Intrusion Detection for Advanced Metering Infrastructures: Requirements and Architectural Directions

Robin Berthier, William H. Sanders, and Himanshu Khurana
Coordinated Science Laboratory, Information Trust Institute, and
Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign
{rgb, whs, hkhurana}@illinois.edu

***Abstract*—The security of Advanced Metering Infrastructures (AMIs) is of critical importance. The use of secure protocols and the enforcement of strong security properties have the potential to prevent vulnerabilities from being exploited and from having costly consequences. However, as learned from experiences in IT security, prevention is one aspect of a comprehensive approach that must also include the development of a complete monitoring solution. In this paper, we explore the practical needs for monitoring and intrusion detection through a thorough analysis of the different threats targeting an AMI.**

AMI, IDS, threat model, network architecture

I. INTRODUCTION

The growth in the number of smart meter deployment initiatives in the world indicates that Advanced Metering Infrastructures (AMIs) are being intensely developed, e.g., in the U.S. [19]. This rapid growth is accompanied by important security concerns about the potential vulnerabilities of the new technologies being introduced. These concerns have been fueled by recent press releases about security flaws found in multiple metering devices [8, 21]. To address the challenge of building a secure AMI, the National Institute for Standards and Technology (NIST) and users' groups, such as the Open Smart Grid [22], have been producing reports and requirements to enable technology and policy makers to include security from the very beginning of the development process. These documents range from risk assessment [23] to security requirements [24] and are completed by additional resources such as academic research publications [2, 3, 4] and attack testing methodology [25].

The different solutions required to build a secure architecture belong to three classes: 1) prevention, 2) detection, and 3) mitigation or resilience. In this paper, we are interested in the detection aspect of the problem, especially in the monitoring of the different communication networks of the AMI. Our objective is to precisely define the requirements for an efficient and complete network and host intrusion detection system. To reach this objective, we address the following questions:

1. What is the threat model of the AMI, and how can attacks manifest in the different communication networks?
2. Which components need to be monitored, and at which layer of the protocol stack?
3. What are the unique constraints of the AMI, which detection technology should be used, and which monitoring architecture should be deployed?

This material is based upon work supported by the Department of Energy under Award Number DE-OE0000097.

The contributions of this study consist in addressing these questions through a review of the literature and a discussion about existing IT solutions and future research areas. This paper is organized as follows. We present an AMI in Section II and detail the related threat model in Section III. We describe the components of an IDS in Section IV, and discuss the monitoring requirements of a comprehensive intrusion detection architecture in Section V. We conclude with a review of the related literature in Section VI and a discussion about the next steps in Section VII.

II. AMI REVIEW

As shown in Figure 1, an AMI includes several communication networks, identified according to their spatial scope:

- The *Wide Area Network* (WAN) serves as a communication link between headends in the local utility network and either data concentrators or smart meters. This network uses long-range and high-bandwidth communication technologies, such as WiMAX, cellular (3G, EVDO, EDGE, GPRS, or CDMA), satellite, Power Line Communication (PLC), and Metro Ethernet. The scale of this network could reach several million nodes.
- *Neighborhood Area Networks* (NANs) ensure communication between data concentrators or access points and smart meters that play the role of interfaces with a *Home Area Network* (HAN). The scale of this network ranges from a few hundred to tens of thousands of nodes.
- Finally, *Field Area Networks* (FANs) allow the utility workforce to connect to equipment in the field.

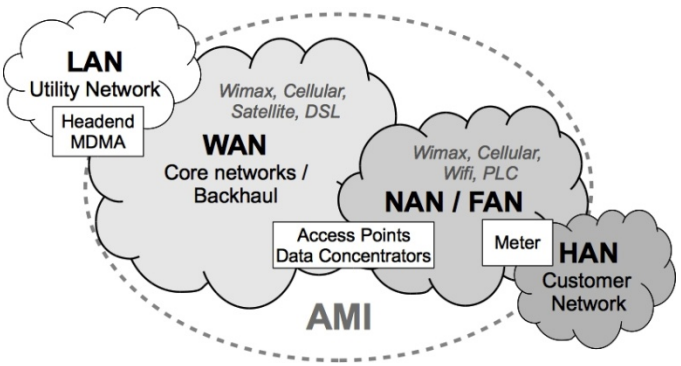

Figure 1: Overview of the AMI Networks

Table 1: Attack Techniques and Consequences

Attack Techniques	Attack Consequences
Network Compromise	
Communication interception and traffic analysis [1, 3, 6]	Integrity of configuration and routing operations, inconsistent traffic origin or destination
Traffic modification, injection, and replay [1, 2, 3]	Integrity of communication traffic, illegitimate network operations, inconsistent traffic origin or destination
System Compromise	
Authorization or authentication violation [1, 5, 6, 7, 8]	Illegitimate network operations, inconsistent traffic origin or destination, illegitimate use of credentials
Spoofing of utility system	Illegitimate network operations, inconsistent traffic origin or destination
Compromise node, spoofing of metering device [1, 2, 3, 4, 5, 6, 7, 8]	Integrity of node software or hardware, illegitimate network operations, inconsistent traffic origin or destination
Denial of Service	
Resource exhaustion [1, 2, 7]	Unresponsive nodes, high bandwidth usage
Signal jamming [2, 5, 7]	Unresponsive nodes, high signal power level
Dropping packets [2]	Integrity of communication traffic

A major challenge in protecting an AMI against malicious activities is to create a monitoring solution that covers the heterogeneity of communication technologies through their requirements (e.g., encryption and real time) and constraints (e.g., topology and bandwidth). It is critical to identify these elements, for two reasons: 1) they can help to define the potential impact of malicious activities targeting the AMI; and 2) they can impose limits on the functionalities and security of a monitoring solution. For instance, the fact that large portions of an AMI network are wireless and use a mesh network topology facilitates network-related attacks such as traffic interception, and the design of the monitoring architecture is more challenging than in a traditional wired network. Moreover, a large number of nodes are deployed in the field or in consumer facilities, which means that attacks requiring physical access are easier to conduct.

III. THREAT MODEL

The threats targeting an AMI can be viewed in different ways: by type of attacker, by motivation, and by attack technique. We explore in this section each perspective through a survey of the literature. Our goal is to understand how the different attacks will likely manifest themselves in an AMI, and, as a result, how we can detect them.

We use the terminology from [26] to list the types of attackers and their motivations. Six types of attackers are considered:

- *Curious eavesdroppers*, who are motivated to learn about the activity of their neighbors by listening in on the traffic of the surrounding meters or HAN.
- *Motivated eavesdroppers*, who desire to gather information about potential victims as part of an organized theft.
- *Unethical customers*, who are motivated to steal electricity by tampering with the metering equipment installed inside their homes.
- *Overly intrusive meter data management agencies*, which are motivated to gain high-resolution energy and behavior profiles about their users, which can damage customer privacy. This type of attacker also includes employees who could attempt to spy illegitimately on customers.
- *Active attackers*, who are motivated by financial gain or terrorist goals. The objective of a terrorist would be to create large-scale disruption of the grid, either by remotely cutting off many customers or by creating instability in the distribution or transmission networks. Active attackers attracted by financial gain could also use disruptive actions, such as Denial of Service (DoS) attacks, or they could develop self-propagating malware in order to create revenue-making botnets.
- *Publicity seekers*, who use techniques similar to those of other types of attackers, but in a potentially less harmful way, because they are more interested in fame and usually have limited financial resources.

Attackers may use a variety of attack techniques to reach their objectives. Based on a survey of the related literature, we categorize these techniques in Table 1. The right column of Table 1 provides information about the attack consequences and will be used in the next section to identify the monitoring mechanisms required for an intrusion detection system.

Table 1 provides a view of threats that is more detailed than the list of attacker types and motivations in [26], but it still offers only a high-level classification of possible malicious activities. At a lower level, vulnerabilities can be found and exploited in the following components:

- For *network compromise* and *denial of service*: flaws or misuses of routing, configuration, name resolution, encryption, or authentication protocols.
- For *system compromise*: software and firmware vulnerabilities, hardware vulnerabilities, read and write access to data storage, or access to encryption keys.

More information about low-level system vulnerabilities and attacks can be found in [7, 8, 25]. In the next section, we describe the components and properties to take into consideration in designing an intrusion detection system that could reliably detect when such vulnerabilities are exploited.

IV. COMPONENTS OF AN IDS

The concept of intrusion detection systems was introduced more than two decades ago and has evolved into an important body of work. A taxonomy has been defined [28], a set of detection techniques and architectures has been created, and caveats and limitations have been carefully studied [29]. Our challenge is to best apply this knowledge to the AMI in order to design a highly efficient monitoring solution that covers our threat model and respects industry-strength requirements while being practical and useful. We believe that the best strategy to achieve this goal is to review existing intrusion detection solutions in the context of the AMI. In doing so, we can leverage possible detection techniques and monitoring architectures with respect to the capabilities and constraints of the AMI. In this process we will identify research areas that need to be investigated further. We discuss these findings at the end of Section V.

A. Definitions and Challenges

Intrusion detection is defined by [27] as the process of monitoring the events that occur in a computer system or network and analyzing them for signs of possible incidents. The components of a traditional IDS are: 1) *sensors* or *agents* to monitor and analyze activity; 2) a *management server* to centralize information collected by the sensors or agents and to manage them; 3) a *database server* to store all the data produced by the IDS; and 4) a *console* to provide an interface to users and administrators so that they can check the status of the system monitored, receive alerts, investigate events, and configure the system.

In the context of the AMI, the major challenges of security solutions are to be robust and to integrate seamlessly with system operations. For an IDS, these challenges translate into two constraints: 1) to be highly accurate in detecting security incidents including unknown attacks; and 2) to have a low overhead on the infrastructure and a minimum impact on management processes.

We review existing detection technologies in the following subsection and explain why we believe a specification-based approach has the potential to meet these constraints.

B. Detection Technologies

The process of detecting malicious activity can be based on three distinct approaches:

- *Signature-based detection* (also known as *misuse detection*) consists in looking for patterns of malicious behavior using a database of predefined attack signatures;
- *Anomaly-based detection* consists in identifying deviations from a normal behavior profile predefined using statistical measures; and
- *Specification-based detection* consists in identifying deviations from a correct behavior profile predefined using logical specifications.

These detection techniques are different for two fundamental reasons. First, signature-based IDS uses a blacklist approach, while anomaly- and specification-based IDSes use a whitelist approach. A blacklist approach requires creation of a knowledge base of malicious activity, while a whitelist approach requires training of the system and identification of its normal or correct behavior. Obviously these two approaches have complementary limitations and advantages: a blacklist-based IDS will not be able to detect unknown attacks and will require frequent updates, while a whitelist-based IDS will often be more expensive to train and to tune. Another limitation of whitelist approaches is that they provide little information about the root causes of attacks.

A second fundamental difference lies in the level of understanding required by each approach. Signature- and anomaly-based IDSes belong to the same group by monitoring activity at a low level, while specification-based IDS requires a high-level and stateful understanding of the activity monitored. For example, in the context of a network IDS (NIDS), signature- and anomaly-based IDSes will monitor network activity at layers 2, 3, and 4 of the OSI model. A signature-based IDS, such as Snort, will also analyze payload information at layers 5 and 7, but in a stateless manner, without tracking the behavior of the underlying applications over time. On the other hand, a specification-based IDS will work by building a state machine of a process, and then monitoring activity to check whether anything escapes from the system boundaries specified previously. This second approach is typically more expensive to develop and less scalable than the first approach. However, it has the strong advantage of being more accurate.

In the context of the AMI, we believe that the best approach is to develop a specification-based IDS. This choice is founded on three arguments. First, the potentially higher accuracy of specification-based IDS compared to signature-based IDS is better suited for the criticality of the AMI. Second, the lack of empirical attack data for the AMI makes the construction of a blacklist of signatures difficult. Third, the limited number of protocols and applications to monitor in the AMI reduces the development cost of specification-based sensors. This last argument is important, because it explains why specification-based IDSes have been applied to specific problems, such as mobile ad hoc networks [12], STP protocols [30], and VoIP protocols [31, 32]. It would be impossible to specify state machines for a large system made of thousands of different protocols, such as the Internet. But the AMI offers a controlled environment in which the task of developing specifications would be cost-efficient.

C. Network and Host Intrusion Detection Systems

Intrusion detection sensors can monitor not only network traffic, but also system information such as file integrity, application logs, and system calls. Compared to NIDSes, these Host Intrusion Detection Systems (HIDSes) have often the advantage of being able to access the root cause of attacks, such as a malicious system process. However, they have two important limitations: 1) they usually add a significant overhead to the local system they monitor, and 2) they are easier to compromise than NIDSes. We note that solutions such as virtualization have been developed to cope with the second limitation [33].

As shown in Table 1, an AMI requires both types of sensor, but to instrument low computation environments such as smart meters with HIDSes is still an open research issue.

V. Proposed AMI Monitoring Architecture

A. Design Considerations

A traditional IDS architecture is made of a collection of lightweight sensors that report to a centralized management server. The core data processing and detection intelligence of this architecture resides on the central component. The main limitation of this approach lies in the difficulty of making it scalable. When dealing with an AMI network that have millions of nodes, this type of architecture is not the optimal way to handle the load.

An alternative solution is to use a distributed architecture in which most of the data processing is handled by the sensors. This type of architecture also includes a central component, but it only has the tasks of coordinating sensors and collecting high-level alerts.

Table 2: Detection Mechanisms and Sensor Locations Classified per Detection Operation

Attack Consequences	Detection Goal and Operation	Type of Sensor, Locations, and References	Protocol Layers (OSI Model)
Stateful Specification-based Monitoring			
Integrity of configuration and routing protocols	Checking of configuration and routing operations against security policy and network configuration	Behavioral/finite state machine monitor, distributed on network nodes [12, 13, 14] or on isolated sensors	3-4
Illegitimate network operations	Stateful checking of protocol operations against security policy and application configurations	Behavioral/finite state machine monitor, centralized on access points [10]	5-7
Stateless Specification-based Monitoring			
Inconsistent traffic origin or destination	Checking of packet header against security policy and network configuration	Firewall log monitor, centralized on access points	3-4
Integrity of communica-tion traffic	Checking of packet payload against protocol specifications	Data validation (e.g., range check), centralized on access points or distributed on isolated sensors [2]	3-7
Illegitimate use of credentials	Checking of system logs against security policy	Authentication log monitor, centralized on access points or distributed on network nodes	5-7
Integrity of node software or hardware	Operating system, application, and file integrity checking	Remote attestation and virtualization, distributed on network nodes [4]	-
Unresponsive nodes	Checking of protocol operations against security policy and application configuration	Physical health report monitor, centralized on access points or distributed on isolated sensors [2]	2-7
Anomaly-based Monitoring			
High bandwidth usage	Traffic monitoring against normal statistical profiles	Threshold monitor (e.g. packet count), centralized on access points or distributed on isolated sensors [2]	3-4
High signal power level	Checking of wireless signal against normal statistical profiles	Threshold monitor, distributed on network nodes or on isolated sensors [2]	1-2

Due to the low computation environment provided by most AMI components, an hybrid architecture consisting of a set of decentralized alert aggregators would be recommended. We believe access points can host the data processing tasks while reporting high level information to a central management server installed in the utility network.

In addition to scalability, a second critical requirement is the reliability of the IDS against accidental failures and malicious attacks. The system needs to operate even if a subset of sensors or even the management server are unavailable or compromised. Failures of, or attacks against, the management server are the most problematic issues, but they can be addressed by the use of redundancy to eliminate single points of failure [34]. Attacks against sensors are greatly limited if sensors are isolated through virtualization or by using their own hardware rather than being embedded on the components they monitor. Techniques to detect compromised nodes include the use of a reputation system to evaluate the level of trust of alert reports [35] or the use of a distributed proof system to prevent a single node from having too much visibility [18]. Finally, the reliability of communications between sensors and management servers can be increased through the use of a separate communication network.

B. Alert Management and Response Mechanism

An important concern with any large-scale monitoring solution is the volume of alerts generated. The lack of a comprehensive alert management process will undermine the usefulness and seamless integration of an IDS. This process consists in two steps: data reduction (including data aggregation) and alert correlation. The first step groups alerts that share similar attributes, while the second step tries to extract high-level information about the intrusions detected. This second step uses correlation rules to map alerts over time, space, and logical sequence [36]. This process allows an IDS to greatly reduce the number of alerts that are communicated to a human operator. Moreover, it can calculate a criticality value for each alert, enabling escalation and prioritization procedures.

An IDS can also be coupled with a set of response mechanisms in order to translate passively generated alerts into automated actions. The possible actions of such Intrusion Prevention Systems (IPS) include blocking a malicious connection, changing the configuration of a firewall, or restoring an application into a clean state [11]. Sophisticated automated response techniques are explored in [20].

C. Proposed Monitoring Solution

Based on the threats described in Section III, we describe in Table 2 the detection mechanisms required to monitor the AMI using a specification-based approach. The third and fourth columns of Table 2 provide information about the type and location of sensors and the protocol layers that need to be monitored. The location includes network nodes (e.g., smart meters), isolated sensors, and access points. An isolated sensor is a security device that can be deployed in a wireless network to monitor the network traffic and report when incidents are detected [2]. Isolated sensors are expensive to deploy and maintain, but they offer higher computation resources and better protection against security compromises than typical network nodes, such as smart meters, do.

We divided Table 2 into three sections: stateful specification-based monitoring, stateless specification-based monitoring, and anomaly-based monitoring. These three categories provide information about the computation resources needed by the different monitoring operations, arranged from high to low. Some monitoring operations work by checking system and network behavior against configurations, policies, and protocol specifications. These resources have to be defined as part of the IDS design process in close collaboration with the developers and users of the AMI systems and networks. We note that techniques have been developed to formally prove that sensors with a local view of the network activity can cover globally defined specifications [14]. For the detection operation highlighted in Table 2, the different resources required are:

- A *network configuration* to provide information about network topology and access control rights;
- *Protocol specifications* to determine correct header and payload formats as well as the behavioral state machines that can capture coherent sequences of request and reply operations;
- *System and network security policies* to specify the allowed behaviors of applications and processes; and
- *Statistical profiles* to determine boundaries of normal network traffic characteristics.

In contrast to configurations, specifications, and policies, profiles of normal behaviors must be tuned through a training phase. The dependency between this training phase and the detection accuracy varies according to the data mining techniques used to perform the detection [9].

To illustrate the monitoring requirements identified in Table 2, we represent them in an AMI in Figure 2. This example assumes that the WAN is a wired network and the NAN is a wireless network where metering devices are organized in a mesh topology and are authenticated by access points between the two networks. These gateways route and control traffic between the utility network and the meters. They have sufficient computing resources and memory to host the core stateful specification-based detection technology. They can also aggregate alerts sent by the metering devices. The embedded hardware of meters cannot support sophisticated monitoring function but they can periodically check the integrity of their software and send health reports to their access points. To prevent a compromised meters from spreading unnoticed within the NAN, a set of IDS sensors with specification-based monitoring capabilities is deployed in key areas. We review in the next section a list of research issues that we think are crucial for the development of this type of monitoring architecture.

D. Open Research Issues

As mentioned in Section II, an AMI requires communication to be encrypted. A key management solution [16, 17] needs to be deployed not only to provide keys to operational components but also to interact with monitoring solution that analyzes encrypted payloads in between nodes.

Another typical constraint of an AMI is the use of a low-bandwidth communication medium. If the IDS uses the same network to communicate, the overhead incurred has to be carefully evaluated, as well as potential denial of service vulnerabilities against the AMI components or the IDS itself.

A third important issue to investigate is the instrumentation of AMI components with specification-based capabilities. Running a stateful specification-based IDS on hardware having limited computation power and memory is a critical challenge.

A final issue related to specification-based IDS is the development of the specifications. Even if the number of protocols used within an AMI is reduced, the complexity of some of them, such as C12.19 and C12.22, incurs a high development cost. We note that solutions exist to partially automate this process [15].

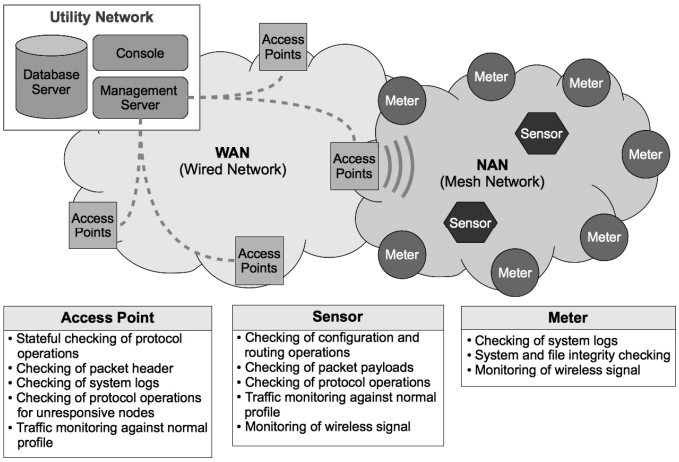

Figure 2: Example of IDS Architecture with Monitoring Operations Detailed per Component

VI. Related Work

In addition to the publications referenced in the previous sections, we detail in this section academic work directly related to the energy industry and to security.

In [1], the author discusses the security requirements and related threats of the four main components of an AMI: smart meters, the customer gateway, the communication network, and the headend. The fact that encryption and authentication alone will not be sufficient to protect the infrastructure is emphasized.

In [2], the authors propose a model-based IDS working on top of the WirelessHART protocol, to monitor and protect wireless process control systems. The hybrid architecture consists of a central component that collects information periodically from distributed field sensors. A set of 8 detection rules working on the physical, data-link, and routing layers cover threats including signal jamming, node compromise, and packet modification.

In [3], the authors focused on the issue of energy theft. They explained, through a detailed threat and security analysis, that the AMI will significantly increase the risk of energy theft. The main reasons are 1) amplification of effort, 2) division of labor, and 3) extended attack surface.

In [4], the authors present an architecture called the Cumulative Attestation Kernel to address the issue of securely auditing firmware updates in embedded systems such as smart meters. The system is designed to be cost-, power-, computation-, and memory-efficient. A prototype is implemented to demonstrate the feasibility of the solution as well as to formally prove that it meets remote attestation requirements.

In comparison to our approach, [1] and [3] focus on describing the security challenges faced by AMIs, while [2] and [4] provide solutions to the specific issues of, respectively, wireless process control systems and remote meter attestation. By describing an IDS, [2] is the closest to our work, but, as shown in Table 2, it does not cover all the monitoring requirements of an AMI.

VII. Conclusion and Future Work

The development of a practical and efficient IDS for an AMI is highly critical. We presented in this paper the requirements and a possible architecture for a comprehensive monitoring solution. Based on a review of the different threats targeting the different AMI components, we surveyed the literature to understand the appropriate detection technologies to deploy and to coordinate. We believe that specification-based detection technology has the potential to meet the industry-strength requirements and constraints of an AMI. However, such technology incurs a high development cost. We are currently investigating techniques to lower this cost and to implement and evaluate this technology on AMI components.

References

[1] F. M. Cleveland, "Cyber Security Issues for Advanced Metering Infrastructure (AMI)," Proceedings of the IEEE Power and Energy Society General Meeting: Conversion and Delivery of Electrical Energy in the 21st Century, pp. 1–5, 2008

[2] T. Roosta, D. K. Nilsson, U. Lindqvist, and A. Valdes, "An Intrusion Detection System for Wireless Process Control Systems," Proceedings of the 5th IEEE International Conference on Mobile Ad Hoc and Sensor Systems (MASS), pp. 866–872, 2008

[3] S. McLaughlin, D. Podkuiko, and P. McDaniel, "Energy Theft in the Advanced Metering Infrastructure," Proceedings of the 4th Workshop on Critical Information Infrastructures Security (CRITIS), 2009

[4] M. LeMay and Carl A. Gunter, "Cumulative Attestation Kernels for Embedded Systems," Proceedings of the 14th European Symposium on Research in Computer Security (ESORICS), pp. 655–670, 2009

[5] W. Sikora, M. Carpenter, and J. Wright, Smart Grid and AMI Security Concerns, InGuardians and Industrial Defender, 2009

[6] M. Davis, SmartGrid Device Security, Presentation at BlackHat 2009

[7] T. Goodspeed, D. R. Highfill, and B. A. Singletary, "Low-level Design Vulnerabilities in Wireless Control System Hardware," Proceedings of the Scada Security Scientific Symposium (S4), 2009

[8] J. Wright (InGuardians), "Smart Meters Have Security Holes," http://www.msnbc.msn.com/id/36055667, 2010

[9] P. García-Teodoroa, J. Díaz-Verdejoa, G. Maciá-Fernándeza, and E. Vázquez, "Anomaly-based Network Intrusion Detection: Techniques, Systems and Challenges," Computers & Security, vol. 28, no. 1–2, pp. 18–28, 2009

[10] F. Hugelshofer, P. Smith, D. Hutchison, and N.J.P. Race, "OpenLIDS: A Lightweight Intrusion Detection System for Wireless Mesh Networks," Proceedings of the 15th Annual International Conference on Mobile Computing and Networking, pp. 309–320, 2009

[11] I. Balepin, S. Maltsev, J. Rowe, and K. Levitt, "Using Specification-based Intrusion Detection for Automated Response," Proceedings of the Recent Advances in Intrusion Detection (RAID), pp. 136–154, 2003

[12] H.M. Hassan, M. Mahmoud, and S. El-Kassas, "Securing the AODV Protocol using Specification-based Intrusion Detection," Proceedings of the 2nd ACM International Workshop on Quality of Service & Security for Wireless and Mobile Networks, p. 36, 2006

[13] R. Gill, J. Smith, and A. Clark, "Specification-based Intrusion Detection in WLANs," Proceedings of the 22nd Annual Computer Security Applications Conference (ACSAC), pp. 141–152, 2006

[14] T. Song, C. Ko, C. Tseng, P. Balasubramanyam, A. Chaudhary, and K. Levitt, "Formal Reasoning about a Specification-based Intrusion Detection for Dynamic Auto-configuration Protocols in Ad Hoc Networks," Journal of Formal Aspects in Security and Trust, pp. 16–33, 2005

[15] N. Stakhanova, S. Basu, and J. Wong, "On the Symbiosis of Specification-based and Anomaly-based Detection," Journal of Computers & Security, vol. 29, no. 2, pp. 253–268, 2010

[16] W. He, Y. Huang, R. Sathyam, K. Nahrstedt, and W. C. Lee, "SMOCK: A Scalable Method of Cryptographic Key Management for Mission-Critical Wireless Ad-Hoc Networks," IEEE Transactions on Information Forensics and Security, vol. 4, no. 1, pp. 140–150, March 2009

[17] R. Sathyam, Analysis of a Key Management Scheme for Wireless Mesh Networks, M.S. thesis, University of Illinois at Urbana-Champaign, 2008.

[18] A. J. Lee, K. Minami, and M. Winslett, "Lightweight Consistency Enforcement Schemes for Distributed Proofs with Hidden Subtrees," Proceedings of the 12th ACM Symposium on Access Control Models and Technologies (SACMAT 2007), Sophia Antipolis, France, pp. 101–110, June 20-22, 2007

[19] Enernex Smart Meter Data for California Energy Commission, http://smartmeterpedia.synthasite.com/Enernex_Map.php, 2010

[20] S.A. Zonouz, H. Khurana, W.H. Sanders, and T.M. Yardley, "RRE: A Game-Theoretic Intrusion Response and Recovery Engine," Proceedings of the 39th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2009), Estoril, Lisbon, Portugal, June 29-July 2, 2009, pp. 439–448

[21] "IOActive's Mike Davis to Unveil Smart Grid Research at Black Hat USA," IOActive press release, July 28, 2009, http://www.ioactive.com/news-events/DavisSmartGridBlackHatPR.php

[22] Open Smart Grid Users Group, http://osgug.ucaiug.org, 2010

[23] AMI-SEC Task Force, "AMI Risk Assessement," http://osgug.ucaiug.org/utilisec/amisec/Shared%20Documents/0.%20AMI%20Risk%20Assessment/, 2010

[24] Smart Grid Interoperability Panel, Cyber Security Working Group, "Draft NISTIR 7628: Smart Grid Cyber Security Strategy and Requirements" (second draft), http://csrc.nist.gov/publications/, NIST (National Institute of Standards and Technology), February 2010

[25] M. Carpenter, T. Goodspeed, B. Singletary, E. Skoudis, and J. Wright, "Advanced Metering Infrastructure Attack Methodology" version 1.0, http://inguardians.com/pubs/AMI_Attack_Methodology.pdf, Jan. 5, 2009

[26] M. LeMay, G. Gross, C. A. Gunter, and S. Garg, "Unified Architecture for Large-scale Attested Metering," Proceedings of the 40th Annual Hawaii International Conference on System Sciences (HICSS), p. 115, 2007

[27] K. Scarfone and P. Mell, I, http://csrc.nist.gov/publications/nistpubs/800-94/SP800-94.pdf, NIST (National Institute of Standards and Technology) special publication 800-94, 2007

[28] H. Debar, M. Dacier, and A. Wespi, A Revised Taxonomy for Intrusion-detection Systems, Annals of Telecommunications, vol. 55, no. 7, pp. 361–378, 2000

[29] T.H. Ptacek and T.N. Newsham, "Insertion, evasion, and denial of service: Eluding network intrusion detection", Technical report, Secure Networks Inc., 1998

[30] P. Jieke, J. Redol, and M. Correia, "Specification-Based Intrusion Detection System for Carrier Ethernet," Proceedings of the International Conference on Web Information Systems and Technologies (WEBIST), 2007

[31] T. Phit and K. Abe, "A Protocol Specification-based Intrusion Detection System for VoIP and Its Evaluation," IEICE Transactions on Communications, vol. 91, no. 12, pp. 3956–3965, 2008

[32] H. Sengar, D. Wijesekera, H. Wang, and S. Jajodia, "VoIP Intrusion Detection through Interacting Protocol State Machines," Proceedings of the International Conference on Dependable Systems and Networks (DSN), pp. 393–402, 2006

[33] M. Laureano, C. Maziero, and E. Jamhour, Intrusion detection in virtual machine environments, Proceedings of the EUROMICRO conference, pp. 520—525, 2004

[34] J.S. Balasubramaniyan, J.O. Garcia-Fernandez, D. Isacoff, E. Spafford, and D. Zamboni, "An Architecture for Intrusion Detection Using Autonomous Agents," Proceedings of the 14th IEEE Computer Security Applications Conference, pp. 13—24, 1998

[35] S. Buchegger and J. L. Boudec, "Performance Analysis of the CONFIDANT Protocol: Cooperation of Nodes — Fairness in Dynamic Ad-hoc Networks." Proceedings of IEEE/ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc), Lausanne, Switzerland, June 2002

[36] H. Debar and A. Wespi, "Aggregation and Correlation of Intrusion-detection Alerts," Proceedings of the Recent Advances in Intrusion Detection (RAID), pp. 85–103, 2001