assessment described above to identify hot spots and the level of effort needed to impact mission success and identify system risks.
- **Mitigation Phase:** Specific mitigation actions are recommended as a result of the analysis. Several attacks have a high mission impact but a low estimated level of effort. The results of this analysis enable the identification and prioritization of capabilities, systems, and science and technology needs to align the system with national space policy and U.S. Department of Defense capabilities documents and functional plans.

4.2 Railroad Control System

The second example application of MIRA is in the area of monitoring and controlling train movements through an initiative called positive train control (PTC). The Rail Safety Improvement Act of 2008 requires the implementation of PTC on certain rail lines by 2015 [1]. The analysis presented in this section is based on publicly-available documentation and should not be interpreted as a complete MIRA analysis of PTC.

The major PTC components are [8]:

- **Back Office Server:** This system stores data on speed restrictions, track geometry and wayside signaling.

- **Onboard System:** This system displays train information to the engineer, monitors and controls train movement if the engineer fails to respond to audible warnings, and uses GPS to report train position.
- **Wayside Signal System:** This system comprises traffic signals located along the track that are connected by cellular modems (primarily 220 MHz) to the back office server and onboard system.
- **Communications Network:** This system comprises a redundant wired and wireless communication network that connects a locomotive, cab car, back office server and wayside interface units

The following components are involved in the analysis:

- **Mission Model:** The mission MOEs of PTC are to reduce fatalities, injuries and the cost of damage due to improper train movements. PTC specifically does not attempt to reduce deaths or injuries when people trespass on railroad tracks or vehicles bypass railroad crossing barriers. PTC does not take human operators out of the system, but it does provide a fail-safe mechanism to stop unsafe train movements should the human operator fail to heed alarms and signals from the onboard system.

To improve reliability, PTC is designed to incorporate multiple redundant capabilities. Wayside signals communicate with both onboard systems and back office servers. This ensures that the onboard system receives signal information even if the back office server does not provide the data. The communications network has wired and wireless network connections. Also, since not all geographic areas support accurate GPS measurements, onboard systems have alternative methods to compute location.

Some of the system MOEs for a PTC system are:

- Accuracy of train position reporting (including the track that the train occupies).
- Accuracy of automated train braking (stopping as close as possible to the target without overrunning it).
- Effective throughput of the communications network to ensure timely delivery of the train consist and other data from the back office server to the onboard system.

Some of the mission-essential functions and required information elements include:

- The onboard system reports train position to the back office server via the communications system.
- The wayside signals report the switch status to the back office server via the communications system.
- The back office server reports the signal status to the onboard system via the communications system.

- The back office server sends train consist data, authorizations and restrictions to the onboard system via the communications system.
 - The onboard system calculates safe braking parameters based on data from the back office server.
 - The onboard system visually and audibly provides information to the engineer.
- **System Architecture Model:** The major components of the PTC model are described above. Additional components of a PTC implementation include wired and wireless network infrastructure components such as radios, base stations, routers and switches; standard information technology components in the back office server; locomotive and cab car components of the onboard system; and handheld terminals for track work crews.
- **Adversary Model:** While PTC was mandated in response to accidental train collisions [8], the centralized train dispatching system is also a potential target for adversaries. Passenger train collisions can cause death and injury to large numbers of people. Freight trains pose a potentially greater threat when they carry toxic chemicals and other hazardous cargo that could harm people near the site of a collision or derailment.
- **Estimating Risk:** According to the system MOEs, the primary concerns are the integrity and availability of the mission functions and information. Confidentiality is a less significant factor in meeting mission objectives, although it is important to protect data regarding future train movements and cargo. For an attacker to achieve the objective of causing a train movement to enter an unsafe area, the train must either violate the assigned authorizations and restrictions, or the data on which the authorizations and restrictions are based must be compromised. In the first case, the back office server would correctly calculate where and when the train should move, but the onboard system would not relay the information to the engineer or execute the required braking actions automatically. In the second case, the back office server would make incorrect decisions because of incorrect train position reporting, train consist reporting or wayside signal reporting.

Further LOE analysis requires information specific to a PTC system. However, the likely targets for attack are the integrity of one or more onboard systems, either directly or through the back office server. The back office server has the potential to provide inaccurate data to an onboard system (causing incorrect braking calculations). It could also impact the integrity of an onboard system through code injection or similar attacks. The attacker LOE would depend on the ability to gain access to an onboard system or back office server. This could be physical access (supply chain, malicious insider or at a train yard) or remote access. Possibilities

for remote access include network attacks on a back office server and malicious data injection into the communications network. Since railroads share track sections, PTC is designed to be interoperable across railroads. An infected onboard system could potentially infect the back office server of another railroad when it operates on the same (compromised) line.

- **Mitigation Phase:** The prioritization of mitigation actions is dependent on the LOEs of specific attacks and the impact of the attacks on mission MOEs. Mitigation actions include:
 - Code integrity monitoring in the onboard system with a failstop on integrity violations.
 - Encrypted and authenticated wired and wireless network communications.
 - Isolating the back office server from non-PTC networks.
 - Standard NIST 800-53 [10] controls to protect the integrity and availability of the back office server.

The mitigation actions must be selected so that they do not impact availability and the timely completion of onboard system functions.

4.3 Potential Improvements

Several potential improvements to MIRA are possible that could aid critical infrastructure system evaluations. Our experience applying MIRA has demonstrated the need for automation to support the analysis, particularly during the scoring process. This is especially important for complex, distributed critical infrastructure systems. As a generic example of the need for automation, MIRA analysts work with mission experts to perform mission impact scoring. Unfortunately, manual scoring does not scale to common instances of MM, especially in the case of large critical infrastructure systems. Suppose, for example, that the five MM components have cardinalities of 5, 7, 7, 10 and 20, there are three attack types (confidentiality, integrity and availability) and there are a total of 50 nodes. Then, a total of 7.35 million distinct mission impact scores must be analyzed and assigned, a task that cannot be performed manually. As mentioned above, we are currently researching an approach for automating mission impact scoring.

Methodologies such as MBA produce results that should be independently validated if at all feasible. We have not as yet done this for MBA in the context of critical infrastructure systems, although it is very important given the nature of critical infrastructure systems. A possible approach, albeit potentially costly to execute, is to have independent red teams attack a critical infrastructure system that was previously analyzed and mitigated using approaches such as MIRA and NMA. The independent red team results could help calibrate how well the MBA methods pinpoint important attack types and mitigate against them in the context of the specific critical infrastructure system and in the

general case. An alternate approach is to emulate portions of a critical infrastructure system in a testbed or use a modeling environment and attempt to validate the analysis results in these settings with red team techniques.

5. Principles and Key Takeaways

The following are some principles and key takeaways related to applying MBA to critical infrastructure systems:

- Mitigation actions for critical infrastructure systems should be very sensitive to performance requirements, especially availability requirements. Timing tolerances in critical infrastructure systems are much stricter than those associated with many traditional information technology systems.
- When analyzing critical infrastructure systems, it is important to vary the mission timeline while holding the other attack variables constant. This can assist in worst case analysis. For example, an attack on the integrity of pump settings in a wastewater treatment plant can be far more damaging before the pumps are programmed compared with attacking the settings after the data has been written to log files.
- Since many critical infrastructure systems are fielded for extended periods during which few major upgrades are permitted, it is important that approaches such as MIRA are applied as early as possible in the system development lifecycle when there is still time to follow through on the mitigation actions that are identified.
- Critical infrastructure systems are very tightly coupled to the missions they support compared with general purpose computing systems. Most assurance methodologies do not prioritize the implementation of countermeasures with respect to mission priorities. With its focus on mission assurance, MIRA enables mitigation efforts that prioritize the preservation of critical mission functions with minimal cost and performance impact.
- The segmentation of networks and systems in critical infrastructure environments must be performed with great care compared with non critical infrastructure environments. The separation of business networks from critical infrastructure systems is essential to limit attacks, as demonstrated by the recent targeting of Aramco computer systems [13]. Analysts applying MBA must keep this in mind during their analysis.

6. Conclusions

As cyber-physical interfaces become more prevalent, especially in critical infrastructure systems, cyber threats are an increasing concern. Mission-based analysis (MBA) is a family of methodologies that allow mission and system owners as well as designers to understand the mission impact of cyber attacks

and targeted mitigation strategies in the context of critical infrastructure systems. MBA, in particular the MIRA variant, advances the state of the art in mission assurance for critical infrastructure systems by expanding the analysis beyond the examination of failures and by creating a framework for integrating techniques for modeling missions, adversaries, threat space and impact of successful attacks on prioritized missions. By utilizing MBA, the critical infrastructure protection community can prioritize resource decisions in fiscally-constrained environments to protect against cyber threats that pose the highest risk, thereby reducing the overall risk to critical systems and missions.

References

- [1] S. Alibrahim and T. Tse, Signal and Train Control, Federal Railroad Administration Research and Development Program Review, Federal Railroad Administration, Washington, DC, 2008.
- [2] C. Burris, J. McEver, H. Schoenborn and D. Signori, Steps toward improved analysis for network mission assurance, *Proceedings of the Second IEEE International Conference on Social Computing*, pp. 1177–1182, 2010.
- [3] P. Katsumata, J. Hemenway and W. Gavins, Cybersecurity risk management, *Proceedings of the Military Communications Conference*, pp. 890–895, 2010.
- [4] M. Keefe, Timeline: Critical infrastructure attacks increase steadily in past decade, *Computerworld*, November 5, 2012.
- [5] T. Llanso, CIAM: A data-driven approach for selecting and prioritizing security controls, *Proceedings of the IEEE International Systems Conference*, 2012.
- [6] T. Llanso, P. Hamilton and M. Silbergliitt, MAAAP: Mission Assurance Analytics Platform, *Proceedings of the IEEE Conference on Technologies for Homeland Security*, pp. 549–555, 2012.
- [7] M. McQueen, W. Boyer, M. Flynn and G. Beitel, Quantitative cyber risk reduction estimation methodology for a small SCADA control system, *Proceedings of the Thirty-Ninth Annual Hawaii International Conference on System Sciences*, p. 226, 2006.
- [8] Metrolink, An introduction to positive train control, Los Angeles, California (www.metrolinktrains.com/agency/page/title/ptc).
- [9] C. Mokkapati, T. Tse and A. Rao, A practical risk assessment methodology for safety-critical train control systems, *Proceedings of the Annual Conference of the American Railway Engineering and Maintenance-of-Way Association*, 2009.
- [10] National Institute of Standards and Technology, Recommended Security Controls for Federal Information Systems and Organizations, NIST Special Publication 800-53, Revision 3, Gaithersburg, Maryland, 2009.

- [11] National Institute of Standards and Technology, Guide for Conducting Risk Assessments, NIST Special Publication 800-30, Revision 1, Gaithersburg, Maryland, 2012.
- [12] P. Ralston, J. Graham and J. Hieb, Cyber security risk assessment for SCADA and DCS networks, *ISA Transactions*, vol. 46(4), pp. 583–594, 2007.
- [13] Reuters, Aramco says cyberattack was aimed at production, *New York Times*, December 9, 2012.