Monitoring Smart Grid Operations and Maintaining Mission Assurance through Adaptive Computer Network Defenses

[Extended Abstract]

Daniel Bilar
Siege Technologies LLC
33 S. Commercial Street
Manchester NH 03101, USA
daniel.bilar
@siegetechnologies.com

James J. Solderitsch
Accenture PLC
11951 Freedom Drive
Reston VA 20190, USA
james.j.solderitsch
@accenture.com

Elvis K. Hovor
Johns Hopkins University
Information Security Institute
3400 North Charles Street
Baltimore MD 21218, USA
ehovor1@jhu.edu

ABSTRACT

We present adaptive computer network defenses with a focus on Smart Grid operations. We model Smart Grid network domains in the MESA-ExtendSim simulator and continuously measure mission assurance indicators (security properties) to gauge the network's mission state. When thresholds are reached, a binary integer optimization problem is solved to identify appropriate host and network defenses (security controls) given parameters (attacker classes) and constraints (cost and mission assurance bounds). The solution is used to actuate defense mechanisms in the simulation. We present mission assurance indicators in an UI before and after defense actuation to gauge defensive effectiveness.

Categories and Subject Descriptors

C.2.3 [**Communication Networks**]: Network Operations—*network management, network monitoring*; D.2.8 [**Software Engineering**]: Metrics—*process metrics, performance measures*; G.1.6 [**Numerical Analysis**]: Optimization—*constrained optimization, integer programming*

General Terms

simulation, optimization, security metrics

Keywords

Optimization, Security Properties, Adaptive Defenses, Mission Assurance, Smart Grid

1. INTRODUCTION

In 2008, IT decision makers at over 80 companies were polled on their use of security metrics. Almost 80% responded that demonstrating IT security effectiveness through metrics to other functional managers helps IT to justify action and budgets [1]. In our work, we use security metrics to guide allocation of best "bang for the buck" (ROI) defense mechanisms against various threat classes in a constrained simulation environment. We specifically address quantification of metrics and experimental validation since they represent methodological lacunae in the literature. A meta-survey of ninety security papers between 1981 and 2008 showed that quantified security was a weak hypothesis because of lack of validation and comparison against empirical data [11].

CSIIRW '11, October 12-14, Oak Ridge, Tennessee, USA

2. RELATED WORK

A taxonomy given by [9] takes "mission assurance" to be synonymous with Mission Essential Functionality (MEF). The conceptual framework's goal, similar to ours, is detection of MEF degradation with runtime support for mission assurance in terms of Quality Of Service and Information Assurance criteria. Our project goes further in that we define metrics, measure them, and implement a quantitative, optimization-based mission assurance framework.

The optimization approach underlying our framework is based on QRAN [2]. QRAN, a quantitative risk analysis and management framework for computer networks, is focused on the risk induced by vulnerabilities present in non-malicious software. It allowed risk managers to get a comprehensive snapshot of the constitutive software on the network, assess its risk with the assistance of a vulnerability database via multi-factor risk metrics, and manage that risk by rank ordering reduction measures subject to cost, functionality and risk tolerance constraints.

The LANL FRNSE (Framework for Responding to Network Security Events) project looks for evidence of threats to the enterprise network and responds with automated steps. One of its goals is substituting scripts for low level security analysts. While FRNSE focuses on detecting and responding to all threats to the enterprise without considering the cost of the response, our work focuses mainly on defensive adaptations that maximize Mission Assurance within the available resource budget. Rather than just responding to events, we try to anticipate attacker next-moves so as to get ahead of them (proactive vs. reactive). However, FRNSE could be one of the defensive tools in our arsenal to make first pass suggestions that are then folded in with our other models and algorithms [3].

3. CONCEPTS

A key concept of our approach is shaped by the mission (or purpose) of a network. Our reference Smart Grid (Fig. 1a) contains multiple networks (clouds), connecting disparate domains via gateway actors (bold rectangles). These networked domains support different missions. Within these network domains, we continuously measure mission assurance indicators (security properties) on actors as we optimize for defenses (security controls) given parameters (attacker classes) and constraints (cost and mission assurance bounds).

3.1 Mission Assurance Properties

We defend, or assure, security properties. In our framework, assuring these properties is akin to assuring the mission. These properties include (but are not limited to) availability (ability to use information/resources), integrity (prevention of unauthorized changes), confidentiality (concealment of information), authenticity (identification and assurance of origin), freshness (non-replay of stale data and commands) and non-repudiation (proof of responsibility).

Mission assurance (MA) decreases as attacker efforts increases. 'Efforts' is meant to denote a general placeholder and includes (but needn't be limited to) person-hours invested, technical resources marshaled and drive.

Effort intensity serves to separate increasingly hardier attack classes. We distinguish between the following five **attack classes**: Opportunistic: quasi-random attacks (likely automated); Hobbyist: low skill human with no strong incentives; Organized crime: medium skill humans with financial incentives; Nation state: highly skilled humans with national security incentives; malicious insider: abuse of privileged insider position with financial, ideological or personal motives.

3.2 MESA Simulation Network

We used the Modeling Environment for SOA Analysis (MESA) (designed by MITRE on top of the ExtendSim framework) to model and simulate parts of the Smart Grid (Fig. 1b). Our first-iteration **prototype focuses just on assuring confidentiality, integrity and availability properties in the Service Provider and Customer domains**.

The **Availability metric** is generated in MESA by increasing service time for polling requests, thereby mimicking degradation of service and modeling DoS. A node outside the domain of interest periodically polls a node in the domain of interest to measure the MA metric.

The **Integrity metric** is generated by increasing both the probability and magnitude of probabilistic error terms added to data. While negative terms model service theft by customers, positive and negative perturbations of control data serves to model integrity attacks against the utility network's control systems [5]. A node periodically samples the readings to measure the metric against threshold values.

The **Confidentiality metric** is modeled in MESA by simulating an information leak (i.e client billing data, power consumption signatures of clients) by increasing the probability of copying 'files' from a node that can only by reached from an internal network to a node that is reachable from outside. An external actor periodically pulls data from that reachable node and upon success or when watermarked data passes by a gateway router/firewall, information has 'leaked'.

Table 1: Optimization parameters

parameter	description
n	# defense responses to optimize over
e	# attack classes of interest
m	# mission assurance metrics of interest
c	# cost factors to be considered

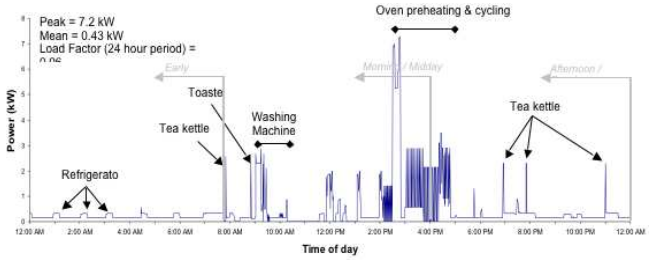

Figure 2: Household power consumption. High variation (load factor = 0.06). Figure drawn from data in [8]

3.3 Optimization Formulation

Table 1 gives an overview of the parameters. Let n be the number of defense responses, let e be number of attack classes, let m be the number of MA curves (e.g. m=3 for Confidentiality, Integrity, Availability) and let c be the number of cost factors associated with defenses (e.g. acquisition, training, cost of ownership). Table 2 gives an overview of the parameters used by the objective function (1) and the constraints (2), (3) of the optimization problem.

Let $\mathbf{D}$ be an m-by-n-by-e matrix whose elements denote the **mission assurance effectiveness of defenses against attacker classes**.

Let $\mathbf{A^{(s)}}$ be an m-by-n subspace of D denoting the **mission assurance effectiveness of defenses against subset of attackers**. A linear function $\mathbf{f} : \mathbf{D} \rightarrow \mathbf{A^{(s)}}$ serves to narrow down the attacker domain of interest. Examples include an average over all attackers, largest attack effort, or a partial subset (e.g. only hobbyist and organized crime).

Let $\mathbf{A^{(s,i)}}$ be the subspace of $\mathbf{A^{(s)}}$ denoting the attacker domain of interest selected by function $\mathbf{f}$. Let $\alpha_\mathbf{i}$ be the **relative weight of the mission assurance metrics** we want to maximize over. Let $\overrightarrow{MA}_\mathbf{rhs}$ denote the **minimum assurance bounds** we need to safeguard for the mission.

Let $\mathbf{C}$ be a c-by-n matrix whose entries denote **the cost associated with defense responses**. Let $\overrightarrow{Cost}_\mathbf{rhs}$ denote the total upper cost limits of the chosen defense responses. The feasible solution(s) are returned in the n-by-1 indicator vector $\vec{x}$ which denotes the set of defense responses (such as watermarking, blackholing, encryption) that should be actuated. The objective function solves for $\vec{x}$ which maximizes (weighed by relative factor α_i) chosen mission assurance metrics.

$$\max_x \sum_s \sum_i \alpha_i A^{(s,i)} \vec{x} \tag{1}$$

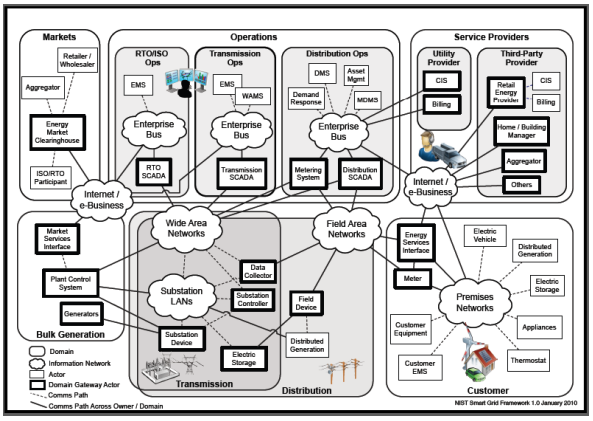

(a) NIST Smart Grid: A domain is a grouping of actors that has similar objectives and rely on similar systems.

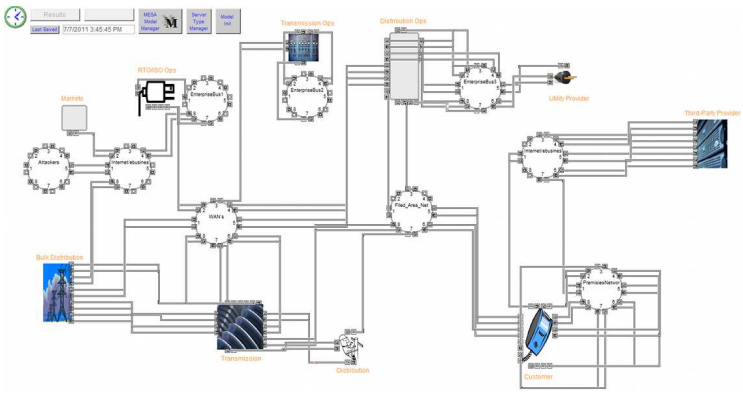

(b) Simplified MESA-ExtendSim model of Smart Grid. Only some actors and domains are shown.

Figure 1: Conceptual (NIST) and MESA-ExtendSim Smart Grid models

Table 2: Optimization Variables

parameter	description
D	mission assurance effectiveness of defenses against attacker classes
$A^{(s)}$	mission assurance effectiveness of defenses against subset of attackers
$A^{(s,i)}$	specific mission assurance effectiveness of defenses against subset of attacker classes
C	costs of defenses
MA_{rhs}	minimum total mission assurance to maintain
$Cost_{rhs}$	maximum total costs allowable
α_i	relative weight of specific mission assurance
x	defense indicator vector

Table 3: Cost items

cost items	description
configuration	deploying defenses
acquisition	delivering defenses
maintenance	updating defenses
opportunity	'road-not-taken'
utilization	negative impact
transition	runtime adoption costs

Table 4: Defense Responses

defense responses	description
watermark	hashtag data
blackhole	change network routes
honeynet	add subnets
encrypt	add network encryption
SESRAA[6]	active denial
signature	enable whitelisting
randomize	change environment
timestamp	timestamp packets
data poison	mutate data
failover	add hosts

subject to

$$A^{(s)}\vec{x} \geq \overrightarrow{MA}_{rhs} \quad (2)$$

$$C\vec{x} \leq \overrightarrow{Cost}_{rhs} \quad (3)$$

such that

$$x \in \{0,1\}^n$$

$$\sum_i \alpha_i = 1$$

We set up a binary integer program (BIP) in MS Excel and solve with a Solver plugin `http://solver.com`. The BIP problem class is generally NP-hard. However, since MA_{rhs} and $Cost_{rhs}$ are plausibly integer valued, the problem can be moved to an efficiently solvable subclass by transforming C and $A^{(s)}$ into totally unimodular matrices [7].

4. IMPLEMENTATION

Our simulation model uses a native feature of MESA's ExtendSim to call out to an Excel macro, either manually before the simulation runs, or automatically as part of the execution. Through this API, MESA sets the initial defense solution indicator vector to 0, executes macros in the Excel sheet to invoke the solver, and reads back the values back into MESA. The optimization response is translated into defense actuations.

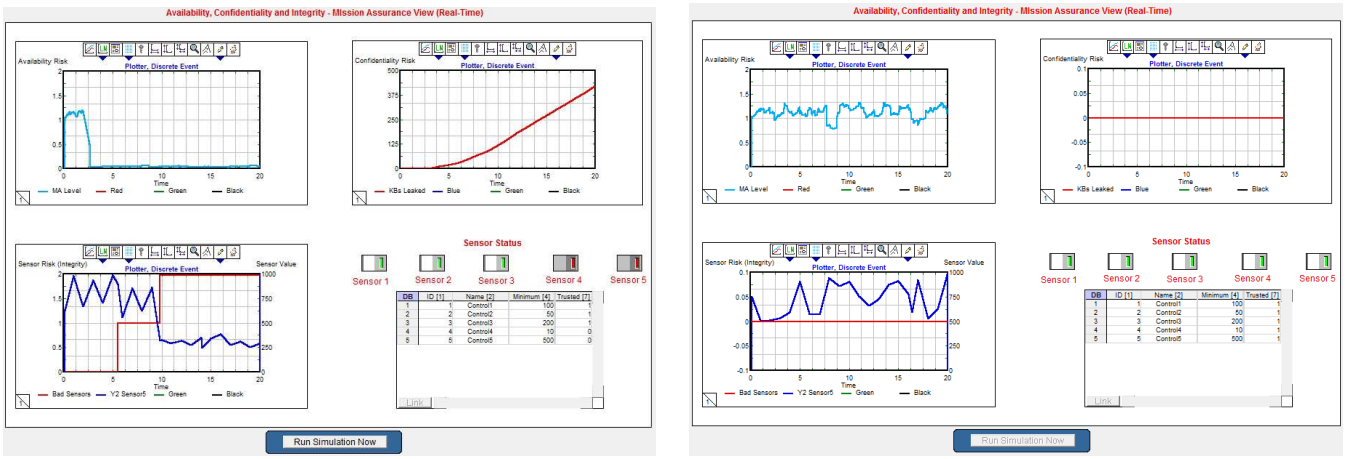

(a) MA measured before defenses

(b) MA measured after defenses

Figure 3: Mission Assurance over time for Availability (upper left), Confidentiality (upper right) and Integrity (lower left)

Simulation traces before and after defense actuation are given in Figs. 3a and 3b. The availability plot shows measured service processing time. In Fig. 3a we see a decrease at time 3. The confidentiality plot shows suspected data leakage overtime (in KBs). The Integrity plot shows sensor readings for a named sensor (blue line) and a count of the number of sensors with abnormal readings (red line). After actuation of three simultaneous defenses (signature, blackholing and watermarking), the dashboard on the right subfigure indicates increased assurance level measurements across the confidentiality, integrity and availability security properties.

5. DISCUSSION

After actuating a single defense for each kind of attack, the measurements under each attack scenario return to normal MA levels. In future work, we hope to show the effects of applying multiple defenses simultaneously, as well as measure the rates of change of Mission Assurance levels in more complex circumstances.

In our prototype, we based defense actuation on deviation from a baseline (effect-based anomaly detection). Our thresholds are set by tapping empirical high resolution (one to five minute intervals) data compiled from residential electricity time use series [12]. Empirical studies show that consumption patterns fluctuate greatly within daily activity periods, with load factors (maximum power compared to mean value) of up to 0.06 (see Fig. 2) [4, 10]. This complicates setting thresholds within acceptable type I and II errors. We will negotiate this problem in the next iteration by identifying anomalies through n-tuple (rather than just one) thresholds and by accurately modeling household consumption through multi-state non-homogeneous Markov chain activity models [13].

6. REFERENCES

[1] R. Ayoub. Analysis of business driven metrics: Measuring for security value. White Paper, 2008.

[2] D. Bilar. *Quantitative Risk Analysis Of Computer Networks*. PhD thesis, Thayer School of Engineering (Dartmouth College), June 2003.

[3] J. Doak and A. Lovato. Framework for responding to network security events (FRNSE). Technical report, Los Alamos National Laboratory, 2008.

[4] B. Fuller, J. Sikora, and J. Lyons. Review of residential electrical energy use data. Technical report, NAHB Research Center, July 2001.

[5] O. Kosut, L. Jia, R. Thomas, and L. Tong. Limiting false data attacks on power system state estimation. In *Information Sciences and Systems (CISS)*, pages 1–6. IEEE, 2010.

[6] Q. Liao, D. Cieslak, A. Striegel, and N. Chawla. Using selective, short-term memory to improve resilience against ddos exhaustion attacks. *Security and Communication Networks*, 1(4):287–299, 2008.

[7] J. F. Maurras, K. Truemper, and M. Akgül. Polynomial algorithms for a class of linear programs. *Mathematical Programming*, 21:121–136, 1981.

[8] M. Newborough and P. Augood. Demand-side management opportunities for the UK domestic sector. In *Generation, Transmission and Distribution, IEE Proceedings*, volume 146, pages 283–293. IET, 1999.

[9] P. Partha, K. Rohloff, M. Atighetchi, and R. Schantz. Managed mission assurance - concept, methodology and runtime support. In *Proceedings of SOCIALCOM*, pages 1159–1164. IEEE, 2010.

[10] S. Taherian, M. Pias, G. Coulouris, and J. Crowcroft. Profiling energy use in households and office spaces. In *Proceedings EECN*, pages 21–30. ACM, 2010.

[11] V. Verendel. Quantified security is a weak hypothesis: a critical survey of results and assumptions. In *Proceedings of NSPW*, pages 37–50. ACM, 2009.

[12] J. Widen, M. Lundh, I. Vassileva, E. Dahlquist, K. Ellegård, and E. Wackelgard. Constructing load profiles for household electricity and hot water from time-use data-modelling approach and validation. *Energy and Buildings*, 41(7):753–768, 2009.

[13] J. Widen and E. Wackelgard. A high-resolution stochastic model of domestic activity patterns and electricity demand. *Applied Energy*, 87(6):1880–1892, 2010.