Available online at www.sciencedirect.com

SciVerse ScienceDirect

Procedia Computer Science 00 (2014) 000–000

Procedia Computer Science

www.elsevier.com/locate/procedia

Conference on Systems Engineering Research (CSER 2014)

Eds.: Azad M. Madni, University of Southern California; Barry Boehm, University of Southern California;
Michael Sievers, Jet Propulsion Laboratory; Marilee Wheaton, The Aerospace Corporation
Redondo Beach, CA, March 21-22, 2014

Mitigating The Risk Of Cyber Attack On Smart Grid Systems

Eric B. Rice[a]*, Anas AlMajali[b]

[a]*Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, USA*
[b]*Information Sciences Institue, University of Southern California, Los Angeles, California, USA*

Abstract

Smart Grid technologies are being developed to upgrade the power grid with networked metrology and controls that can improve efficiency and provide new methods to manage the system. While these technologies offer great benefits, they also introduce new classes of risk, most notably creating new attack vectors that can be exploited by cyber attack. To assess and address risks in cyber-physical systems like these, the system designer's toolset needs to include concepts drawn from cyber security, reliability, and fault tolerance design, integrated into a common methodology. In this paper, we discuss the fragmented landscape of studies into the risk of cyber attack on smart metering systems, and then draw on concepts from systems engineering and fault tolerance design to organize and unify the pieces.

Keywords: Smart Grid, Systems Engineering, Cyber Security, Cyber-physical Systems

* Corresponding author. Tel.: +1-818-393-6643.
E-mail address: Eric.B.Rice@jpl.nasa.gov

1. Introduction

Smart Grid technologies are being developed to upgrade the national power infrastructure with networked metrology and control infrastructure. These efforts aim to improve operational efficiency in the near term, and create new opportunities for advanced technologies in the future. Smart Grid solutions may greatly improve the reliability of the grid and reduce costs of power delivery, but they also introduce new reliability and cyber security risks. Demand Response (DR) technologies provide the utility a path to manage system load in addition to the available supply, creating control loops that may couple loosely with existing controls. Advanced Metering Infrastructure (AMI) deployments are intended to improve the efficiency of meter reads, but also provide other monitor and control capabilities, such as the ability to connect or disconnect service to a customer by remote command. In both examples, new control paths are added that may affect the individual and aggregate behavior of loads on the power grid in ways that have not previously been experienced by grid operators.

The risks introduced by AMI deployments are of particular interest because of the immediacy of AMI deployments. Commercial AMI networks are being deployed worldwide, with penetration rates in the U.S. reaching as high as 30.2% in 2012[2]. These commercial deployments commonly follow current best practices for security, which include protections for the core data security concerns of integrity, availability, and confidentiality[3]. Although these security designs can provide assurance that cyber attacks will be less likely to succeed, the consequence of a successful attack may still be high. A compromise of metering networks may allow adversarial access to control functions that, if corrupted, have consequences above and beyond the integrity or availability of data in the system, including risks to human safety and system integrity. As market penetration increases, so do the potential risks associated with the security of these controls.

Over the last several years, there have been numerous publications that attempt to characterize risk of cyber attack directed at smart grid technologies. Some of these efforts have focused on breadth-first analysis of risk, with only a cursory discussion of mitigations. Others have focused on one element of the system, or on a limited scenario, demonstrating robustness of that part of the system to a given attack. Studies that address mitigation of attacks tend to focus on a single feature or component, with little discussion of the properties of the end-to-end system that make the mitigation viable.

In this study, we intend to tie together some of the independent threads in the literature, by applying systems engineering and fault management concepts to an example cyber-physical scenario. In our previous work[1], we built a simple model that allowed us to study the inter-system interactions of a metering network with the power system during a load-drop attack, in which an adversary gains access to control functions of the AMI to send a sequence of service disconnect commands to customer meters. In this discussion, we apply broader systems and fault management concepts to expand that analysis to characterize the range of potential behaviors and outcomes, and to derive constraints for detection and response techniques that may help reduce the consequence of such an attack.

2. Related Works

In approaching this analysis, we draw on three threads from related disciplines. We aim to contribute to the existing body of work in smart grid attack analyses by applying concepts from fault management systems design within a broader framework of systems engineering for security.

2.1. Smart Grid Cyber Attack Analyses

A number of publications in recent years have attempted to address the risk of cyber attacks against metering infrastructure. Broader papers have performed analyses of the threats to power infrastructure in general and AMI systems in particular, while more focused works have attempted to characterize one part of the system during an attack. Sridhar et al[5] performed an end-to-end assessment of the risks the power system faces from malicious attack. While the analysis includes concepts of control loop integrity and load drop attacks directed at AMI systems, it stops short of describing the mechanisms by which an attack would impact the grid, or estimating the potential outcome of a particular scenario. Clements et al.[6] provide more detailed analysis of the Western Electricity Coordinating Council (WECC) service area response to specific load-drop and oscillatory load control attacks. Results showed that the WECC is able to tolerate the conditions of the scenarios studied, however, this analysis is done with such a limited set of criteria that neither the results nor the method provide insight that can be generalized to the protection of power infrastructure. Temple et al.[7] take an additional step to propose a potential mitigation – addition of a delay to the execution of a service disconnect command. A command delay provides a built-in window for detection and response, but Temple's work stops short of specifying the necessary performance of such functions. Anderson and Fuloria[8] further discuss architectural considerations for protecting meter communications, including the use of encryption.

Taken individually, these publications each tell an incomplete story, but taken together, they each tell a different piece of a larger story. The study of smart grid cyber security risks can benefit from a broader perspective that can tie each of these efforts into a cohesive mitigation design. A systems approach would not only provide a broader context for studies like the above, but would also provide concepts and methods for evaluating them together and against each other when adopting the individual recommendations.

2.2. Systems Engineering and Cyber Security

Bayuk[9, 10, 11] provides a wealth of foundational thought on the application of systems engineering methods to security analysis and design, and a recent issue of INCOSE's Insight publication discusses more detailed topics in system security design. The organizing principles and process in these publications can help organize the above piecemeal studies into a cohesive risk management design for AMI systems. Of particular interest to our work, Dugan and Snell[12] discuss the derivation of robust security requirements. Applied to the above smart grid studies, requirements analysis can provide performance specifications that can help size the protective features (e.g. how much command delay is sufficient, compared to the performance of detection and response?) and manage their impact on overall system function (e.g. how much overhead for encryption can the AMI network tolerate?).

Such a requirements process does not stand alone, however. A more robust and resilient system design must be derived from a clear understanding of the acceptable and unacceptable behaviors of the system. Woody[13] describes thread analysis, which is the specification of desired end-to-end behaviors of a system. While Temple's analysis of a delay feature demonstrates its viability as risk mitigation, it relies on the assumption that the capability exists to detect disconnect commands and send cancellation messages, without specifying the performance of that capability. Applying Woody's mission thread analysis or another behavior specification method, it would become clear what the system behavior must be for the delay feature *and the related capabilities* to sum to a viable mitigation.

2.3. Fault Management and Cyber Security

The disciplines of fault management and dependable computing provide some useful formal concepts that can be applied in analysis of cyber attacks and the design of mitigations. In all of the above work, the unifying concept is one of *cyber attack as a cause of system failure*. This idea is expressed well in the context of system dependability in the taxonomy developed by Avizienis et al.[14], where intentional, malicious, human-induced faults in both software and hardware are discussed alongside other causes of system failure. This insight also opens cyber security design up to many of the concepts in fault management and related areas of study.

In a pair of papers, Johnson and Day[15,16] present a systems framework for the analysis and management of failures. For this paper and future work, we intend to apply Avizienis's definition of *attack*, and Johnson & Day's definitions of *fault*, *failure*, *anomaly*, and *root cause* to describe events. The concepts of *critical failure effect* and *time to criticality* are key concepts to assess behaviors.

Integrating cyber attacks that target system function into this conceptual framework is simple. All of the definitions and methods are broad enough to be inclusive; they do not discriminate along lines of system properties (they can be applied equally well to systems of varying properties and physics), and they can be used to describe causes and effects that are either intentional or random in nature. Fovino et al.[17] use similarly general language to integrate fault and attack analysis into a single fault tree structure that provides a formal method for analyzing intentional and random events side-by-side.

Definitions[14,15,16]

Failure	The unacceptable performance of intended function.
Anomaly	The unexpected performance of intended function.
Fault	A physical or logical cause, internal to the system, which explains a failure.
Attack	Malicious external fault.
Critical Failure Effect	The first failure effect to compromise a critical system objective.
Time To Criticality	The time it takes for a failure to propagate to a critical failure effect.

3. Analysis of The Load Drop Attack

In our prior work in this area, we constructed a cyber-physical model of an AMI system connected to a power system, and measured the integrated system's sensitivity to a load drop attack. Our goals were three-fold: to identify a critical failure effect, at which point the ability of the system to maintain power flow was compromised; to find the worst-case attack scenario that was executable given the properties of the AMI network; and to identify the time to criticality for the scenario, as the worst-case (shortest) time that an attacker may induce the critical failure effect. The following discussion uses the same model of the system to apply a more general systems analysis.

3.1. AMI-Power Cyber-Physical Model

The model used to assess system behaviors that are representative of (though not directly applicable to) real-world system was constructed in four parts, illustrated in Fig 1.

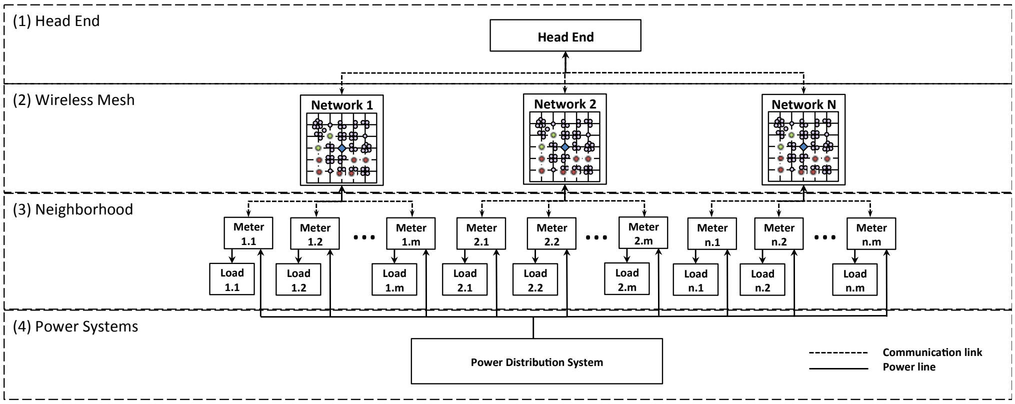

Fig. 1. Smart Grid System Model

1. Head End – The back-end utility control system (referred to as the head-end) has monitor and control capabilities for the meter network. The functional properties of the head end were not modeled in any significant detail, as it was assumed that its contribution to the behavior of the overall scenario is not driving.

2. Wireless Mesh - The wireless communications infrastructure that connects the head-end to smart meters consists of a collection of smaller sub-networks containing a subset of meters.

3. Neighborhood Model - The topology of the meter wireless mesh and the connection of smart meters to loads was determined by what we termed the Neighborhood Model, which drew on unit density for the 90057 zip code and distribution of customer demand for the Los Angeles Department of Water and Power service area. With this model, we were able to derive the stimulation applied to the power system from the execution times of load disconnect commands.

4. Power System Model - Finally, the power system was modeled using a WECC-approved configuration of the IEEE 9-bus model, consisting of 3 generators and 3 loads..

Further details of this model can be found by referencing the paper, “A Systems Approach to Analysing Cyber-Physical Threats in the Smart Grid.”

3.2. Failure Effects

The failure effects of interest for a load-drop attack can be derived from applicable regulatory standards for power system stability and reliability. In order of consequence, failure effects are listed in Table 1.

Table 1. Failure Effects

	Failure Effect	Description
1	System shutdown	Measured as zero power output by all generators at the end of the simulation. In the scenario discussed here, generator shutdown is affected by violating the operating limits for frequency. The Electric Power Research Institute (EPRI) tutorial[18] defines this as a frequency exceeding 61.8 Hz for at lest 0.25 seconds.
2	Power quality violation	IEEE 1159-2009[19] defines limits for acceptable power delivery to customers. Delivered voltage is required to stay in the range 0.9 pu < V < 1.1 pu. System frequency is required to stay within +/- 0.10 Hz of nominal. The consequences of operating outside of power quality requirements include damage to or degradation of customer and utility equipment.
3	Anomalous operation	A set of cases that fall outside normal operational behavior, but that do not rise to violations of standards or requirements.
4	Normal operations	A set of cases that resemble normal operations. May still contain unauthorized actions, but consequences are not measurable at the system level.

3.3. Fault/Attack Analysis

A subset of an integrated fault/attack analysis is included in Fig 2. The highlighted boxes (1-9) show the attack vector of interest to this exercise. The primary function of power control is to maintain generation that is matched to the demand in the system. Attack scenarios that disrupt enough load or supply quickly enough will be able to overcome the system’s ability to compensate. Power operations contingency plans tend to focus on cases where generation is insufficient to meet demand[4], so it is interesting to consider cases where supply far exceeds load in the system, captured as item 4 in the tree.

When load is removed from the system, there are two immediate technical consequences: at the generators, frequency increases proportionally to the difference between load and generation; and voltage increases proportionally to the error in reactive load. Because our simulations included reductions to the purely active, or resistive, part of load, we did not measure an impact to voltage that could otherwise be expected.

Moving from right to left in the tree, it becomes apparent how smart grid technologies introduce risks at a high level of abstraction, and how cyber-physical risks occur. With the ability to influence load in the system, smart grid technologies are capable of directly impacting the power balance. By exploiting these digital technologies, an attacker would have direct influence over quality and integrity of the power delivered to customers.

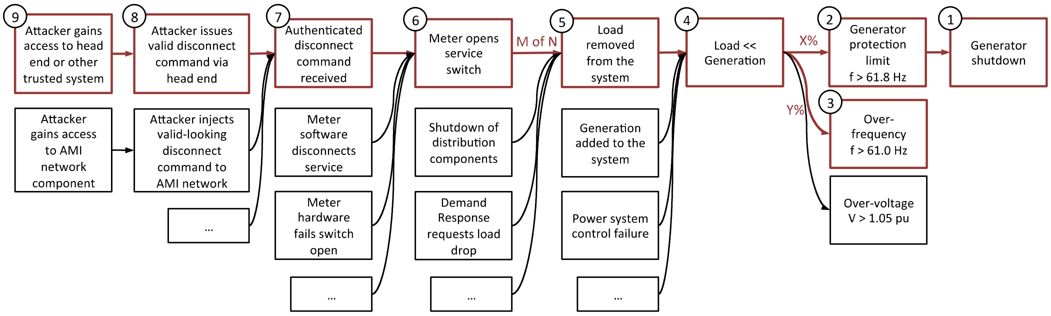
9
Attacker gains access to head end or other trusted system
8
Attacker issues valid disconnect command via head end
7
Authenticated disconnect command received
6
Meter opens service switch
M of N
5
Load removed from the system
4
Load << Generation
X%
2
Generator protection limit f > 61.8 Hz
1
Generator shutdown
Y%
3
Over-frequency f > 61.0 Hz
Over-voltage V > 1.05 pu
Attacker gains access to AMI network component
Attacker injects valid-looking disconnect command to AMI network
…
Meter software disconnects service
Meter hardware fails switch open
…
Shutdown of distribution components
Demand Response requests load drop
…
Generation added to the system
Power system control failure
…

Fig. 2. Fault/Attack Tree

3.4. Load Drop Sensitivity

In our simulations, we stimulated the power system with load drops that varied in both magnitude, from 0 percent to 100 percent of system load, and in duration, from 0 seconds to 100 seconds, to characterize the sensitivity of the system to load drops. This analysis characterizes the propagation paths 4 → 2 → 1 and 4 → 3 in the tree, by determining the percentage of load that needs to be dropped to result in the failure effects 2 and 3 (X% and Y%). For each case that resulted in a total loss of generation, we extracted the time and load at which the system shut down, and found that the system behavior was well structured. The risk of system shut down clustered above a clear boundary, and there were no cases of non-convergent or unstable behavior below it.

Extending that assessment, we then looked for the load and time at which any of the above failure effects occurred in. The results, in Fig 3a, reveal two clear boundaries. The lower boundary, starting around 2% load drop, is the point at which the system reaches an over-frequency of 61.0 Hz. The upper boundary is determined primarily by the generator frequency protection limit at 61.8 Hz, and most cases occur above 50% of system load. Over- and under-voltage cases did show up in the data, but primarily as the system response to a generator shutdown.

The clear boundaries in this view of the system allow us to reduce the problem significantly. We can discuss the system in terms of clearly defined regions of consequence, as depicted in Fig 3b, rather than estimating the consequences of each individual case. The regions are numbered in accordance with the list of failure effects in Section 3.2 above. At a glance, the system is very tolerant of such an event – even 2% of a metropolitan service area is a large number of customers to be addressed in a short window to have an impact on power quality.

3.5. AMI Disconnect Command Execution

In our AMI network simulations, we simulated the execution of disconnect commands as they are sent by the head end and traverse the wireless mesh network to characterize the behavior of the path 8 → 7 → 6 in the above tree. These simulations rely on two key related assumptions: 1. Disconnect commands are sent individually to each meter; and 2. They are sent in the most efficient pattern possible, so that the load of the execution is distributed to all branches of the network equally. We then varied the spacing between commands to find the performance limits of the network that affect the behavior of the attack. Detailed results can be seen in our earlier paper[1].

The design of this experiment provided some key insights to the viability of potential attack scenarios. Most importantly, the scenarios bracketed the performance limitations of the network – we were able to find the fastest possible execution of a stream of commands given the properties of the network. Above a particular rate, in this case command spacing smaller than 100 ms, the network became clogged and commands were delayed. This places an upper bound on the rate of command delivery, which in turn puts an upper bound on the slope of the load drop curve.

By overlaying a load drop curve on top of the behavioral boundaries, we are able to estimate the magnitude of event that will have a corresponding consequence. By overlaying the worst-case attack execution, derived from the network simulations, we are able to derive fundamental properties that can be applied as requirements for protective measures for this scenario. The Time to Criticality for the two consequence levels is annotated in Fig 3c. TTC_1 is the time to criticality for failure effects in region 1 (system shut-down), and TTC_2 is the time to criticality for failure effects in region 2 (violation of power quality requirements). This application of the concept of TTC differs from that proposed by Johnson and Day, in that they measure the TTC as the time to the *first* intolerable failure effect, and we are applying it to *each* failure effect of interest.

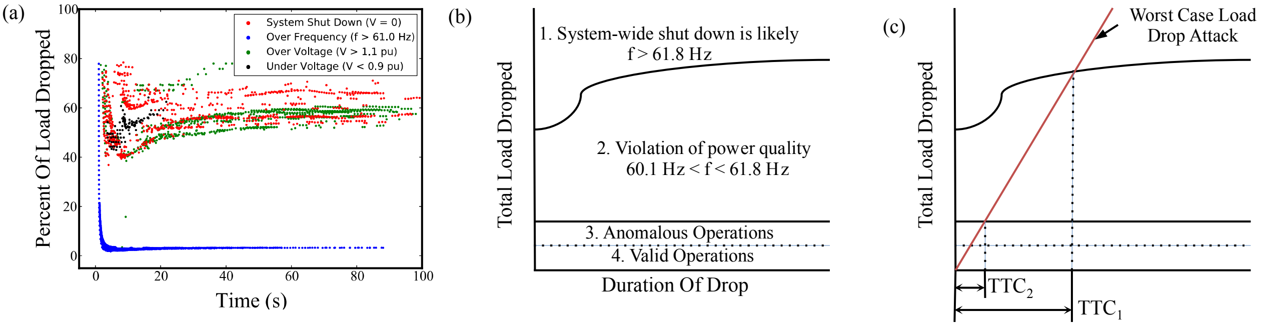

Fig. **3**. (a) Power System Sensitivity to Load Drop, (b) Regions of Consequence, (c) System Parameters

4. Mitigations

Mitigations for this scenario can be categorized by their effect on the behavior of the scenario. By taking a risk management approach, we can divide mitigations into categories of *prevention* (mitigations that reduce the likelihood of an event) and *tolerance* (mitigations that reduce the consequence of an event)[16]. While a simple distinction, this is an important one in the context of cyber events, where most controls are preventative in nature. A few common example cyber controls are listed in Table 1, with an assessment of their impact on the scenario.

Table 2. Example Cyber Controls

	Feature	Applied At	Notes	Likelihood	Consequence
1	Access Controls	Head End	Reduces likelihood of malicious access to HE	↓	•
2	Encryption	HE-Meter interface	Reduces attacker's ability to generate unique valid commands independent of HE	↓	•
3	Command Source Authentication	HE-Meter interface	Ensures meter will not accept commands from an unauthorized 3rd party	↓	•

As with most cyber controls, if the control is subverted, a malicious actor potentially gains unfettered access to create unwanted effect, and the potential consequences have not been addressed. The cyber security community has put significant effort into detecting subversions of controls (like unauthorized access), but in the case of a cyber-physical application like this one, there is an opportunity to implement deeper features that address risk by limiting allowable actions in the system to those that are safe for the system.

Top-level performance requirements for mitigations to the scenario can be derived from the above analysis. Features that detect and respond to an ongoing scenario must act quickly enough to complete before a failure effect is reached. To avoid failure effect 1, detection and response must execute within TTC_1; to avoid failure effect 2, detection and response must execute within TTC_2; and so on. Taken in this context, Temple's proposed delay feature can be viewed as one feature in a system design. A delay to command execution does not preclude any given failure effect (the disconnect commands are still executed), so by itself, the delay does not mitigate the risk. A random delay will decrease the slope of the load drop, and therefore increase the time to criticality for all failure effects. Table 3 identifies potential features that may reduce consequences of the scenario in a more complete way.

Table 3. Mitigation Features

	Feature	Applied At	Notes and Effect	Likelihood	Consequence
1	Random Delay	Meter	Increases TTC. Enables detection	•	•
2	Throttle network	AMI Network	Reduces command rate, increases TTC.	•	•
3	Detect load drop	Meter or power system	Detects ongoing event.	•	•
4	Detect commanding	Head end or meter network	Detect load drop commands as they are sent. Detects ongoing event.	•	•
5	Restrict traffic	Meter network	Stop commands from reaching meters. Truncates load drop.	•	↓
6	Cancel commands	Head end/Meter	Cancel received disconnect commands. Truncates load drop.	•	↓

Features 1 and 2 in Table 3 will increase the TTC, which in turn provides extra time for an event to be detected and for a response to execute. Features 3 and 4 represent potential methods for detecting a load drop attack, and 5 and 6 are finally methods by which the attack can be directly addressed. Although the implementation of any controls will involve further characterization of attack scenarios, this provides an initial trade space to address attacks of this class in AMI networks, and puts earlier work in a more organized context.

5. Conclusions And Future Work

In this work, we applied a set of methods and concepts from systems disciplines to a cyber attack scenario in the smart grid. This work demonstrates an application of systems engineering to cyber-physical security problems, with the result that we were able to derive useful performance metrics for the design and evaluation of risk reductions, in a systematic way. Future directions will include expanding the analysis to other failure effects of interest, and developing the mitigations in more detail. While the results of the simulations performed for this work are not directly applicable to the protection of deployed smart grid systems, it is our hope that the methods and concepts discussed here provide some new guidance in approaching the risks that utilities face with their smart grid upgrades.

Acknowledgements

We would like to acknowledge our colleagues at USC, JPL, and LADWP for discussions and support in the work discussed here. We would like to specifically acknowledge Kymie Tan, Arun Viswanathan, and Clifford Neuman for their close collaboration and support.

This research was carried out at the Information Sciences Institute, University of Southern California, and at the Jet Propulsion Laboratory, California Institute of Technology, through an agreement with the National Aeronautics and Space Administration.

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. N66-001-10-C-2018. Neither the United States Government nor any agency thereof, the Los Angeles Department of Water and Power, nor any of their employees make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, complete- ness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or any agency thereof. Figures and descriptions are provided by the authors and used with permission.

References

1. Almajali A, Rice E, Viswanathan A, Tan K, Neuman C, A Systems Approach to Analysing Cyber-Physical Threats in the Smart Grid, *IEEE International Conference On Smat Grid Communcations*, 2013
2. Federal Energy Regulatory Commission, *2013 Assessment of Demand Response and Advanced Metering Staff Report*, 2013
3. Cisco White Paper, *Cisco Connected Grid Security for Field Area Network*, 2012
4. Western Electricity Coorinating Council, *NERC/WECC Planning Standards*, 2012
5. Sridhar S, Hahn A, Govindarasu M, Cyber–Physical System Security for the Electric Power Grid, *Proceedings of the IEEE*, Vol. 100, No. 1, January 2012, DOI 10.1109/JPROC.2011.2165269
6. Clements SL, Kirkham H, Elizondo M, Lu S, Protecting the Smart Grid: A Risk Based Approach, Power and Energy Society General Meeting, 2011 IEEE, DOI 10.1109/PES.2011.6039120
7. Temple WG, Chen B, Tippenhauer NO, Delay Makes a Difference: Smart Grid Resilience Under Remote Meter Disconnect Attack, *IEEE International Conference On Smat Grid Communcations Symposium - Smart Grid Cyber Security and Privacy*, October 2013
8. Anderson R, Fuloria S, Who Controls the Off Switch?, *IEEE International Conference On Smat Grid Communcations*, 2010 DOI 10.1109/SMARTGRID.2010.5622026
9. Bayuk J, Systems Security Engineering, *IEEE Security and Privacy*, Vol 9, Issue: 2, March/April 2011, DOI 10.1109/MSP.2011.41
10. Bayuk J, Horowitz B, An Architectural Systems Engineering Methodology for Addressing Cyber Security, *Wiley Online*, DOI 10.1002/sys.20182
11. Bayuk J, Barnabe D, Goodnight J, Hamilton D, Horowitz B, Neuman C, Tarchalski S, *Systems Security Engineering*, Systems Engineering Research Center Techical Report No. SERC-2010-TR-005
12. Duggan R, Snell M, Uncertainty in Security: Using Systems Engineering Approaches for Robust System Security Requirements, *INCOSE Insight*, Vol 16 issue 2, p. 28-30
13. Woody C, Mission Thread Security Analysis: A Tool for Systems Engineers to Characterize Operational Security Behavior, *INCOSE Insight,* Vol 16 issue 2, p. 37-40
14. Avizienis A, Laprie JC, Randell B, Landwehr C, Basic Concepts and Taxonomy of Dependable and Secure Computing, IEEE Transactions On Dependable and Secure Computing, Vol 1, No 1, January-March 2004
15. Johnson S, Day J, Conceptual Framework for a Fault Management Design Methodology, *AIAA Infotech Conference,* 2010.
16. Johnson S, Day J, System Health Management Theory and Design Strategies, *AIAA Infotech Conference*, 2011.
17. Fovino I, Masera M, De Cian A, Integrating Cyber Attacks Within Fault Trees, Reliability Engineering and System Safety 94 (2009), DOI 10.1016/j.ress.2009.02.020
18. Electric Power Research Institute Technical Report, *EPRI Power Systems Dynamics Tutorial*. EPRI, Palo Alto, CA: 2009. 1016042.
19. IEEE Standards Board, IEEE Recommended Practice for Monitoring Electric Power Quality, IEEE1159-2009, 2009