Security Concerns and Disruption Potentials Posed by a Compromised AMI Network: Risks to the Bulk Power System

M. M. Olama[1], J. J. Nutaro[1], V. Protopopescu[1], and R. A. Coop[2]
[1]Oak Ridge National Laboratory, Oak Ridge, Tennessee, USA
[2]University of Tennessee, Knoxville, Tennessee, USA

***Abstract*— *The advanced metering infrastructure (AMI) of a smart electrical grid is seen as both a network for improving the efficiency of the electrical power system and as a potential target for cyber-attackers bent on disrupting electrical service. In this paper, we examine how a hijacked AMI network might be used to instigate widespread blackouts, and the physical barriers that the electrical system itself poses to such an attack. To this end, we present a simple, but potentially useful, model for gauging the quantity of load that an attacker must control for an attack to be successful. Conversely, the model suggests a scheme for mitigating the attack, but at the cost of decreasing the usefulness of smart meters as devices for the legitimate regulation of electrical load.*

Keywords: Advanced metering infrastructure, smart meters, energy management systems, swing equation, cyber security

1 Introduction

Smart electric meters, capable of two-way communications and having software and hardware to enable energy management in real-time, are a major part of the advanced metering infrastructure (AMI) of a smart electrical grid. These meters have tremendous potential to improve the efficiency and reliability of the national power system. For example, a washing machine in a household with a smart meter could be set to run only when energy is cheap. This reduces energy costs for the power consumer and reduces the size of demand peaks, which are served with expensive, relatively inefficient sources of energy. Current plans call for nearly 17 million smart meters to be installed in U.S. homes and businesses over the next few years. While proponents of a smart grid and, in particular, an AMI have touted the potential to improve the electricity system, critics have expressed concerns about the susceptibility of AMI meters to cyber-attacks.

Cyber security practitioners [1]-[5] have claimed that smart meters are susceptive to hacking, and thereby are a potential enabler of unauthorized access to command and control processes that could be abused to disrupt electrical service. Such claims are often disputed by engineers that operate large electrical power systems. Significantly, the vulnerabilities found in smart meters have not been put to such a use.

However, if such an attack were to be carried out, it might unfold as follow. The attacker gains control over a substantial quantity of load by hijacking a large number of smart meters. Using these meters, the attacker creates a large imbalance between power used and power supplied by switching off the load that he controls. This causes a large and sudden change in the frequency of the power system, thereby forcing some generators to disconnect from it. By repeating this attack several times in the course of a few minutes, large numbers of generators may be forced to disconnect, thereby instigating a large-scale blackout.

This scenario is often used as evidence of a systemic risk posed by vulnerabilities in smart meters and their attendant infrastructure for communication and control. Opponents of these claims cite the intrinsic robustness of a power system to sudden changes of load. Indeed, the inertia of the power system's generators and the presence of automatic controls designed specifically to deal with imbalances of supply and demand are a physical barrier to successful attack.

In this paper, we estimate the physical requirements that must be met by an attack seeking to disable a large power system by the sudden and simultaneous manipulation of numerous electrical loads. We further propose that a random delay imposed on energy management actions can mitigate the most disruptive effects of such an attack. Towards this end, we examine two key quantities in such an attack: the amount of load controlled by an attacker and the swiftness with which it is switched. Our method of analysis is illustrated by its application to data published for the power grid that serves the western United States.

2 Motivation and method of attack

Among the motives for instigating a computer-based, rather than physical, attack against electrical infrastructure is the relatively small risk of getting caught. Cyber-attacks are notoriously difficult to attribute to a specific source (see, e.g., the discussion by Libicki [6]), and so the risks faced by a cyber-attacker are typically much less than those faced by the instigators of a physical attack. This aspect of a cyber-attack may make it attractive to malicious actors who want to disrupt electrical service but have a low tolerance for risk; this may be particular true for potential attackers who would never consider carrying out a physical attack.

There are two basic avenues for a cyber-attack on electrical loads: 1) through the meter hardware or 2) through a computer that controls metering or other electrical management, functions. Attacking the meters would perhaps be the most effective. However, such an attack requires both known, exploitable flaws in the meter software and a capability to exploit that flaw via remote operations on a very large number of meters; e.g., by the use of a botnet-style worm (see, e.g., [7], [8]).

The other avenue for an attack is via computers (e.g., home computers) that interact with and control some actions of a smart meter or load. For example, a home energy management system that is operated by a personal computer may be susceptible to attack via the Internet. By hijacking this computer, the attacker may be able to modify the power consumption of electrical appliances under its control (see, e.g., [9]). It is likely in this case that the computer controls only part of the load at a home or business, e.g. the PC may be able to remotely turn the furnace on or off, and so such an attack vector may be relatively less effective than one that targets the meter itself.

In any case, a significant hurdle that a cyber-attacker must overcome is to ensure that the infected devices are physically collocated. For an attack to be effective, it must be able to disconnect (or otherwise change the consumption of power by) loads within a specific electrical power system. Hence, all of the hijacked load must reside within the geo-physical region served by the targeted power system.

3 Effect of a sudden change of load

The sudden disconnection of an electrical load is felt by a generator as an acute easing of the torque which opposes its turning rotor. Automatic controls act to bring all forces back into balance by adjusting the production of power at the generator to match demand. A sudden reconnection of the same load will have the opposite effect, causing the generator to feel an acute tug in opposition to its spinning rotor. This causes it to slow and, again, automatic controls act to bring supply and demand into balance. If, however, the rotating speed of the machine drifts too far from normal, then automatic protection devices disconnect it from the electrical network. This can turn an abnormal event into a cascading failure.

It is conceivable that an attacker in control of a sufficiently large number of smart meters will use this physical phenomenon to instigate a large blackout. The number of meters required, and the precision with which their switches must be operated, are determined by the physical properties of the generators and the settings of their controllers. The main method in this attack is a change in load that happens more quickly than the automatic controls can respond.

We use a model derived from the swing equation and a simplified speed governor (see, e.g., the power system model presented in [10]) to approximate the system-wide, average change in frequency of a large electrical power system following a sudden change in the real electrical load. This model has two parameters for characterizing the electrical generation system, both of which can be measured (see, e.g., [11], [12]). These are its inertia M and the rate τ of its response to frequency excursions. This model also has one parameter for describing the change in load; this is the fraction α of the base electrical load P_e that is shed.

The change ω of the frequency of the power system during the event is related to the changing power P_m output by the generators and power P_e demanded by the loads by

$$\dot{\omega} = \frac{1}{M}\left(P_m - P_e\left(1 - \alpha u(t)\right)\right)$$

$$\dot{P}_m = -\frac{1}{\tau}\omega \qquad (1)$$

where $u(t)$ is the step function and the model begins in steady state with $\omega(0) = 0$ and $P_m(0) = P_e$.

Though simple, this model captures the three salient features that act unavoidably to oppose an attack relying on the manipulation of load. These features are (i) the inertia of the generators, which opposes sudden changes in frequency; (ii) the speed governors that act to correct an imbalance before dangerous changes in frequency are realized; and (iii) the rate at which those governors must act, which is determined in turn by the generation inertia and the size and rate at which the imbalance forms.

Only the frequency excursion is of interest here; solving Eqn. (1) we get

$$\omega(t) = \alpha P_e \sqrt{\frac{\tau}{M}} \sin\left(\frac{t}{\sqrt{M\tau}}\right) \qquad (2)$$

which has a maximum amplitude $\omega_{\max}$ at

$$\omega_{\max} = \alpha P_e \sqrt{\frac{\tau}{M}} \qquad (3)$$

This derivation shows that the maximum frequency deviation is proportional to the fraction of load that is removed by the step change. Using

$$\epsilon = P_e\sqrt{\frac{\tau}{M}} \tag{4}$$

Equation (3) can be written as

$$\omega_{\max} = \epsilon\,\alpha \tag{5}$$

and given measurements of α and $\omega_{\max}$ for an observed event, ϵ can be calculated.

If $\omega_{\max}$ is greater than the maximum excursion tolerable by the power system (again, a quantity which is known or can be estimated) then the system is at risk. Such a sudden change in load can be prevented in at least three ways.

First, hardware in the meter itself can impose a random delay and thereby force a ramped response from the population of meters. If the rate of response is sufficiently slow, then it will give automatic controllers sufficient time to safely adjust the power output of their generators. If implemented in trustworthy hardware, this protection mechanism is effective regardless of how the attack is instigated; that is, it cannot be disabled by software faults or computer-based attacks.

Second, business logic, implemented in software, can monitor for unsafe load changes and refuse to execute them, force the change into safe limits, ask the operator for confirmation, or all three. This kind of process control supplements other security measures which seek to ensure that requests come only from authorized operators, that worms and viruses do not infect the smart meters and enable malicious operation of the electrical switch, and to prevent other similar kinds of contingencies. However, the software that implements the protective business logic is itself subject to cyber-attack, and so its effectiveness as a security measure cannot be guaranteed.

Third, AMI installations could be deliberately limited in their scope. Large-scale penetration of AMI would be much more difficult given isolated AMI networks that employ different types of metering hardware, operating systems, and networking protocols. This is probably infeasible as a long term security solution, but could serve as a risk mitigation strategy while pilot deployments are rolled out. A staged deployment will provide opportunity for a utility to discover security problems, gain valuable operating experience in this respect, and to do so before the population of advanced meters reaches a threshold sufficient for causing wide spread disturbances.

In the next section, we address the first approach as a practical means for reducing the risks of an AMI being misused as a tool for the widespread disruption of the electrical service. Specifically, we show with an analytical model how imposing a random delay in the meter mitigates the impact of a cyber-attack on the power system.

4 Mitigating strategy

The maximum frequency deviation of the power system, $\omega_{\max}$, due to a sudden change in load can be reduced by enforcing a random time delay between a request that the smart meter opens or closes its switch and the moment at which the action is actually carried out. The switching delay must be determined at random for each request, and the delay implemented by a device not accessible in any way via remote access to the meter. If a hardware delay can prevent a sudden change in electrical load, automatic speed controls in the generators can prevent a dangerous jump in ω.

This security control is modeled in a way similar to the attack in (1), but with the step $u(t)$ replaced with the ramp $\beta(t)$ as

$$\dot{\omega} = \frac{1}{M}\Big(P_m - P_e\big(1 - \alpha\beta(t)\big)\Big)$$

$$\dot{P}_m = -\tfrac{1}{\tau}\omega \tag{6}$$

where $\beta(t)$ is the ramp function

$$\beta(t) = \begin{cases} t/\gamma & t < \gamma \\ 1 & t \geq \gamma \end{cases} \tag{7}$$

and $\gamma > 0$ is the duration over which the load change occurs. This is illustrated in Fig. 1.

Solving Eqns. (6) and (7) for ω gives

$$\omega(t) = \begin{cases} \frac{P_e\alpha\tau}{\gamma}\left(1 - \cos\left(\frac{t}{\sqrt{\tau M}}\right)\right) & t < \gamma \\ \frac{P_e\alpha}{\gamma M}\left(\cos\left(\frac{t-\gamma}{\sqrt{\tau M}}\right) - \cos\left(\frac{t}{\sqrt{\tau M}}\right)\right) & t \geq \gamma \end{cases} \tag{8}$$

If the power system is stable in the sense that it will ultimately damp out any excursion (and this must be the case while the power system is operating) then it is sufficient to consider just the first swing (i.e., half period); subsequent swings will have decreasing amplitudes.

From this perspective, there are two possibilities. If the interval γ is sufficiently large, then the first swing occurs while $t < \gamma$. Its magnitude is

$$\omega_{max}(t < \gamma) \leq 2\alpha\tau P_e/\gamma \tag{9}$$

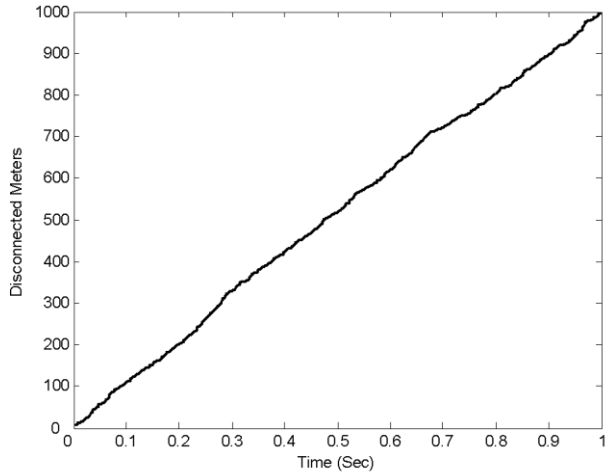

Fig. 1. Simulated ramp response of 1000 meters with enforced delays selected at random from [0,1].

and this limit is reached at

$$t_{max}(t < \gamma) = \pi\sqrt{\tau M} \tag{10}$$

In this case, the rate of the response of the generators to the imbalance counteracts the magnitude and suddenness of the change in load.

More desirable for the attacker is that γ be small, in which case inertia carries the system through its first swing. This is the case of practical interest, for which ω is bounded by

$$\omega_{max}(t \geq \gamma) \leq \frac{P_e\alpha}{\gamma M}\left(\cos\left(\frac{\pi-\frac{\gamma}{\sqrt{\tau M}}}{2}\right) - \cos\left(\frac{\pi+\frac{\gamma}{\sqrt{\tau M}}}{2}\right)\right) \tag{11}$$

and this limit is reached at

$$t_{max}(t \geq \gamma) = \frac{\pi\sqrt{\tau M}}{2} \tag{12}$$

Indeed, if γ is very small, then Eqn. (11) can be approximated by

$$\omega_{max}(t \geq \gamma) \approx \frac{P_e\alpha}{M\sqrt{\tau M}} \tag{13}$$

Thus, a conservative limit for ω_{max} is given by

1. Eqn. (11) while γ is less than the t_{max} in Eqn. (12),
2. Eqn. (11) from γ equal this t_{max} until Eqn. (11) is less than Eqn. (9), and
3. Eqn. (9) afterwards.

5 Illustration

To illustrate this method of analysis, we estimate the combinations of γ and α required to cause particular frequency excursions in the western United States. The data used for these calculations is derived from a June 14 event in the western U.S. (as governed by WECC; see [12], especially Figures 1 and 2). In that event, 4.589 GW of generation was lost resulting in a 0.4 Hz frequency excursion. Chassin *et. al.* measured the system inertia during this event as 17.8 GW.sec^2. The base load P_e for this specific event is not reported in their paper, but 90 GW is the midpoint for the range of values reported and using this for P_e gives a notional 5% change in load. With this data and Eqn. (13), the control constant τ is calculated as 0.0224.

Fig. 2 illustrates the conservative limit for ω_{max} using the above data. When the ramping time is small, the size of the change in load is the dominant factor. This observation is consistent with the approximation in (13). The size of the excursion can be controlled, however, by lengthening the ramping interval; using the data above, the magnitude of the excursion falls quickly for $\gamma > 1$ second.

This observation suggests that the impact of hijacked loads on the power system may be mitigated with an enforced ramping time. One place to enforce this is in the meters themselves; an inexpensive circuit may be introduced into the electronics that selects a delay uniformly in [0, γ] for that meter's switch. The aggregate effect of this delay is described by (7), and if γ is large enough then the population of smart meters adjusts the total load P_e at a tolerable rate. This scheme prevents an attacker from causing an unacceptably large change in frequency by placing a physically enforced limit on his actions. Moreover, if this delaying circuit is installed when the meter is manufactured then its protective function comes at a very small price.

Though this scheme prevents rapid changes in load at the meters, it does not prevent a similar type of attack at other points in the electrical system; e.g., sub-stations typically have a capability to disconnect large amounts of load for the purposes of emergency load shedding. Moreover, it is not known whether the risk posed by an attack on smart meters justifies the proposed limit on their use: rapid control of load via smart meters is desirable when it is used to regulate (rather than disrupt) frequency.

6 Conclusions

Risk that may be posed by smart meters is of particular concern because these meters, unlike a control center or substation, are readily accessible to an attacker. Smart meters are easily purchased, giving a potential attacker the opportunity to study these devices in detail; and smart meters that are installed in a home or business may be

linked directly to the home or business computer network. Indeed, this feature is being used in some energy management systems, and it raises the possibility of hijacking loads via the Internet.

The model developed in this paper is a tool for estimating the impact that hijacked loads may have on an electrical power system. Though the model is simple and the calculated estimate necessarily rough, the resources required to construct this estimate are minimal and can be obtained relatively easily (again, see [11], [12]). The model makes concrete the argument that inertia and speed controls are a barrier to causing widespread disruption of electrical service by the use of hijacked loads. The technological sophistication and engineering resources required to overcome this barrier remains a topic for future analysis.

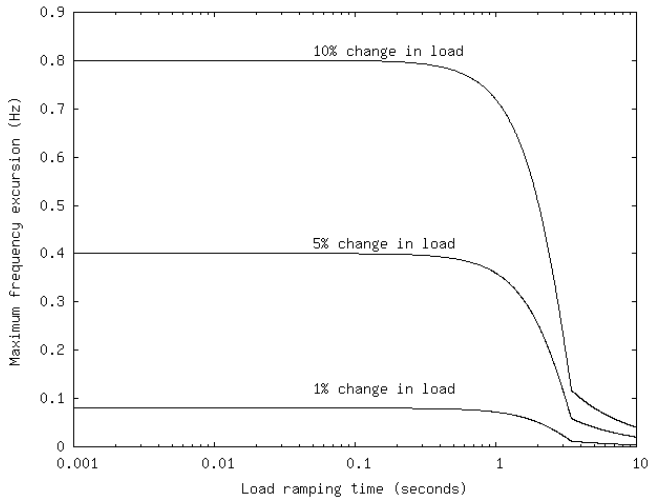

Fig. 2. Estimate of the maximum frequency excursion as a function of the load ramping time.

7 References

[1] A Risk-based Approach to Determining Electronic Security Perimeters and Critical Cyber Assets, Technical Report, IOActive Inc., 2009.

[2] Study of Security Attributes of Smart Grid Systems – Current Cyber Security Issues, Technical Report, Office of Electricity Delivery and Energy Reliability, U.S. Department of Energy, April 2009.

[3] S. Harris, "Chinese Hackers Pose a Clear and Present Danger to U.S. Government and Private-Sector Computer Networks and May be Responsible for Two Major U.S. Power Blackouts," *National Journal Magazine*, May 31, 2008.

[4] 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. 2009.

[5] R. Berthier, W. H. Sanders, and H. Khurana, "Intrusion Detection for Advanced Metering Infrastructures: Requirements and Architectural Directions," *Proc. of the IEEE International Conference on Smart grid Communications*, pp. 350 – 355, 2010.

[6] Martin C. Libicki, Cyberdeterrence and Cyberwar, RAND Corporation, 2009.

[7] B. Stone-gross, M. Cova, L. Cavallaro, B. Gilbert, M. Szydlowski, R. Kemmerer, C. Kruegel, and G. Vigna, "Your Botnet is My Botnet: Analysis of a Botnet Takeover," *Proc. of the 16th ACM Conference on Computer and Communications Security (CCS'09)*, pp. 635-647, Nov. 2009.

[8] S. Staniford, D. Moore, V. Paxson, and N. Weaver, "The Top Speed of Flash Worms," *Proc. of the 2004 ACM Workshop on Rapid Malcode (WORM'04)*, pp. 33–42, New York, NY, USA, 2004.

[9] M. LeMay, G. Gross, C. A. Gunter, and S. Garg, "Unified Architecture for Large-scale Attested Metering," *Proc. of the 40th Annual Hawaii International Conference on System Sciences (HICSS)*, pp. 115-124, 2007.

[10] J. Arrillaga, C. P. Arnold, and B. J. Harker, *Computer Modeling of Electrical Power Systems*, New York, Wiley 1983.

[11] T. Inoue, H. Taniguchi, Y. Ikeguchi, and K. Yoshida, "Estimation of Power System Inertia Constant and Capacity of Spinning-Reserve Support Generators Using Measured Frequency Transients," *IEEE Transactions on Power Systems*, vol. 12, no. 1, pp. 136-143, Feb. 1997.

[12] D.P. Chassin, Z. Huang, M.K. Donnelly, C. Hassler, E. Ramirez, and C. Ray, "Estimation of WECC System Inertia Using Observed Frequency Transients," *IEEE Transactions on Power Systems*, vol. 20, no.2, pp. 1190-1192, May 2005.