

Final Technical Report

Project Title: Secure, Scalable Control and Communications for Distributed PV

Project Period: 12/1/15-11/30/18

Submission Date: 1/15/19

Recipient: Sandia National Laboratories
Address: 1515 Eubank SE
Albuquerque, NM 87123

Website: www.sandia.gov

Award Number: DE-EE0001495-1593

Project Team: Sandia National Laboratories
DNK Consulting
Montana Tech

Principal Investigator: Jay Johnson
Phone: (505) 284-9586
Email: jjohns2@sandia.gov

Business Contact: Shannon Boynton
Phone: (505) 284-2631
Email: sboynto@sandia.gov

Technology Manager: M. Kemal Celik, Ph.D.

Project Officer: Thomas Rueckert



Sandia National Laboratories

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Executive Summary

An increasing number of jurisdictions are adopting Distributed Energy Resource (DER) interconnection standards which require photovoltaic (PV) inverters, energy storage systems, and other DER to include interoperable grid-support functionality. These functions provide grid operators the tools to support local and bulk power system operations with DER equipment, but the associated grid operator-to-DER communications networks must be deployed with appropriate cybersecurity features. In some situations, additional security features may prevent control system scalability or increase communication latencies and dropouts. These unintentional consequences of the security features would therefore hinder the ability of the grid operator to implement specific control algorithms. This project evaluated the tradeoffs between power system performance and cybersecurity metrics for several grid services. This was conducted in two parts. First, the following distributed control algorithms were studied to determine the power system performance impact from latency, dropouts, and availability:

- **Distribution:** Voltage regulation with a volt-var shift algorithm
- **Transmission:** Synthetic inertia
- **Transmission:** Communication enabled fast acting imbalance reserve
- **Transmission:** Communication enabled frequency droop

Using parametric power system simulations, critical thresholds for communication delays, equipment availability, and packet loss were determined.

Second, the team developed multiple cybersecurity reference architectures for DER communication networks that included network segmentation, bump-in-the-wire encryption, and moving target defense. These architectures were built in a communication network and power co-simulation environment called SCEPTRE at the Distributed Energy Technologies Laboratory (DETL) at Sandia National Laboratories. These simulations included actual protocol traffic through simulated remote terminal units (RTUs) and other Linux and Windows Virtual Machines (VMs) representing networking equipment, DER, and the utility. Cyber resilience was quantified with an adversary-based assessment methodology for each of the architectures containing the security features (enclaving, encryption, moving target defense). The communication latency of these networks were also measured.

The team was then able to compare cybersecurity metrics to network latency and power system performance. It was found that the distribution voltage regulation technique was robust to long DER communication latencies and therefore any of the security features could be implemented with no impact on the results. The transmission-level grid services were more sensitive to network latency, however, the latency associated with each of the security features (<10 ms increase) would not impact the power system performance substantially. In fact, the low-level DER Modbus reads/writes were by far the greatest percentage of the total communication time, which cannot be reduced by stripping any security features.

In contrast, adding the network security features, drastically improved confidentiality, integrity, and availability of the DER network. The network segmentation was effective at preventing an adversary from freely traversing the DER network and controlling the entire DER aggregation. Encryption provided confidentiality of the network so man-in-the-middle attacks could not be conducted readily. And the moving target defense was effective at preventing replay attacks.

Based on these findings, applying all the security features to DER networks will help protect DER communication networks and power system operations with minimal impact to communication latency or power system performance. It is recommended that utilities, grid operators, and DER aggregators employ network segmentation, encryption, and moving target defense to provide secure distribution and transmission grid-support services.

Contents

1	Background	5
2	Project Structure	6
3	Distributed DER Control Algorithms	8
3.1	Communication Enabled Synthetic Inertia	8
3.1.1	Introduction	8
3.1.2	Model and Simulation	9
3.1.3	Effects of Gain on System Response	10
3.1.4	Effects of Latency on System Response	12
3.1.5	Effects of Availability on System Response	13
3.1.6	Conclusions and Future Work	16
3.2	Communication Enabled Fast Acting Imbalance Reserves	17
3.2.1	Introduction	17
3.2.2	Description of CE-FAIR	18
3.2.3	Results and Discussion	19
3.2.4	Conclusions	20
3.3	Communication Enabled Frequency Droop	20
3.3.1	Introduction	20
3.3.2	CE Droop vs Conventional Droop	21
3.3.3	Effect of Variations in Time Delay (τ_d)	23
3.4	Distribution: Hierarchical Control of Volt-VAr Function for Steady-State Voltage	25
3.4.1	Introduction	25
3.4.2	Test Setup	26
3.4.3	Hierarchical Voltage Control	27
3.4.4	Simulation and Communication Results	29
3.4.5	Conclusions	32
4	Cybersecurity Reference Architectures and SCEPTRE Deployment	33
4.1	Power Simulations	33
4.2	SCEPTRE	34
4.2.1	Power Simulation Interaction	36
4.2.2	Remote Terminal Units	36
4.2.3	Security Mechanisms	37
4.2.4	Topologies	42
5	Communication Latency	47
5.1	Communication Latency (Emulated)	51
5.1.1	Network Segmentation	51
5.1.2	Encryption	53
5.1.3	Moving Target Defense	53
5.2	Communication Latency (Physical)	55
5.2.1	Geographic Separation	55
5.2.2	Smart Inverter Read and Write Times	55
5.3	Latency Observations	57

6	Red Team Assessments	58
6.1	Scope and Rules of Engagement	58
6.2	Methodology	58
6.3	Tools	59
6.4	Emulytics Challenges	59
6.5	Threat Catalog	60
6.6	Results	61
6.6.1	Reconnaissance	61
6.6.2	Packet Replay	64
6.6.3	Denial of Service	67
6.6.4	Man in the Middle	67
6.6.5	Flat network topology without encryption	67
6.6.6	Flat network topology with encryption	67
6.6.7	Segmented-Network Topology without Encryption	68
6.6.8	Segmented-Network Topology with Encryption	68
6.6.9	Segmented-Network Topology with HIL and without Encryption	68
6.6.10	Moving Target Defense Network Topology	68
6.7	Comments	69
6.8	Summary	69
7	Challenges with Operating the Power System with No Communications	70
7.1	Problem Conceptualization	70
7.1.1	Dispatch Drift without Communications	71
7.2	Recommendations	74
8	Significant Accomplishments and Conclusions	77
9	Inventions, Patents, Publications, and Other Results	78
10	Path Forward	80
	References	81

1 Background

There is ample evidence from the last few years that many power system networks in the US [1, 2, 3, 4, 5, 6] and abroad [7] are the target of active cybersecurity reconnaissance and attacks. The most widely publicized attacks are those that caused widespread blackouts in Ukraine in 2015 and 2016 [8, 9], but there have been several other disconcerting trends including: the increase in operation technology (OT)-focused malware, e.g., Crash Override and Black Energy [10, 11], deep reconnaissance into power system networks [12, 13, 14], and growing willingness to deploy powerful cyber weapons that are affecting critical infrastructure [8, 9, 15]. Attackers often use myriad techniques to gain footholds in information technology (IT) networks and then pivot to other computers, servers, and networks to exfiltrate sensitive information, monitor operations, or plan for sophisticated attacks [16].

At the same time, penetrations of Distributed Energy Resources (DER)—e.g., Photovoltaics (PV) and Energy Storage Systems (ESS)—continue to grow rapidly on distribution and subtransmission systems [17, 18]. Over the last decade, an increasing number of inverter vendors and aggregators have provided monitoring portals for their customers. Like many other Internet of Things (IoT) devices, customers can monitor or control their equipment via proprietary communication protocols. However, this IoT equipment now controls a substantial portion of the total power production in certain jurisdictions, like Hawaii and California [19, 20].

In 2018, a revision to the US interconnection and interoperability standard, IEEE Std. 1547, required DER equipment to have either an IEEE 2030.5 (SEP 2), IEEE 1815 (DNP3), or SunSpec Modbus communication interface [21]. New California Public Utility Commission (CPUC) Electric Rule 21 regulations that go into effect in early 2019 define IEEE 2030.5 [21] as the default application protocol for Investor Owned Utilities (IOUs) communications to DER [22, 23]. The adoption of standardized communication protocols is a critical step toward interoperability between power system operators and DER equipment, but a comprehensive national approach to DER cybersecurity is absent. At this point, the network architecture and cybersecurity requirements are not fully defined for CA, and other states will also have to make similar decisions as utilities and other grid operators start interacting with interoperable DER. Herein, the team researched communication requirements for different grid applications in order to generate reference architectures and cybersecurity recommendations for discussions in California and elsewhere regarding the best methods to take advantage of increasing PV penetrations and advanced inverter capabilities.

There are many security requirements for operators of critical infrastructure in the U.S. Power system operators are required to adhere to the North American Electric Reliability Corporation (NERC) Critical Infrastructure Protection (CIP) standards which cover—among other things—training, security and information management, perimeter defenses, and incident reporting [24]. NREC requirements are reserved for bulk power equipment operating at or above 100 kV, so DER equipment and associated networks are exempt from these requirements. The solar industry and national government understand this gap in power system security and are working to address the requirements by reviewing and updating security requirements in the DER communication protocols [25, 26], standing up DER cybersecurity working groups [27], and seeking new security standards for DER devices and networks [28].

There are a wide range of R&D areas that may improve the national DER cybersecurity posture [29]. In this work, three network defense techniques were analyzed with respect to power system performance and security tradeoffs. Specifically, network segmentation, encryption, and moving target defense (MTD) were deployed in a virtualized environment to (A) calculate the additional communication latencies associated with these features, (B) determine the impact these would have for distribution- and transmission-level grid services (e.g., voltage regulation, frequency reserves, protection, etc.), and (C) evaluate any security improvements in the broad areas of confidentiality, integrity, and availability by conducting adversary-based (red team) assessments. This work produced power system performance and cybersecurity metrics to advise the solar and power system industry on best cybersecurity practices for DER networks.

2 Project Structure

The ultimate goal of the project was to enable high penetrations of solar generation on the grid (greater than 100% of peak load), while maintaining or improving grid performance, reliability, and security. With any system, there are tradeoffs. Dithering any of the above metrics will result in a different level of performance, reliability, cost, and security. One of the advantages of widely distributed solar generation is the inherent reduction in risk from the loss of a single asset. While there will clearly be large plants, and many smaller sources (e.g., behind-the-meter solar), the benefit of a large distributed system is that the loss of some fraction of plants can easily be tolerated, especially if the remaining sources are dispatchable (e.g., via headroom provided by curtailment or energy storage), or load can be modulated.

To understand these tradeoffs, the project was divided into four components. The first piece of work took a closer look at the DOE SETO communications target metrics (below) by clearly articulating the impact of communication Quality of Service (QoS) on grid operations using distributed control and communications architectures. Depending on the application, some communication metrics may be relaxed significantly, resulting in significant cost savings. For other applications, the metrics are not sufficient to maintain or improve the stability and security of the power grid with very high penetrations of PV generation.

- Scalability: up to 5,000,000 nodes
- Availability: $> 99.999\%$
- Response time: < 1 second
- Cost: LCOE < 6 cents/kWh by 2020
- Interoperability: compliance with open standards

These metrics were studied for the applicable, aforementioned distributed controls algorithms are discussed in Section 3. Results from these power system simulations provided detailed insight into how power system operations are affected by QoS (i.e., availability, packet loss, and communications latency). Generally, the distribution controls were resilient up to 20 s of communication delay, but the transmission-level controllers were impacted between 110-400 ms.

The second component of the project was to develop and deploy multiple cybersecurity reference architectures for DER networks. The reference architectures used three different cybersecurity features: network segmentation, encryption, and moving target defense. The DER networks currently being deployed in the U.S include encryption, but generally do not include segmentation or moving target defense. However, additional security features are being used in these 'real-world' deployments including Public Key Infrastructure (PKI)/node authentication (see details on trust mechanisms in [25]). These reference architectures were constructed in a networking and power co-simulation environment called SCEPTRE in the Distributed Energy Technologies Laboratory (DETL) at Sandia National Laboratories in Albuquerque, NM. SCEPTRE was updated to work with OpenDSS, a distribution system simulation program, so the voltage regulation algorithm could be simulated accurately. PowerWorld simulation test cases were also created for the transmission control algorithms, but those were not used for any of the red team assessments.

The third component of the research was to conduct adversary-based assessments of the cybersecurity reference architectures in SCEPTRE to quantify the cybersecurity resilience. Determining resilience metrics is challenging, and a research field of it's own, but the team scored the results of different attacks on the impact to confidentiality, integrity, and availability (the cybersecurity CIA triad, which represents the fundamental needs of a communication system). Specifically, the red team conducted network reconnaissance of each of the experimental architectures and launched a range of attacks on this network, including:

- Data Compromise: alteration or access of confidential of data by unauthorized users.

- Code Injection: malicious code injection, such as inputs to the human-machine interface (HMI).
- Local Exploits: exploiting vulnerabilities in user applications, such as the SunSpec Dashboard and SunSpec System Validation Platform (SVP).
- Remote Exploits: exploiting vulnerabilities in network services.
- Privilege Abuse: exploiting existing privileges for authorized users on the system.
- Interception: man-in-the-middle or eavesdropping of authenticated communications.
- Denial of Service: rendering the system unusable to authorized users, such as overloading the RTU processors.
- Policy: exploiting flaws in policy, such as firewall security settings.
- Insider Threat: exploiting authorized user knowledge or access for malicious purposes.

The results of these attacks were used to score the architecture. SCEPTRE was also used to determine the additional latency associated with adding these security mechanisms. Fig. 1 shows conceptually how SCEPTRE was used to generate power system and latency data, along with the cybersecurity metrics for a given cybersecurity reference architecture.




Figure 1: Conceptualization of the SCEPTRE Environment Providing Cybersecurity Metrics and Power System Performance Data.

The final component of the project was to study the difficulties associated with operating the power system under low or no communications scenarios. As the U.S. moves toward a smarter and more interoperable power system, grid operations become more efficient, economical, and autonomous. But if communications are lost, grid operators must revert to safer modes of operation in order to provide power to customers. It was found that if communications were lost at minimum load in the Western Interconnect, generators would increase their output to match load without consideration of the security-constraints on the transmission systems. This would overload certain paths, tripping them and leading to a cascading blackout.

3 Distributed DER Control Algorithms

There has been limited research on the impact of communications delay on power system stability. The majority of papers have focused on small signal stability in the face of communications delay [30, 31, 32]. Likewise, there has been relatively little research on the impact of cybersecurity algorithms on latency, and thus grid stability. One paper addresses expected latencies for secure communications for wide area control [33], although it does not address the stability impact. There have been numerous publications on cybersecurity for the “smart grid”, but none address the impact on stability [34, 35]. Communication needs vary greatly depending on the service provided, location, and topology of the power system. Therefore, this project developed communications and cybersecurity requirements for various applications (e.g., voltage support, renewable firming, transient response, small signal stability, etc.) that provide context to the communications tradeoffs. This can aide DOE and power system industry in specifying the design of networked control systems to meet the needs of the grid in 2030.

While some of the control system and communication analysis can be performed using simplified models and arrive at closed form results, this project focused on detailed analysis performed with simulations of representative systems. For this effort, both transmission and distribution applications were studied. The transmission level simulation and analysis was performed using PSLE and MATLAB. The distribution level simulation and analysis were performed using OpenDSS and MATLAB. The impact of latency and availability is presented in the following sections for the following use cases:

- **Transmission:** synthetic inertia (SI)
- **Transmission:** communication enabled fast acting imbalance reserve (CE-FAIR)
- **Transmission:** communication enabled frequency droop (CE-Droop)
- **Distribution:** hierarchical control of Volt-VAR function for steady-state voltage (VV shift)

It should be noted that there have previously been a number of studies investigating latency-impacts for power systems. A study of implementing optimal fixed structure control to improve small signal stability of the western North American power system (wNAPS) was performed in [36]. Evaluation of the similar target metrics for a central control algorithm using advanced inverter functions at the distribution system level was done in [37]. Other works have also looked at the impact of communication latency on controls for power systems. A study of latency effects on inter-area oscillation damping controls in the wNAPS was done in [38]. A study on characterizing latencies in a wide-area measurement system was performed in [39].

3.1 Communication Enabled Synthetic Inertia

3.1.1 Introduction

One use case of interest is synthetic inertia (SI). Designed to replace the inertial response of traditional generators, this concept may also be applied to photovoltaic generation. Examples of SI in the context of wind generation are found in [40, 41]. Similarly, synthetic inertia provided by energy storage is proposed in [42]. The power electronics interfacing of PV generation allow for a type of rapidly responding imbalance reserve for primary frequency response. In the proposed modification featured in this paper, global system frequency information is used instead of local measurements in the synthetic inertia control law. This concept, referred to as communication enabled synthetic inertia (CE-SI), uses this system frequency for every PV plant in the entire system. The utilization of communication among distributed energy resources could provide the frequency sharing mechanism.

Section 3.1.2 describes the PV plant model and synthetic inertia implementation, the test system that is used, and the simulation performed to analyze controller performance. Section 3.1.3 analyzes SI and

CE-SI performance as a function of controller gain. Section 5 examines how CE-SI performance varies with increased latency. The impact of communication availability is investigated in Section 3.1.5. Finally, conclusions and possible future directions of this research are discussed in Section 3.1.6.

3.1.2 Model and Simulation

This study was conducted using GE's Positive Sequence Load Flow (PSLF) platform. While primarily done with PV generation in mind, the analysis in this study can be easily applied to any form of converter-interfaced generation source. PV generation was modeled using a custom current injection model in PSLF. The conceptual block diagram for CE-SI is shown in Figure 2. For classical implementations of SI, a local frequency with minimal latency is employed. For CE-SI, system frequency measurement information is relayed to all PV plants. The details of the communication scheme (e.g., centralized aggregator or peer-to-peer) are abstracted away and it is assumed that this shared information may come with time delay and is subject to communication intermittency. Furthermore, it is assumed that the plants are operating in a curtailed mode and have the capacity to increase their power output to respond to frequency dips caused by a loss of generation.




Figure 2: Conceptual block diagram representation of CE-SI. The frequency weights k_i are the proportion of the associated generator's inertia to the total system inertia; this summation creates a weighted system frequency over N system generators.

The control law for synthetic inertia, ignoring the low-pass filters, is given by

$$\Delta P = -k_{in} \frac{df}{dt}, \quad (\text{MW}) \quad (1)$$

where f is the frequency employed in the calculation. The units of k_{in} are $\text{MW} \cdot \text{s}^2$. In CE-SI, system frequency is used; this quantity is derived from a weighted average of generator machine speeds where generators with a larger proportion of total system inertia are more heavily weighted. The effect of this averaging is to smooth out frequency changes that tend to be volatile during sudden power imbalances. In either case, the primary goal of synthetic inertia is to emulate the inertial response of synchronous

generators that slows down the rate of change of frequency (RoCoF), particularly the deceleration of machine speeds after generation loss.

To evaluate the effectiveness of CE-SI, this paper analyzes the system response to generation drop events for different configurations. These responses are compared to a case with no control implemented as a baseline. The case with no PV generation, the base case, is used to determine if all of the replaced inertia can be successfully emulated. Finally, the responses for classical SI are analyzed to determine if CE-SI can provide potential improvements.

The test system for this study was based on a representation of the Northeast Power Coordinating Council (NPCC) region, consisting of 140 buses and 48 generators with a total of 28.042 GW of generation. The modified case with PV generation replaces 26 traditional generators with approximately 13.729 GW of generation; this results in a PV penetration level of approximately 49%. The event used to perturb the system is the loss of the "Monroe" unit: a loss of 655 MW or 2.3% of total system generation, occurring at five seconds into the simulation.

3.1.3 Effects of Gain on System Response

In this section, the performance of CE-SI and SI are evaluated for a wide range of k_{in} (gain) values. The system frequency responses for CE-SI with no time delay are shown in Fig. 3. The first second following the generation drop event is shown in Fig. 4 to illustrate how increasing k_{in} mitigates average machine speed deceleration. These results show enhancement in the frequency response compared to the base case.




Figure 3: System response for 100 different values of k_{in} and no latency.

In order to properly evaluate this relationship between latency and gain, the maximum rotor angle separation among machines in the entire system is introduced as a proxy metric. To compute this metric, the rotor angle for each conventional generator in the system is recorded. Using this information, the maximum angle difference among all machines is computed at every point in time and recorded as a time signal, $M_{rot}(t)$. Finally, the maximum value of this signal $M_{rot}(t)$ is determined for each simulation.

$$\text{Max. Rotor Angle Separation} = \max_t \max_{i,j} |\delta_i(t) - \delta_j(t)| \quad (2)$$

Maximum rotor angle separation is useful for evaluating rotor angle stability [43]. Its primary purpose in this analysis is to indicate that a particular simulation encountered a loss of synchronism. The maximum




Figure 4: Zoom-in of Fig. 3 to illuminate changes in machine speed deceleration.

rotor angle separation for the data sets are shown in Fig. 5. Two cases with CE-SI are shown, with one including time delay, τ_d .





Figure 5: Maximum rotor angle separation for three implementations of synthetic inertia at different values of k_{in} .


The maximum rotor angle separation was typically observed to occur in the period of steady state simulation prior to the generation drop event (i.e., before five seconds); for the no control case, this fixed value was found to be 72.48° . For SI at all gains, it was observed that the maximum rotor angle separation was unchanged across all tested values of k_{in} . Thus, SI did not encounter any stability issues; this is likely due to the assumption that frequency measurements were assumed to be instantaneous. For CE-SI with no latency, there were marginal increases in this metric with increments in gain. However, for CE-SI with latency, rotor angle separation indicative of loss of synchronism was observed for values of k_{in} that exceed 350.

To determine the effectiveness of SI and CE-SI on the inertial response of the system, the maximum and average RoCoF, in the two seconds following the disturbance, are computed. The results are shown in Fig. 6. Data points that correspond to cases deemed to be unstable are omitted from the figures. Cases where rotor angle separations exceed 180 degrees are considered unstable. As seen in the frequency response time series, increased gain is effective in reducing the maximum RoCoF. For cases with no latency, CE-SI is the

superior choice. Note that both SI and CE-SI are easily able to exceed emulating the inertial response of the base case at appropriate gain levels. However, as latency is added to CE-SI, performance degrades and it is uncertain if CE-SI can do better than SI.



(a) Maximum value of RoCoF.



(b) Time average of RoCoF.


Figure 6: Maximum and time average of RoCoF during the two seconds following the generation drop event for different gain values.

Solely with respect to improving the maximum RoCoF of the system response, CE-SI is successful and even more so with increased gain. However, this benefit is not without drawbacks such as the increased settling time and the possible introductions of oscillations which can be seen at higher gain values. Additionally, as latency increases for CE-SI, instabilities are introduced when the gain is high. As latency increases, the upper limit on k_{in} decreases. This relationship is explored further in the following section.


3.1.4 Effects of Latency on System Response

As observed in the results presented in Section 3.1.3, communication latency plays a role in limiting the effectiveness of CE-SI's emulated inertial response. To explore this relationship in more detail, the system response using CE-SI was simulated for a range of time delay values with fixed values of k_{in} . The RoCoF metrics for the simulations are shown in Fig. 7. Cases determined to be unstable are omitted.

The observed trend is that communication time delay decreases the effectiveness of the emulated inertial response. The examined CE-SI cases easily replicate or improve upon the original inertial response (i.e.,



(a) Maximum value of RoCoF.



(b) Time average of RoCoF.

Figure 7: Maximum and time average of RoCoF during the two seconds following the generation drop event for different time delay values.

that of the base case), especially as gain is incremented. While increased gain values expand the range of latencies where CE-SI is an improvement over the base and no control cases, the results in Section 3.1.3 are a reminder that stability is a concern for higher gain scenarios. The maximum rotor angle separations for the same data is shown in Fig. 8. For the gain values selected, stability issues were observed for the $k_{in} \geq 400$ simulations. Finally, to relate back to SI, the maximum and average RoCoF for the highest gain tested for SI ($k_{in} = 1000$) were 6×10^{-3} Hz/s and 4.1×10^{-3} Hz/s, respectively. Based on the results, it appears that the communication latency must be low for CE-SI to improve over SI. However, in practice, a dedicated communication network is unlikely to have latencies exceeding 50 ms on average [38, 39]. Since these results vary on the disturbance tested, the PV penetration level, and even the system itself, the criteria of CE-SI being a clear improvement over SI is subject to change.

3.1.5 Effects of Availability on System Response

In addition to time delays in communication, availability is a challenge for distributed control schemes. Availability, in this scope, is defined as the probability of a device successfully receiving external communication and having the opportunity to act on the received information. In the case of CE-SI, availability




Figure 8: The maximum rotor angle difference between any two machines during each simulation. Higher k_{in} incurs instability for smaller time delay values.

is the probability of the device successfully receiving system frequency information and computing how much power should be modulated. The communications link was designed for a 16 Hz update rate. In the event of a lost message, the plant maintains its present power output until communication is successfully restored. For cases of low availability, it is possible that consecutive dropouts may occur causing the plant to freeze its output for extended durations.




Figure 9: System responses for availability values ranging from 1-100%.

Sweep of availability levels In Fig. 9, the system response for one realization each of availability levels ranging from 1 to 99% is shown. The k_{in} value for all simulations was fixed at 200 and τ_d was set to zero. It can be observed that the system response is qualitatively similar for availability levels from approximately 40% and up. As availability drops below this threshold, significant variation in system response can be seen.

Fig. 10 shows the active power output of the PV plant installed at the Millstone bus in New England. As with the system frequency, the power output demonstrates very little variation until availability declines




Figure 10: Sample PV plant output for different availability values.

to less than 50%. For these availability levels, extended periods of communication loss are readily apparent as the power output levels off for seconds at a time. Since it is improbable for high availability scenarios to have consecutive occurrences of communication loss, the system response is similar for values above a certain level of availability. The saturated maximum RoCoF curve in Fig. 11 demonstrates this relationship.




Figure 11: Maximum RoCoF value for instances of different availability.

Monte Carlo simulations While the simulations performed for Section 3.1.5 are useful for discerning availability trends at a glance, they are unlikely to capture the expected behavior of the randomness involved. To address this, Monte Carlo simulations were performed at selected availability levels. Fifty simulations each were performed for availability levels of 99%, 75%, 25%, and 1%. Fig. 12 compiles the maximum RoCoF for each run into histograms for each availability level. The bin range and width is fixed for each subplot for consistency. As mentioned in the previous subsection, reduction in availability leads to larger variations in system response. However, increases in variance do not tend to manifest until relatively low availability levels and such levels may be deemed unacceptable for communication standards. Fig. 13




Figure 12: Histograms of maximum RoCoF value using fifty instances each at four different availability levels.

combines the system response time series for each data set using different colors for each instance. As in the previous result, the system frequency responses tend to remain consistent until very low availability levels.

Because the communication messaging rate is much faster than the system dynamics, the system performance consistency can be explained by the "catch-up" mechanism in the synthetic inertia control law. The frequency derivative used is slow, smooth, and continuous except during the instantaneous generation drop event. Dropouts in communication are compensated for when communication is restored and mismatches in pre- and post-dropout frequency information can be proportionally accounted for. Thus, availability is likely to be more of a factor when the RoCoF is highest, e.g., shortly after a generation drop event. On the other hand, if the communication rate were to be significantly reduced, then response variation would likely manifest sooner at higher availability levels; longer periods between communications would lead to more deviations in measured frequency. Based on these results, the standard of communication reliability in terms of availability does not need to be set very high as long as communication to the PV plants occurs at a reasonable rate.

3.1.6 Conclusions and Future Work

This work proposes a novel approach of implementing synthetic inertia in PV plants and other converter-interfaced generators and explores its potential benefits and drawbacks. The proposed approach uses communicated system frequency in place of locally measured frequency. The primary costs of this proposed modification are those associated with the communication necessary to convey this information. To that end, the effects of latency and availability on CE-SI were analyzed. Since the main goal of SI is to emulate the inertial response of traditional generators to arrest machine speed decelerations in generation




Figure 13: Fifty sample paths (using same colormap as previous figures) of system response at four different availability levels.

loss events, the effects of communication hurdles on the system response RoCoF were investigated. The primary findings were that CE-SI has the potential to improve RoCoF-based metrics over SI provided that the gain is set sufficiently high and the latency values are small. Results show that communication latency introduces system instability and is a limiting factor on how high the gain can be set. Furthermore, it was observed that communication availability standards do not have to be strict as long as the communication rate is sufficiently faster than the system dynamics.

Future research directions include implementation of this control approach in a real setup and combination with other control loops to improve other aspects of the frequency response.

3.2 Communication Enabled Fast Acting Imbalance Reserves

3.2.1 Introduction

Primary frequency regulation or frequency response is the response a power system undergoes following a large power imbalance event such as a loss of generation or a sudden disconnection of load [44]. Traditionally, the frequency response of a power system is determined by the combined effort provided by the governing action of individual generators. This action is driven by a feedback control loop that, based on local frequency measurements, adjusts the machine power output level [44, 45]. The governor actuation is typically slow as it involves the movement of the large inertias of the synchronous machines [46]. Contrary to these slow acting dynamics of synchronous generation, converter interfaced generators (CIGs) can adjust power levels almost instantaneously.

Taking advantage of the fast actuation of CIGs, a new control scheme for these devices to participate in the frequency response of power systems is presented. In the proposed method, CIGs are required to change their power levels depending on a feed-forward command signal which is constructed through

active monitoring of power imbalances in the system. Because power imbalances are predictive of frequency disturbances, this approach enables the system to respond before frequency errors grow large. A communication infrastructure is necessary to implement this approach to monitor and communicate power imbalances from the place they occur to the actuating CIGs. Communication-based methods have been proposed before to help with power system frequency regulation but are usually based on frequency measurements [47, 48].

3.2.2 Description of CE-FAIR

The proposed control scheme is named Communications Enabled – Fast Acting Imbalance Reserve (CE-FAIR) and is explained as follows.

Implementation The implementation of CE-FAIR in a centralized or hybrid approach can be observed in Fig. 14 and is explained as follows:

1. A power imbalance disturbance occurs in the system. The imbalance is detected through direct monitoring mechanisms in the system such as protection devices in generators and/or indirect methods such as PMUs.
2. The monitoring system sends a power imbalance signal to a server or an aggregator.
3. Based on knowledge of the system, the aggregator decides how much power each CIG should contribute to mitigate the power imbalance (see Section 3.2.2 below).
4. The server initiates communication with all the actuators predetermined to respond to the power imbalance.
5. The converter interfaced generators start adjusting (either increasing or decreasing) their power output as soon as they receive the message from the aggregator. The rate of adjustment in power is limited by the resource and the power electronics capabilities. As a result of different communication latencies, each CIG may have a unique start time.




Figure 14: Stages on the CE-FAIR (centralized implementation).

CE-FAIR can also be implemented in a distributed scheme, where the communications of power imbalances occur in a peer-to-peer fashion. CIGs must periodically communicate their available power range so that in the event of an imbalance each CIG has the necessary information to calculate the desired change in output power.

Redispatch determination The redispatch of each CIG in CE-FAIR is explained in this section. Each CIG receives a power command to modify its power output as follows

$$\Delta P_i = K_{FF}^i P_{imb} \quad (3)$$

where P_{imb} is the power imbalance determined by the monitoring system. The proposed scheme is a feed-forward controller that employs a communication infrastructure to inform of a power imbalance. A key element in defining how much each CIG participates is K_{FF}^i , defined as

$$K_{FF}^i = \eta \frac{P_i}{P_{available}}, \quad \text{with } P_{available} = \sum_{j=1}^N P_j \quad (4)$$

which is the *feed-forward* proportional gain of the i^{th} CIG and is determined by how much power the i^{th} CIG produces out of the total power production of all the CIGs. P_j is the available power level of the j^{th} device and variable η is defined as the fraction of the power imbalance that will be replaced by the CE-FAIR action. An η value of 1 implies that the CIGs will collectively adjust their power by as much as P_{imb} and is the case where (ignoring redispatch losses) the system frequency will be restored to its nominal value. The computation of K_{FF}^i is done at the aggregator level (in the centralized implementation) and communicated to each CIG at the moment they are required to act. The time elapsed from the moment the event occurs and the moment a particular CIG starts responding is named T_{FF}^i and is critical for the performance of the proposed control scheme.

3.2.3 Results and Discussion

A WECC-developed model of the western North American Power System (wNAPS) is used in this work to study the efficacy of the proposed control scheme in Section 3.2.2. The model chosen is the heavy summer 2016 case which has 20910 buses, 3033 generators and is operating at roughly 178 GW of power. Around 40 GW of power, being produced by conventional generators, is replaced by a developed model for a CIG. In total, 216 conventional machines were substituted by CIGs for a penetration level of $\sim 23\%$. The CIG model corresponds to a power dispatchable, controllable current source with a first-order approximation for the inverter. The disturbance for this analysis is the loss of the Columbia Generating Station (CGS) unit which is the largest generator in the Pacific Northwest. This corresponds to a loss of 1.15 GW occurring at 2 s. For this event, CE-FAIR was tested for $\eta = 35\%$ and 100% , and 5 different actuation latencies ($T_{FF} = 0, 0.25, 0.5, 1$ and 2 s) which are kept the same for all CIGs in the system. Two additional cases, one where there are no CIGs in the system (No CIG), and another where they have no control (No Control), are also considered. The simulations for this study were performed using the GE Positive Sequence Load Flow (PSLF) platform.

Figs 15 (a), (b) show the frequency response of the system for the loss of the CGS unit for all the cases mentioned above. The effects of including CE-FAIR to the CIG of the system can be summarized as follows: (i) The proposed controller improves both the frequency nadir and the settling frequency of the system. (ii) Increases in the actuation latency are reflected in more pronounced frequency nadirs. The frequency response of the system when CE-FAIR is active is the same as in the no control case for time intervals below the actuation latency value. The communication latency must be less than the time to frequency nadir in order for the scheme to provide a benefit. (iii) Increases in the power compensation level η improve both the frequency nadir and the settling frequency. It should be highlighted that when $\eta = 1$, the settling frequency is higher than the nominal of 60 Hz. This is the result of the new generation experiencing less losses than the generation lost. The distributed nature of the CIG generation (as compared to the centralized loss) results in it being on average closer to the loads. This result is however system and event dependent. (iv) Frequency zeniths (or overshoots) higher than the nominal frequency are also observed with the use of CE-FAIR in particular for this level of CIG penetration.




Figure 15: Frequency response for different values of η and time delays with 40 GW of power coming from CIG in the wNAPS model.

3.2.4 Conclusions

A new frequency control strategy that takes advantage of communications and fast responding resources is proposed. The approach is based on monitoring power imbalances and informing, through a communication infrastructure, commanded power levels to CIGs in a feed-forward control fashion. Because the system relies on communications to redispatch the CIG, its performance may be impacted by time lags and delays in the transfer of information. The efficacy of the proposed control scheme was demonstrated using simulation results for a model of the WECC.

3.3 Communication Enabled Frequency Droop

3.3.1 Introduction

Frequency needs to be kept within a tight range close to the nominal for power systems to operate adequately. The real-time control action that ensures frequency fluctuations are always within an acceptable range is the primary frequency regulation of the system. Traditionally, the primary frequency regulation is mainly defined by a droop governing action installed in particular units in the system. This action is effectively a proportional controller as

$$\Delta P_j = k_R(f_{\text{ref}} - f_{\text{eq}}) \quad (5)$$

where f_{ref} is the nominal frequency (60 Hz in North America). In the typical implementation of this controller, installed in conventional synchronous machines, signal f_{eq} correspond to the machine rotor speed. The feedback signal is hence a local signal.

With the advent of renewable energy and more generally of converter interfaced (CI) components and their interconnection to grid, the frequency response of the system is affected. In particular, it has been determined that power electronic interfaces tend to reduce the inertia of the system and they generally do not participate in frequency regulation. As a result of these facts, frequency response in the system is deteriorated with the inclusion of CI resources. To alleviate the deleterious effect of CI integration to frequency regulation, a governing-like function such as the one in (5) is implemented to CI devices. In such cases f_{eq} is typically chosen as the frequency at the bus where the CI device is installed. This frequency

tends to be volatile and noisy as is typically obtained from electrical quantities such as the phase angle. In this work a slight modification on the selection of the f_{eq} is proposed and evaluated. The proposed signal to use is,

$$f_{eq} = \sum_{j=1}^N k_j f_j, \quad (6)$$

as the feedback signal for the governing action of CI devices. This signal is a weighted average of machine speeds which tends smooth out the variations of local frequency measurements. Using this approach effects such as local oscillations and wide-area phenomena such as inter-area oscillations are attenuated; only the global trend of the variation in frequency is retained. The weights, k_j , in 6 are usually defined as,

$$k_j = \frac{H_j}{\sum_{i=1}^N H_i} \quad \text{for all } j \in \{1, \dots, N\} \quad (7)$$

where N is the total number of generators in the system and H_j is the inertia constant of the j^{th} machine. In the proposed approach the feedback frequency signal is computed using system-wide machine speed (or frequency) information. It is then assumed that this information is transmitted through communication channels from the place they occur to a central location, an aggregator, that computes the feedback signal and then transmits it to the CI devices. An alternative approach is to have each CI device compute the feedback signal which can be interpreted as a distributed implementation of the proposed approach. The mechanism of computation of the feedback signal is beyond the scope of this work which rather analyzes the advantages and disadvantages of using such type of signal. The proposed modification to use a global frequency measurement for governing control is referred to Communication-Enabled Droop (CE-Droop) and its benefits are evaluated in this work and compared to the conventional Droop approach that uses local frequency measurements.

3.3.2 CE Droop vs Conventional Droop

Test System In this work the NPCC system was used to test the proposed CE-Droop approach. The disturbance under consideration was the tripping of the Monroe generating unit in the MISO region at 2 seconds. This event causes a loss of roughly 655 MW (nearly 2.3% of the total power production). The particular system under consideration is a version of the NPCC with approximately 50% of solar PV penetration, obtained by substituting roughly half of the conventional generators with PV plants. Note that the installation of the solar plants was performed evenly through the system.

Effect of Variations in Droop Gain (k_R) In this Section, we present the effects of increasing the gain k_R to every PV plant in the system. Two scenarios are considered, one where each PV plant has a CE-Droop controller and another where the conventional droop (with local measurement only) is used. In the case of CE-Droop it is assumed that the global frequency information is available instantaneously (i.e. with no time delay). For each of this scenarios k_R was varied from 0 to 1000. In addition the no solar case is also studied for comparison and a case where the PV plants are unresponsive to frequency fluctuations. Note that this latter case correspond to the case where $k_R = 0$.

Fig. 16 (a) shows the frequency response of the system for the scenario where PV plants have conventional droop only and Fig. 16 (b) show the same results when the PV plants used the proposed CE-Droop approach. The no solar case shows that including PV definitely deteriorates the frequency response of the system as both the frequency nadir and the settling frequency are lower than in the no solar case. The results in Fig. 16 also show that implementing either Droop or the proposed CE-Droop to the PV plants in




Figure 16: System frequency response for the loss of the Monroe generating unit. Comparison of conventional and communication enabled droop approaches.

the system restores and even improves the frequency response of the system. As anticipated, higher gains tend to provide better results for both scenarios.




Figure 17: Comparison in the frequency response of the system for different values of gain k_R for the two scenarios of CE and conventional droop.

Fig. 17 compares the conventional droop method against the proposed CE-Droop approach for $k_R = 40, 100, 200, 1000, 2000$. These results show that for the higher gain case ($k_R = 200$) CE-Droop shows an undamped oscillation with frequency of 1.31 Hz.

The frequency nadir of the system and the time the system takes to reach it were determined for the results in 16. These metrics were used to quantify the effect that increasing the value of k_R has on the frequency response of the system. Fig. 18a shows the variations in the frequency nadir of the system as a function of k_R for both the CE-Droop and conventional Droop scenarios. These results show that the nadir is reduced by increases in k_R and they also show that a small gain is enough to restore the nadir of the system from the solar no control case to the no solar case. The results also show that for the same k_R value the resulting nadir using CE-Droop is always above the nadir when using the conventional droop. Regarding the time to frequency nadir metric, it is observed in Fig. 18b that CE-Droop also reaches the frequency nadir faster than the conventional droop approach for all levels of k_R . Combined these results show that the frequency nadir of the system is enhanced using CE-Droop compared to using only conventional droop. The settling frequency of the system is also affected by k_R variations however there is no difference between the CE-Droop and conventional droop approaches as shown in Fig. 17.

To analyze the effect that both CE-Droop and the conventional droop approaches have on the transient




Figure 18: Frequency nadir and the time it takes to reach it with variations in the effective gain k_R .




Figure 19: Maximum rotor angle separation as a function of the gain k_R .

stability of the system, the maximum separation of rotor angles is determined for each of the scenarios and cases in Fig. 18. This quantity was calculated by first computing the maximum rotor angle separation among all the generators in the system as a function of time $\theta_{\max}(t)$ and then computing the maximum value taken by this function as

$$\Theta_{\max} = \max_{t \in [0, T_f]} (\theta_{\max}(t)) \quad (8)$$

The maximum rotor angle separation as a function of gain k_R ($\Theta_{\max}(k_R)$) is presented in Figure 19 for the CE-Droop and conventional droop methods. They show that $\Theta_{\max}$ is unaffected by changes in k_R when conventional droop is implemented while it mildly increases as k_R augments in the CE-droop scenario.

3.3.3 Effect of Variations in Time Delay (τ_d)

Computing the global frequency in (6) for the proposed CE-Droop approach is not instantaneous as it involves both the transfer of information and a computation processing time. The feedback signal f_{eq} then reaches each PV plant after a time lag. In this Section the effect of this time lag, noted τ_d , is analyzed with the assumption that the delay experienced by each individual PV plant is the same¹.

¹In an actual implementation each PV plant experiences a different time delay. However it is expected that the variations among time delays for different PV plants to be on average smaller than the considered feedback delay.




Figure 20: System frequency response for the loss of the Monroe generating unit. Comparison of the effect of time delay on CE-Droop for different gains values.

Fig. 20 show the frequency response of the system following the loss of Monroe for with increments in the time delay τ_d value for k_R values of 100, 200, 400. The results when $k_R = 100$ are shown in Fig. 20a and they show that as τ_d increases oscillations in the system start to appear. Results in Figs. 20 (a), (b) for k_R equal to 200 and 400 respectively, show that time delay values of 0.4 and higher cause the system to lose synchronism.

The analysis of transient stability of the system was again carried out using the maximum separation of rotor angles $\Theta_{\max}$. Fig. 21 show variations in $\Theta_{\max}$ with increases in τ_d for different values of k_R as well as for the no control and no solar cases. The results in Fig. 21 complement those in Fig. 20 and they show that for values of k_R of 200 and 400 the system becomes unstable with time values above 350 ms and 110 ms respectively.

The impact of the time delay τ_d on the frequency nadir and the time the system takes to reach it are also investigated. Fig. 22 (a) show how increases in the time delay are reflected in the frequency nadir of the system for different values of gain k_R . These results show that the frequency nadir of the system is deteriorated with increases in τ_d , however they also show that the decrease is not significant and in all (stable) cases considered the CE-Droop perform better than the no control and the no solar case. In Fig. 22 (b) the impact that τ_d has on the time to frequency nadir is presented; it shows that this time is






Figure 21: Maximum rotor angle separation as a function of the time delay τ_d .



(a) Frequency Nadir as a function of the time delay τ_d .



(b) Time to Frequency Nadir as a function of the time delay τ_d .

Figure 22: Frequency nadir and the time it takes to reach it with variations in the time delay τ_d for different gain values k_R .

only slightly affected by increments in the time delay.

3.4 Distribution: Hierarchical Control of Volt-VAr Function for Steady-State Voltage

3.4.1 Introduction

High penetrations of PV interconnected to the distribution system can cause over-voltages at the point-of-interconnection (POI) beyond ANSI Range A [49]. Given the increased deployment of renewable and distributed energy resources (DERs), innovative strategies for grid modernization and control are required. With the emergence of smart grid and advanced inverter functions, such as fixed PF, constant reactive power (VAr), and volt-VAr (VV), there is an increasing interest in solutions for utilizing advanced inverter functions to mitigate PV impacts and increase PV hosting capacity [50]. Many of the advanced inverter functions, especially those that allow manipulation of the VAr generation/absorption, lend themselves to assisting with voltage issues [51].

A simple hierarchical control was developed to dispatch PV VV settings to provide distribution system voltage regulation. The effectiveness of the controller depends on a suitably fast and reliable communication

infrastructure. The focus of the research was to evaluate the potential communication requirements when utilizing intelligent VV dispatch to mitigate over-voltages. The necessary communication infrastructure was tested by evaluating the effectiveness of the hierarchical control with varying communication intervals, reliability, and delays.

3.4.2 Test Setup

A rural 12 kV distribution feeder serving a highly commercial load area was chosen as the test feeder. The feeder model consists of 215 buses and 39 service transformers. The feeder has a peak load of 3.98 MW. The feeder voltage is regulated via the substation transformer load tap changer (LTC); there are no line voltage regulators or switching capacitors.

The load data for the week ending on the date the minimum daytime load occurred, October 25th, was selected as the simulation week for all studies. The minimum daytime load was defined as the lowest load level that occurred between 10:00-14:00 when the solar power output is high. The minimum daytime load was found to be 1.51 MW (36% of peak). The measured substation supervisory control and data acquisition (SCADA) data at 15-minute resolution was used to model the load variation.

Quasi-static time series (QSTS) power flow analysis [52] was performed at 5-second resolution by linearly interpolating the load data. The analysis was performed in OpenDSS using the GridPV toolbox [53]. A map showing the layout of the feeder topology and the simulated PV locations is shown in Figure 23. Two PV systems were simulated as shown by the yellow stars in Figure 23. Each PV system is 750




Figure 23: Map of the test feeder with the PV test scenario.

kW, with an aggregate total of 1.5 MW of PV, just slightly less than the minimum daytime load level. To create the scenario of interest, the AC-to-DC ratio of PV1 was set to 1.12 and PV2 to 1.05. Under peak solar output, PV1 had a maximum VAR support capacity of ± 378 kVar, and PV2 had ± 240 kVar.

Similar to the motivation for selecting the minimum daytime load period for simulation, the maximum possible PV production was assumed in order to create the worst-case scenario for highest feeder voltages. The maximum possible solar production is modelled using a clear-sky global horizontal irradiance (GHI) profile generated using the Ineichen clear sky model via the GridPV toolbox [53]. The GHI was then converted to a plane-of-array (POA) irradiance, assuming a 30° surface tilt, using the PV_LIB toolbox [54].

3.4.3 Hierarchical Voltage Control

Distribution system voltages are defined by ANSI C84.1 Range A to be within $\pm 5\%$ of nominal voltage with respect to 10-minute average voltages [49]. Because of the 10-minute average definition in the standard, feeder voltage regulation has an extended period to detect and correct voltage issues. The hierarchical voltage control is composed of two layers. At the local layer, the PV systems monitor and attempt to regulate their local voltage using the VV function. The system control layer communicates with all PV systems to dispatch new VV settings for additional system benefits.

Local Control Layer Each PV system is controlled locally using the VV advanced inverter function. This research focused on a scenario where PV inverters on a distribution feeder were all assigned a “default” VV curve [55], shown with the blue line in Figure 24. The PV system is measuring the local voltage on the output of the inverter and reacting in real-time to change its reactive power output.

The VV curve is designed to adjust the PV inverter VAR level depending on the voltage, producing VARs as the voltage goes outside the deadband. Note that at high voltage, the PV system has negative reactive power (absorbing), which will work to bring the system voltages lower. The inverse is true when the voltage is low, with reactive power being injected into the distribution system. VV is able to provide local voltage regulation, but it does not have any additional information about the system conditions to regulate voltages throughout the feeder.

The y-axis on the VV curve in Figure 24 is dependent on the size of the PV inverter relative to the DC rating of the PV system. It is common to set the maximum and minimum values of the curve to equal the reactive power headroom of the inverter during full PV real power output conditions. For example, for PV1 with an AC-to-DC ratio of 1.12, at full real power output, there is reactive power headroom of 0.504_{PU} , or 0.450_{PU} headroom with the inverter AC rating as the base.

Centralized Control Layer Due to the limitations of local voltage control, a centralized system layer was added to the hierarchical control. It was assumed that the hierarchical controller received measurements from the PV systems and could send dispatch control signals back to the PV systems. Figure 24 illustrates the VV curve and the logic behind the curve shifting strategy for voltage regulation.

Since the system controller has information about the voltages around the feeder, it can dispatch new VV curves to the PV systems to modify their reactive power output to help the system voltages even when their local voltages are fine. For a given voltage, by shifting the VV curve to the left, absorption of more VARs is induced, lowering the system voltage as a result. The inverse is true for a right shift in the curve. An example motivating the need for the centralized system control layer is shown in Figure 25. Figure 25 shows the feeder voltage profiles at 11:00 AM on 10/25/09 for the default VV curves (red) and with controller shifted VV curves (blue). With the local default VV control during light load and high irradiance conditions, PV2 near the end of the feeder was unable to mitigate its local over-voltage after exhausting its VAR support capabilities (240 kVar). PV1 was only absorbing a portion of its available reactive power (378 kVar capacity), since it was below 126 V. The profile resulting from the centralized controller shifting the VV curve of PV1 left, to absorb more reactive power and keep the system voltages within ANSI Range A, is shown in blue. Figure 26 illustrates the overall logic of the hierarchical voltage




Figure 24: VV curve shifting logic.




Figure 25: Feeder voltage profiles during daytime minimum with default VV curves (red) and with PV1's VV shifted left (blue).

controller. The centralized control receives the voltage measurements at the PV inverters to modify the VV curves, i.e. the VAr levels, to mitigate any voltages outside ANSI Range A (0.95 to 1.05_{PU}, 114 and 126 on a 120 V base) [49]. The hierarchical controller dispatches new VV settings to the PV inverters only if the voltage was out of range.

The adjustable controller voltage deadband was set slightly tighter than the ANSI limits, 116.4 to 125.5 on a 120 V base, to prevent voltages outside of ANSI Range A. If a PV voltage was outside the controller deadband, all PV systems, except those reporting a voltage outside the deadband, were assigned a new VV curve shifted by an adjustable shift interval (set to 0.001_{PU} or 0.12 on a 120 V base) in the appropriate direction, depending on whether it was an over-voltage or under-voltage. The centralized controller deployed shifted VV curves to all inverters with additional VAr capabilities, per the communication interval, until all voltages are within ANSI Range A.

The logic provides voltage regulation for the feeder using existing measurements, and it also has the benefit of monitoring and regulating the voltage at the PV locations. At night when the PV inverters are off, the VV curves were reset to the default settings. During periods between the communication intervals, the local VV control reacts to regulate the local voltage based on its VV settings in a hierarchical control framework.

Due to VV modeling limitations, there were only two PV systems simulated in this example, but the control strategy described could be extrapolated to scenarios with more PV systems interconnected. In the case of many PV deployments on several phases, the centralized controller would dispatch the new VV curve settings based on the phase with voltage issues and the phases to which each PV system is interconnected. An intelligent, cascading proximity prioritization would also be prudent in a case with many PV systems, but was not applicable in this example.

3.4.4 Simulation and Communication Results

The hierarchical voltage control and communication network model were implemented in MATLAB, with OpenDSS running the distribution system model and power flow. All voltage issues are defined by the 10-minute moving-average voltage of any bus being outside the ANSI Range A. No under-voltage issues were observed in this case, so only the over-voltages and time above ANSI Range A (1.05_{PU}) were analyzed.

Voltage Regulation Results Figure 27 shows the 10-minute moving-average voltage results for the minimum daytime load week for the test PV scenario. When the PV is at unity power factor, there are 42.5 hours that the feeder is outside ANSI during the week. The total hours with ANSI violations is decreased to 24.9 hours when the default VV curves are added to the PV systems. Finally, Figure 27 shows how the hierarchical voltage control removes all ANSI violations and keeps the maximum feeder voltage below 1.05_{PU}.

The initial testing assumed that the controller requested and obtained parameters from all PV systems, processed calculation of dispatch values, dispatched values, and inverters implemented dispatched commands all within a 5-second simulation interval. The communication infrastructure requirements were tested by investigating three components of the communication network: communication interval, reliability, and delay. The sections below demonstrate how the effectiveness of the VV control was dependent on the communication network by quantifying the amount of time outside ANSI Range A observed.

Communication Interval Results The communication interval, i.e. how frequently communication must occur, was studied by varying the communication between the inverters and controller from every 5 seconds to every 15 minutes. As the measured data and dispatched settings from the centralized controller were exchanged less often, the effectiveness of the hierarchical controller decreased. Figure 28 shows the results




Figure 26: Controller diagram for voltage regulation using VV.

for different communication intervals during the simulation week for the VV controller simulations with 100% reliable communication and no network delays.

For all intervals less than the ANSI metric of 10 minutes, the controller mitigated all over-voltages by modifying the PV VV curves to keep the voltage within the deadband. At communication intervals




Figure 27: Maximum feeder voltage comparisons for simulation week.

longer than 10 minutes, voltage deviations that persisted for more than 10 minutes between communication intervals resulted in voltage violations. In this case, any interval greater than 10 minutes would result in violations per ANSI Range A.




Figure 28: VV based controller results for the simulation week with different controller communication intervals.

Communication Reliability Results The communication reliability was tested by implementing a random probability of successful communication. For example, when simulating a system with 99% reliability, each communication signal had a 99% probability of being received without errors. This stochastic model was applied to both the measurement signal coming into the centralized controller and the dispatch signal going

to the inverters.

The QSTS simulation was run three times for each test condition so that the stochastic results could be averaged. The results in Figure 29 show an extreme sweep of different communication network reliabilities for a few communication intervals. The 7-minute interval was the longest one observed to not have any issues all the way down to 80% reliability, and the 11-minute interval was the shortest one observed to have issues even at 100% reliability, as was also observable in Figure 28.




Figure 29: Controller results for different communication intervals and network reliability.

Communication Delay Results In the previous simulations, it was assumed that there were no communication delays and the time from measurement to new setting implementation was performed within a 5-second simulation interval. In reality, there are many communication delays that may extend the process beyond one second [51]. While this would generally be achievable in a few seconds, the effect of the communication delays was tested by increasing the delay up to 20 seconds from the time of measurement to implementation of new settings.

Figure 30 shows the time above ANSI during the simulation week for different communication intervals and communication delays. Note that even extremely large communication latencies of 20 seconds do not have an impact on the effectiveness of the hierarchical voltage controller.

3.4.5 Conclusions

A hierarchical control algorithm was developed to utilize photovoltaic system advanced inverter functions, specifically VV, to provide distribution system voltage regulation and to mitigate voltages outside ANSI Range A by using voltage measurements at the PV inverters. The controller was developed and demonstrated on a week-long analysis of a simple two-PV scenario on an actual 3-phase distribution system model. The necessary communication infrastructure for effective control was evaluated by testing three different communication aspects: 1) interval, 2) reliability, and 3) delay.

Based on this specific test system, the hierarchical VV controller mitigated any 10-minute average voltages above ANSI Range A up to a 10-minute communication interval assuming 100% communication reliability and no delay. These results were synonymous with other similar research focused on mitigating voltage regulator tap operations with a time delay setting of 30 seconds [51, 56], i.e. the communication timeframe required is directly correlated with the application time-urgency.




Figure 30: Time outside ANSI (min) during the simulation week with different communication intervals and network delays.

The reliability of the communication network did not have an impact on the controller for communication intervals of 7 minutes or less, even all the way down to 80%. At 9-minute communication intervals, the communication network must be at least 95% reliable, and at 10-minute communication intervals, the communication network must be at least 99.5% reliable for the hierarchical controller to be fully effective at mitigating all over-voltages.

A communication delay of up to 20 seconds, despite being a very high delay assumption in reality, had no impact on the effectiveness of the hierarchical voltage controller. From a distribution perspective, the ANSI requirement of 10-minute averages represents one of the least time-sensitive metrics where centralized control of advanced functions is applicable.

4 Cybersecurity Reference Architectures and SCEPTRE Deployment


This section talks about the cybersecurity reference architectures that were created to study different security features along with the software tool, called SCEPTRE, that was used to perform the assessments.

4.1 Power Simulations

To demonstrate and evaluate the cybersecurity implications of communication for DER support of the grid, an EmulyticsTM environment was created and utilized to combine power simulation with virtual networking using the SCEPTRE platform developed by Sandia National Laboratories. In describing this environment, first the underlying power model and simulation will be described in Section 4.1. Then the components of the SCEPTRE platform utilized in this research will be discussed in Section 4.2. To examine the cybersecurity implications for communications within a control network and its impact on the complete system, analysis of the networking is insufficient. Rather, a combination of the networking and an underlying power system simulation is required. This section will discuss two power system models that were created for the red team assessments.

The first model represented the 12 kV distribution feeder with two 750 kW PV sites from the volt-var shift use case, as shown in Figure 23. This model was anonymized from a utility partner and developed for voltage regulation studies [57, 58]. The feeder included 215 buses and 39 transformers, and had a peak load of 3.98 MW. The GridPV toolbox in OpenDSS [59] was used to conduct quasi-static time series (QSTS)

A transmission model was also created using the 42-bus transmission model with 26 PV systems, as shown in Figure 31. While this model was implemented in SCEPTRE as a proof-of-concept for transmission co-simulations with high penetration PV scenarios, it was not used in any of the red team assessments.



In Figure 32, three domains are shown which represent the SCEPTRE network emulation (top), Minimega-control virtual machines and RTUs (middle), power system simulation provider (bottom). In the upper layer, the network directs actual communication packets through emulated routers, switches, and

other networking components. Devices in this network (utility DERMS, DER, etc.) are built on Windows and Linux VMs in the middle layer. The power simulation at the bottom, updates values represented in the VMs/RTUs in the middle layer, but also updates the power simulation based on changes to those devices through an unexposed back-end network. On the right, the HIL DER is shown. For the HIL SCEPTRE experiments in this project, the PV inverter was connected to a real PV array but the AC power was provided with a grid simulator. As changes in the power system occurred in the provider (e.g., a new voltage level), those changes were communicated to the Ametek RS180 grid simulator to change the terminal voltage on the inverter. The active and reactive power of the inverter were calculated using a LabVIEW data acquisition system (DAS) connected to current and voltage probes. From those measurements, the power factor of the inverter was calculated and updated in the OpenDSS simulation.




Figure 32: Overview of the SCEPTRE environment with a HIL DER.

The rest of this section will discuss these components in greater detail. Section 4.2.1 will discuss the use of a power simulation within SCEPTRE and Section 4.2.2 will discuss the RTUs representing DERs in our models. The implementation of encryption and other potential security measures within SCEPTRE are presented in Sections 4.2.3 and 4.2.3. Finally, Section 4.2.4 will present the various control network topologies represented and utilized for the Red Teaming activities of Section ??.

4.2.1 Power Simulation Interaction

The interface between a power simulation and SCEPTRE is done through what is known as a “Provider” in SCEPTRE. This code uses any Application Program Interface (API) available, or develops one, to connect to a simulation platform such as OpenDSS. The provider then connects back into SCEPTRE through a standard communication interface with a Publish/Subscribe model using ZeroMQ.

In the case of the power hardware-in-the-loop (PHIL) experiments, a physical PV inverter was connected to the power simulation as shown in Fig.32. The OpenDSS simulation was run and power system parameters (voltage, frequency, DER current, etc.) were passed through ZeroMQ to the virtualized devices to update their internal Modbus maps with the latest data. In the case of the HIL inverter, the voltage was sent to the Ametek RS 180 grid simulator over a TCP/IP connection using IEEE-488.2 (SCPI) commands to update the AC voltage applied to the inverter. A Windows computer running a LabVIEW program connected to current and voltage probes on the AC terminals of the PV inverter acted as the Data Acquisition System (DAS). The LabVIEW program broadcast UDP packets containing active and reactive power measurements to the SCEPTRE environment every second. This data was captured by the SCEPTRE and used to update the active and reactive powers of the PV inverter in the OpenDSS power simulation with measurements from the physical inverter. Therefore, the OpenDSS power simulation was resolved every second with new data from the physical inverter. Interestingly, the DAS measurements needed to be used as opposed to those in the Ametek, because the Ametek measurements of active and reactive power were not accurate—likely because the probes were not located at the terminals of the inverter and this was a small PV inverter compared to the grid simulator capacity.

4.2.2 Remote Terminal Units

Within SCEPTRE, the RTUs are virtual machines (VMs) that perform the following functions.

1. Represent field devices in an ICS or SCADA network
2. Connect and interact transparently with the underlying system simulation
3. Perform simple control logic
4. Communicate with the control network using a standard protocol, such as Modbus/TCP or DNP3

To represent a field device in a SCADA network, SCEPTRE RTUs are generally operated as pared down versions of Linux with select capabilities. This better represents the limited functionality of a field device that is only completing a few specific tasks. Furthermore, it allows for an increased number of RTUs to be created within a SCEPTRE experiment without requiring more hardware resources to be allocated. The RTU communicates with the power simulation using a ZeroMQ messaging framework on the backend of SCEPTRE. This process uses a Publish/Subscribe model to both pull data from the power simulation and push control settings and commands back into the simulation. RTUs are programmed with simple control logic representative of its role in the emulated SCADA network. This control logic and processing should be equivalent to that which would be on a field device in the corresponding real-world system.

The RTU, as well as any other components in the ICS network such as Programmable Logic Controllers (PLC), Front End Processors (FEP), or an Open Platform Communications (OPC) Server, can communicate commonly used control protocols on the ICS network such as Modbus/TCP or DNP3. This means that the communications used within the control network should be represented and utilized to communicate information to the appropriate locations. These communications and network connections are analyzed in the Red Teaming activities of Chapter ??.

For this work, measurement and grid-support functionality was added to SCEPTRE RTUs to represent the capabilities of a SunSpec-compliant smart inverters using SunSpec Information Models 1, 101, 123, and 126, which represent the Common, Inverter (Single Phase), Immediate Controls, and Static Volt-VAR data

[61]. The associated connections between the inverter and OpenDSS and PowerWorld were constructed across the ZeroMQ network.

4.2.3 Security Mechanisms

To perform cybersecurity analyses of the impact of the security control mechanisms on both the security of the DER control system and the corresponding impact to the grid, the implementation of various security measures was required. This section discusses those mechanisms examined in this work and their implementation within the emulated control network.

Network Segmentation and Enclaving Network segmentation is a known and commonly used strategy for providing additional security to a control network. The extreme example of such a control mechanism is the “Air Gap”, where a control network is completely physically separated from any other network. In reality, maintenance of a control network commonly requires crossing such an air gap. To address that need, control networks may be logically or physically segmented to reduce the impact of a security compromise to a subset of the system. This can be done through physical segmentation of the control network, or through mechanisms such as Virtual Local Area Networks (VLANs).

Network segmentation is a technique to minimize common-mode vulnerabilities, whereby enclaves are isolated with firewall rules, VPNs, proxies, or other networking technologies so that traffic between them is only allowed by exception. Extensive research on segmentation for military microgrids has been completed previously [62, 63], and the enclaving concepts represented here were derived from [62, 64], but have been modified for an ADMS application. The downside of this approach is the additional network administration and communication latency. There are technical challenges to segment DER networks because networking equipment will not necessarily be owned and operated by a single entity. It may be possible to enclave the devices if communications are passed directly to the DER through networks that are owned by the grid operator, e.g., through an advanced metering infrastructure (AMI) mesh radio or dedicated SCADA network to DER systems. However, in the majority of commercial and residential PV systems, communications will be established through wired or wireless networks via the public internet, as shown in Fig. 33. In those cases, it is more difficult to enclave the networks because internet service providers (ISPs) control the network routing and firewall rules cannot be implemented easily without assistance from the ISPs. Therefore, the use of VPNs, proxies, or some other technology would be required to logically isolate the enclaves. Three options and their pros and cons are presented in 1. These include using firewall rules—which will only work when operating over a private network, using hardware proxies to hide traffic, and using encrypted VPN tunnels.

Firewall rules have been used in the past for military systems [63], but these are not effective when operating with internet connected devices because the network traffic channels are not consistent and ISP systems are designed for speed, not security. This method becomes an option when the utility or other grid operator is communicating to DER equipment through a network that they own. In that case, they may apply specific firewall rules to create enclaves. Blocking all connections initially and then allowing specific ones is considered a best practice with firewall rules. When the network is privately controlled, this approach will allow a utility to whitelist traffic from a DERMS to each respective HAN (or commercial/utility DER LAN or enclave), and all other traffic would be dropped. This can be easy to manage if the number of HANs is small; however, if that number is extensive and continually changing, it would become a difficult operational management issue. Constantly changing firewall rules can also introduce a greater chance for the dropping of legitimate connections by mistake thereby causing significant ongoing support concerns. The advantage—almost necessity—of using a firewall on a private network is to ensure that only traffic desired from the DERMS to a specific HAN can transit the network, and all other traffic (potentially malicious) is dropped.




Figure 33: Different DER control network architectures in which DER data is exchanged over public internet or AMI radio networks.

Table 1: Methods of enclaving DER networks.

Enclave mechanism	Pros	Cons
Firewall rules (whitelist DER-to-headend connections) for grid operator or aggregator-owned networks (e.g., AMI networks)	• Private network • Extends grid operator or aggregator LAN to FAN or HAN	• More costly • More management • Complex • Potentially less data bandwidth or speed
Use hardware proxy, which monitors for DER/utility traffic and exchanges it	• Works well for aggregators • Traffic send via ISPs using RESTful HTTP or TLS connections	• Relies on 3rd party (ISP) to manage network (could have more latency if QoS is an issue) • Need maintenance contracts • Privacy concerns (for unencrypted traffic) • Less flexibility
Virtual private networks (VPNs) between DER and grid operator	• Direct connections between DER and utility • Reduced latency • Grid operator controls and easily changes segmentation	• VPN management and maintenance difficult • Could burden facility/home bandwidth

Another option would be for the utility to provide each physical site with a hardware proxy between the ISP connection and HAN or facility LAN. A hardware proxy is simply a small device like a cable/DSL modem that would have two primary connections. One connection would be to receive the general ISP connection, and the other would be to output to the HAN. Additional connections would be required if the device were intended to connect directly to an inverter—in that case, the hardware proxy would route traffic identified as intended for the inverter before passing it off to the HAN. This proxy would monitor for traffic specific to the inverter and pass that traffic directly to it; all other traffic would be passed unmonitored to the HAN. Controlling specific traffic between the DERMS and an individual (or group) HAN would be similar to the firewall option over a private utility network. There would be potential privacy concerns if the proxy were compromised by an adversary who could then manipulate the network traffic for their benefit. A challenge with the hardware proxy is to install and support a physical hardware device at each site; this could present additional support and maintenance cost to the utility. It may be possible to provide each facility with an ISP-friendly switch/modem in place of what the ISP has provided (most markets have few ISPs to pick from, and so it might be easy to provide an ISP-friendly hardware device). The hardware proxy would also allow for priority traffic specific to the inverters, i.e., priority over regular HAN traffic. Finally, the hardware proxy would need to update the utility through a dynamic-DNS-like service so that the utility was always aware of the (potentially changing) publicly routable IP address of the home or facility.

Alternatively, the utility could maintain an ongoing virtual private network (VPN) connection directly to the inverters through the existing ISP network and corresponding switch/modem. A VPN is an encrypted tunnel for communication between two systems over a network. This would provide the utility with a direct, secure connection between the DERMS and each HAN over a public network based on well-established open standards. Communication encryption prevents eavesdropping or manipulation by an outside party. Traffic specific to each HAN could be communicated through the VPN tunnel with the assurance that it remains secured from any malicious actors along the communication path (similar to traffic sent over a private network). Additional support and maintenance requirements would be necessary from the utility, similar to the hardware proxy, but additional support from a HAN's ISP would not be necessary as most ISPs support VPN tunnels for their customers without any additional service changes. To initiate a new connection, each inverter would initiate the VPN to a known utility IP address providing a "plug-and-play" deployment. An alternative to an inverter initiated connection would be to deploy a facility gateway where the VPN connection could be originated to each respective HAN.

Once the method of generating the enclaves is selected, the specific rules for cleaving the devices must be decided. There must be a balance: too many enclaves mean higher likelihood of mistakes in the firewall, VPN, proxies, or router configurations, slower communication times, and more difficulty deploying more DER; but at the same time, there must be a certain number of enclaves to prevent control of a critical magnitude of generation. Here, we offer two basic approaches:

- A segmented network with DER placed in one of three enclaves at random or convenience, e.g., Fig. 34
- A critically segmented network with no more than 20% of the total capacity in a single enclave, e.g., Fig. 35

The placement of DER in a specific enclave could be done based on geography, power system topology, nameplate capacity, or other metadata. More sophisticated methods of determining the number of enclaves and which DER should be placed in each should be considered an area of future research. These approaches were initially explored in an earlier Virtual Power Plant project [64].

Denial of Service (DoS) risk is reduced with network segmentation. Breaking up a network reduces the overall target space by transferring risk from a single network segment to multiple enclaves. As a result, an adversary would need a foothold inside each enclave to launch a DoS attack on the entire DER fleet.




Figure 34: Segmented network with DER placed in random enclaves.




Figure 35: Segmented critical network with no more than 20% of the generation placed in one enclave.

This may be relatively trivial for sophisticated adversaries, but it will require additional work on their behalf. Unfortunately, a DoS attack against the utility server as a single target could result in stopping all inbound and outbound DER-related traffic. Using whitelisted firewall rules in the DER infrastructure would be a useful—but time-consuming—solution on the utility side to further reduce the risk from DoS attack related noise. Network segmentation would also reduce the risk of the DER equipment becoming a part of a botnet used for distributed DoS (DDoS) attacks because it increases the effort required to reach the devices. To assess this approach co-simulations of a distribution power system and DER control network were created and run. Adversary-based assessments were then conducted on these environments to investigate the effectiveness of enclaving techniques to a range of attacks including DoS, as described in ??.

Encryption As SCEPTRE utilizes known and common control protocols to communicate, it is also capable of integrating common encryption schemes into those communications. The SunSpec-compliant DER inverter RTUs communicate Modbus TCP to the utility ADMS (SVP) VM. Since Modbus is passed in plaintext, it was desired to encrypt these communications to make the environment more realistic. The most typical way to accomplish this, and the method used in this research, is to incorporate Transport Layer Security (TLS) to secure communications using a common implementation such as OpenSSH [65]. Therefore, in some of the environments SSH components were added to the environments to pass the Modbus traffic between the utility and the DER subnets through an encrypted tunnel. This additional security mechanism was performed for all three enclaving strategies.

Moving Target Defense (MTD) Moving target defense (MTD) is class of technologies that dynamically modify a system environment to create uncertainty for adversaries. Chavez et al. developed Artificial Diversity and Dynamic Security (ADDSec) as a MTD tool that leverages software defined networking (SDN) to randomize network parameters and communication paths. ADDSec has the ability to randomize IP addresses and port numbers both in anticipation of and in response to detected network activity. This is meant to thwart an adversaries ability to conduct reconnaissance and establish communications between devices on the network, and has been proven to be effective at increasing the resilience of grid wide area networks against certain types of attacks [66]. ADDSec is comprised of several components:

1. Dynamic reconfiguration of networking and routing parameters, using pseudo-random number generators as a source of entropy, including randomization of IP addresses and ports, to thwart reconnaissance and prevent unwanted connections.
2. Generation of unique application binaries within a system to raise the difficulty of producing software exploits.
3. An ensemble of machine learning algorithms that analyze host statistics, networking information, and network traffic to autonomously detect and trigger reconfiguration of the systems in real-time.

For more information on how ADDSec operates, such as its use of software defined networking (SDN) to enable transparency to end-devices, please refer to [66, 67]. Intuitively, the use of dynamic configurations to decrease predictability for attackers seems reasonable as a means for enhancing cyber security, but techniques to measure the resilience benefits of MTDs to-date has been primarily survey- and opinion-based [68]. Integrating ADDSec in our emulated system topologies allows us to evaluate its effectiveness against a controlled baseline. It is also worth noting that while ADDSec has been proven effective against reconnaissance attacks, it does not necessarily provide protection against a persistent threat that has previously been introduced to the network or prevent an attacker from carrying out an exploit on a previously compromised host.

An example of this technology is shown in Figure 36. On the left is a utility subnet consisting of an Advanced Distribution Management System (ADMS), Geographical Information System (GIS), and DER

management system (DERMS). On the right, is a collection of DER in a campus or utility/commercial site on a single switch. There is a “IP Generator” computer in the bottom that sends the new IP addresses to the switches in front of actual DER or computation devices. The MTD changes the IP addresses of these switches but the utility-owned and DER nodes retain static IP addresses. Actual implementation would likely require multiple MTD subsystems that independently reconfigure the IP addresses of the utility subnet and DER devices. Since this technology requires a sperate network to be overlaid on the publicly-addressable one, it is likely that residential DER will not have the ability to be included in the MTD/SDN overlay, but commercial and utility owned sites could employ this technology.




Figure 36: Implementation of Moving Target Defense on a DER communication network.

4.2.4 Topologies

This section enumerates the different cybersecurity reference architectures used for the red team evaluations. The topologies specify the connections and network information for the virtual machines and networking components within the emulation experiment. For this research, a total of seven topologies were created vary the security control mechanisms for the virtual DER network:

1. Flat network with plaintext Modbus traffic (no encryption)
2. Flat network with encrypted Modbus traffic using SSH tunnels
3. Segmented network with plaintext Modbus traffic (no encryption)
4. Segmented network with encrypted Modbus traffic using SSH tunnels
5. Critically Segmented network with plaintext Modbus traffic (no encryption)
6. Critically Segmented network with encrypted Modbus traffic using SSH tunnels
7. Flat network with plaintext Modbus traffic (no encryption) with a moving target defense overlay that randomized IP addresses every 3 seconds

These topologies use a combination of the various security measures mentioned in Section 4.2.3, with different enclaving strategies, encrypted tunnels, and MTD approaches between the utility and DER subnets.

In the case of topology #3, this environment was deployed with simulated inverters and also replacing one of the inverters with a physical 3.0 kW residential-scale PV inverter with a HIL feedback loop.

Specifically, a VM running the SunSpec Validation Platform (SVP) [69] was placed in the utility network and was monitoring and controlling 20 DERs. The SVP was issuing voltage regulation setpoints according to the volt-var shift algorithm described in [57, 58].

Besides the various enclaving strategies, Each topology also simulated some extraneous traffic on the network through a tool in Minimega called Protonuke, which merely added some extra load by simulating internet traffic related to web browsing, mail, and SSH connections between extra VMs acting as servers and clients. Finally, a Kali virtual machine was deployed inside each topology for the red team to use in their assessment.

Flat Distribution Clear The flat topology is a simple flat network for all twenty DERs in our system, meaning that each DER can reach and communicate with every other DER. Two routers were used to connect a utility network through a Wide Area Network (WAN), such as the internet or any other large distributed network. This configuration is shown in Figure 37.

The flat distribution clear topology was intended to depict connection through a traditional ISP network, or similar, such as could be found in a large neighborhood or a public campus. The DER inverters were connected directly to the public network and shared that network with any other devices that would be connected to it. In this topology, the SVP virtual machine connected directly to each DER inverter using TCP Modbus with no additional protections in place. Meanwhile, the Protonuke client shared the larger public network and connected to the Protonuke server on a separate network to generate some extraneous traffic as would be case on a used network. The Protonuke server was the sole device on this separate network and is only present to represent connections and traffic to and from other networks.

Flat Distribution Encrypted The flat distribution encrypted topology is very similar to the clear version described in the last section, but adds in an SSH tunnel as shown in Figure 38. This SSH forwarding tunnel was established to encapsulate the Modbus TCP traffic between the SVP virtual machine and each DER. This is performed by taking Modbus traffic in the clear on the unsecured sides of the “SSH CORP” and “SSH WILD” virtual machines and securing them using OpenSSH. In short, traffic within the flat control network for the DERs and within the “CORP network” are plain Modbus with no security, but across the public network these messages have been encrypted. This setup is meant to represent a utility providing an SSH server at the DER sites to be used to forward traffic through an encrypted tunnel. This limits exposure of the Modbus traffic but does not protect it all the way to each DER inverter. For instance, if an adversary could position themselves between the “SSH WILD” server and a DER inverter, they could capture and manipulate the Modbus traffic to the inverter with ease. An alternative to this topology would place a SSH server directly next to each DER inverter. The SVP system could then communicate over the public network to the SSH server at the DER, which would then communicate to the DER inverter directly. In that instance, an adversary would have to be on the private network to capture or manipulate the Modbus traffic, but this extra security comes with significantly increased costs due to the need for extra devices for SSH tunneling. This can be minimized by building security functionality, such as SSH tunneling, directly into smart inverters themselves.

Segmented Distribution Clear In the segmented clear topology, the flat control network for the DERs are broken into three separate segments. The inverters themselves are split equally between these segments, with 6-7 inverters per segment as shown in Figure 41.

By splitting the inverters into separate segments, the difficulty of impacting all twenty is increased. However, as this is a “clear” topology, each segment is directly routable and the Modbus traffic is passed




Figure 37: Flat Distribution Clear




Figure 38: Flat Distribution Encrypted




Figure 39: Segmented Distribution Clear

in the clear and unsecured.

The segmented distribution clear topology was meant to break the DER inverters into separate, publicly routable networks. This offers a couple things to the utility. First, it would remove the reliance on a single network to maintain traffic to all DER inverters. If there was a network outage on one segment, a portion of the DER inverters would still be available for communication with the ADMS. Additionally, this segmentation reduces the target surface for an adversary, who now requires access to each segment to intercept or manipulate Modbus traffic for that segment. In other terms, it increases the steps required to impact all the inverters. The Kali virtual machine and the Protonuke server were placed on a fourth network segment connected to the publicly routable network.

Segmented Distribution Encrypted The segmented distribution encrypted topology was similar to the segmented distribution clear topology, with the main difference being the inclusion of SSH servers on each segment as shown in Figure 40. The SVP VM passed TCP Modbus traffic through the SSH CORP system and then through the encrypted tunnels with each of the “SSH Wild” servers (one per segment). Those SSH servers at each network segment then communicated in the clear to each DER inverter within their segment. As noted for the flat encrypted topology, this will add protections to the Modbus traffic over the public network but in this case also comes with the additional measure of network segmentation to limit impact to grid support when some of the DERs are compromised.

Segmented Critical Distribution Clear The segmented critical clear topology is very similar to the segmented clear topology, with the only difference being that the number of segments is increased to five, as shown in Figure 41. This means that each segment has only four inverters and thus the impact of a single segment being compromised is reduced. In other words, four inverters has less generation capacity than six or seven, so the impact of losing a segment is less than would be observed in the segmented topology.

One other aspect that a segmented critical network topology brings into play is the ability to specify various assets as various levels of criticality. That is, some enclaves could be protected at higher levels than others depending on the corresponding system impact from loss of that enclave. For example, one segment could be prioritized for a hospital where another segment may not be as critical to maintain grid support capabilities. Again, since this is a clear topology the SVP control messages communicated with each DER inverter using TCP Modbus with no additional protections.

Segmented Critical Distribution Encrypted Similar to the other encrypted topologies, the segmented critical distribution encrypted topology uses SSH tunnels to forward TCP Modbus to each inverter through an SSH server. Since there are now five network enclaves, there are five fielded SSH servers—one for each network segment as shown in Figure 42.

5 Communication Latency

When cybersecurity features are added to control networks, there is an increase in communication latency from processing data, exchanging keys, binding certificates, performing encryption, or reconfiguring the system. These operations have the risk of adversely affecting real-time grid operations if the delays are significant. Several experiments were conducted to determine the communication latency associated with adding security features to DER networks. These studies were conducted using SCEPTRE and physical DER and Phasor Measurement Units (PMUs). Section 5.1 provides assumptions and caveats for the latency measurement experiments. Section 5.1 presents the impact to communication times as measured in several emulation experiments and Section 5.2 discusses the latency observed in several cases with real PMUs and smart inverters.




Figure 40: Segmented Distribution Encrypted




Figure 41: Segmented Critical Distribution Clear

Secure, Scalable Control and Communications for Distributed PV Sandia National Laboratories




Figure 42: Segmented Critical Distribution Encrypted

Before continuing, it is important to note how the latency results obtained from a system emulation are useful and what the limitations are [70]. For instance, the absolute latency values for the various security methods studied will likely not be representative of hardware implemented in the field. Various implementations of the protocol stack and firmware and hardware variations may lead to very different results. However, relative impacts from applying additional security mechanisms are illustrative and help to provide understanding on the scale of the additional time required. This information is useful in determining whether system performance will or will not be significantly degraded. Field testing with components in the environment they will be utilized should still be conducted to verify operation is as desired.

While most of this chapter is presenting and examining results of communication latencies for various security mechanisms in a control network, this section quickly discusses limitations of communication latency experimentation. It is well known from the Nyquist–Shannon sampling theorem that when measuring transient behavior in a system, the sampling rate of the system measurements must occur at least twice as fast as the fastest behavior of interest. When sampling at rate greater than this limit, it is possible to reconstruct the behaviors of interest. The importance of this limit in this context is that it gives a bound for the sampling rate needed to observe power system dynamics. When communication latencies come into play, they should fall within those bounds for the behaviors that are being observed or controlled, such as for analyzing transient stability. Likewise, in analyzing feedback control systems for DER, the sampling rate is important for the purpose of analyzing closed-loop stability of the system. Analysis of the closed-loop stability for various control mechanisms and behaviors of interest can include study of such as delay margin or the location of the discrete system poles (eigenvalues) demonstrate the regions of stability. Furthermore, power system concerns and constraints feed into the limits developed for the timing required for various controls in grid control.

As noted in DOE’s 2017 report on the Modern Distribution Grid: Volume III, the communication timing requirements for DER are on the order of seconds, with typical bandwidth and latency requirements of 10 kbps and 5 seconds, respectively [71]. These communications requirements represent generalized limits on how much latency can be tolerated between the utility and smart inverters. Prior work on three transmission-level and one distribution-level distributed DER control algorithms provided a far more detailed view of the relationship between communication latency and performance DER control algorithms. It was found that the hierarchical volt-var shift distribution algorithm was effective with latencies up to 20 seconds [58], whereas the transmission services were severely impacted with lower latencies. Synthetic Inertia experienced a loss of machine synchronism defined by rotor angle separation with latencies between 200-400 ms (depending on the gain) [72]; communications-enabled fast acting imbalance reserve was ineffective if the delay is longer than the time to the frequency nadir (e.g., 1-10 seconds depending on system inertia) [73]; and communications-enabled DER droop control experienced oscillations with latencies of 110-400 ms (depending on the gain) [74]. These findings all indicate the control algorithm will lose effectiveness with increasing latency, leading to a range of potential problems.

5.1 Communication Latency (Emulated)

5.1.1 Network Segmentation

As explained above, the cybersecurity implications for power system performance depends on the grid-support service being provided. Certain communications-enabled DER services are robust to latency and other quality of service characteristics (dropped packets, DER availability, etc.), while others are not. In this section, a SCEPTRE experiment was created to calculate the increased latency associated with adding network segmentation. Notably, the main difference in these topologies is the extra hop from breaking the DER control network into multiple segments. A round trip time (RTT) for a segmented DER network

and a flat network were calculated by pinging the DER from the utility Windows VM. The results for more than 10,000 individual measurements is shown in Fig. 43. If we assume these results are normally distributed, the mean and variance of the distribution is shown in Table 2, along with the minimum, maximum, and median values. Overlayed on top of the histogram is are plots of the normal distributions fit to the datasets. These distributions are scaled virtually to match the histogram and, thus, are not probability density functions (PDFs) because they are not continuous over time. Basically, the y-axis of this graph is a probability of a measurement falling into a bin of the histogram, not a probability of a measurement being a certain value. Another important note is these distribution results contain long tails on the right due to several outliers from instances when packets are dropped. This causes the probability density functions (PDF) shown in Fig. 43 to appear as though they don't match the results well.




Figure 43: Differences in Communication Times for Flat and Segmented Networks using Modbus/TCP with no Transport Security

Table 2: Network Segmentation Latency using Modbus/TCP

Case	Mean, μ (ms)	Standard Deviation, σ (ms)	Min (ms)	Max (ms)	Median (ms)
Flat Distribution Clear	1.5605	0.5396	0.7861	16.8277	1.5192
Segmented Distribution Clear	1.8234	0.2834	1.0188	11.2763	1.8024

5.1.2 Encryption

Secure Shell (SSH) cryptography protocols were used to wrap unsecured Modbus communications at the transport layer. To show the impact of the encryption, a SCETPRE environment was constructed with multiple SSH tunnels using common encryption cyphers and modes of operation. The network topology is shown in Figure 44. This experiment measured the communication time required for Modbus/TCP packets to traverse the network when wrapped in transport security using TLS. If we assume these results are normally distributed, the mean, standard deviation, min, max, and median of the distribution are as shown in Table 3. Note that due to the outliers (large maximum values), the standard deviation is far larger than one would expect from Figure 45.




Figure 44: Topology for Testing Communication Latency of Encryption Ciphers and Cipher Modes with Transport Security

5.1.3 Moving Target Defense

Previous research on an emulated grid wide area network (WAN) has shown that ADDSec moving target defense can be beneficial to system security during a reconnaissance or denial of service type attack in which an attacker is sending packets over the network. In a study on ADDSec resilience [66], latency measurements in the form of round trip time were taken for a WAN in which one device has been compromised by a self-propagating worm. Without ADDSec, the network hosts were infected within minutes, leading to a doubling of latency. With ADDSec in operation, fewer network hosts were infected within the same timeframe, and latency was increased to a much lesser extent. Moreover, the study demonstrated that




Figure 45: Differences in Communication Times for Common Ciphers and Cipher Modes for Transport Security of Modbus/TCP

Table 3: Encryption Latency using Modbus/TCP

Case	Mean, μ (ms)	Standard Deviation, σ (ms)	Min (ms)	Max (ms)	Median (ms)
AES128-CTR	4.0526	0.8295	2.0698	81.1382	4.0604
AES192-CTR	4.0662	1.5339	2.1778	206.7507	4.0748
AES256-CTR	4.3728	1.5879	2.0645	206.9327	4.3957
AES128-GCM	4.1056	1.5665	2.2905	205.8982	4.0985
AES256-GCM	4.4290	1.5858	2.2220	205.7683	4.4418
ChaCha20-Poly1305	4.0496	0.6043	2.1506	45.0614	4.0565
Clear	2.3834	0.4236	1.0010	15.1254	2.3847

the overhead to network latency introduced by the ADDSec SDN controller during normal operations was minimal in comparison to the latency increase during an attack.

In applying MTD to a DER communication system, one must be careful to consider network constraints and the environment in which it is operating. Although modern DER grid-services do not have strict latency or timing requirements, this could change with the integration of more sophisticated transmission or distribution grid-support services. This said, the additional latency from MTD is nearly negligible. In prior work, communication latencies for various MTD modes were determined with different randomization time periods; it was found that MTD increased the average latency by less than 1 ms but caused slightly higher dropout rates (approx. 1 dropout per 33.3 seconds with IP randomization every 3 seconds) [66]. Other approaches to MTD, like path randomization, may increase latency more. A 11.73 ms increase in RTTs for path randomization was reported by Chavez [67]. However, even though MTD does not significantly increase latency on the system, it does potentially introduce other forms of system overhead that need to be considered.

5.2 Communication Latency (Physical)

The following sections discuss results captured using physical hardware. First, the communication times for PMU messages between ABQ and several geographically distributed in locations within the continental United States are discussed to better understand latency impacts of distance. Then, timing measurements for several smart DER are discussed based on tests conducted at the Distributed Energy Technologies Laboratory (DETL) at Sandia National Laboratories.

5.2.1 Geographic Separation

The results of Figure 46 show the transit times for one-way messages from the respective PMUs to Sandia National Laboratories in Albuquerque, NM. The PMU transit times to Albuquerque are calculated using the GPS timestamp and the GPS time at the receiver. The connection to TX is over a dedicated fiber line and has minimal network hops. Conversely, the NM PMU has numerous routers and switches in the communication path which slow down the packets. In general, these results show that the architecture (switch and router hops) and communication medium (copper vs. fiber) is more important to data-in-flight times than geographic separation. This is important to keep in mind when developing the control network architecture to ensure that the number of network hops is does not impair control system operations.

Table 4: Encryption Latency using Modbus/TCP

Case	Mean, μ (ms)	Standard Deviation, σ (ms)	Min (ms)	Max (ms)	Median (ms)
PMU-1 (ABQ $\leftrightarrow$ NM)	78.9117	8.9063	61.0000	105.0000	79.0000
PMU-2 (ABQ $\leftrightarrow$ WA)	67.1551	1.5846	63.9999	86.0002	68.0001
PMU-3 (ABQ $\leftrightarrow$ TX)	36.2080	3.2368	30.9999	66.9999	36.0000

5.2.2 Smart Inverter Read and Write Times

1000 Modbus read and write times were collected for two commercially available residential-scale DER devices and one CHIL device [75] in DETL using the SunSpec System Validation Platform. The results are shown in Fig. 47 and Table 5. Inverter 1 has an large standard deviation for both read and write times. It is not clear if there are internal communication checks or other inverter processes that could be slowing some of the responses. This connection also includes a large number of outliers which significantly affect




Figure 46: Differences in Communication Time from Geographically Separated Phasor Measurement Units to ABQ

the distribution. Similar results are reported in [64]. Like Inverter 1, Inverter 2 had a direct connection of Modbus/TCP over 1 network hop, but responds much faster to both read and write requests. The connection to Inverter 3 included an Ethernet-to-Serial converter in the path to translate Modbus/TCP to serial Modbus. This added an additional delay due to the processing required to perform that conversion—possibly accounting for some of the larger average communication times for reads and writes with that device. It is believed that the variations observed in these results are not primarily due to various in the networking architecture. It is more likely that the inverters include different implementations of the protocol stack, processor hardware, and scheduling differences of processing tasks for I/O to and from memory. Further analysis would be required to determine the precise reason for the variations and draw generalization about the expected DER read and write times.




Figure 47: Differences in Communication Times for Several Common Smart Inverters

5.3 Latency Observations

Based on the results for network segmentation, encryption, MTD, geographical separation, and DER read/write times, some observations can be made about the impact to the control system when adding security features. In general, large geographic distances have the possibility of adding 50-100 ms of latency for utility-to-DER communications due to the additional networking equipment (routers and switches) between endpoints. DER read and write times vary widely; they can be 1 second or larger in some situations. In contrast, network segmentation adds less than 1 ms, encryption adds on the order of 3-5 ms of additional latency, and MTD adds 1 ms. Therefore, for the proposed cybersecurity features, it is

Table 5: Round-trip Communication Time for Modbus/TCP with Smart Inverters in DETL

Case	Mean, μ (ms)	Standard Deviation, σ (ms)	Min (ms)	Max (ms)	Median (ms)
Inverter-1 (Read Time)	163.0757	26.1437	44.9998	1145.9999	161.0000
Inverter-2 (Read Time)	3.0319	0.9801	0.9999	7.0000	3.0000
Inverter-3 (Read Time)	165.8618	1.0560	162.9999	168.0002	165.9999
Inverter-1 (Write Time)	168.3799	133.6979	38.0001	1435.0002	148.0000
Inverter-2 (Write Time)	1.9383	0.9110	0.9999	12.0001	2.0001
Inverter-3 (Write Time)	33.7298	0.6583	31.9998	36.0000	33.9999

not believed they will impact the grid-support service performance since they add only contribute a minor percentage of the total latency between the utility and DER.

6 Red Team Assessments

Red team assessments are authorized, adversary-based, cyber assessment conducted to strengthen defenses through awareness and exploitation of the system’s potential vulnerabilities. The primary objectives for each of the topologies assessed was to identify and compromise the DER devices (power inverters) by modifying communication or grid-support functions (Freq-Watt, Volt-Var, Power-Factor settings, etc.) or disrupting network communications.

6.1 Scope and Rules of Engagement

The security assessment focused on the communications between the emulated DER devices on the network, the simulated corporate and provider networks and any hardware-in-the-loop (HIL) devices. Rules of Engagement were defined as:

- Limited to the SCEPTRE experiment network
- HITL devices are in-scope
- Realistic code that exists on field devices (inverters, etc.) is in-scope
- SCEPTRE, Phenix, Minimega, OpenDSS, PowerWorld are out-of-scope

6.2 Methodology

The assessment incorporated elements from Sandia National Laboratories’ Information Design Assurance Red Team (IDART), NIST’s Guide to Industrial Control Systems (ICS) Security Guidelines, Department of Homeland Security’s ICS-CERT Recommended Best Practices, and collective expertise of PV inverter systems. The red team assessed each of the topologies with a specialized methodology based on this guidance.

The assessment team developed impact metrics based around the triad of confidentiality, integrity, and availability, also known as CIA, which serve as the core attributes for many cybersecurity risk evaluation frameworks, including the NIST Common Vulnerability Scoring System (CVSS) Impact ratings [76]. These attributes are prioritized according to the system environment and mission, with relative importance levels

captured in a system critical matrix (SCM). In this set of experiments, the DER network mission is constant, while the environmental parameters change in accordance with the network topologies. The following criticality levels were adapted from the NIST CVSS:

- HIGH [Value: 3.33] - The system MUST NOT compromise CIA for the respective critical information type; a compromise of CIA will have grave consequences. Items within the SCM are designated HIGH when a critical information compromise creates Denial of Service (DoS) conditions, severe system degradation, or loss of trust system-wide.
- MEDIUM [Value: 2] - The system should not compromise CIA for the respective critical information type; a compromise of CIA will have serious consequences. Items within the SCM are designated MEDIUM when a critical information compromise results in system degradation, but is not system-wide or not grave enough to be designated as HIGH.
- LOW [Value: 1] - The system compromise of CIA for the respective critical information type is unlikely to have serious or grave consequences.

6.3 Tools

Reconnaissance of each network topology began by actively probing from the Kali machine on the sub-networks provided. The use of Nmap and OpenVAS provided IP identification, host fingerprinting and vulnerability assessments of the devices on the network. TCPdump and Wireshark were utilized on the flat networks to capture packets on the wire for use in a replay attack as well as identification of communication protocols in the encrypted topologies. The use of COTS tools, SunSpec Dashboard, vendor-specific applications, and ModbusPal were used to craft vendor specific protocol traffic to the devices on the networks.

6.4 Emulytics Challenges

EmulyticsTM environments enable rapid security prototyping and red teaming exercises. The networks were designed to represent realistic network topologies and passed real DER protocol and encryption packets. To further improve the fidelity of the derived results, a power hardware-in-the-loop (HIL) PV inverter was added to allow the attackers to work with a physical device. However, while EmulyticsTM environments do well in faithful simulations of cyber-physical systems, the attack surfaces are typically reduced; human elements are removed, hardware, software, and firmware diversity are decreased, and overall system complexity is limited. In some instances, the discovered vulnerabilities may be artifacts of the testbed setup itself, as they may be introduced by the emulation and not present in the field. Due to these limitations, emphasis is placed on the rules of engagement which define the scope of assessment and were designed to have the red team concentrated on realistic vulnerabilities.

Although the red team methodology defined criticality levels and quantitative scoring values, they are still largely subject to the prior experience and priorities of the assessor. Assessors with varying levels of familiarity with a particular type of DER device, protocol, or system architecture may assign varying criticality levels to the same information compromise based on their interpretation of the potential impact, and subject matter expertise in DER operations is needed to accurately grade consequences of compromise to the system itself. Moreover, resources are measured in the form of time to compromise, which is subject to variability with human actors in the loop and variation in the system setup. This adds variability and uncertainty in the results that are difficult to remove. Even autonomous red teaming systems, which are meant to produce reproducible baselines, are currently still reliant on feedback from human expertise. Given these drawbacks, the red team's quantitative risk analysis would be informed based on the frequency and ease of common attacks and the goals of this assessment.

6.5 Threat Catalog

The red team developed a threat catalog to categorize the types of vulnerabilities that a threat actor may seek to exploit. By tailoring each assessment against this catalog, the red team ensured their methodology was reproducible and applicable across a wide range of systems and threat models.

- Internet Security
 - Network Surveying
 - Port Scanning
 - System Identification
 - Services Identification
 - Vulnerability Research
 - Internet Application Testing
 - Router Testing
 - Firewall Testing
 - Intrusion Detection System Testing
 - Trusted Systems Testing
 - Password Cracking
 - Denial of Service Testing
 - Containment Measures Testing
- Information Security
 - Document Grinding
 - Competitive Intelligence Scouting
 - Privacy Review
- Physical Security
 - Access Controls Testing
 - Perimeter Review
 - Monitoring Review
 - Alarm Response Testing
 - Location Review
 - Environment Review

The red team modeled a threat from an attacker equipped with specialized knowledge of DER system protocols with considerations for potential insider access, but not to the level of a nation-state actor with customized exploits. The following threats were examined:

- Data Compromise: alteration or access of confidential of data by unauthorized users.
- Code Injection: malicious code injection, such as inputs to the HMI.
- Local Exploits: exploiting vulnerabilities in user applications, such as the SunSpec Dashboard and SunSpec SVP.
- Remote Exploits: exploiting vulnerabilities in network services.
- Privilege Abuse: exploiting existing privileges for authorized users on the system.
- Interception: man-in-the-middle (MITM) or eavesdropping of authenticated communications.
- Denial of Service: rendering the system unusable to authorized users, such as overloading the RTU processors.
- Policy: exploiting flaws in policy, such as firewall security settings.
- Insider Threat: exploiting authorized user knowledge or access for malicious purposes.

```
# Nmap 6.47 scan initiated Thu Nov 8 10:39:19 2018 as: nmap -sP -T5 --min-parallelism 100 -oG output.file.txt 75.75.0.0/16
Host: 75.75.128.10 () Status: Up
Host: 75.75.128.11 () Status: Up
Host: 75.75.128.12 () Status: Up
Host: 75.75.128.13 () Status: Up
Host: 75.75.128.14 () Status: Up
Host: 75.75.128.15 () Status: Up
Host: 75.75.128.16 () Status: Up
Host: 75.75.128.17 () Status: Up
Host: 75.75.128.18 () Status: Up
Host: 75.75.128.19 () Status: Up
Host: 75.75.128.20 () Status: Up
Host: 75.75.128.21 () Status: Up
Host: 75.75.128.22 () Status: Up
Host: 75.75.128.23 () Status: Up
Host: 75.75.128.24 () Status: Up
Host: 75.75.128.25 () Status: Up
Host: 75.75.128.26 () Status: Up
Host: 75.75.128.27 () Status: Up
Host: 75.75.128.28 () Status: Up
Host: 75.75.128.29 () Status: Up
Host: 75.75.128.101 () Status: Up
Host: 75.75.128.132 () Status: Up
Host: 75.75.128.250 () Status: Up
Host: 75.75.128.251 () Status: Up
Host: 75.75.129.101 () Status: Up
Host: 75.75.129.132 () Status: Up
# Nmap done at Thu Nov 8 11:02:17 2018 -- 65536 IP addresses (26 hosts up) scanned in 1377.50 seconds
```

Figure 48: Nmap host discovery scan

6.6 Results

To execute the assessment, the red team used Kali Linux equipped with network reconnaissance and network attack tools. The red team conducted assessments for two scenarios:

- **Network Attacker** An intruder who has access to a subnet where the inverters are deployed. This adversary has no access to the DER device but does have access to one of the network switches (through public networks or as ISP operator).
- **Homeowner Attacker** The intruder is on the DER home area network (HAN) with physical access to the DER.

For each topology, the team conducted reconnaissance and active attacks including Packet Replay, Denial of Service, and MITM.

6.6.1 Reconnaissance

Network scans using Nmap and OpenVAS discovered and fingerprinted devices. Nmap was used to discover devices and networks that existed in the topology. Figure 48 shows part the results from an nmap scan. Nmap was then run at a different level of granularity to discover open ports and determine basic OS fingerprinting. Figure 49 shows the open ports on the Windows SVP machine and a linux network services machine on the corporate network.

OpenVAS was then used to probe the open ports and test for vulnerabilities. Figures 50 and 51 show the scan result from an inverter and details on a vulnerability, respectively.

SunSpec Dashboard application is designed to communicate with Modbus SunSpec RTUs. Access parameters required to connect are IP address, IP port, slave ID, and timeout period. The application displays available registers on the device and, depending on the parameter, may be writeable. Figure 52 shows the connection to a SCEPTRE RTU inverter.

```

root@      :~# nmap -O 192.168.0.100

Starting Nmap 6.47 ( http://nmap.org ) at 2018-09-10 08:03 MDT
Nmap scan report for 192.168.0.100
Host is up (0.0024s latency).
Not shown: 991 closed ports
PORT      STATE SERVICE
135/tcp   open  msrpc
139/tcp   open  netbios-ssn
445/tcp   open  microsoft-ds
49152/tcp open  unknown
49153/tcp open  unknown
49154/tcp open  unknown
49155/tcp open  unknown
49156/tcp open  unknown
49157/tcp open  unknown
Device type: general purpose
Running: Microsoft Windows 7|2008
OS CPE: cpe:/o:microsoft:windows_7:- cpe:/o:microsoft:windows_7::spl cpe:/o:microsoft:windows_server_2008::spl cpe:/o:microsoft:windows_8
OS details: Microsoft Windows 7 SP0 - SP1, Windows Server 2008 SP1, or Windows 8

OS detection performed. Please report any incorrect results at http://nmap.org/submit/ .
Nmap done: 1 IP address (1 host up) scanned in 15.38 seconds

root@      :~# nmap -O 75.75.129.101

Starting Nmap 6.47 ( http://nmap.org ) at 2018-09-10 08:10 MDT
Nmap scan report for 75.75.129.101
Host is up (0.00079s latency).
Not shown: 996 closed ports
PORT      STATE SERVICE
22/tcp    open  ssh
25/tcp    open  smtp
80/tcp    open  http
443/tcp   open  https
Device type: general purpose
Running: Linux 3.X
OS CPE: cpe:/o:linux:linux_kernel:3
OS details: Linux 3.11 - 3.14
Network Distance: 2 hops

OS detection performed. Please report any incorrect results at http://nmap.org/submit/ .
Nmap done: 1 IP address (1 host up) scanned in 15.21 seconds

```

Figure 49: Nmap host fingerprinting results

Greenbone

Security Assistant

Logged in as Admin

admin

| Logout

Fri Sep 14 13:49:01 2018 UTC

Scan Management

Asset Management

SecInfo Management

Configuration

Extras

Administration

Help

▼ Report: Results

1 - 9 of 9 (total: 9)

PDF

Done

Filter:

sort-reverse=severity result_hosts_only=1 min_cvss_base= levels=hmlg

Vulnerability	Severity	Host	Location	Actions
TCP Sequence Number Approximation Reset Denial of Service Vulnerability	5.0 (Medium)	75.75.128.14	general/tcp	
TCP timestamps	2.6 (Low)	75.75.128.14	general/tcp	
CPE Inventory	0.0 (Log)	75.75.128.14	general/CPE-T	
ICMP Timestamp Detection	0.0 (Log)	75.75.128.14	general/icmp	
Record route	0.0 (Log)	75.75.128.14	general/icmp	
OS Detection Consolidation and Reporting	0.0 (Log)	75.75.128.14	general/tcp	
Traceroute	0.0 (Log)	75.75.128.14	general/tcp	
NTP read variables	0.0 (Log)	75.75.128.14	123/udp	
LDAP Detection	0.0 (Log)	75.75.128.14	5502/tcp	

(Applied filter: sort-reverse=severity result_hosts_only=1 min_cvss_base= levels=hmlg autofp=0 notes=1 overrides=1 first=1 rows=100 delta_states=gn)

1 - 9 of 9 (total: 9)

Figure 50: OpenVAS scan result on an inverter

Medium (CVSS: 5.0) NVT: TCP Sequence Number Approximation Reset Denial of Service Vulnerability
Summary The host is running TCP services and is prone to denial of service vulnerability.
OID of test routine: 1.3.6.1.4.1.25623.1.0.902815
Vulnerability Detection Result Vulnerability was detected according to the Vulnerability Detection Method.
Impact Successful exploitation will allow remote attackers to guess sequence numbers and cause a denial of service to persistent TCP connections by repeatedly injecting a TCP RST packet.

Figure 51: OpenVAS vulnerability description




Figure 52: SunSpec Dashboard connected to an inverter

By default, Modbus slaves listen on port 502. From the results of the above scans, Modbus was identified to be running on a custom port. The inverters were easily accessible with the SunSpec Dashboard application on this port and inspecting the connection revealed the SunSpec ID number, 0x53756e53, identifying the SunSpec Modbus Map for the Modbus traffic on the non-standard port 5502, shown in Figure 53.

Wireshark was used to analyze and reverse-engineer the communications using SunSpec Dashboard. For example, the mappings of SunSpec Model 101 (0x65), length 50 (0x32) are shown in Wireshark in Figure 54.

The settings for each Modbus register address were mapped by capturing packets for each value read, as seen in the packet capture. For example, the Modbus Nameplate registers were read using SunSpec Dashboard and packets were captured similar to Figure 54 above, and Inverter Model 101 bytes are shown in Figure 55.

6.6.2 Packet Replay

Packet replay is an attack in which data transmission is resent or repeated in a manner that causes undesired results. Utilizing the mapping discovered using Wireshark and the SunSpec Dashboard, packets captured in the reconnaissance could be resent via netcat, a simple network communication tool. With some value modification and scripting, the packet replay attack could be carried out on all inverters and multiple register locations autonomously. Shown in Figure 56 below is the packet capture of cycling all inverters connection register via a script.

Other replayed packets included modification of inverter phase voltages, DC voltages, current, and power. Unauthorized actions demonstrate that an adversary could easily transmit fraudulent data to




Figure 53: SunSpec ID number on port 5502 using SunSpec Dashboard




Figure 54: SunSpec Table 101 (0x65), length 50 (0x32)




Figure 55: Mapping of Modbus Nameplate registers




Figure 56: SunSpec Dashboard showing the connection register state cycling controls on two inverters

falsify the inverter's state and disrupt network communications.

6.6.3 Denial of Service

A Denial of Service (DoS) attack is an interruption in an authorized communication on a network. Reconnaissance of the inverters indicated that they were susceptible to TCP SYN Flood attacks. Tools in the Kali suite called floodrouter6 and hping3 were used to send spurious router advertisements and TCP SYN requests from random IP addresses to devices on the network, respectively, causing network communication outages. The DoS attacks were successful on inverters and network devices between inverters and the SVP, in most cases.

6.6.4 Man in the Middle

Man-in-the-Middle (MITM) is the act of eavesdropping, dropping, delaying or altering communications while in transit from source to destination. For this attack, a tool in the Kali suite called Ettercap was used. The tool ARP spoofs via ARP cache poisoning the two devices of which the communications are desired to be intercepted. In all cases, MITM attacks worked for any devices on the same subnet. In unencrypted topologies or segments, MITM attacks between an inverter and its destination (gateway or SSH gateway), packets were captured and visible in plaintext.

6.6.5 Flat network topology without encryption

- **Observations** Reconnaissance showed the Red Team's position was on the same subnet as the inverters. A router separated the utility's DERMS system into a separate network. The router and DERMS were susceptible to Denial of Service attacks. Man-in-the-Middle attacks were possible between each inverter and the router. Packet replay was possible directly to the inverters.
- **Challenges** As a baseline for the assessment, no challenges were found.

6.6.6 Flat network topology with encryption

- **Observations** Reconnaissance showed that encryption was added via a bump-in-the-wire (SSH server) technique, and as a result of a flat network with the Red Team presence being on the subnet of the inverters, the traffic on the subnet was unencrypted until it passed through the SSH server. Tools were able to reach directly to the inverters, which allowed for register changes using netcat and SunSpec Dashboard without challenge. In this, and all encryption enabled topologies in this assessment, the SSH gateway machines had a password-less root login enabled for SSH, an oversight for the Blue Team setup. The Red Team was able to log in, pull SSH encryption keys and fingerprints, capture traffic from the inverters before encryption, and pivot onto the corporate segment of the network through the SSH tunnel. This would have been a valuable pivot if devices on the corporate network were not already routable and reachable by the Red Team's position on the Kali machine. On this topology, DoS, MITM and packet replay were all successful.
- **Challenges** On a bump-in-the-wire encryption setup, an attacker positioned outside the local network will only see encrypted traffic across any potential attacker-controlled parts of the network, preventing an attacker from reading or modifying the traffic passing through. After obtaining the SSH keys from the SSH tunnel hosts, decryption of the tunnel traffic was investigated. Even the older version of SSH found in the topologies utilizes a Diffie-Hellman key exchange encryption suite so an attacker who has decrypted traffic from the SSH key negotiation cannot decrypt the rest of the conversation.

6.6.7 Segmented-Network Topology without Encryption

- **Observations** The Red Team was provided two access points, one on an empty subnet and one on one of the subnets with inverters. On the empty subnet, or outsider access, MITM was unsuccessful because there were no hosts susceptible to an ARP poisoning attack. Any attempts to pivot and deploy MITM tools were unsuccessful due to Linux package dependencies on an air-gapped network. In both positions, DoS and packet replay to the inverters were successful.
- **Challenges** From an outsider position on an emulated network, it is not a target-rich environment. Pivoting into subnets with targets is difficult when hosts do not have the human element and the OS vulnerabilities seen in the real-world.

6.6.8 Segmented-Network Topology with Encryption

- **Observations** The addition of encryption from the segmented unencrypted network added bump-in-the-wire SSH gateway hosts on a subnet basis. While MITM was still an available attack when on the same subnet, an outsider without the ability to pivot and deploy tools remains excluded from this attack vector.
- **Challenges** No unique challenges were introduced in this topology.

6.6.9 Segmented-Network Topology with HIL and without Encryption

- **Observations** The addition of a physical inverter in the topology provided a target from which the Red Team previously conducted an assessment. In that assessment, the team conducted successful reconnaissance, vulnerability scans, packet replay, MITM, sniffed passwords, DoS attack, and validated proper bookkeeping (logs) of the device during a security event. In contrast with the previous assessment, the DER device was not on the same subnet, and thus the vendor software and DER Connection Assistant tools were unable to discover the device. SunSpec Dashboard was able to successfully connect to the DER device, and showed some minor differences from the emulated inverters in the topology—most notably the lack of a connection register previously used to disconnect the inverter’s communications.

It was shown that flipping the DER Volt-VAR curve about the reactive power-axis caused the device to sink power when phase voltages were low. This is not a desired operating mode for the equipment because it will drive the power system away from nominal voltage. The effect can be seen in Figure 57.

Except for the addition of the HIL, the attacks conducted on the topology were the same for the Segmented-Network Topology without Encryption.

- **Challenges**

6.6.10 Moving Target Defense Network Topology

- **Observations** The MTD environment was difficult to conduct reconnaissance because the IP addresses of the equipment regularly changed. DoS attacks on the router were still successful because it prevented traffic between the utility and the DER equipment.
- **Challenges** This environment was far more challenging to craft MITM because the target’s IP address kept changing.




Figure 57: SunSpec Dashboard reading DER with inverted VV curve

6.7 Comments

Throughout the course of testing and assessment, a few issues were found with experimental setup and design. These were corrected as found and are listed below:

- A Modbus packet reading 10 holding register values causes the SCEPTRE inverter process to quit unexpectedly.
- The SVP process representing the utility DERMS stopped periodically, ceasing Modbus traffic. The fix was to automatically restart the provider process.
- Modification of an inverter's Volt-Var settings caused all other inverters' settings to be modified (unintended).
- The MTD mechanism is sensitive to communication loss during DoS attacks, and does not restart, leaving the networking in a fail-visible state.

6.8 Summary

Based on the assessment results, it is clear that adding each of the cybersecurity features improved the security posture of the DER network. Adding enclaving increased difficulty for the Red Team to access all

Topology	Encryption	Access	Attacks			Results			t_1	t_{all}	Score
			DoS	Replay	MITM	C	I	A			
Flat	None	Homeowner	x	x	x	0	0	0.2	0.5	1.0	0.2
Flat	RFC 7539	Homeowner	x	x	x	0	0	0.2	0.5	1.0	0.2
Segmented	None	Outsider	x	x	x	0	0	0.2	0.5	1.0	0.2
Segmented	None	Homeowner	x	x	x	0	0	0.2	0.5	1.0	0.2
Segmented	RFC 7539	Outsider	x	x	x	0	0	0.2	0.5	1.0	0.2
Segmented	RFC 7539	Homeowner	x	x	x	0	0	0.2	0.5	1.0	0.2
Segmented	None	Outsider	x	x	x	0	0	0.2	0.5	1.0	0.2
MTD	None	Homeowner	x	x	x	0	0	0.2	0.5	1.0	0.2

- t_1 is the estimated time to access and control 1 DER device in hours
- t_{all} is the estimated time to access and control all the DER devices in the simulation in hours
- RFC 7539 is the IETF Protocol for the ChaCha20 stream cipher and Poly1305 authenticator

Table 6: Quantified security scores for different DER communication networks based on red team assessments.

the DER devices to coordinate a cyber attack that affected a large number of DER. The use of encrypted tunnels between the utility and the DERs prevented the adversary from launching attacks from the public internet, although, an error in the SSH implementation exposed an enclave to control. This is an important result, as it reinforces the risks associated with poor network management practices. The moving target defense technique is not widely used because of its added expense and capital/configurational overhead but provided a substantial improvement in the cyber resilience of the DER network.

7 Challenges with Operating the Power System with No Communications

The last component of the project explored the reliance on communication systems for grid operations and considered possible design changes to minimize communication needs. These low-communication strategies provided operational alternatives for power systems to increase cyber resilience and potentially act as fallback operating modes for power systems that are impacted by cyber or physical events. For each of the designs, the dependencies between the communication system and electric power grid are described and analyzed. Arguably, for an increased cost, the system can retain some level of preparedness for a loss of communications and control in partnership with standardized cyber attack mitigation and recovery measures. Ultimately, the preservation of system load requires a broad view of subsystem interconnectivity within the grid and should be evaluated holistically. Far greater detail on the history of the power system and issues that appear with communications loss, refer to [77].

7.1 Problem Conceptualization

Operating the electric power grid without communications systems is possible, but doing so would require significant and costly modifications to existing infrastructure and major changes to operating procedures. A more prudent approach is to focus on technologies and procedures necessary to transition temporary operation without communication, when needed, preserving the significant economies afforded by communication systems during their availability. Over the last century, physical control systems such as flywheel governors have largely been replaced by increasingly complex, wide-area digital control and communications systems. Modern systems afford higher performance and economy, can cover greater distances, and

have increased the reliability of our grid. These systems allow higher resource utilization of our grid infrastructure. This not only lowers costs, but also it defers the need to expand transmission and distribution infrastructure which has growing public opposition. Regulatory effects, environmental issues, and 'Not In My Backyard' public attitudes have greatly minimized the expansion of transmission infrastructure in recent decades. This has resulted in efficient, higher transfer capacities but with lower actively controlled margins and sharper edges to instability. Although these margins are well managed using advanced communications and controls, a loss of communications and control requires immediate actions by system operators to maintain sufficient operating margins.

For the North American grid, as currently designed, a pervasive loss in communications would require an immediate reduction in power transfer to ensure stability, which implies the selective tripping of loads and/or the redispatch of generation. Alternatives to this scenario would likely require significant changes to the grid system architecture and an increase in cost, but there are key themes that will emerge if a large grid is made to safely and reliably operate void of communications. These include:

1. Grid operations must be coordinated during a loss of communications.
2. A transition between operation with and without communications should be able to occur in real time.
3. Find a solution for today's grid that will also work for tomorrow's grid consisting primarily of inverter-based resources.

In order to analyze the impact of a loss of communications would be to the power system, two scenarios were created based on a model of the Western North American Power System shown in Fig. 58 (see details of the mini-WECC model in [78]). In the first, constrained dispatch coordination is lost at low load levels, generators collectively increase their production to match increasing load but a pathway is overloaded and trips. In the second scenario, the communications required to enact a remedial action scheme (RAS) are severed and a fault in the southwest causes an overload and trip of the same north-south pathway. The scenarios exhibit different, related grid problems, i.e., the first scenario compounds the effects of the second. Such compounding effects are difficult to predict and result in increased uncertainty with respect to harmful consequences. Details of the first scenario is provided in the following section. The second scenario is presented in [77].

7.1.1 Dispatch Drift without Communications

In the WECC there are about 6000 industrial generators, only some of which are in service and dispatched at a given time, depending on the season, time of day, and maintenance schedules. A coordinated dispatch is called a Security Constrained Economic Dispatch (SCED), which means that the most economic choice of generation is selected, as long as it does not impact system security (reliability). Security can affect the dispatch because electricity flows across the network according to the laws of physics (from sources to loads, through a network), without a means of direct control. However indirect control of electrical flows can be created by changing the source power injected into the network at different locations—different generators. Doing this is often necessary to ensure that transmission lines do not become overloaded, hence the economic dispatch becomes constrained to ensure grid security.

Between dispatches, which typically occur at one-hour intervals, generating reserves are standing by in case a contingency occurs, such as the loss of a large power plant. Random deviations in load and renewable energy are managed by regulating reserves using an automatic generation control system (AGC). The AGC system depends on communication systems and will be assumed to be out of service for the purposes of this discussion. This leaves system balance to be managed using a combination of generator droop control and generator plant operator intervention.




Figure 58: Western North American Power System model used for the power system studies.

In the event of a loss of communications, the use of security constrained dispatch will not be available since it relies on topology data obtained using a state estimation system, which is informed by SCADA data. Although generation and load schedules will be available based on forecast system conditions, these schedules will not be modifiable based on system topology changes, load changes, generation outages, or contingencies. Therefore, the schedules will be semi-valid.

To better understand the implication of an uncoordinated dispatch, the system was assumed to have lost communications at the lowest daily load, about 0400 hours. As load increased, the generator plant operators are assumed to follow the guidelines of *Western Interconnection Generator Frequency Operating Guide During a Loss of Complete Communications* [79], increasing plant output in response to frequency decline as explained above. The maximum deviation between an actual generator's output and its schedule

is expected to occur at the highest load value, corresponding to about 1700 hours. This can be assumed since each generating plant that is operating per [79] will not operate in a coordinated manner. Some plant operators may be quick to respond, resolving the system frequency problem before other plants have a chance to increase theirs. The result is that over time, all plant outputs are not equally increased, nor are they following a schedule that ensures system security constraints are met. This randomness in generator output is attributed to human reaction time.

It is reasonable to assume that before a failure in communications occurs, all generating plants have a predetermined schedule from which they will attempt to operate throughout the day. It's also fair to assume that the load forecast for the day (from which the schedules are based) are highly accurate—this assumes the loss of communications does not influence societal load patterns. However, during normal operation various power plants are replaced by others at scheduled intervals due to economic operations. The increase in power output of one plant is time synchronized with the decrease in power output of another. In this manner, generation and load remain balanced during changes to dispatch (ignoring diurnal load increases and regulation requirements). However, for the loss of communications scenario, the generating plants are not expected to decrease their output, even if schedules dictate. So generators that are scheduled to come online during the event will do so, but as they try to increase their power output to their scheduled amount, they will find system frequency increasing as a result, and will back off their power in order to avoid causing an over frequency condition per [79]. These generators will likely stay online, providing capacity if a drop in system frequency indicates that it is needed, but increases in generator power outputs will not be according to schedule, thus they will not necessarily be secure dispatches. This means that the power flows across the grid network may or may not cause transmission overloads or other reliability problems.

Using the mini-WECC model, power flows were determined (in percent, normalized to the day's peak load condition) for select WECC paths. Path flow limits, in MW, exist but are not available for public dissemination. Each path identifies power flow over a specific set of power lines. Generally, paths provide a means of observing system power flows at a higher level than individual transmission lines. For each path, limits exist, which must be enforced during operations. Enforcement of path limits does not guarantee that individual transmission lines will not be overloaded (causing a violation of a security limit) but generally, the path limits are defined in a manner that prevent overloads.

We see multiple WECC paths in Fig. 58. These paths are simulated on the mini-WECC model, a reduced order model, and some individual transmission lines have been consolidated. Although the mini-WECC model omits some of the specific transmission lines used to define WECC paths, the existing lines are sufficient to provide a good indication of WECC path flows. The paths of greatest interest are those which represent the highest consistent flows. The UL75 metric is defined by the percentage of time each path is loaded beyond 75%. For this assessment, we've selected paths 19, 36, 48, 49 and 66. Although some of these paths are typically less than 75% loaded, some of them have a large impact on overall system stability. Others, such as Path 10 are essentially radial, largely unaffected by the system dispatch, except for the Colstrip units supplying them, thus they provide little interest to this analysis.

A Monte Carlo analysis was conducted using the mini-WECC model. One thousand random dispatches were evaluated, with each generator's power output uniformly distributed between $\pm 15\%$ of its nominal output power for the system peak load. For each dispatch, the path flows were calculated and stored, as shown in Fig. 59. The histograms in Figure 60 provide insight into the possible transmission path loads during a severe loss of communications. The most interesting result is the effects of Path 66, the California-Oregon Intertie (COI). This particular path is critically relied upon during heavy loads, transferring excess hydro generation into California from the Pacific Northwest. There are several Remedial Action Schemes, which rely upon wide-area communication systems, and which permit the reliable transfer of power along this path. As can be seen from Figure 60, there are conditions in which Path 66 has more than doubled its flow. Sustaining an extreme condition like this is not possible, and would result in line trips to interrupt

flow, very likely before achieving the high flows shown. The consequences of unanticipated line trips could range from tens of thousands of MW of load dropped in California, to generation overspeed and subsequent trip in the Pacific Northwest, any of which could lead to a cascading outage.




Figure 59: Generator plant dispatch ranges for a case corresponding to generation dispatch of 110.6 GW and max load of 106 GW.

This example has served to explain one aspect of our dependency on communications systems for grid operations. The example assumed very extreme conditions regarding a complete and comprehensive loss of communication. Although these results provide a bounding of what could happen, they do not reflect what is expected to happen, since the pervasive loss of communications is an extremely unlikely scenario. But the results do provide some insights for even a more limited loss of communications. As seen in upper left histogram of Figure 60, a complete loss of communications results in path flows far exceeding system capability, implying that local protection relays would likely trip, such as Zone 3 impedance relays which emulate overcurrent protection) well before the condition shown was reached. Once a transmission line trips due to overloading, other transmission lines in the network are forced to carry the power using a different path to the load, thus increasing their load, placing them at risk of tripping and a potential cascading scenario.

7.2 Recommendations

The example above shows just one of many challenges that will occur if communications are lost in power system operations. While this is highly unlikely and other methods of communications (e.g., cell phones) will be used in a situation like this, it is illustrative of problems that will emerge with communication failures. Fortunately, there are ways to provide secure operation of the grid without communications, but they imply:




Figure 60: Histograms of select path flows in the WECC when the system dispatch is randomly (uniformly) adjusted within $\pm 15\%$ from its dispatch during peak load conditions.

1. Changes in control system design, especially wide area controls including automatic generation controls, wide area protection, and some local protection such as the use of transfer tripping and blocking.
2. Changes in operating procedures; these not only affect the cases with and without communications, but the difficult task of coordinating a transition from the former to the latter in real-time.
3. Changes in system planning

Based on the Dispatch Drift analysis above, it is clear that a drop in power delivery would be required to maintain system reliability when communication is lost because otherwise transmission path overloading will occur, resulting in line tripping, and even more wide-spread load shedding. There is a need to develop strategies, procedures and apply technologies which can make a loss of communications to our power grid less costly financially and socially. These strategies should be incorporated into ongoing infrastructure planning and operating processes. This approach should be used as a reinforcement to enhance cybersecurity response and recovery plans.

In the future, analysis of a partial loss of communications would be interesting. It would require (A) a significant model development endeavor which combines power system models with communications models, and (B) identifying and collection communication system models which interact with the grid, major control systems, and RAS Systems. Additionally, topic areas for managing system operations with limited to no communications, include:

1. **The development of operating procedures** to coordinate and manage system operation without communication systems in place. This work would focus on establishing methods to allow secure operation in terms of frequency stability, line and path loading, voltage control, redispatch and

changing dispatch schedules, real-time balancing, and voltage, transient and small signal stability management. The work would also recognize that the most probable cases will include only a partial loss of communication, and procedures should allow for different scenarios in this respect.

2. **The development of resilience communication systems** using multiple modes of communication, which are anticipated to have different failure modes. For example, if one mode of communication is vulnerable to an Electro Magnetic Pulse (EMP), consider a redundant mode of communications that is not.
3. **The use of advanced demand response** has a significant potential to help the grid maintain stability. Although allowable, grid operation procedures prevent the use of load shedding except in extreme circumstances. If that practice is used by grid operators to prevent a disaster, and if they are successful, they will have difficulty proving that it was necessary and may be criticized for it. The use of advanced demand management will be employed in stages of lowest to highest criticality.
4. **The integration of communications into system planning studies** is only partially conducted. For example, the NERC standard TLP-001-4 requires system performance to be maintained under specific contingency conditions. However, these conditions do not include ensuring system performance while modeling communication contingencies. Although NERC communication standards do exist which aid the overall system reliability such as NERC COM-001 and COM-002, communication systems are not integrally modeled with power assets to determine system performance, but by exception. Overall system planning studies should be considered comprehensively with both communications and power systems infrastructure. Doing such, especially at the interconnection scale, would be a major undertaking.

8 Significant Accomplishments and Conclusions

This project produced a number of contributions to the power system and cybersecurity scientific and engineering communities. Three transmission-level and one distribution-level distributed control algorithms were developed and studied, along with their resistance to networking quality of service parameters. It was found that the hierarchical volt-var shift distribution algorithm was effective with latencies up to 20 seconds, whereas the transmission services were severely impacted with lower latencies. Synthetic Inertia experienced a loss of machine synchronism defined by rotor angle separation with latencies between 200-400 ms (depending on the gain); CE-FAIR is ineffective if the delay is longer than the time to the frequency nadir (e.g., 1-10 seconds depending on system inertia); and CE-Droop experienced oscillations with latencies of 110-400 ms (depending on the gain). These findings provide grid operators information regarding appropriate settings and resistance to communication failures or latencies.

This project then constructed the first DER communication network simulation connected to a distribution system power simulation. To do this, a number of new devices and interfaces had to be created. The SCEPTRE interface to OpenDSS was specifically created for this project in order to provide realistic cybersecurity assessments of distribution-connected equipment. A first-of-its-kind virtualized SunSpec-compliant RTU was assembled and interfaced the power simulation and incorporated into the virtual DER communication network. This project was the first to conduct adversary-based cybersecurity assessments on a virtualized DER network. The results were enlightening in terms of the changes to the cybersecurity posture of the system with different security features.

Based on these experiments, the marginal decrease in communications speed and bandwidth are justified to significantly increase the security posture of OT networks. It is recommended that DER communication networks are segmented and employ encryption. Moving target defense is more challenging to implement in the field because of the required out-of-band communication network, although many DER devices have cell modems or other communication modes to reach the OEM or aggregators. Future work should be conducted on this promising technology to determine if possible deployment on DER networks is financially and technically practical. Additionally, the use of SCEPTRE to virtualize DER equipment, communication networks, and power systems was found to be highly effective at comparing security methodologies in terms of QoS performance and resilience to cyberattacks. Continued development and use of the co-simulation platform is recommended to (a) assess security features for DER and other communication networks and (b) evaluate the impact of these technologies on communications and control system performance metrics.

9 Inventions, Patents, Publications, and Other Results

In the first year, two conference papers were presented at the IEEE Innovative Smart Grid Technologies (ISGT-2017) conference. A paper was accepted to IEEE Power Engineering Letters. A presentation was made at two working group meetings in November: the WECC Modelling and Validation Working Group (MVWG) and the WECC Renewable Energy Modeling Task Force. A technical advance (patent application) was submitted in early November for Communication-Enabled Fast Acting Imbalance Reserve (CE-FAIR).

In the second year, the team focused on the impact to the power system from cybersecurity attacks and submitted a impact study to IET Cyber-Physical Systems: Theory & Applications. Similar work was presented at a DOE CEDS Outreach webinar in 2018. The team also produced a cybersecurity primer for DER vendor, aggregators, and utilities. The SCEPTRE environment the team was constructing was also introduced and demonstrated at the Workshop on Co-Simulation Platforms for the Power Grid at LBNL.

In the final year, the team produced an IoT Journal manuscript on the findings of the SCEPTRE red team assessments, along with a more detailed SAND report on the methodology and results. A cumulative list of publications, patents, and awards is listed below.

- R. Byrne, R. Elliott, F. Wilches-Bernal, R. Concepcion, J. Neely, O. Lavrova, and J. Quiroz, "Small signal stability of the western North American power grid with high penetrations of renewable generation," in proceedings of the 43rd IEEE Photovoltaic Specialists Conference, Portland, OR, June 2016.
- M. Reno, J. Quiroz, O. Lavrova, and R. Byrne, "Evaluation of Communication Requirements for Voltage Regulation Control with Advanced Inverters," in proceedings of the IEEE North American Power Symposium 2016, Denver, CO, September 2016.
- F. Wilches-Bernal, R. Concepcion, J. Neely, R. Byrne, and A. Ellis, "Communication Enabled – Fast Acting Imbalance Reserve (CE-FAIR)," in IEEE Transactions on Power Systems, vol. 33, no. 1, pp. 1101-1103, Jan. 2018.
- J. Johnson, B. Richardson, K. Schwalm, I. Onunkwo, P. Cordeiro, B. Wright, N. Jacobs, C. Lai, "Assessing DER Network Cybersecurity Defenses in a Power-Communication Co-Simulation Environment," IEEE Internet of Things Journal (in preparation).
- J. Johnson, J. Quiroz, R. Concepcion, F. Wilches Bernal, M. Reno, "Power System Effects and Mitigation Recommendations for DER Cyber Attacks," IET Cyber-Physical Systems: Theory & Applications (accepted).
- R. Concepcion, F. Wilches-Bernal, R. Byrne, "Effects of Communication Latency and Availability on Synthetic Inertia," 2017 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT), Washington, DC, pp. 1-5. April 23-26, 2017.
- R. Guttromson, M. Donnelly, and D. Trudnowski, "Widespread loss of communications in grid systems," Tech. Rep., Sandia National Laboratories, 2018.
- J. Quiroz, M. Reno, O. Lavrova, R. Byrne, "Communication Requirements for Hierarchical Control of Volt-Var Function for Steady-State Voltage," 2017 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT), Washington, DC, 2017, pp. 1-5.

- I. Onunkwo, B. Wright, P. Cordeiro, N. Jacobs, C. Lai, J. Johnson, T. Hutchins, W. Stout, A. Chavez, B. T. Richardson, K. Schwalm, "Cybersecurity Assessments on Emulated DER Communication Networks," Sandia Technical Report, Dec 2018.
- C. Lai, N. Jacobs, S. Hossain-McKenzie, C. Carter, P. Cordeiro, I. Onunkwo, J. Johnson, "Cyber Security Primer for DER Vendors, Aggregators, and Grid Operators," Sandia Technical Report, SAND2017-13113, Dec 2017.
- J. Johnson, "The Potential Impact of DER Cybersecurity Vulnerabilities," DOE OE CEDS Outreach Webinar Series, 27 June 2018.
- N. Jacobs, J. Johnson, "SCEPTRE: Power System and Networking Co-Simulation Environment," Workshop on Co-Simulation Platforms for the Power Grid, LBNL, Berkeley, CA, 21 May 2018.
- R. Byrne, "Secure, Scalable Control and Communications for Distributed PV," WECC Renewable Energy Modelling Task Force (REMTF), Los Angeles, CA, November 16, 2016.
- F. Wilches-Bernal, "Secure, Scalable Control and Communications for Distributed PV," WECC Modelling & Validation Working Group (MVWG), Los Angeles, CA, November 17, 2016.
- F. Wilches-Bernal, J. Neely, R. Byrne, R. Concepcion, A. Ellis, "Communication Enabled Fast Acting Imbalance Reserve (CE-FAIR)", technical advance, submitted November 3, 2016.
- O. Lavrova, "Communication Enabled Fast Acting Imbalance Reserve," SunSpec Alliance Industry Meeting, Las Vegas, NV, September 13, 2016.
- R. Byrne, elevated to IEEE Fellow effective January 1, 2017, notified November 22, 2016.

10 Path Forward

There is substantial follow-up work that would easily build on the contributions from this project. The general approach of simulating power system control algorithms to determine sensitivities to latency, dropouts, and availability should be expanded to other control methods. This will help inform grid operators of their operation margins for different systems and further warn them if a control algorithm is unsuitable when using unresponsive devices or slow communication networks. Additionally, this work showed a wide range of round-trip communication times. It is important to have a better understanding of these latencies to estimate power system operations.

This project also constructed the first DER network and distribution system co-simulation environment. This SCEPTRE environment is extremely powerful and could be used to study a range of cybersecurity features or approaches. And while the team was able to construct and evaluate network segmentation, SSH encryption, and moving target defense, this system should also be used to determine appropriate techniques for:

- Architecting different communication topologies (e.g., utility-to-DER, utility-to-aggregator-to-DER, etc.) with sufficient security features at each level of the communication system.
- Building effective Intrusion Detection Systems (IDSs) for DER networks.
- Experimenting with different encryption and trust mechanisms (symmetric and asymmetric/PKI encryption), various key sizes, etc.
- Defining role-based access controls for DER.
- Researching different DER communication protocols (IEEE 2030.5, IEEE 1815, SunSpec Modbus, IEC 61850-7-420, OpenADR, etc.) and associated security features.
- Selecting firmware update/patch rates to minimize risks of malware/botnets without significantly impacting vendor or aggregator operations.
- Choosing perimeter defenses, intrusion prevention systems (IPSs), and firewall rules for utilities and aggregators.

Clearly, there are many experiments that could be constructed using the virtualized infrastructure to study the DER security status quo and design improvements for adoption by industry. Results from these assessments can be fed to codes and standards development organizations to provide the greatest standardized security to the solar and power industry. In fact, the SunSpec/Sandia DER Cybersecurity Workgroup continues to research these topics and regularly provides input to standards development processes, and this team is working on many of these topics through DOE CEDS- and SETO-funded projects.

References

- [1] Idaho National Laboratory, “Cyber threat and vulnerability analysis of the U.S. electric sector, mission support center analysis report,” Tech. Rep. INL/EXT-16-40692, Idaho National Laboratory, August 2016.
- [2] R. Smith, “How a U.S. utility got hacked,” *The Wall Street Journal*, December 2016.
- [3] R. Smith, “Russian hackers reach U.S. utility control rooms, homeland security officials say,” *The Wall Street Journal*, June 2018.
- [4] P. Caine, “Russian-backed hackers infiltrating US power grid,” *WTTW*, August 2018.
- [5] D. Q. Wilber, “Russian malware found on vermont electric utility laptop,” *Los Angeles Times*, January 2017.
- [6] D. Bradbury, “Staff dust off their typewriters after malware attack,” *Naked Security*, August 2018.
- [7] D. Kirkpatrick, “British cybersecurity chief warns of russian hacking,” *The New York Times*, November 2017.
- [8] K. Zetter, “Inside the cunning, unprecedented hack of ukraine’s power grid,” *Wired*, March 2016.
- [9] A. Greenberg, “How an entire nation became russia’s test lab for cyberwar,” *Wired*, June 2017.
- [10] A. Greenberg, “‘Crash Override’: The malware that took down a power grid,” *Wired*, June 2017.
- [11] R. J. Campbell, “Electric grid cybersecurity,” *Congressional Research Service*, September 2018.
- [12] B. G. M. Riley, J. Dlouhy, “Russians are suspects in nuclear site hackings, sources say,” *Bloomberg*, July 2017.
- [13] N. Perlroth, “Hackers are targeting nuclear facilities, homeland security dept. and F.B.I. say,” *The New York Times*, July 2017.
- [14] N. Perlroth, “Cyberattacks put russian fingers on the switch at power plants, U.S. says,” *The New York Times*, March 2018.
- [15] A. Greenberg, “The untold story of notpetya, the most devastating cyberattack in history,” *Wired*, August 2018.
- [16] E-ISAC, “Analysis of the cyber attack on the ukrainian power grid: Defense use case,” tech. rep., E-ISAC, March 2016.
- [17] Solar Energy Industries Association and GTM Research, “U.s. solar market insight q3 2018,” tech. rep., SEIA, September 2018.
- [18] U.S. Energy Information Administration, “U.S. battery storage market trends,” tech. rep., SEIA, May 2018.
- [19] S. Page, “Hawaii will soon get all of its electricity from renewable sources,” *Think Progress*, May 2018.
- [20] I. Penn, “California invested heavily in solar power. now there’s so much that other states are sometimes paid to take it,” *LA Times*, June 2017.

- [21] IEEE Std. 1547-2018, “IEEE standard for interconnection and interoperability of distributed energy resources with associated electric power systems interfaces,” tech. rep., Institute of Electrical and Electronics Engineers, Inc., Feb 2018.
- [22] IEEE Std. 2030.5-2013, “IEEE adoption of smart energy profile 2.0 application protocol standard,” tech. rep., Institute of Electrical and Electronics Engineers, Inc., Nov 2013.
- [23] SunSpec Alliance, “Common smart inverter profile: Ieee 2030.5 implementation guide for smart inverters, version 2,” tech. rep., SunSpec Alliance, March 2018.
- [24] North American Electric Reliability Corporation, “Critical infrastructure protection standards.” <https://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx>. Accessed: 11-14-2018.
- [25] J. J. P. Cordeiro, J. Obert, “Recommendations for trust and encryption in der interoperability standards,” tech. rep., Sandia National Laboratories, December 2018.
- [26] C. Lai, N. Jacobs, S. Hossain-McKenzie, C. Carter, P. Cordeiro, I. Onunkwo, and J. Johnson, “Cyber security primer for der vendors, aggregators, and grid operators,” Tech. Rep. SAND2017-13113, Sandia National Laboratories, December 2017.
- [27] SunSpec Alliance, “Sunspec der cybersecurity workgroup.” <https://sunspec.org/sunspec-cybersecurity-workgroup>. Accessed: 11-14-2018.
- [28] C. C. D. Saleem, “Certification procedures for data and communication security of distributed energy resources,” tech. rep., NREL Technical Report, December 2018.
- [29] J. Johnson, “Roadmap for photovoltaic cyber security,” Tech. Rep. SAND2017-13262, Sandia National Laboratories, December 2017.
- [30] H. Jia, N. Guangyu, S. T. Lee, and P. Zhang, “Study on the impact of time delay to power system small signal stability,” in *Proceedings of the 2006 IEEE Mediterranean Electrotechnical Conference, MELECON*, pp. 1011–1014, May 2006.
- [31] Y. Xiaodan, J. Hongjie, and Z. Jinli, “A LMI based approach to power system stability analysis with time delay,” in *TENCON 2008 - 2008 IEEE Region 10 Conference*, pp. 1–6, Nov 2008.
- [32] J. Hongjie and Y. Xiaodan, “A simple method for power system stability analysis with multiple time delays,” in *Proceedings of the 2008 IEEE Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century*, pp. 1–7, July 2008.
- [33] K. Gajrani, A. Bhargava, K. G. Sharma, and R. Bansal, “Cyber security solution for wide area measurement systems in wind connected electric grid,” in *Proceedings of the 2013 IEEE Innovative Smart Grid Technologies - Asia (ISGT Asia)*, pp. 1–5, Nov 2013.
- [34] G. N. Ericsson, “Cyber security and power system communication - essential parts of a smart grid infrastructure,” *IEEE Transactions on Power Delivery*, vol. 25, pp. 1501–1507, July 2010.
- [35] M. Braendle, “Cyber security for power systems - a closer look at the drivers and how to best approach the new challenges,” in *2011 64th Annual Conference for Protective Relay Engineers*, pp. 322–327, April 2011.

- [36] R. H. Byrne, R. J. Concepcion, J. Neely, F. Wilches-Bernal, R. T. Elliott, O. Lavrova, and J. E. Quiroz, "Small signal stability of the western north american power grid with high penetrations of renewable generation," in *Photovoltaic Specialists Conference (PVSC), 2016 IEEE 43rd*, pp. 1784–1789, IEEE, 2016.
- [37] M. J. Reno, J. E. Quiroz, O. Lavrova, and R. H. Byrne, "Evaluation of communication requirements for voltage regulation control with advanced inverters," in *IEEE North American Power Symposium (NAPS)*, 2016.
- [38] J. W. Stahlhut, T. J. Browne, G. T. Heydt, and V. Vittal, "Latency viewed as a stochastic process and its impact on wide area power system control signals," *IEEE Transactions on Power Systems*, vol. 23, pp. 84–91, Feb 2008.
- [39] F. Zhang, Y. Sun, L. Cheng, X. Li, J. H. Chow, and W. Zhao, "Measurement and modeling of delays in wide-area closed-loop control systems," *IEEE Transactions on Power Systems*, vol. 30, pp. 2426–2433, Sep 2015.
- [40] Y. C. Zhang, V. Gevorgian, E. Ela, V. Singhvi, and P. Pourbeik, "Role of wind power in primary frequency response of an interconnection," in *International Workshop on Large-Scale Integration of Wind Power Into Power Systems as Well as on Transmission Networks for Offshore Wind Power Plants*, Oct 2013.
- [41] J. Ekanayake and N. Jenkins, "Comparison of the response of doubly fed and fixed-speed induction generator wind turbines to changes in network frequency," *IEEE Transactions on Energy Conversion*, vol. 19, pp. 800–802, Dec 2004.
- [42] G. Delille, B. Francois, and G. Malarange, "Dynamic frequency control support by energy storage to reduce the impact of wind and solar generation on isolated power system's inertia," *IEEE Transactions on Sustainable Energy*, vol. 3, pp. 931–939, Oct 2012.
- [43] P. Kundur, N. Balu, and M. Lauby, *Power System Stability and Control*. Discussion Paper Series, McGraw-Hill Education, 1994.
- [44] P. Kundur, *Power System Stability and Control*. New York, NY: McGraw-Hill, 1994.
- [45] V. Vittal and A. R. Bergen, *Power Systems Analysis*. Upper Saddle River, NJ: Prentice Hall, 2nd ed., 1999.
- [46] NERC, Princeton, NJ, *Balancing and Frequency Control: A Technical Document Prepared by the NERC Resources Subcommittee*, 2009.
- [47] C. Zhang, T. Liu, and D. J. Hill, "Distributed load-side frequency regulation for power systems," in *Power Systems Computation Conference (PSCC), 2016*, (Genoa, Italy), pp. 1–7, 2016.
- [48] C. Wu, S. Kar, and G. Hug, "Enhanced secondary frequency control via distributed peer-to-peer communication," in *European Control Conference, 2016*, (Aalborg, Denmark), pp. 1–6, 2016.
- [49] American National Standards Institute (ANSI), "Electric power systems and equipment - voltage ratings (60 Hz)," Tech. Rep. C84.1-2011, American National Standards Institute (ANSI), 2011.
- [50] J. Seuss, M. J. Reno, R. J. Broderick, and S. Grijalva, "Improving distribution network PV hosting capacity via smart inverter reactive power support," in *2015 IEEE Power Energy Society General Meeting*, pp. 1–5, July 2015.

- [51] M. J. Reno, J. E. Quiroz, O. Lavrova, and R. H. Byrne, "Evaluation of communication requirements for voltage regulation control with advanced inverters," in *2016 North American Power Symposium (NAPS)*, pp. 1–6, Sept 2016.
- [52] R. J. Broderick, J. E. Quiroz, M. J. Reno, A. Ellis, J. Smith, and R. Dugan, "Time series power flow analysis for distribution connected PV generation," Tech. Rep. SAND2013-0537, 2013, Sandia National Laboratories, Albuquerque, NM, 2013.
- [53] M. J. Reno and K. Coogan, "Grid integrated distributed PV (GridPV) version 2," Tech. Rep. SAND2014-20141, Sandia National Laboratories, Albuquerque, NM, 2014.
- [54] P. P. M. Collaborative, "PV_LIB toolbox for matlab." https://pvpmc.sandia.gov/applications/pv_lib-toolbox/, 2016.
- [55] International Electrotechnical Commission (IEC), "Object models for power converters in distributed energy resources (DER) systems," Tech. Rep. TR 61850-90-7, International Electrotechnical Commission (IEC), February 2013.
- [56] M. J. Reno, M. Lave, J. E. Quiroz, and R. J. Broderick, "PV ramp rate smoothing using energy storage to mitigate increased voltage regulator tapping," in *2016 IEEE 43rd Photovoltaic Specialists Conference (PVSC)*, June 2016.
- [57] J. E. Quiroz, M. J. Reno, O. Lavrova, and R. H. Byrne, "Communication requirements for hierarchical control of volt-var function for steady-state voltage," in *IEEE ISGT, Arlington, VA*, April 2017.
- [58] M. Reno, J. Quiroz, O. Lavrova, and R. Byrne, "Evaluation of communication requirements for voltage regulation control with advanced inverters," in *IEEE North American Power Symposium, Denver, CO*, September 2016.
- [59] M. J. Reno and K. Coogan, "Grid integrated distributed pv v) version 2," Tech. Rep. SAND2014-20141, Sandia National Laboratories, 2014.
- [60] Minimega, "A distributed vm management tool." minimega.org. Accessed: 11-14-2018.
- [61] SunSpec Alliance, "Specifications & information models." <https://sunspec.org/about-sunspec-specifications>. Accessed: 11-20-2018.
- [62] J. Stamp, C. Veitch, J. Henry, et al., "Microgrid cyber security reference architecture (v2)," Tech. Rep. SAND2015-9711, Sandia National Laboratories, November 2015.
- [63] J. Stamp, et al., "Design tradeoffs and cyber security for microgrids," in *Energy Exchange: Federal Sustainability for the Next Decade*, Aug 2016.
- [64] J. Johnson, et al., "Design and evaluation of a secure virtual power plant," Tech. Rep. SAND2017-10177, Sandia National Laboratories, 2017.
- [65] OpenSSH, "Openssh 7.9 release notes." <http://www.openssh.com/txt/release-7.9>. Accessed: 10-19-2018.
- [66] A. R. Chavez, J. R. Hamlet, and W. Stout, "Artificial diversity and defense security (addsec) final report," Tech. Rep. SAND2018-4545, Sandia National Laboratories, April 2018.

- [67] A. R. Chavez, W. M. S. Stout, and S. Peisert, “Techniques for the dynamic randomization of network attributes,” in *2015 International Carnahan Conference on Security Technology (ICCST)*, Taipei, 2015.
- [68] S. Hossain-McKenzie, C. Lai, A. Chavez, and E. Vugrin, “Performance-based cyber resilience metrics: An applied demonstration toward moving target defense,” in *44th IECON*, Washington DC, 2018.
- [69] SunSpec Alliance, “Sunspec system validation platform (svp).” <https://sunspec.org/sunspec-svp>. Accessed: 11-14-2018.
- [70] S. Jones, K. Gabert, and T. Tarman, “Evaluating emulation-based models of distributed computing systems,” Tech. Rep. SAND2017-10634, Sandia National Laboratories, August 2017.
- [71] US DOE OE, “Modern distribution grid: Decision guide volume iii,” tech. rep., US DOE, June 2017.
- [72] R. Concepcion, F. Wilches-Bernal, and R. Byrne, “Effects of communication latency and availability on synthetic inertia,” in *2017 IEEE ISGT*, Washington DC, April 2017.
- [73] F. Wilches-Bernal, R. Concepcion, J. Neely, R. Byrne, and A. Ellis, “Communication enabled – fast acting imbalance reserve (ce-fair),” *IEEE Trans. Power Systems*, vol. 33, pp. 201–213, January 2018.
- [74] F. Wilches-Bernal, et al., “Impact of communication latencies and availability on droop-implemented primary frequency regulation,” in *49th NAPS*, Morgantown, WV, 2017.
- [75] J. Johnson, R. Ablinger, R. Bruendlinger, B. Fox, and J. Flicker, “Design and evaluation of sunspec-compliant smart grid controller with an automated hardware-in-the-loop testbed,” *Technology and Economics of Smart Grids and Sustainable Energy*, vol. 2, December 2017.
- [76] NIST, “Common vulnerability scoring system calculator version 3.” <https://nvd.nist.gov/vuln-metrics/cvss/v3-calculator>. Accessed: 11-14-2018.
- [77] R. Guttromson, M. Donnelly, and D. Trudnowski, “Widespread loss of communications in grid systems,” tech. rep., Sandia National Laboratories, 2018.
- [78] R. Byrne, D. Trudnowski, J. Neely, R. Elliott, D. Schoenwald, and M. Donnelly, “Optimal locations for energy storage damping systems in the western north american interconnect,” in *IEEE PES General Meeting, Conference & Exposition, 27-31 July 2014*, pp. 1–5.
- [79] North American Electric Reliability Corporation, “Generating Unit Operations During a Complete Loss of Communications, Draft NERC Reliability Guideline,” tech. rep., 2018.