

AMI Threats, Intrusion Detection Requirements and Deployment Recommendations

David Grochocki, Jun Ho Huh, Robin Berthier, Rakesh Bobba, Alvaro A. Cárdenas and Jorjeta G. Jetcheva and William H. Sanders

Information Trust Institute, Coordinated Science Laboratory,
and Electrical and Computer Engineering Department
University of Illinois at Urbana-Champaign
{dgrocho2,jhhuh,rgb,rbobba,whs}@illinois.edu

Fujitsu Laboratories of America, Inc
{alvaro.cardenas-mora,jjjetcheva}@us.fujitsu.com

Abstract—Advanced Metering Infrastructures (AMI) facilitate bidirectional communication between smart meters and utilities, allowing information about consumption, outages, and electricity rates to be shared reliably and efficiently. However, the numerous smart meters being connected through mesh networks open new opportunities for attackers to interfere with communications and compromise utilities’ assets or steal customers’ private information.

The goal of this paper is to survey the various threats facing AMIs and the common attack techniques used to realize them in order to identify and understand the requirements for a comprehensive intrusion detection solution. The threat analysis leads to an extensive “attack tree” that captures the attackers’ key objectives (e.g., energy theft) and the individual attack steps (e.g., eavesdropping on the network) that would be involved in achieving them. With reference to the attack tree, we show the type of information that would be required to effectively detect attacks. We also suggest that the widest coverage in monitoring the attacks can be provided by a hybrid sensing infrastructure that uses both a centralized intrusion detection system and embedded meter sensors.

I. INTRODUCTION

The introduction of a new metering infrastructure to energy delivery systems is a significant change that requires a tremendous amount of planning. The importance of this upgrade is defined by its magnitude (millions of meters have to be replaced) and advanced capabilities (e.g., two-way communications for all devices). Among the planning efforts required, design of the right security foundation is a critical one to ensure that the infrastructure will reach an acceptable level of resiliency against a wide array of threats. In particular, it is essential to identify the requirements for a comprehensive monitoring solution that would enable utilities to gain situational awareness over the security state of their infrastructure.

Several initiatives to understand the threat landscape of Advanced Metering Infrastructures (AMIs) have been conducted over the past few years, addressing topics that range from high-level threat models [1] to specific attack scenarios [2] and even experimentally tested attack techniques [3]. However, there is still a gap between the identification of threats and the specification of a comprehensive monitoring solution. Utilities need to understand the risks of AMI deployments and the requirements for intrusion detection before they choose the monitoring architecture in which to invest.

Fleury et al. [4] explore a comprehensive set of threats against energy control systems, but not threats specific to the AMI and mesh network settings. Berthier and Sanders [1] cover threats targeting AMI but at a high level. They provide preliminary insights into requirements for IDSes, but they do not detail the types of information that would be needed for detection or discuss the monitoring-coverage tradeoffs inherent in different types of intrusion detection systems.

This paper presents an extensive survey of AMI-specific threats (focusing on wireless in mesh networks) and a detailed mapping to the *information required* for accurate attack detection. Our “attack tree” presented in Section III-C captures attackers’ high level objectives and breaks them down into more fine-grained attack steps, demonstrating how they might be achieved. Based on these findings, we discuss a few possible deployment schemes for IDS, and how effective each might be in detecting the attack steps. In particular, we find that a hybrid sensing infrastructure, whereby a central sensor is instrumented together with distributed meter sensors, would provide the widest detection coverage. The detailed discussion is in Section V.

II. AMI OVERVIEW

The role of an AMI is to enable communication between utility companies and electricity meters, including remote electricity usage readings (on-demand and periodic), sending of updated price information to the meters, transmission of alerts about outages, and upgrades of meter firmware, among other communications. Some messages require real-time delivery, while others can be buffered and delayed without negative consequences. In addition, AMIs have security and privacy requirements, since sensitive customer information is frequently exchanged, and some of them provide a remote disconnect feature. To accommodate the aforementioned requirements and also a wide range of meter deployment topologies, e.g., from dense urban settings to sparse rural environments, meter manufacturers have designed highly flexible network architectures that can include different communication media. Those architectures usually follow the same network hierarchy, such that a wide area network (WAN) connects utilities to a set of gateways in the field, and then neighborhood area networks (NANs), also called field area networks (FANs),

connect gateways to meters. A WAN uses long-range and high-bandwidth communication technologies, such as long-range wireless (e.g., WiMAX), cellular (e.g., 3G, EVDO, EDGE, GPRS, or CDMA), satellite, or Power Line Communication (PLC). NANs typically have shorter range requirements and can be deployed using wireless (e.g., IEEE 802.11, IEEE 802.15, or proprietary) or PLC-based technologies. In some cases, meters can directly include cellular capabilities or even use the customer's home Internet connection to bypass the need for separate WANs and LANs. In this paper, we focus on NANs that use a wireless mesh network. The mesh topology brings robustness to the network, since communication routes can automatically adapt when failures occur. However, they also represent a challenge for the deployment of an efficient security monitoring solution.

III. THREATS ANALYSIS AND ATTACK TREE

The addition of a communication infrastructure and the new computational capabilities of smart grid devices adds a significant attack surface to traditional energy delivery systems. For example, cyber intrusions that would previously have required physical access to the utility network may now be possible through a remote exploit. In the context of AMIs, the fact that smart meters are not only connected to the utility network but also directly accessible by customers enables new attack vectors. Indeed, field area networks in which meters are deployed appear to be an attractive target for adversaries, because they consist of large numbers of physically accessible devices and have limited or no security monitoring capabilities.

The goal of this section is to review the threats and attack objectives that are specific to AMI networks and tie them to individual attack steps. A number of representative case studies are explored to connect the attackers' objectives with more fine-grained, individual attack steps. The results lead to an extensive attack tree and to the identification of the information required for detecting such attacks. Note that our analysis was performed within the scope of AMI networks; access to Home Area Networks (HANs) may yield additional attacker motivations and involve additional attack steps; however, methods of compromising HANs are beyond the scope of this paper.

A. Survey of Previous Literature

The key characteristics of an AMI that could attract malicious activity are 1) access to a communication infrastructure other than the Internet, 2) access to millions of low-computation devices, 3) access to sensitive customer information, 4) high visibility and high impact in the case of disruption (e.g., power outage), and 5) financial value of energy consumption data. Consequently, attackers could be motivated to abuse the communication infrastructure, reduce their energy bills, steal information from targeted customers, remotely disconnect targeted customers or large regions, or create denial-of-critical-services.

A large set of attack techniques can be combined to reach those objectives. We conducted a thorough survey of previous literature from 11 different universities and independent corporations to identify fundamental attack steps. A first category of research we studied did not specifically cover AMIs but were useful in clarifying the threats common to wireless networks.

For instance, [5] examines attacks on wireless networks to motivate solutions to address the privacy issue, [6] developed a threat model to guide the design of a secure WLAN architecture, [7] and [8] study threats on mobile ad-hoc networks (MANET), [9] focuses on sensor networks, and [10] and [11] investigate threats specific to mesh networks.

In the category of publications focusing on the smart grid, [12] presents the design of a firewall to secure wireless communication in energy delivery systems. [2] examines attacks targeting energy theft in AMIs; the authors later used that analysis to motivate a new methodology for penetration testing in AMIs [3]. While those efforts have been important in shedding light on the security issues surrounding wireless mesh networks and AMIs, to the best of our knowledge, there has not been an AMI threat survey detailed and extensive enough to guide the design of a comprehensive security monitoring solution.

Our next step was to combine the attacks discussed in the literature in order to build a holistic view of the AMI attack ecosystem. From an initial list of 5 attack motivations and 30 unique attack techniques, we first filtered out those irrelevant to the AMI environment, and then worked on decomposing the remaining ones into individual attack steps. The motivation for the decomposition was to understand the fundamental pieces of information required by a monitoring solution to detect any combination of those attack steps, including combinations that we did not cover in our threat model. We illustrate the decomposition through the following three case studies, and present the results in a set of attack trees shown in Figure 1.

B. Case Studies

1) *Distributed Denial of Service (DDoS) Attack Against the Data Collection Unit (DCU)*: The attacker's motivation in performing a distributed denial of service (DDoS) attack [13] is to compromise the DCU and prevent relays deployed between the WAN and NANs from communicating or functioning (see Figure 2). Assuming that the entry point of the attack is in the smart meters, the following are typical individual attack steps that would be involved: (1) installation of malware on the meters through physical tampering or exploitation of a network vulnerability; (2) coordination of a DDoS campaign among the compromised meters; and (3) crafting and sending of a large number of malicious network packets to the DCU.

2) *Stealing Customer Information*: The motivation of the second attack is to collect customer information and learn about customer behavior by eavesdropping on the incoming and outgoing network traffic of the meters. Considering that the AMI traffic may be encrypted, this attack may involve the following individual steps: (1) theft of the decryption keys (or the master seed number that is used to generate the keys) accomplished by physically accessing the meters or performing brute-force attacks on the crypto system; (2) eavesdropping on the AMI traffic to intercept the messages; and (3) decryption of the messages and collection of the message contents.

3) *Sending Remote Disconnect Commands Through the Data Collection Unit*: Here, the attacker wishes to disconnect a large number of customers by exploiting the "remote disconnect" functionality on the meters (see Figure 3). The

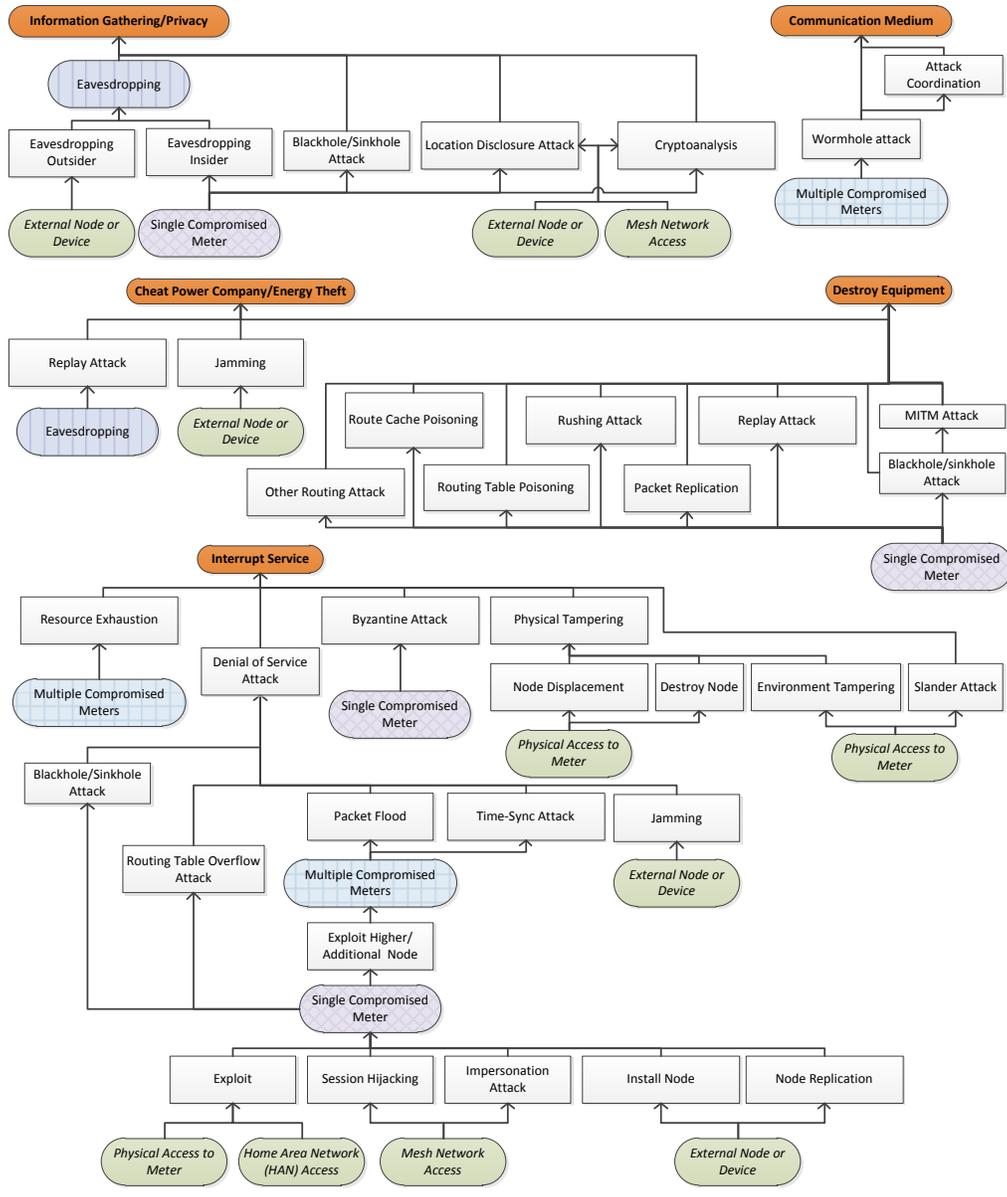


Fig. 1. Attack Trees

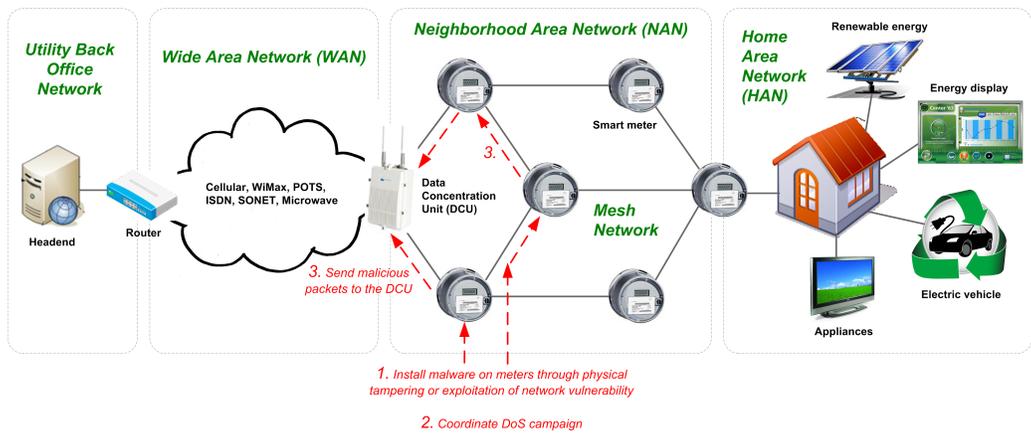


Fig. 2. Distributed Denial of Service Attack Against the DCU

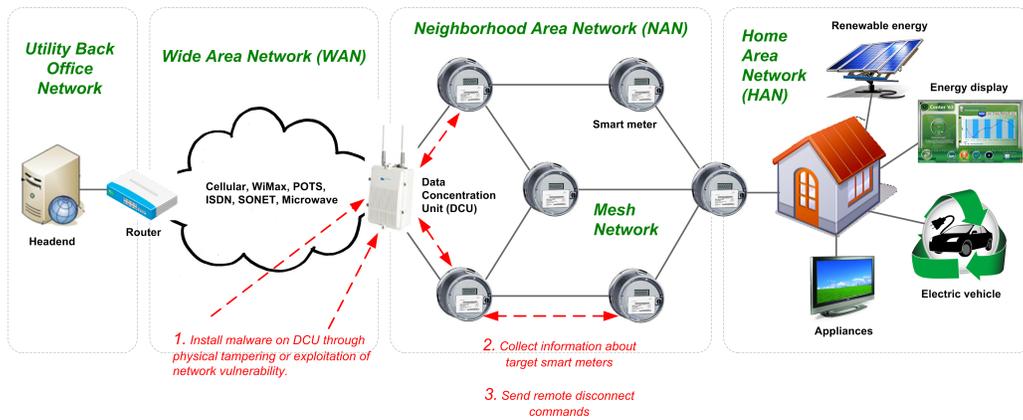


Fig. 3. Remote Disconnect Command Attack

DCU is very likely to be the point for launches of these attacks as it is one of the more suitable devices for triggering the remote disconnect command for many customers without being detected by the utility. The attack steps involved are (1) installation of a malware on the DCU through physical tampering, exploitation of a network vulnerability, or abuse of insider privileges; (2) identification of the meters and collection of information about them (e.g., IP addresses); and (3) sending of remote disconnect commands to the targeted meters.

C. Attack Tree

Based on the results collected from the previously discussed case studies (and more that are not covered in this paper), we created “attack trees” (see Figure 1) that provide an overview of the attackers’ key objectives and the smaller attack steps that would be involved in achieving them. The root nodes (orange) represent the objectives, and the child nodes represent the individual attack steps. A child node may comprise multiple attack steps and can reappear in a different branch. For example, “eavesdropping” (information gathering/privacy) consists of “eavesdropping outsider” and “eavesdropping insider”; this node reappears under “replay attack” (cheat power company/energy theft) but we show only the parent node. Leaf nodes (green) represent various ways an adversary may gain access to the network.

To illustrate how one may go about constructing an attack that utilizes the tree, take for example, the goal of “interrupting service.” One attack can be constructed by gaining physical access to a meter and then using a buffer overflow attack (exploit) to gain root access, thus compromising a single node. Once a single meter has been compromised, a network exploit could be used to compromise multiple meters. Once the adversaries have multiple meters under their control, they can coordinate a distributed denial of service attack and use a packet flood on the target, thereby interrupting service against the target. If the target is core component to network communications, this attack might result in the utility being cut off from large segments of the network.

Next, we look at the types of information that are needed to effectively detect those attack steps.

IV. INFORMATION REQUIRED FOR DETECTION

The basic attack steps, which were decomposed from the attack scenarios in the threat survey, and the information required to detect those attack steps are presented in Table I. Each line of the table is an individual low-level attack technique that can be used alone or in combination with other techniques to build complex attack scenarios. As explained in the previous section, our goal is to identify fundamental attack techniques at the lowest level and to identify the core information required for their detection. Once acquired, that information, associated with detection technology, ensures that any combination of attacks could be detected.

The information required for detection can be organized into three categories:

- *System information:* health reports from meter, and gateways (CPU, battery consumption), firmware and software integrity of AMI devices, clock synchronization.
- *Network information:* NAN collision rate, packet loss, node response time, traffic rate, health and integrity of routing table, associations between physical addresses and node identity.
- *Policy information:* Authorized AMI protocols, authorized AMI devices, authorized traffic patterns, authorized route updates, authorized firmware updates.

The knowledge extracted from the mapping between attacks and information required for detection is crucial to the design of a comprehensive and cost-efficient monitoring solution. Indeed, the above categorization reveals that data must be collected from different locations in the infrastructure. For example, the need for information on health and integrity of routing tables requires routers (in this case, meters) to be instrumented so that they can send periodic health reports or at least be remotely queried for health and integrity checks. However, instrumentation of all routers in the network may be too expensive, and a better solution, from a cost point of view, could be to rely on attack manifestations at other locations in the system instead of routers for detection. Next, we will review those tradeoffs by investigating different intrusion detection architectures.

V. ARCHITECTURE AND DISCUSSION

Having identified the information required to detect common attacks, we are now in a position to sketch possible

TABLE I
LIST OF INDIVIDUAL ATTACK TECHNIQUES AND INFORMATION REQUIRED TO DETECT THEM

Category	Attack technique	Target	Information required
DoS	Collision in Packet Transmission	NAN Link Layer	NAN collision rate, node response time
DoS	Packet Flood	Node in NAN (Meter/DCU)	CPU and memory usage of target incoming network traffic to target, authorized network protocols, network health information, packet-per-second rate, node response time
DoS	Jamming	NAN Physical Layer	NAN signal level, node response time
DoS	Alter Routing Table	Routing Protocol	Routing table health, node response time
DoS	Drop Packets	NAN Traffic	Packet loss among nodes in mesh network
DoS	Destroy Node	Node in NAN (Meter/DCU)	Node availability / response time
DoS	Time-Desynchronization	Node in NAN (DCU)	Time-synchronization traffic among nodes or time configured on nodes
DoS	Resource Exhaustion (Battery, Bandwidth, or CPU)	Node in NAN (Meter/DCU)	Traffic among meters, valid traffic profile or node health (CPU, battery consumption), network health (bandwidth usage)
Spoofing	Impersonate Regular Node	Node in NAN (Meter)	Associations between physical addresses and node identity
Spoofing	Impersonate Master Node	Node in NAN (DCU)	Associations between physical addresses and node identity, associations between regular and master node registrations
Spoofing	Man-in-the-Middle	NAN Traffic	Associations between physical addresses and node identity
Spoofing	Wormhole	NAN Traffic	Associations between physical addresses and node identity, routing table integrity/update
Spoofing	Slander	Distributed Detection System	Integrity of trust and reputation system
Eavesdropping	Passively Listen to Traffic	NAN Traffic	N/A (undetected)
Eavesdropping	Active cryptanalysis	NAN Traffic	Traffic among meters
Physical	Compromise Meter	Node in NAN (Meter)	integrity of meter firmware, memory contents of meter, meter firmware upgrade policy, meter status, information about bandwidth and wireless signal
Communication	Attack Coordination	Traffic in NAN	network protocols that are authorized for use, network traffic among the meters, network characteristics of legitimate traffic

IDS deployment schemes. This section explores four different approaches and discusses how effective each would be in detecting attacks.

A. Centralized IDS

The most cost-effective solution would be a centralized deployment scheme in which a single IDS sensor is deployed at the head-end, monitoring all the traffic that flows to and from the AMI. This IDS sensor would have access to the traffic reaching the utility network, to maintenance and upgrade policies, and to system logs from AMI appliances. Thus, it could detect systemic attacks that target the utility network and insider attacks that leave traces in access logs, and it could also analyze anti-tampering alerts sent by smart meters.

While the set of information required for detection shows that a central sensor at the head-end is necessary, our analysis in the previous section shows that it is not sufficient. Indeed, there will also be a significant number of attacks performed within the AMI that the central sensor would miss. For instance, attack techniques such as “installing malware on the meter” or “eavesdropping on NAN traffic” through an active cryptanalysis (i.e., by injecting traffic to force nodes to generate encrypted packets) would be undetected by the central sensor because it would not have access to information such as the integrity of the meter firmware, the memory contents of the meter, the NAN traffic among meters, the network bandwidth usage, or the routing table integrity.

B. Embedded Sensing Infrastructure

To effectively monitor the traffic among meters and to get access to meter-specific information, it would be sensible to

place sensors within the meters themselves. Shin et al. [14] discuss an IDS architecture that selects a subset of meters as sensors while minimizing the number of meters that need to be instrumented. Those sensors would have complete visibility over meter-specific information such as health reports, firmware and software integrity, and memory contents. Access to that information would allow attack techniques such as “meter compromise” to be detected more effectively.

Since in the mesh network, meters function not only as end points but also as relays, a collection of embedded sensors would have complete visibility over the traffic that flows within the AMI. Such redundancy in traffic monitoring would greatly improve accuracy and increase the trustworthiness of alerts that are generated. Moreover, the meter sensors would be effective in detecting any attack that originates from the HAN (e.g., through a compromised appliance); such attacks may target the meter first, and then extend to the DCU or the utility assets. The meter sensors would inspect incoming messages and detect any unusual or prohibited commands.

However, if only meters are instrumented, attacks performed directly on the DCU would not be covered. Data required to detect physical tampering with the DCU or exploitation of insider privileges, for instance, would be out of reach. Moreover, since the AMI might be encrypted, meter sensors would need access to multiple *decryption keys* to be able to properly inspect the packet contents originating from other meters, increasing the impact of a compromised meter sensor. The centralized IDS, on the other hand, would have access to traffic that has already been decrypted. It is also worth noting that most meters have limited processing power, storage, and

communication capabilities. Deploying a resource-intensive IDS sensor on a meter, whether it is in the form of software or hardware, might be a detriment to the meter's daily operations and require hardware upgrades that increase costs.

C. Dedicated Sensing Infrastructure

An alternative deployment scheme would be the dedicated sensing infrastructure, in which a small number of dedicated sensors monitor networks in the field. The key advantage of these sensors is the high availability of processing power and storage, which would allow complex IDS functions (e.g., full specification-based detection) to be performed; the daily meter operations would not be affected. This solution offers an interesting trade-off between network visibility and deployment cost, since the number of dedicated sensors needed to cover NANs would be far less than the total number of meters.

Much like embedded sensors, those dedicated sensors will have access to the AMI network data and be capable of monitoring the traffic flowing within the AMI. Distribution of the decryption key for decrypting the packet contents would be more manageable, since there would be a smaller number of sensors sharing the keys; this also means that there would be a smaller attack surface for key compromises. Nevertheless, dedicated sensors will not be able to monitor attacks performed directly on the meters or the DCU, as they will not have access to meter-specific or DCU-specific data. Attacks that start from the HAN and try to compromise the meters would also go undetected. In addition, from a practical point of view, figuring out where to deploy the dedicated sensors is not trivial; a potential site (e.g., a pole top) would need to be surveyed and rented, and permits would need to be obtained.

D. Hybrid Sensing Infrastructure

A combination of the central sensor and embedded meter sensors might provide the widest coverage in detecting attacks. Attacks that could not be covered by the centralized IDS (e.g., attacks performed directly on the meters, or malicious packets that flow within the AMI) would be covered by the meter sensors. The meter sensors would also cover attacks that originate from the HAN. Nevertheless, it might be hard to convince meter vendors to embed the sensor capabilities, as they could push costs up where margins are already small.

Alternatively, dedicated sensors could also be used together with the centralized sensor, monitoring the traffic that flows within the AMI and managing complex IDS operations. There should be more financial incentives for security companies to build them and utilities to deploy them, as fewer of them might be needed (compared to meters), especially in dense urban areas. Attacks that originate from the HAN or attacks performed directly on the meters would be missed, thus reducing the monitoring coverage.

VI. CONCLUSION

This paper provides a detailed look at the threats facing an AMI, the kind of attacks that might be performed, and the information that would be required to detect them effectively. Our attack trees capture the attackers' high-level objectives and show the individual attack steps that may be performed to achieve them. Our analysis of the individual attack steps

identifies the types of information that would be required to detect those attack steps. The types of information needed to detect attacks provides guidance for the design of an effective sensing infrastructure. Specifically, our analysis suggests that a hybrid sensing infrastructure, whereby a centralized IDS is orchestrated together with embedded meter sensors, provides the widest coverage in monitoring attacks. In the future, we plan to simulate different IDS deployment schemes and work out the optimal number of meter sensors or dedicated sensors that would be needed to monitor the entire AMI as well as perform a cost-benefit analysis of the proposed architectures.

VII. ACKNOWLEDGMENT

This material is based upon work supported by the Department of Energy under Award Number DE-OE0000097. The authors would like to thank Jenny Applequist for her editorial assistance.

REFERENCES

- [1] R. Berthier, W. Sanders, and H. Khurana, "Intrusion detection for advanced metering infrastructures: Requirements and architectural directions," in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, oct. 2010, pp. 350–355.
- [2] S. McLaughlin, D. Podkuiko, and P. McDaniel, "Energy theft in the advanced metering infrastructure," in *Proceedings of the 4th international conference on Critical information infrastructures security*, ser. CRITIS'09. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 176–187. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1880551.1880566>
- [3] S. McLaughlin, D. Podkuiko, S. Miazvezhanka, A. Delozier, and P. McDaniel, "Multi-vendor penetration testing in the advanced metering infrastructure," in *Proceedings of the 26th Annual Computer Security Applications Conference*, ser. ACSAC '10. New York, NY, USA: ACM, 2010, pp. 107–116. [Online]. Available: <http://doi.acm.org/10.1145/1920261.1920277>
- [4] T. Fleury, H. Khurana, and V. Welch, "Towards a taxonomy of attacks against energy control systems," *Critical Infrastructure Protection II*, pp. 71–85, 2009.
- [5] X. Wu and N. Li, "Achieving privacy in mesh networks," in *Proceedings of the fourth ACM workshop on Security of ad hoc and sensor networks*, ser. SASN '06. New York, NY, USA: ACM, 2006, pp. 13–22. [Online]. Available: <http://doi.acm.org/10.1145/1180345.1180348>
- [6] D. Welch and S. Lathrop, "Wireless security threat taxonomy," in *Information Assurance Workshop, 2003. IEEE Systems, Man and Cybernetics Society*, june 2003, pp. 76–83.
- [7] A. Rai, T. R.R., and S. Upadhyay, "Different Types of Attacks on Integrated MANET-Internet Communication," *International Journal of Computer Science and Security*, vol. 4, no. 3, pp. 265–274.
- [8] B. Kannhavong, H. Nakayama, Y. Nemoto, N. Kato, and A. Jamalipour, "A survey of routing attacks in mobile ad hoc networks," *Wireless Communications, IEEE*, vol. 14, no. 5, pp. 85–91, october 2007.
- [9] A. A. Cardenas, T. Roosta, and S. Sastry, "Rethinking security properties, threat models, and the design space in sensor networks: A case study in scada systems," *Ad Hoc Netw.*, vol. 7, no. 8, pp. 1434–1447, Nov. 2009. [Online]. Available: <http://dx.doi.org/10.1016/j.adhoc.2009.04.012>
- [10] T. Gamer, L. Voelker, and M. Zitterbart, "Differentiated security in wireless mesh networks," *SECURITY AND COMMUNICATION NETWORKS*, vol. 4, no. 3, pp. 257–266, MAR 2011.
- [11] H. Redwan and K.-H. Kim, "Survey of security requirements, attacks and network integration in wireless mesh networks," in *New Technologies, Mobility and Security, 2008. NTMS '08*, nov. 2008, pp. 1–5.
- [12] X. Wang and P. Yi, "Security framework for wireless communications in smart distribution grid," *Smart Grid, IEEE Transactions on*, vol. 2, no. 4, pp. 809–818, dec. 2011.
- [13] D. Jin, C. Lee, D. Nicol, I. Shin, and H. Zhu, "Simulation-based Study of Distributed Denial-of-Service Attacks in Advanced Metering Infrastructure," in *INFORMS Annual Meeting*, Charlotte, NC, USA, November 2011.
- [14] I. Shin, J. H. Huh, C. Lee, and D. M. Nicol, "A Monitoring Architecture for Smart Meter Mesh Networks in the Smart Grid," (submitted for review), 2011.