Delay Makes a Difference: Smart Grid Resilience Under Remote Meter Disconnect Attack

William G. Temple, Binbin Chen, Nils Ole Tippenhauer
Advanced Digital Sciences Center, Singapore
{william.t, binbin.chen, nils.t}@adsc.com.sg

***Abstract*—Modern smart meters commonly provide a service switch which allows remote connection or disconnection (RCD) of electrical service over a utility's communication network. While this feature is valuable for utilities, researchers have raised concerns about possible (ab)use by malicious attackers, noting the high economic cost of blackouts, as well the potential for controlled on-off switching of meters to affect power grid stability, for example by disturbing its frequency. However, while security concerns have been identified, little work has been done to develop and assess concrete countermeasures that are specific to these attacks. In this paper, we design novel randomized time delay countermeasures for smart meter RCD attacks, and demonstrate their effectiveness under sophisticated attack scenarios. We show that even if an attacker successfully issues malicious RCD commands, a well-designed time delay countermeasure makes the smart grid more resilient by: 1) preventing rapid changes in overall system load; and 2) providing time for a utility to potentially detect and stop an attack in progress. In particular, we demonstrate that a geometric delay mechanism can greatly reduce the magnitude of an attack with little impact on a utility's day-to-day operations.**

I. Introduction

Smart meters are one of the most visible elements of the transition to a modern smart grid. Across the United States and Europe, and increasingly in other parts of the world as well, millions of smart meters have been deployed as utilities invest in Advanced Metering Infrastructure (AMI) to enable dynamic pricing, remote service switching, and other services. Smart meters are typically controlled and queried through wireless or power-line communication. However, this communication and remote control capability also introduces potential attacks with severe consequences for consumers and asset owners.

In particular, the service switch and associated remote connect/disconnect (RCD[1]) capability of smart meters has caught the attention of the security community in recent years [1], [2], [3]. An RCD attack could cause a widespread blackout (or blackmailing of such a blackout) [1], or could potentially harm the power network or other loads by causing voltage or frequency deviations [2]. In either case, a successful attack would have severe economical and political consequences.

While security measures like data encryption and intrusion detection systems (IDS) offer some level of protection for AMI systems, they provide little recourse if an attacker *is* able to compromise the system and issue malicious disconnect commands to hundreds of thousands (or millions) of meters. In this work, we consider the use of random execution delays for RCD commands to mitigate such attacks. Even if an attacker compromises a system, enforcing a random time delay in every meter before executing RCD commands makes a smart grid more resilient by: 1) preventing rapid changes in load which may destabilize the power system; and 2) providing time for a utility to detect and stop an attack in progress.

Such delay mechanisms are technically feasible today: some meters support a configurable time delay when restoring service after a disruption [4]. While delayed RCD has been mentioned in the context of smart grid system resilience in [5] and [6], these documents do not provide specific delay mechanisms, nor do they analyze how delay would improve the overall security posture of smart grids. To the best of our knowledge, we are the first to provide in-depth design and assessment of time delay countermeasures for RCD attacks.

Our contributions are threefold: we investigate the effect of time delay on utilities' RCD value proposition, we design novel randomized time delay countermeasures that mitigate attacks regardless of the specific attacker strategies employed, and we use detailed simulations to demonstrate the effectiveness of these countermeasures. Our analysis and simulation results show that a geometric delay mechanism can significantly enhance a smart grid's ability to withstand, detect, and recover from an attack, with RCD delays of up to two hours.

The remainder of this paper is organized as follows: In Section II we present a framework for modeling attacks on RCD-enabled smart meters. In Section III we identify time delay as a critical countermeasure against such attacks. In Section IV we analyze the effects of time delay on utility RCD use cases. We make use of this information to design randomized RCD delay mechanisms in Section V. In Section VI we simulate these delay mechanisms and assess their impact on a smart grid's resilience. Finally, we conclude in Section VII.

II. Modeling Remote Disconnect Attacks on AMI

Threats posed by remote meter disconnection have been identified in the literature [1], [2], but little work has been done to develop and assess concrete countermeasures that are specific to these attacks. In this section, we describe the system and attacker models used to study time delay countermeasures for RCD attacks in the remainder of the paper.

[1]We refer to this feature as RCD throughout the paper, although we will focus our discussion on remote *disconnection*.

A. System Model

We consider a simplified smart grid system (see Fig. 1), where a utility company uses an AMI head end unit to control a set of smart meters at their customers' homes. Communication between the head end and the meters is routed via intermediate data concentrator units (DCUs), which are located in neighborhood area networks (NANs). The head end can issue requests (through a DCU) for any meter to disconnect. Our attacker model allows us to abstract away from the exact protocols, communication technology (e.g., 3G, power line communication), and security features available.

In addition to this AMI component, the system model includes an outage management system (OMS) at the utility, which can be reached by customers via phone or (mobile) internet. This OMS allows the utility to aggregate and analyze user reports. Additional utility services which may play a role in outage management (e.g., customer information system, geographical information system) are not explicitly modeled.

B. Attacker Model

We consider an attacker with two possible goals: 1) denial-of-electrical-service for a utility's customers during a certain time window, and 2) physical disturbance of power grid frequency via load shedding.

The impact of the first attack, a large-scale disconnect attack on a major city, can be devastating and has been discussed in detail in [1]. Outside of this worst-case scenario, smaller-scale RCD attacks (e.g., neighborhood level) can also have serious repercussions, and they may be less challenging for attackers to carry out [3]. The second attack, a targeted connection/disconnection of meters to affect physical power system parameters, has been discussed in [2]. Such an attack would likely require a high level of knowledge about the system, as well as a large number of compromised devices.

We model a strong attacker for both scenarios as follows: we assume an attacker with perfect knowledge of the target utility's infrastructure (including network topology, smart meter specifications, and control strategies) as well as knowledge of countermeasures employed in the system. We assume this attacker can transmit disconnect messages to any meter, and that these commands will be accepted as authentic. This abstraction encompasses a number of possible attack vectors; for example, compromise of the AMI head end's master key, exploitation of a weak protocol, and insider attacks from an AMI operator. We assume the attacker can prevent transmission of messages on the main communication channel, if needed (e.g., through jamming).

This attacker is limited to interact with the head end, DCU, and smart meters. The attacker cannot directly attack the OMS, the utility, or secondary communication channels.

III. Delayed Disconnect as a Countermeasure

There are several countermeasures available to utilities concerned about RCD attacks. Well-designed authentication and key management schemes [1], [7] are essential components of the solution, but in the absence of strong detection and response mechanisms, they leave little recourse if an attacker successfully gains access to the system. Intrusion detection systems are an active area of research [3], [8], but an attacker may be able to evade or disable the system as well.

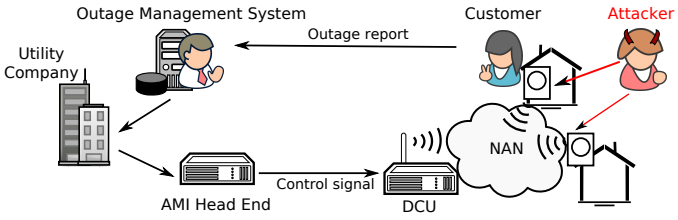

Fig. 1. Simplified smart grid environment with AMI, an outage management system and an attacker sending remote disconnect commands to meters.

For this reason, we focus on a more fundamental countermeasure in this paper—a random time delay for all RCD operations. This delay should be implemented within each smart meter, and configuration of delay parameters should ideally require physical presence (e.g., by turning a knob or requiring a local connection to a dedicated component) to prevent an attacker from changing them remotely. Even if the IDS or authentication countermeasures fail, this delay mechanism makes a smart grid more resilient by: 1) preventing rapid changes in load which may destabilize the power system (see Section VI-B); and 2) providing time for a utility to potentially detect and stop an attack in progress.

The specifics of these detection and recovery mechanisms depend on the nature of the compromise and on the utility's system. In this paper we analyze one possible scenario, and leave more detailed discussion for future work. In particular, since our modeled attacker may interfere with AMI network traffic, we consider a detection mechanism involving a utility's outage management process (see Fig. 1).

If an attacker successfully disconnects a customer's smart meter, the homeowner will likely contact the utility company and report the loss of power within a reasonable amount of time. If the attacker targets a large number of customers, a utility's OMS will begin to receive a significant number of reports, which would trigger an investigation. Once the investigation concludes that the outages are malicious, the utility can respond. Here, we assume that the utility is able to trigger a *fail-safe* mode remotely, canceling all pending meter disconnect requests. This trigger could be sent through the primary communication channel, or if this is fully controlled by the attacker, a secondary channel could be used. One possibility is the use of secondary or reconfigurable radios [9]. Another is manual activation by the user—similar to the manner in which pre-pay meters can provide a small amount of "reserve" energy if the user presses a button [1].

The presence of humans in the loop makes this countermeasure challenging for an attacker to circumvent—provided the delay period is long enough for this detection/intervention process to unfold. This raises the question of how a delay affects the day-to-day operations of utilities, who use RCD for many different applications. We investigate these issues in the following sections by discussing the impact of RCD delay

TABLE I
SMART METER REMOTE CONNECT/DISCONNECT USE CASES

No.	Use Case	Utility Benefit	Time Requirement	Alternatives
1	Routine service switching	cost savings	hours	Manual switch
2	Non-paying customers	employee safety, revenue assurance	hours	Manual switch
3	Demand limiting	demand-side management	minutes–hours	Dynamic pricing
4	Load shedding	lower-level load control	minutes–hours	Substation breakers

on utility operations, designing delay mechanisms that make the best use of the allowable time flexibility, and assessing how these countermeasures improve the security posture of a smart grid system under sophisticated attack scenarios.

IV. IMPACT OF DELAY ON RCD VALUE PROPOSITION

In this section, we discuss four utility RCD use cases (summarized in Table I) and estimate their sensitivity to timing delays. Using these estimates, we develop delay countermeasures in Section V which improve system resilience with little impact on utilities' day-to-day operations.

Routine service switching. When a customer moves in or moves out, the utility needs to do a routine service toggling. Utilities can significantly reduce operating costs by carrying out such operations remotely, rather than dispatching service technicians. While this remote service switching is typically carried out in a matter of minutes [10], [11], it is not particularly time-critical. If a residence has been vacated, there will be very little residual power demand. Furthermore, while move in/out dates are generally known to a utility days in advance, a precise time of day is not. Therefore, it is reasonable to expect an RCD window of up to several hours (e.g., between 12–4pm) for this application.

Non-paying customers. Without remote disconnect capability, the utility has to deploy service people to disconnect non-paying customers. Sometimes, these service people are then physically threatened when they attempt to access the meter. Hence, remote disconnect capability provides greater revenue assurance for a utility, as well as an important safety benefit. If a customer with an RCD-enabled smart meter doesn't pay the bills, the utility can use either a "hard" switch-off, or a "soft" switch to pre-payment metering [1]. As with routine service switching, a transition to pre-pay mode (or a service shut-off) is not particularly time-critical. Without RCD, it could easily take hours (or even days) to identify a non-paying customer, schedule a manual disconnect, and dispatch a service truck to that particular neighborhood. Therefore, this use case should also allow an RCD window of a few hours.

Demand limiting. Many meters have a demand limiting mode, which penalizes customers (either by temporary disconnection, or by enforcing a higher electricity price) who exceed a predefined maximum demand level. When assessing the RCD timescale for demand limiting, it is important to understand that *demand* is defined as the customer's load averaged over a specified amount of time [12]. In practice, this time interval would likely be between 5 and 30 minutes. If a smart meter detects a demand limit violation in one time interval, it should be permissible to enforce the appropriate consequence at any time within the next interval. However, longer delays before shutoff may be possible in certain cases, such as for a government enforcing a mandatory per-household demand cap [1]. In such scenarios, the meter disconnect acts more as a deterrent than a real-time response mechanism.

Load shedding. Although utilities have always been able to shed load by opening substation circuit breakers, the presence of RCD-enabled smart meters allows for a more precise response. The execution time requirement for load shedding differs depending on the application. For example, in emergency situations the utility must be able to shed load in near real-time. In a system with supply shortages necessitating rolling blackouts, RCD-enabled smart meters enable more precise load control with less stringent time constraints.

V. DESIGNING DELAY COUNTERMEASURES

Our study in Section IV suggests that introducing delay before executing RCD commands would have little impact on the intended use cases, as long as the delay is below a suitable threshold. We denote this maximum allowable RCD time delay as d_{max}, which is likely to vary from tens of minutes to 1 or 2 hours based on the application scenarios. In this section, we investigate how to use this allowable period of delay to improve the resilience of the smart grid.

A. Delay Countermeasure Model

In our analysis, we consider a utility company with n RCD-enabled smart meters. The utility's goal is to minimize the total number of disconnected meters by designing an appropriate time delay distribution over $[0, d_{max}]$. As discussed in Section III, we consider a detection process involving a utility's outage management system, which receives outage reports from customers. We develop two models for the OMS detection process, with different levels of abstraction: a simplified model for analytical evaluation (Section V-B), and a more detailed model for simulation (Section VI-A).

In the simplified model we assume that the utility is able to detect and stop an attack in progress within d minutes ($d < d_{max}$) once τ *meters* have been disconnected in a certain time window. In this model, d includes the time necessary for affected (i.e., disconnected) customers to report the outage, and the time needed for the utility to investigate, conclude an attack is taking place, and send out fail-safe commands (see Section III) stopping all pending disconnections.

In the more detailed model, we assume the attack is detected once τ' *outage reports* have been received from customers via the OMS during a certain time window (individual customer reporting times are sampled from a specified distribution).

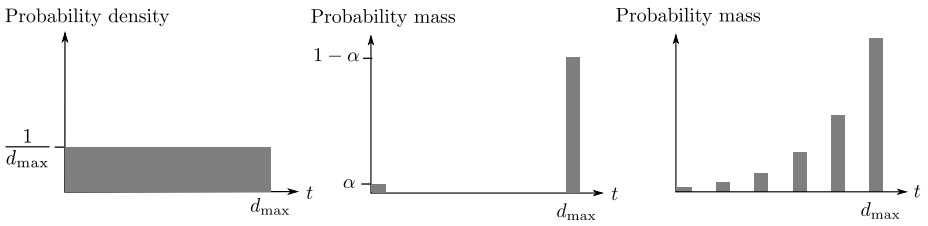

(a) Uniform distribution (b) Bernoulli distribution (c) Geometric distribution

Delay mechanism	Number of disconnected meters	
	Analytical bound	Example
Uniform	$\tau + dn/d_{\max}$	$\approx 100,000$
Bernoulli	$2\sqrt{\tau n} - \tau$	$\approx 48,000$
Geometric (with $\beta = 3$)	$4\tau + \frac{2(n-\tau)}{3^{\lceil d_{\max}/d \rceil}-1}$	$\approx 5,600$

(d) Summary of analysis with an example setting of d=10 min, $d_{\max}$ = 60 min, τ=1000, and n=600, 000

Fig. 2. Three RCD delay mechanisms and the summary of their performance (our analysis assumes $d < d_{max}$).

Once this condition is met, the utility takes an additional time, d', to investigate before sending out fail-safe commands.

In either case, the value of τ (or τ') should correspond to a small fraction of the total number of meters, but it must be large enough to exclude most non-security-related outages. The response time d (or d') should be around the same order of d_{max} (i.e., tens of minutes or hours). We set the time window for detection according to the attacker's target time window, which is assumed to be on the order of hours.

B. Delay Mechanisms

The performance of the RCD time delay countermeasure depends on the delay distribution, which must be designed against an attacker who can select the number of meters to target and schedule when to send disconnect commands. In this section, we assess three delay distributions—uniform, Bernoulli, and geometric—and develop analytical bounds for the number of meters an attacker can successfully disconnect under her optimal strategy (proofs are omitted due to space limitations). Fig. 2 (a)–(c) shows the probability density / mass functions of the three delay mechanisms.

Uniform delay mechanism. Under this basic delay mechanism, once a meter receives an authenticated disconnect command, it selects a random back-off delay uniformly from $[0, d_{max}]$. Each meter selects its back-off delay independently and disconnects itself after its selected delay.

Examine the moment when τ meters have just disconnected themselves. Under our detection and response model, the utility would be able to stop the attack within another d minutes, canceling all disconnect commands that are still outstanding by then. During these d minutes, the attacker can continue to disconnect an additional nd/d_{max} meters on expectation, assuming that an attacker has already sent out disconnect commands to all meters. As there is no better attack strategy, this gives us an analytical bound on the total number of disconnected meters at $\tau + dn/d_{max}$, which grows linearly with n given a constant d/d_{max}.

Bernoulli delay mechanism. An alternative delay distribution to reduce the number of disconnected meters is the Bernoulli distribution. Basically, each meter tosses a biased coin to choose between two possible delay intervals, each spanning at least tens of seconds to avoid stability issues (see Section VI-B). Specifically, with probability α (a system parameter), the meter disconnects itself almost immediately. Otherwise, it postpones its actual disconnection toward the end of the allowed period (i.e., delay $\approx d_{max}$). The intuition behind this design is simple: if a large number of meters are attacked, the meters that disconnect themselves almost immediately would likely trigger OMS detection, which in turn stops the ongoing attack for the larger fraction of meters that postpone their disconnections.

Here the best strategy of an attacker is no longer to send disconnect commands to all n meters in one shot: sticking to this strategy would allow the utility to set α to a very small value and thus reduce the number of disconnected meters. Instead, an attacker can maximize the number of disconnected meters by launching a two-batch attack: the attacker chooses the total number of meters in a first batch such that the subset of immediately disconnected meters in this batch will not trigger the utility response. Around the moment when the subset of meters in the first batch that choose to postpone begin to disconnect themselves (and will soon trigger the utility response), the attacker sends out commands to all the other meters. Under such an attack, all meters in the first batch and on expectation a subset of α fraction of meters in the second batch would have disconnected themselves before pending disconnect commands are called off. Summing up these two terms gives an analytical bound on the total number of disconnected meters at $\tau/\alpha + (n - \tau/\alpha)\alpha$. The utility could strategically choose $\alpha = \sqrt{\tau/n}$ to minimize this analytical bound at $2\sqrt{\tau n} - \tau$, which grows at a sub-linear rate with n.

Geometric delay mechanism. If the maximum RCD delay d_{max} is a few times larger than d, a utility can further reduce the number of disconnected meters by applying a geometric distribution over the allowable delay period. Let $k = \lceil d_{max}/d \rceil$. Under this delay mechanism, a meter will choose among k possible delay intervals, with the ith ($i = 1, \ldots k$) delay interval spanning around $(i-1)d/d_{max}$ and chosen with a probability of $(\beta^i - \beta^{i-1})/(\beta^k - 1)$. Here, β is a system parameter that can be optimized according to k and other parameters (i.e., n and τ). Note that when $k = 2$, the optimization of β would degenerate the geometric delay mechanism to the Bernoulli delay mechanism.

With this geometric distribution and our model, regardless of what strategy an attacker has chosen, on expectation less than $\tau(1+\beta) + (n-\tau)(\beta-1)/(\beta^k - 1)$ meters would be disconnected. The first term bounds the number of disconnected meters that have been issued disconnect commands before τ meters are disconnected, and the second term bounds the number of meters that are issued disconnect commands after

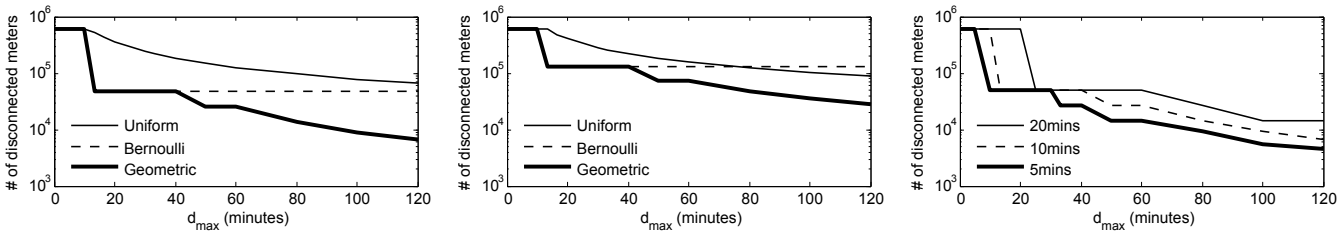

(a) Default simulation setting (b) A larger detection threshold ($\tau' = 9600$) (c) Varying d' (with geometric delay mechanism)

Fig. 3. Number of meters disconnected by an attacker, varying the maximum allowable RCD delay time d_{max}. We utilize logarithmic scales to illustrate a decrease of several orders of magnitude with increasing d_{max}.

that. With a constant β, the first term grows linearly with τ, which is a small fraction of n, and the second term decays geometrically with k, which will also become a small fraction of n even for a small value of k (e.g., $k = 6$).

Analysis Summary. Fig. 2 (d) summarizes our analytical results for the three delay mechanisms. For each, the number of disconnected meters grows at a different asymptotic rate with respect to n. As illustrated in the given example setting, under the geometric delay mechanism, as long as d_{max} is a few times (e.g., $6\times$) larger than d, the number of disconnected meters would be small ($< 1\%$ of n). In comparison, under the uniform delay mechanism, a significantly larger number of meters ($> 10\%$ of n) would be disconnected. For the Bernoulli delay mechanism, the number of meters disconnected is on the order of $\sqrt{\tau n}$, thus it highly depends on the detection threshold τ. When τ is around 0.2% of n (as in the example setting), the number of disconnected meters can still reach around 8% of n. A comparison of these three delay mechanisms suggests: 1) if possible, a utility company should make d_{max} a few times larger than d, and use the geometric delay mechanism; 2) otherwise, if $\lceil d_{max}/d \rceil = 2$ and the geometric distribution mechanism degenerates to the Bernoulli one, a utility needs to make its detection threshold τ as small as possible to promote its resilience under RCD attack.

VI. Resilience with RCD Time Delay

In this section, we assess the resilience of the smart grid to an RCD attack, in terms of both the capability of the utility to detect and stop an attack in progress, and the ability of the underlying power system to absorb such attacks.

A. Resilience to Denial-of-Electrical-Service Attack

We have used a simplified detection and response model to help design the delay mechanisms in Section V-B. We now use a more detailed simulation model to validate our design and analytical results, and to gain further insight into the effects of model parameters τ' and d'.

Simulation setting. In our default simulation setting, we consider an attacker who can control all the $n = 600,000$ meters of a utility. She adopts different attack strategies according to the deployed delay mechanism, with an intent to maximize the total number of meters she can disconnect (see Section V). The random outage reporting time for affected homeowners is modeled using a shifted exponential distribution with a cumulative distribution function of $F(x) = (1-e^{(1-x)/10})H(x-1)$, where x is measured in minutes and H is the Heaviside step function. The 1 minute time shift here captures the minimum delay incurred when a homeowner carries out initial investigation and accesses her reporting channel (e.g., phone). A similar reaction delay has been observed in studies of earthquake reporting through social media [13].

For the utility, its default detection threshold, τ', is set to 1200 (or 0.2% of n) according to the trace in [14]. This threshold would prevent 95% of non-security-related outages in the trace from triggering false alarms. The additional utility response time, d', is set to 10 minutes by default.

Simulation results. While our simulation model captures more operational details than our analytical model, Fig. 3(a) reveals similar findings: a maximum RDC delay d_{max} of up to 2 hours can significantly reduce the number of disconnected meters, by providing time for a utility to detect and stop an attack in progress. Furthermore, the geometric delay mechanism performs the best among the three studied. When d_{max} is below 40 minutes, the geometric delay distribution degenerates to the Bernoulli delay distribution. However, further increasing d_{max}, allows the geometric delay mechanism to outperform the latter, and for $d_{max} = 120$ minutes, the number of disconnected meters under the geometric delay mechanism is only around 1% of n.

In Fig. 3(b), we change the default setting by having a larger detection threshold of $\tau' = 9600$, which would prevent 99% of non-security-related outages in [14] from triggering false alarms. The geometric distribution still performs the best among the three. A comparison of Fig. 3(a) and Fig. 3(b) shows that a larger detection threshold τ' results in a larger number of disconnected meters. For the geometric distribution, reducing the false positive rate from 5% to 1% results in about four times as many affected meters. This suggests that the nature of a utility's non-security-related outages affects its resilience to RCD attack through τ'.

Finally, Fig. 3(c) evaluates the geometric delay mechanism under varying utility response delay d'. It shows that the number of disconnected meters becomes smaller with a smaller value of d'. The reduction is most significant for some relatively small value of d_{max} (e.g., when $d_{max} = 20$).

In summary, we find that random time delay countermeasures greatly reduce the impact of an RCD attack. For all simulated delay distributions and utility parameters (τ', d'), the presence of time delay limited the number of realized meter disconnections to less than 22% of the $600,000$ meters

targeted. In the best case, a geometric delay with $d_{max} = 120$ minutes, around 1% of targeted meters are disconnected. As discussed in Section IV, a delay of this magnitude would have little effect on a utility's typical RCD use cases.

B. Resilience to Attacks on Power System Stability

So far, we have discussed smart meter RCD attacks targeted at denial-of-electrical-service, for economic or other reasons. However, another potential attacker goal is the manipulation of physical power system characteristics, such as frequency, via targeted meter on/off switching [2].

For an RCD attack to cause transient stability problems, a large amount of load must be shed instantaneously. The authors of [15] simulate malicious load shedding and re-application (without delay countermeasures) in the Western Electricity Coordinating Council region and find that—while the system remains stable—reliability criteria can be violated. Implementing randomized time delay countermeasures for smart meter RCD can enhance the power system's resilience to these attacks in two important ways. First, delay mechanisms reduce the scale of load shedding attacks, as we show in Section V-B. Second, the random delay makes attacks less precise, as the attacker cannot control in detail when certain meters will disconnect. However, for the proposed geometric and Bernoulli delay mechanisms, care must be taken to spread the load shedding in each discrete interval over a small period of time. A study in [16] shows that, in the IEEE 14-bus system, spreading load drop over tens of seconds to a minute is sufficient to ensure stability.

Based on the simulation from the previous section, we can analyze the power system stability impact of our delay countermeasures (we leave a more detailed analysis for future work). Implementing any of the delay mechanisms in Section V-B extends the timescale of an RCD load-shedding attack from seconds to minutes (or hours), which gives system operators a chance to re-balance generation and load through automatic generation control and/or re-dispatch. Consider the results shown in Fig. 3(a). With a geometric delay distribution, and d_{max} of only 20 minutes, an attacker is limited to shutting off roughly $47,450$ out of $600,000$ meters. Assuming each meter's demand was a conservative 5kW at disconnection (see [12], [17] for typical demand traces), the total load loss is under 240MW. Modern Combined Cycle power plants can ramp as fast as 50MW per minute [18], and hydroelectric plants can respond faster [19]. Therefore, a large power system can adequately respond to such attacks within minutes.

VII. Conclusion

Smart meters with RCD capability introduce a potentially serious cyber threat, which must be properly understood and defended against. In this paper, we show that a properly designed random time delay countermeasure makes a smart grid more resilient by: 1) preventing rapid changes in overall system load, and 2) providing time for a utility to potentially detect and stop an attack in progress. We design three delay distributions and show that time delays of two hours or less have little impact on utilities' day-to-day operations, but "make a difference" during an attack. In particular, our simulation results show that a geometric delay distribution over $[0, 2]$ hours can reduce the number of realized meter disconnections to around 1% of the number of meters being targeted.

Acknowledgements

This work is supported by Singapore's Agency for Science, Technology, and Research under the Human Sixth Sense Programme. We thank W. H. Sanders for his helpful feedback.

References

[1] R. Anderson and S. Fuloria, "Who controls the off switch?" in *Proc. of the Conference on Smart Grid Communications (SmartGridComm)*, 2010.

[2] M. Costache, V. Tudor, M. Almgren, M. Papatriantafilou, and C. Saunders, "Remote control of smart meters: friend or foe?" in *Proc. of the European Conference on Computer Network Defense (EC2ND)*, 2011.

[3] D. Grochocki, J. H. Huh, R. Bobba, W. H. Sanders, A. A. Cardenas, and J. G. Jetcheva, "AMI threats, intrusion detection requirements and deployment recommendations," in *Proc. of the Conference on Smart Grid Communications (SmartGridComm)*, 2012.

[4] "GE I-210+ remote disconnect FAQs," http://www.dbmss.ca/PDF%20files/I210+_and_I210+c_Remote_Disconnect_FAQ_2.pdf.

[5] U.K. Gov. Dept. of Energy & Climate Change, "Government response to the consultation on the second version of the smart metering equipment technical specifications," 2013.

[6] U.S. Resilience Project, "Smart grid cyber security & interoperability requirements," http://www.usresilienceproject.org/workshop/participants/pdfs/Appendix_B_Smart_Grid_Security_Requirements.pdf.

[7] M. Nabeel, S. Kerr, X. Ding, and E. Bertino, "Authentication and key management for advanced metering infrastructures utilizing physically unclonable functions," in *Proc. of the Conference on Smart Grid Communications (SmartGridComm)*, 2012.

[8] R. Berthier, W. H. Sanders, and H. Khurana, "Intrusion detection for advanced metering infrastructures: Requirements and architectural directions," in *Proc. of the Conference on Smart Grid Communications (SmartGridComm)*, 2010.

[9] A. Ghassemi, S. Bavarian, and L. Lampe, "Cognitive radio for smart grid communications," in *Proc. of the Conference on Smart Grid Communications (SmartGridComm)*, 2010.

[10] "Avista utilities update to idaho public utility commission staff on remote reconnect/disconnect pilot," http://www.puc.idaho.gov/fileroom/cases/elec/AVU/AVUE0709/company/20130211UPDATE%20ON%20REMOTE%20RECONNECT%20DISCONNECT.PDF.

[11] "Texas-new mexico power company's request for approval of an advanced metering system (ams) deployment and ams surcharge," http://www.smartgrid.gov/sites/default/files/doc/files/TexasNew_Mexico_Power_Company_Request_For_Approval_Advance_201005.pdf.

[12] W. H. Kersting, *Distribution system modeling and analysis*. CRC Press, LLC, 2012.

[13] A. Crooks, A. Croitoru, A. Stefanidis, and J. Radzikowski, "#earthquake: Twitter as a distributed sensor system," *Transactions in GIS*, vol. 17, no. 1, pp. 124–147, February 2013.

[14] "PEPCO tracker outage summary," http://pepcotracker.com/summary.csv.

[15] S. Clements, H. Kirkham, M. Elizondo, and S. Lu, "Protecting the smart grid: A risk based approach," in *IEEE Power and Energy Society General Meeting*, 2011.

[16] X. Lou, D. Yau, H. Nguyen, and B. Chen, "Profit-optimal and stability-aware load curtailment in smart grids," *IEEE Transactions on Smart Grid, to appear*.

[17] S. Barker, A. Mishra, D. Irwin, P. Shenoy, and J. Albrecht, "SmartCap: Flattening peak electricity demand in smart homes," in *PERcom*, 2012.

[18] "Flexefficiency 50 combined cycle power plant," http://www.ge-energy.com/products_and_services/products/gas_turbines_heavy_duty/flexefficiency_50_combined_cycle_power_plant.jsp.

[19] N. Jaleeli, L. S. VanSlyck, D. N. Ewart, L. H. Fink, and A. G. Hoffmann, "Understanding automatic generation control," *IEEE Transactions on Power Systems*, vol. 7, no. 3, pp. 1106–1122, August 1992.