SANDIA REPORT
SAND2017-10177
Unlimited Release
Printed September 2017

Design and Implementation of a Secure Virtual Power Plant

Jay Johnson, Jack Flicker, Anya Castillo, Clifford Hansen, Mohamed El-Khatib, David Schoenwald, Mark A. Smith, Russell Graves, Jorden Henry, Trevor Hutchins, Jason Stamp, Derek Hart, Adrian Chavez, Mitch Burnett, Jose Tabarez, Casey Glatter, Boqi Xie, A. P. Meliopoulos, Phuc Huynh, Hao Zhu, Katherine Davis

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and 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-NA-0003525.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: http://www.osti.gov/scitech

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: http://www.ntis.gov/search

SAND2017-10177
Unlimited Release
Printed September 2017

Design and Implementation of a Secure Virtual Power Plant

Jay Johnson, Clifford Hansen
Renewable and Distributed Systems Integration
Sandia National Laboratories

Jack Flicker
Advanced Microelectronics and Radiation Effects
Sandia National Laboratories

Anya Castillo
Mission Analytics Solutions
Sandia National Laboratories

Mohamed El-Khatib, David Schoenwald
Electric Power Systems Research
Sandia National Laboratories

Mark A. Smith, John Eddy
Systems Readiness and Sustainment Technology
Sandia National Laboratories

Russell Graves
Cyber Physical Security
Sandia National Laboratories

Jorden Henry, Trevor Hutchins, Casey Glatter
Critical Infrastructure Systems
Sandia National Laboratories

Jason Stamp
Special Cyber Initiatives
Sandia National Laboratories

Derek Hart
Emulytics Initiatives
Sandia National Laboratories

Adrian Chavez
Networked System Surveillance & Assurance
Sandia National Laboratories

Mitch Burnett
Brigham Young University
Electrical Engineering

Jose Tabarez
New Mexico State University
Electrical and Computer Engineering

Boqi Xie, A. P. Meliopoulos
Georgia Institute of Technology
School of Electrical and Computer Engineering

Phuc Huynh, Hao Zhu, Katherine Davis
University of Illinois Urbana-Champaign
Electrical and Computer Engineering

Abstract

For three years, Sandia National Laboratories, Georgia Institute of Technology, and University of Illinois at Urbana-Champaign investigated a smart grid vision in which renewable-centric Virtual Power Plants (VPPs) provided ancillary services with interoperable distributed energy resources (DER). This team researched, designed, built, and evaluated real-time VPP designs incorporating DER forecasting, stochastic optimization, controls, and cyber security to construct a system capable of delivering reliable ancillary services, which have been traditionally provided by large power plants or other dedicated equipment. VPPs have become possible through an evolving landscape of state and national interconnection standards, which now require DER to include grid-support functionality and communications capabilities. This makes it possible for third party aggregators to provide a range of critical grid services such as voltage regulation, frequency regulation, and contingency reserves to grid operators. This paradigm (a) enables renewable energy, demand response, and energy storage to participate in grid operations and provide grid services, (b) improves grid reliability by providing additional operating reserves for utilities, independent system operators (ISOs), and regional transmission organization (RTOs), and (c) removes renewable energy high-penetration barriers by providing services with photovoltaics and wind resources that traditionally were the jobs of thermal generators. Therefore, it is believed VPP deployment will have far-reaching positive consequences for grid operations and may provide a robust pathway to high penetrations of renewables on US power systems. In this report, we design VPPs to provide a range of grid-support services and demonstrate one VPP which simultaneously provides bulk-system energy and ancillary reserves.

ACKNOWLEDGMENTS

The Virtual Power Plant project was supported by the Energy and Climate Laboratory Directed Research and Development Program at Sandia National Laboratories.

CONTENTS

FIGURES

TABLES

NOMENCLATURE

AGC	Automatic Generation Control
AGF	Advanced Grid Function
AMI	Advanced Metering Infrastructure
ANSI	American National Standards Institute
ARIMA	Autoregressive Integrated Moving Average
ARMA	Autoregressive Moving Average
ARP	Address Resolution Protocol
ARPA-E	Advanced Research Projects Agency-Energy
BEMS	Building Energy Management System
BESS	Battery Energy Storage System
CAISO	California Independent System Operator
CEC	California Energy Commission
CEI	Comitato Elettrotecnico Italiano
CLI	Command Line Interface
CONUS	Continental U.S.
CPUC	California Public Utilities Commission
CVaR	Conditional Value-at-Risk
DER	Distributed Energy Resource
DERMS	Distributed Energy Resource Management System
DERP	Distributed Energy Resource Provider
DETL	Distributed Energy Technologies Laboratory
DOE	Department of Energy
DOY	Day of Year
DR	Demand Response
DS-DQSE	Distribution System Distributed Quasi-Dynamic State Estimator
EBP	Estimation-Based Protection
ESS	Energy Storage System
ERCOT	Electric Reliability Council of Texas
EU	European Union
FERC	Federal Energy Regulatory Commission
FW	Frequency-Watt
GMLC	Grid Modernization Laboratory Consortium
HIL	Hardware-in-the-Loop
HMI	Human-Machine Interface
HRRR	High Resolution Rapid Refresh
IDS	Intrusion Detection System
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IOU	Investor-Owned Utility
IP	Internet Protocol
ISO	Independent System Operator
LDRD	Laboratory Directed Research and Development
LMP	Locational Marginal Price
LTC	Load Tap Changing Transformer

MAC	Media Access Control
MdS	Mesa del Sol
MOL	Minimum Operating Level
MOS	Model Output Statistics
NAM	North American Mesoscale
NOAA	National Oceanic and Atmospheric Administration
NODES	Network Optimized Distributed Energy Systems
NREC	North American Electric Reliability Corporation
NREL	National Renewable Energy Laboratory
NYISO	New York Independent System Operator
OLE	Object Linking and Embedding
OPC	OLE for Process Control
OPF	Optimal Power Flow
PF	Power Factor
PID	Proportional-Integral-Derivative
PJM	PJM Interconnection LLC
PMU	Phasor Measurement Unit
PNM	Public Service Company of New Mexico
P/Q	Active and Reactive Power
PV	Photovoltaic
PW	PowerWorld
pvlib	PV Library
PVRPM	Photovoltaic Array Performance Model
QSTS	Quasi Static Time Series
RE	Renewable Energy
RMT	Reserve Magnitude Target
RMVT	Reserve Magnitude Variability Tolerance
RTO	Regional Transmission Organization
RTU	Remote Terminal Unit
SCADA	Supervisory Control and Data Acquisition
SIRFN	Smart Grid International Research Facility Network
SIWG	Smart Inverter Working Group
SMB	Server Message Block
SNL	Sandia National Laboratories
SOC	State-of-Charge
SPF	Specified Power Factor
TCP	Transmission Control Protocol
UNM	University of New Mexico
VaR	Value-at-Risk
VDE	Verband der Elektrotechnik, Elektronik und Informationstechnik
VIU	Vertically Integrated Utility
VPP	Virtual Power Plant
VV	Volt-Var

1 INTRODUCTION

Virtual Power Plants convert variable renewable energy systems into monolithic dispatchable resources which provide electric utilities/ISOs/RTOs with mechanisms to perform frequency regulation and respond to grid disturbances more quickly and efficiently than large fossil power plants. This will increase renewable energy penetrations above 100% of peak load by eliminating the need for dedicated ancillary services because VPPs can dispatch identical services in accordance with utility and Independent System Operator (ISO)/Regional Transmission Organization (RTO) requirements. This capability is established with a resilient real-time controller interconnected to thousands of distributed energy resources (DERs). The VPP aggregator optimizes the individual behavior of distributed power electronics-based converters, small generators, and demand response units to dispatch frequency control reserves, thereby improving:

1. *Grid operations* by dispatching solar photovoltaic (PV) power at appropriate times to minimize the effects of displacing traditional generation—achieved with autonomous and commanded advanced grid functions which mitigate bulk system frequency deviations.
2. *Grid reliability* by providing additional operating reserves to utilities/ISOs/RTOs and distributing the controls to many DERs so single points of failure (e.g., large generators) no longer significantly disrupt the grid.

Real-time VPP optimization transforms variable renewable energy systems, demand response units, and other DER into dispatchable aggregations that provide system operators with mechanisms to provide grid services more quickly than large thermal plants. Due to the sheer number of DERs and their small sizes, it is not practical for bulk system operators to optimize and control individual DERs. In that regard, a VPP represents a framework for cohesive optimization and control of large numbers of small DERs which are then seen as a single entity by grid operators. VPPs provide grid support services using robust communications, robust control, and efficient optimization of large and diverse sets of DER; and ultimately, this functionality may eliminate the need for dedicated ancillary services entirely.

The national and international trend of increasing renewable energy penetrations is a worst-case scenario for bulk system reliability as grid inertia and governor control are displaced and frequency deviations from RE variability are increasingly common[1]. Therefore, instituting frequency response reserves with DERs in accordance with utility, ISO/RTO, and NERC requirements are critical for future grid resiliency. In general, VPPs could be composed of grid operator-owned assets or privately-owned DER that are controlled under a legal agreement. In the case of operating in regions with vertically integrated utilities (VIUs), the VPP would be scheduled and dispatched as part of an operating plan in which the VIU manages its generation, transmission, and distribution services centrally[2]. In market-based jurisdictions, the VPP would submit offers into the day-ahead and real-time markets.

---

[1] C. Martinez, S. Xue, and M. Martinez, "Review of the recent frequency performance of the Eastern, Western and ERCOT Interconnections," Lawrence Berkeley National Laboratory, Tech. Rep., December 2010.

[2] N.P. Padhy. Unit commitment – a bibliographical survey, IEEE Transaction on Power Systems 19(2):1196–1205, 2004.

Researchers have been investigating the advantages of establishing aggregations of distributed generators for some time. Algorithms have been developed for optimizing VPP aggregations[3,4], VPP control mechanisms[5], communication methods[6], market mechanisms[7], scheduling/ forecasting power levels for reserve markets[8], and techniques of maintaining grid stability[9,10,11]. There have also been demonstrations in the deregulated markets of Europe: Kassel University incorporated DERs into a large, geographically dispersed VPP[12]; the European Union (EU)-funded FENIX program investigated market scenarios[13,14,15]; and an operational VPP in Denmark is using a Distributed Energy Market[16]. Additionally, two German utilities began large (150 MW) VPP demonstrations consisting of various DERs using Siemens hardware[17,18].

In the United States, the Investor-Owned Utilities (IOUs) in California have stated within the California Public Utilities Commission (CPUC)/California Energy Commission (CEC) Smart Inverter Working Group (SIWG) meetings that they are not interested in controlling residential-scale DER equipment. Instead, they prefer to communicate to aggregators to update the Electric Rule 21 or IEEE 1547 advanced grid-support functions. In practice, many companies like Enphase, SolarCity, Vivint Solar, and SunPower already have communication networks to their

---

[3] N. Capodieci, G. Cabri, "Managing Deregulated Energy Markets: An Adaptive and Autonomous Multi-agent System Application," 2013 IEEE Systems, Man, and Cybernetics (SMC) Conference, pp. 758 – 763.
[4] L.L. Pfitscher, D.P. Bernardon, L.N. Canha, V.F. Montagner, L. Comasseto, M.S. Ramos, "Studies on parallelism of feeders for automatic reconfiguration of distribution networks," 47th International Universities Power Engineering Conference (UPEC), pp.1-5, 2012.
[5] S. Lukovic, I. Kaitovic, M. Marcello, U. Bondi, "Functional requirements of embedded systems for monitoring and control structure of Virtual Power Plants," Environmental, Energy, and Structural Monitoring Systems, pp. 19–26, 2009.
[6] S, Sucic, A. Martinic, D. Francesconi, "Utilizing SOA-ready devices for virtual power plant control in semantic-enabled Smart Grid Analyzing IEC 61850 and OPC UA integration methodology," Smart Grid Communications (SmartGridComm), 17-20 Oct. 2011.
[7] I. Praca, C. Ramos, Z. Vale, M. Cordeiro, "MASCEM: a multiagent system that simulates competitive electricity markets," IEEE Intelligent Systems, vol.18, no.6, pp.54,60, Nov-Dec 2003.
[8] M. Vasirani, R. Kota, R.L.G. Cavalcante, S. Ossowski,N. R. Jennings, "An Agent-Based Approach to Virtual Power Plants of Wind Power Generators and Electric Vehicles," IEEE Transactions on Smart Grid, vol. 4, no. 3, pp. 1314-1322, Sept. 2013.
[9] J. F. Baalbergen, V. Karapanos, M. Gibescu, L. van der Sluis, "Emergency voltage control with decentralized generation," IEEE PES Innovative Smart Grid Technologies (ISGT Europe), pp. 1–10, 2011.
[10] H. Morais, T. Sousa, P. Faria, Z. Vale, "Reactive power management strategies in future smart grids," Power and Energy Society General Meeting (PES), pp. 1–5, 2013.
[11] D. Pudjianto, C. Ramsay, G. Strbac, "Virtual power plant and system integration of distributed energy resources," Renewable Power Generation, vol.1, no.1, pp.10-16, March 2007.
[12] P. Fairley, Real Electricity Flows from Virtual Power Plants, MIT Technology Review, 16 April, 2012.
[13] A.v.d. Welle, C. Kolokathis J. Jansen, C. Madina, A. Diaz, FENIX deliverable D3.3 Report: Financial And Socio-Economic Impacts Of Embracing The Fenix Concept, Final Report, 30 September 2009.
[14] L. Nikonowicz, J. Milewski, "Virtual power plants—general review: structure, application and optimization," Journal of Power Technologies, vol. 92, no. 3, pp. 135–149, 2012.
[15] K. Kok, "Short-term economics of virtual power plants," 20th International Conference on Electricity Distribution, 2009.
[16] L. Nikonowicz, J. Milewski, "Virtual power plants—general review: structure, application and optimization," Journal of Power Technologies, vol. 92, no. 3, pp. 135–149, 2012.
[17] B. Müller, "The Internet of Things: Virtual Power Plants", URL: http://www.siemens.com/, Accessed: April 2015.
[18] Siemens and RWE, "Siemens offers cloud-based Web service for virtual power plants," Joint Press Release, 13 Feb 2014.

end devices for monitoring. In the future, the IOUs and other utilities will send their needs to these entities to update their fleet of devices. While the CAISO and other ISO/RTO markets do not support DER aggregation bids, there are a number of indications that this may be changing. California Electric Rule 24 permits demand response providers (such as EnerNOC or EnerNex) to solicit customers to participate in their demand response programs and bid into the CAISO wholesale electricity market. In November 2016, the Federal Energy Regulatory Commission (FERC) published a Notice of Proposed intended "to remove barriers to the participation of electric storage resources and distributed energy resource (DER) aggregations in the capacity, energy, and ancillary service markets"[19]. There is increasing discussion among industry leaders on this topic from the regulatory standpoint[20].

Despite significant research in this area and multiple VPP field demonstrations, there remains a number of technical research gaps, including:

1. Prior VPP field demonstrations were not fully automated or did not respond quick enough to provide synthetic inertia or frequency droop/primary frequency reserves. (Most demonstrations provide following, secondary, or tertiary control reserves.) No VPPs simultaneously provide multiple reserve services with a hierarchical, hybrid, or multi-level optimization and control architectures.
2. Previous VPPs only use large (generally commercial-scale) DER assets, leaving smaller residential-scale solar, wind, and battery systems out of the aggregation. No VPPs incorporated residential PV systems or securely communicated to a large number of smaller devices over public Internet channels using forthcoming DER communication standards (SunSpec Modbus, IEC 61850-90-7, IEEE 1815, or IEEE 2030.5).
3. No VPP demonstrations dynamically adjusted the DER dispatch when resources are unavailable due to communication failures or cyber security compromise.

This project investigated solutions to these challenges and the VPP design presented herein could act as the basis for distribution-level DERs aggregation, optimization and control for multiple ancillary service markets. This project tackles a major component of U.S. grid modernization by establishing a robust methodology for integrating large penetrations of renewable energy on the current U.S. electric grid without making large investments in new transmission lines or reserve generators. This improves grid resilience by making critical infrastructure more robust to attacks, safeguards U.S. energy infrastructure with cyber-resilient communications to distributed assets, and includes the ability to respond quickly to targeted or natural disruptions to the grid. VPPs also reduce the U.S. dependence on oil and gas, combats climate change by building a renewable energy and DER control system to provide grid services, and secures the prosperity of the country by transforming the U.S. to more efficient and cleaner energy technologies. This technology will be an invaluable asset for the U.S. smart electricity grid of the future.

---

[19] FERC, Electric Storage Participation in Markets Operated by Regional Transmission Organizations and Independent System Operators, 18 CFR Part 35, Docket Nos. RM 16-23-000; AD16-20-000, Nov 2016.

[20] K. Shallenberger, DER aggregation: Sector experts identify emerging trends in a nascent market, Utility Dive July 24, 2017.

2 BACKGROUND INFORMATION

The Sandia virtual power plant is designed to provide grid services through the use of interoperable, controllable DER equipment. There is a wide range of grid services that are required for the safe, secure, reliable operation of the power system and many of these services can be provided with appropriately-designed DER control mechanisms which communicate modes of operation via commanded and autonomous DER grid-support functions.

2.1 Grid Services

Conventional power is generated by large power plants and transported through a high-voltage transmission system to a low voltage distribution system to serve end loads but this paradigm is shifting with the increased of distributed generation. Grid operators regulate this power system to provide dependable power to customers using automated and manual ancillary services which ensure the power is reliably transported from generation to end loads. The name and capability of the grid services depend on how they are defined for a given region (e.g., ISO/RTO regions use different names for similar services[21]) which often causes confusion when discussing grid services, but they fall into a few general categories of balancing, frequency support, and voltage support[22]. FERC described essential reliability services as including load and resource balance, voltage support, and frequency support (which are then subdivided into many more categories)[23]—but there are also dozens of proposed new services. A recent DOE Grid Modernization Laboratory Consortium (GMLC) effort enumerated such grid services[24]; a portion of that list is provided below with the headings organized by Group, Category, Subcategory, and Grid Service Name:

1. Scheduling
    1.1. Production Coordination
        1.1.1. Production Optimization
            1.1.1.1. Unit Commitment
            1.1.1.2. Energy Scheduling
        1.1.2. Wheeling
    1.2. Dispatch/Balance
        1.2.1. Generation-Load Matching/Meet Energy Imbalance by Generation Change
            1.2.1.1. Redispatch
            1.2.1.2. Wind and Solar Curtailment
        1.2.2. Demand/Modification

---

[21] J.F. Ellison, L.S. Tesfatsion, V.W. Loose, R.H. Bryne, "Project Report: A Survey of Operation Reserve Markets in U.S. ISO/RTO-managed Electric Energy Regions," Sandia National Laboratories Technical Report, SAND2012-1000, Sept 2012.

[22] E. Hirst, B. Kirby, "Electric-Power Ancillary Services," Oak Ridge National Laboratory Technical Report ORNL/CON-426, Feb 1996.

[23] NERC Essential Reliability Services Task Force, "Concept Paper On ERS that Characterize Bulk Power System Reliability," Draft, September 2014.

[24] N. Samaan, S. Widergren, R. Melton, A. Somani, R. Pratt, J. Taft, "Grid Services Master List," PNNL-26599, June 2017.

1.2.2.1. Peak Demand Limiting
1.2.2.2. Energy Shifting (Net Load Adjustments)
1.3. Energy Settlements
1.3.1. Interchange
1.3.2. Loss Compensation
1.3.3. Uninstructed/Instructed Deviation of Generation
2. Regulation and Stabilization
2.1. Frequency Support
2.1.1. Inertial Support and Damping
2.1.1.1. Primary System Inertia
2.1.1.2. Inertia Augmentation
2.1.1.3. Inertia Modulation
2.1.1.4. Damping
2.1.2. Primary
2.1.2.1. Droop/Governor Response
2.1.2.2. Load Modification
2.1.3. Secondary
2.1.3.1. Regulation UP
2.1.3.2. Regulation Down
2.1.3.3. Load Modification
2.1.4. Tertiary
2.1.4.1. Replacement
2.1.4.2. Market Responses
2.1.4.3. Load Modification
2.2. Voltage Support
2.2.1. Voltage Adjustment
2.2.1.1. Bulk System Voltage Regulation
2.2.1.2. Distribution-Level Voltage Regulation
2.2.1.3. Edge-Device Voltage Regulation Support
2.2.2. Reactive Power Control
2.2.2.1. Bulk System Var Regulation
2.2.2.2. Distribution-Level Var Regulation
2.2.2.3. Edge Device Var Regulation Support
2.3. Ramping
2.3.1. Ramp Response
2.3.1.1. Generation Dispatch
2.3.1.2. Storage Dispatch
2.3.1.3. Load Modification

The details of these services are explained in detail in the GMLC report, but at a high level, there are multiple markets or VIU operations which commit generators to balance energy or provide reserves for specific periods of time. This is typically done a single day ahead of time, although there are intra-day adjustments as well. The generators operating reserves are either non-event reserves or contingency reserves, further subcategorized based on their speed and magnitude[25].

[25] E. Ela, M. Milligan, B. Kirby, Operating Reserves and Variable Generation, NREL/TP-550-51928, August 2011.

As shown in Figure 1, within non-event reserves there are regulating reserves which absorbs fast, smaller power imbalances and following reserves which absorbs the slower, larger power imbalances. Depending on the region, the Automatic Generation Control (AGC) signal for regulation is issued every 4-8 seconds. Load following reserve signals are issued at much slower rates – typically one hour.

The primary concern of the transmission system operator is to maintain bulk system stability by matching generation to load. Small deviations in this balance are common and compensated for by the inertial energy of the online generators. When the load is greater than generation, the bulk system frequency decreases; and when there is more power production than load, the frequency increases. If the frequency deviates from the nominal frequency by a large enough magnitude, automated protection mechanisms will initiate. As frequency drops, first load shedding (disconnecting consumer regions) occurs in multiple tiers and eventually power plants are disconnected to ensure they are not damaged. Both protection mechanisms cause customer blackouts. Similarly, transmission and distribution systems are designed to disconnect from voltage deviations caused by faults, or excessive active or reactive power flows to protect equipment.

When there is a fault on the system (e.g., sudden generator, load center, or transmission line loss) contingency reserves prevent bulk system collapse. Initially, the inertial response of the operating generators will prevent the frequency from changing instantaneously. As soon as the change in frequency is measured, primary (~1-30 seconds from the fault), secondary (~30 sec-10 min from the fault), and tertiary (~10-30+ min from the fault) reserves are employed to arrest the frequency change and restore the system to nominal grid frequency. These reserves are initiated in a cascading fashion to ensure the frequency is restored reliably.

Voltage regulation on the transmission and distribution systems is necessary for the proper operation of the power equipment and end loads. Many devices are used to maintain the voltage such as load tap changing transformers (LTCs), voltage regulators, and capacitor banks.

It was not the intention of this work to redefine these services, define new services, or investigate the markets for grid services; but rather design the technology for DER aggregations to participate as a single entity in the services. The ancillary services that this project targeted are shown in Figure 1. Ultimately, the team investigated two scenarios:

1. Co-optimizing the participation of a VPP in an energy market and a tertiary reserve market.
2. Providing voltage regulation on distribution circuits with a VPP of DER located on the same feeder.

but only programmed the former VPP. The simulation results of the voltage regulation VPP are presented in a second report.[26]

[26] A.P. Meliopoulos, G. Cokkinides, B. Xie, C. Zhong, J. Johnson, "Full State Feedback Control for Virtual Power Plants," Sandia Technical Report, September 2017.

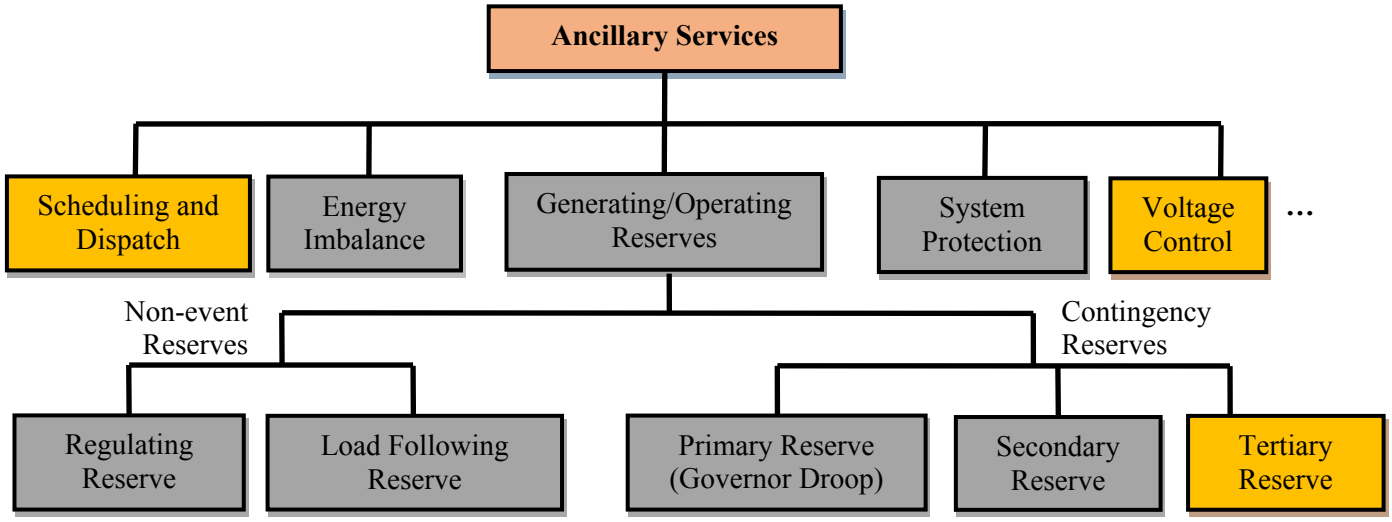

Figure 1. Target ancillary services for the VPP project were scheduling (bidding into energy markets), tertiary contingency reserve, and distribution voltage regulation.

2.2 DER Grid-Support and Interoperability Capabilities

In the last decade, DER interconnection and interoperability codes and standards in the US and around the world have been rapidly changing to offer grid operators additional resources to provide ancillary services[27]. The impetus to change the DER requirements arose from increasing penetrations of grid-connected, distributed, inverter-based renewable energy systems which reduced system inertia and replaced it with variable, non-dispatchable generators. Since the mid-2000s, the U.S. has relied on the IEEE 1547 series[28] of standards to harmonize the interconnection requirements across the country. However, very high penetrations of photovoltaic installations—driven by state renewable portfolio standards, favorable economics, and consumer preference—led regulators to conclude there is an eminent need to have DER grid-support capabilities for reliable power system operations. As a stop-gap measure, IEEE 1547a[29] was drafted in 2014, which permitted—but did not require—DER to actively regulate voltage and frequency with agreement of the Area Electric Power System operator.

Taking the regulation process further, the Hawaiian and Californian state Public Utility Commissions changed their interconnection standards to require grid-support functions for inverter-based DER. Currently, HI Rule 14H[30] and CA Rule 21[31] have requirements for wide

---

[27] D. Rosewater, J. Johnson, M. Verga, R. Lazzari, C. Messner, R. Bründlinger, K. Johannes, J. Hashimoto, K. Otani, International development of energy storage interoperability test protocols for renewable energy integration, EU PVSEC, Hamburg, Germany, 14-18 Sept, 2015.
[28] IEEE Standard 1547-2003, Standard for Interconnecting Distributed Resources with Electric Power Systems, 2003.
[29] IEEE Standard 1547a-2014, Standard for Interconnecting Distributed Resources with Electric Power Systems: Amendment 1, 2014.
[30] Hawaiian Electric Company, Inc., Rule No. 14, Service Connections and Facilities on Customer's Premises, Section H: Interconnection of distributed generating facilities with the company's distribution system, effective

frequency and voltage ride-throughs, and active and reactive power functions. IEEE 1547 is currently undergoing a full revision to add these functions to the interconnection standard and re-harmonize the requirements across the nation, but this standard is not expected until late 2017 or 2018.

As a result of the new US interconnection requirements, and similar requirements in Europe, DER vendors have added many grid interactive functions to their equipment. Many of these capabilities are defined in the International Electrotechnical Commission (IEC) Technical Report (TR) 61850-90-7[32] and encoded in the SunSpec Alliance Modbus information models[33]. Many of these same functions are defined in the German VDE[34], Italian technical rule CEI 0-21[35] for low voltage and CEI 0-16[36] for medium and high voltage, and others[37]. In the vast majority of cases, modern commercial DER equipment can configure a range of DER advanced grid functions through interoperable communication tools. Examples of these capabilities are provided in prior work by the Smart Grid International Research Facility Network (SIRFN)[38] and Sandia National Laboratories[39].

In this work, the team choose to focus on the challenges of providing (a) energy and tertiary contingency reserves by co-optimizing the DER resources and (b) providing distribution voltage regulation with a VPP isolated on a single feeder. The active power headroom in renewable energy DER devices was created through the curtailment function described in IEC 61850-90-7, SunSpec Models, and the IEEE 1547 full revision. By curtailing the PV or other renewable energy resources through power electronics commands, the active power can be increased and decreased to respond to the needs of the VPP, in addition to any production generated by storage technologies. One distinct feature of a VPP serving balancing and frequency support services is that the VPP does not need to have a single point of connection to the grid, but instead is composed of an aggregation of different DER sources that connect to the grid at geographically diverse points of common coupling. Therefore, the VPP could be used to aggregate DGs, energy

---

October 21, 2015.

[31] Pacific Gas and Electric Company, Electric Rule No. 21, Generating Facility Interconnections, Filed with the CPUC on 20 Jan, 2015.

[32] IEC Technical Report 61850-90-7, "Communication networks and systems for power utility automation–Part 90-7: Object models for power converters in distributed energy resources (DER) systems," Edition 1.0, Feb 2013.

[33] SunSpec Alliance Interoperability Specification, SunSpec Inverter Models, Document #12020 Version 1.5, released 4-14-2015.

[34] VDE Application Guide VDE-AR-N 4105: Generators in the low voltage distribution network. Application guide for generating plants' connection to and parallel operation with the low-voltage network, 1/08/2010.

[35] CEI Reference Technical Rules for the Connection of Active and Passive Consumers to the HV and MV Electrical Networks of Distribution Company, CEI 0-16 and 0-16, 2014.

[36] CEI Reference Technical Rules for the Connection of Active and Passive Users to the LV Electrical Utilities, CEI Reference 0-21, December 2013.

[37] J. Johnson, S. Gonzalez, A. Ellis, "Sandia DER Interoperability Test Protocols; Relationship to Grid Codes and Standards," IEEE International Conference on Standards for Smart Grid Ecosystems, Bangalore, India. 6-7 Mar, 2014.

[38] J. Johnson, R. Bründlinger, C. Urrego, R. Alonso, "Collaborative Development Of Automated Advanced Interoperability Certification Test Protocols For PV Smart Grid Integration," EU PVSEC, Amsterdam, Netherlands, 22-26 Sept, 2014.

[39] J. Hernandez-Alvidrez, J. Johnson, "Parametric PV Grid-Support Function Characterization for Simulation Environments," IEEE PVSC, Washington, DC, 25-30 June, 2017.

storage systems, entire microgrids, demand response units, electric vehicles and even entire distribution stations across a state or even the entire interconnection. The voltage regulation functionality is created with the fixed power factor function which is defined in IEC 61850-90-7, SunSpec Models, and the IEEE 1547 full revision.

2.3 VPP Economics and Regulatory Challenges

While the technical aspects of VPPs are the focus of this work, it is worth a short discussion of the economic, legal, and regulatory challenges. While it is likely that capital costs (and regulation barriers) associated with creating VPPs will temporarily prevent them from being formed, DER interoperability and interconnection requirements will slowly reduce this upfront cost as more devices come 'out-of-the-box' with the ability to provide grid services. Already many DER installers, vendors and aggregators have the ability control large fleets of DER equipment. Once the VPPs are created, they must demonstrate they are cost competitive with traditional thermal generators when bidding into ancillary markets or operating under the direct or indirect supervision of utilities/ISOs/RTOs. Some of the prior work in the EU has looked at techno-economic analysis of VPPs, but further studies which include fuel costs of the traditional generators, opportunity cost of curtailing the PV output to provide headroom, operations and maintenance costs of the equipment, thermal losses incurred in network transmission of power, and other VPP operations costs must be completed. Ultimately, if VPPs are allowed to participate in ISO/RTO markets, it clear that many aggregators will bid into these markets. Then it is up to the markets to decide who is responsible to provide power.

There are also multiple legal and regulatory hurdles preventing virtual power plants in the U.S. First, while demand response and energy storage is allowed to participate in energy, capacity, and ancillary service markets[40] (see FERC Order 719[41]), ISO/RTO regulations historically do not allow participation by aggregations of the DER resources[42] but this is changing. CAISO is beginning to recognize aggregations of small DER are useful assets and should be included as market participants. In fact, CAISO currently allows Non-Generator Resources (like energy storage systems) to bid in wholesale markets and they are considering adding a Distributed Energy Resource Provider (DERP) asset class to markets in the future.[43] This project did not investigate regulations preventing VPP participation in ISO markets, nor determine regulatory obstacles to broad VPP deployment in American reserve markets. Yet, it is only through sound research and field demonstrations—such as those started in this project—that regulators will be incentivized to change these requirements.

---

[40] J.B. Eisen, "Distributed Energy Resources, 'Virtual Power Plants', and the Smart Grid," Environmental & Energy Law & Policy Journal, pp. 191-213, 2012.
[41] Federal Energy Regulatory Commission 18 CFR Part 35 Wholesale Competition in Regions with Organized Electric Markets, October 17, 2008.
[42] A. Zurburg, "Unlocking Customer Value: The Virtual Power Plant," worldPower, 2010.
[43] CAISO, "Energy Storage and Aggregated Distributed Energy Resource Education Forum," April 16 & 23, 2015, accessed 14 September 2017, http://www.caiso.com/Documents/Presentation-EnergyStorageandAggregatedDistributedEnergyResource-EducationalForum.pdf

Further, energy and financial accounting for each of the DER users is also legally murky. For example, how would a VPP prove the reserve metrics were met when there is no single meter showing the net effect of the disparate devices? While control of DER would be worked out with a contract between the VPP and the DER owners, there are a number of difficulties regarding balanced payment for access to the DERs. Proper metering of each of the assets is necessary to compensate owners correctly, but in the case of PV and other renewable sources when the inverter is curtailed to a percentage of the nameplate rating, the waived power is unknown and therefore the lost opportunity cost of not producing that power is unknown. Similarly, when demand response is taken, it is unknown how much power is saved over the duration of the DR event; therefore, it is difficult to report a firm number as load reduction back to the ISO/RTO and compensate the owner. In the case of demand response, there may also be contractual agreements and/or penalties if the owner disables the controllability of the resource. Lastly, there are questions about the fairness of the VPP controls: for example, if Battery 1 has a better communication link to the VPP and is used more heavily than Battery 2, how is the rate structure designed so that the DER are compensated appropriately? For storage systems, preparing the DER for a reserve event will likely to require arbitrage, which is illegal in many regions. The regulatory and accounting questions must continue to be enumerated so roadblocks are identified; from there, detailed options for navigating through these obstacles can be established to demonstrate VPP viability and assist with industry commercialization.

3 VPP DESIGNS

The end goal is a VPP software system which provides utilities, balancing authorities, and third parties a portfolio of services to adjust the operation of a DER fleet. Virtual power plants can be used for a range of grid services. These services could be changed day to day or even hourly, and it is possible to provide multiple grid services simultaneously (like in the case of microgrid applications). These services can be selected depending on the needs of the grid operator at any given time. Sandia has been creating DER control and optimization mechanisms for a range of services for many years, including voltage regulation[44,45,46] , small signal stability[47] and wide-area damping for transmission systems[48], frequency regulation[49], and frequency reserves (e.g., synthetic inertia[50], frequency-droop (primary reserve)[51], secondary[52] or tertiary contingency reserves[53]). Each of these services will require different market/utility/ISO/RTO interactions, optimization, and control, as well as studies of DER interoperability/dispatchability, communication limitations, protocol assessment, market profitability/influences, control stability, forecasting, cyber-resilience, and effective utilization of different DER resources. This knowledge will ultimately be integrated into a monolithic software program that will provide each of these ancillary services. At this stage of development, the team has built a fundamental framework for a VPP that contributes to energy and reserve markets (Section 3.1), which can then be extended to simultaneously provide a variety of grid services based on operational needs. The following sections include VPP designs for different grid services; note that these designs include many of the same core VPP components which will be discussed in further detail in Section 4.

It should also be noted that providing fleets of DER with autonomous functionality in some ways represents VPP operations, but this does not include transactive energy concepts or account for

---

[44] J. Seuss, M.J. Reno, R.J. Broderick, R.G. Harley, "Evaluation of reactive power control capabilities of residential PV in an unbalanced distribution feeder," 2014 IEEE PVSC, pp. 2094-2099, 8-13 June 2014.
[45] 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.
[46] J. Quiroz, M. Reno, O. Lavrova, R. Byrne, "Communication Requirements for Hierarchical Control of Volt-VAr Function for Steady-State Voltage," IEEE ISGT 2017, Arlington, VA, April 23-26, 2017.
[47] 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.
[48] J. Neely, J. Johnson, R. Bryne, R. T. Elliott, Structured optimization for parameter selection of frequency-watt grid support functions for wide-area damping, DER Journal, vol. 11, no. 1, pp. 69-94, 2015.
[49] J. Johnson, J. Neely, J. Delhotal, M. Lave, "Photovoltaic Frequency-Watt Curve Design for Frequency Regulation and Fast Contingency Reserves," IEEE Journal of Photovoltaics, vol. 6, no. 6, pp. 1611-1618, Nov. 2016.
[50] R. Concepcion, F. Wilches-Bernal, R. Byrne, "Effects of Communication Latency and Availability on Synthetic Inertia," IEEE ISGT 2017, Arlington, VA, April 23-26, 2017.
[51] F. Wilches-Bernal, R. Concepcion and R. H. Byrne, "Impact of Communication Latencies and Availability on Droop-Implemented Primary Frequency Regulation," Proceedings of the IEEE North American Power Symposium, Morgantown, WV, September 17 -19, 2017.
[52] F. Wilches-Bernal, R. Concepcion, J. Neely, R. Byrne, and A. Ellis, "Communication Enabled Fast Acting Imbalance Reserve (CE-FAIR)," IEEE Trans. Power Systems, 2017.
[53] J. Neely, S. Gonzalez, J. Delhotal, J. Johnson, M. Lave, "Evaluation of PV Frequency-Watt Function for Fast Frequency Reserves," IEEE Applied Power Electronics Conference (APEC), Long Beach, CA, March 20-24, 2016.

the financial components of the operations. Here we focus on current or future markets which could provide revenue streams for the VPP. In the first VPP design, the VPP interacts with the ISO/RTO market directly to generate revenue. In the second VPP, the VPP can provided primary frequency reserves which is not a current market, but has been considered in CAISO[54] and other regions[55]. The final VPP design was for voltage regulation, in which it was envisioned that a utility would be faced with purchasing voltage regulation equipment (e.g., a capacitor bank), or using the DER on the feeder to perform the voltage regulation. The utility in this case could pay the VPP (and DER owners) for this service to avoid the cost of purchasing the new equipment. This scenario is perhaps most likely contextualized in a VIU setting.

Many of the VPP designs for this project are based on a proposed aggregation of DER assets at the Sandia National Laboratories Distributed Energy Technologies Laboratory (DETL), the $10M University of New Mexico-operated Mesa del Sol Aperture Center microgrid, and the DOE-funded Public Service Company of New Mexico (PNM) Prosperity Site. The sites are shown in Figure 2 and the combined collection of DER assets at these locations is:

- PNM 500 kW Battery Energy Storage System (BESS)
- PNM 500 kW PV
- MdS Site Controller (BEMS) connected to:
    - 240 kW natural-gas-powered reciprocating engine
    - 80 kW natural-gas-powered phosphoric acid fuel cell
    - 50 kW PV array mounted as a parking lot shading structure
    - 163 kWh / 50 kW advanced lead-acid battery bank
- DETL 30 kW battery
- DETL 225 kW diesel genset
- DETL SunSpec Gateway (managing 10 residential inverters)

These sites have also been (and continue to be) a central part of several research ventures. For instance, PNM research at the Prosperity Site has demonstrated PV smoothing[56], load shifting[57], and peak smoothing[58]. This team also conduced PV-smoothing control demonstrations using the prosperity battery and the MdS genset[59,60].

---

[54] R. Mullin, CAISO Seeks Primary Frequency Response Market, RTO Insider, December 26, 2016. URL: https://www.rtoinsider.com/caiso-frequency-response-36087/

[55] W. Li, P. Du, N. Lu, "Design of a New Primary Frequency Control Market for Hosting Frequency Response Reserve Offers from both Generators and Loads," IEEE Transactions on Smart Grid, vol. PP, no.99, pp.1-1.

[56] S. Willard, O. Lavrova, B. Arellano, J. Hawkins, A. Mammoli, B. McKeon, “Smoothing and shifting PV – Applying energy storage to enhance the benefits of renewable energy.” 2012 World Renewable Energy Forum; ASES 2012.

[57] A. Ellis, D. Schoenwald, J. Hawkins, S. Willard, B. Arellano, “PV output smoothing with energy storage.” 38th IEEE Photovoltaic Specialists Conference (PVSC), 2012.

[58] O. Lavrova, F. Cheng, S. Abdollahy, H. Barsun, A. Mammoli, D. Dreisigmayer, S. Willard, B. Arellano, C. van Zeyl, “Analysis of battery storage utilization for load shifting and peak smoothing on a distribution feeder in New Mexico.” 2012 IEEE PES Innovative Smart Grid Technologies (ISGT).

[59] J. Johnson, A. Ellis, A. Denda, K. Morino, T. Shinji, T. Ogata, M. Tadokoro, “PV Output Smoothing using a Battery and Natural Gas Engine-Generator,” 39th IEEE Photovoltaic Specialists Conference, Tampa Bay, Florida, 16-21 Jun, 2013.

[60] J. Johnson, K. Morino, A. Denda, J. Hawkins, B. Arellano, T. Ogata, T. Shinji, M. Tadokoro, A. Ellis, “Experimental Comparison of PV-Smoothing Controllers using Distributed Generators,” Sandia Technical Report SAND2014-1546, Feb 2014.

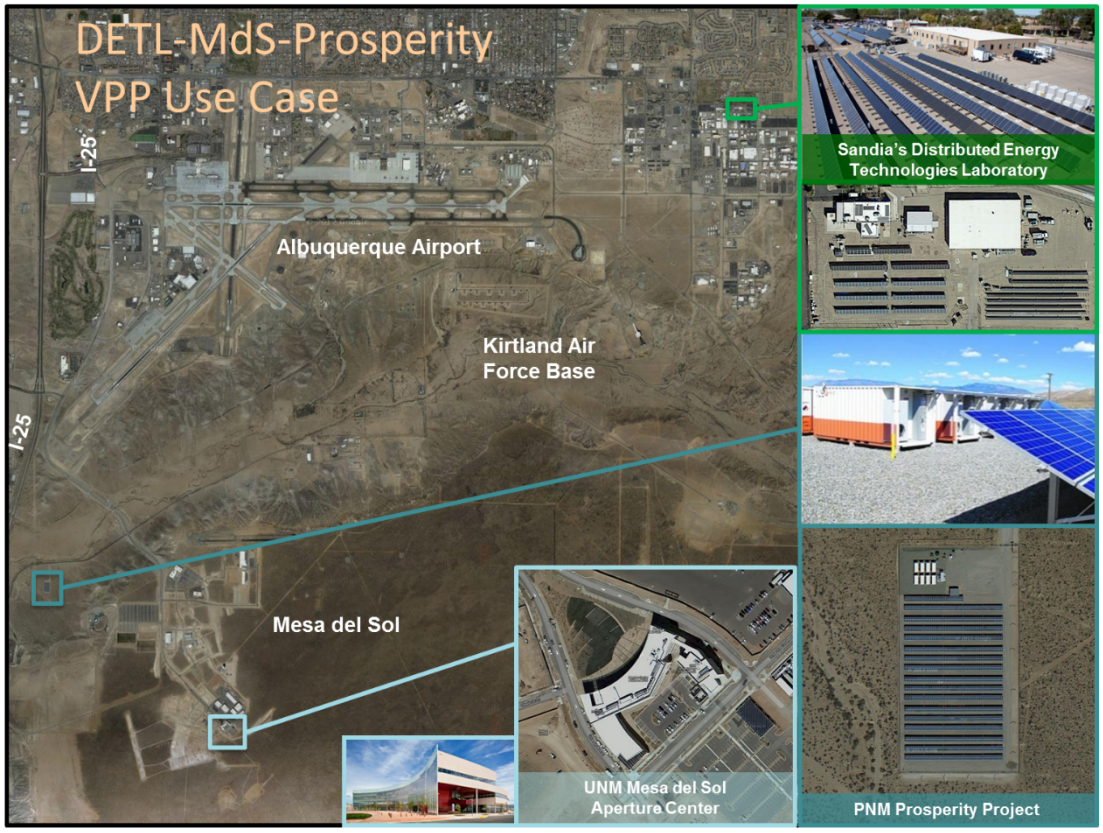

Figure 2. The sites in the UNM-PNM-SNL VPP use case.

3.1 VPP Design for Energy and Reserve Markets

A modular VPP architecture was designed to co-optimize bids into energy and tertiary contingency reserve markets, as shown in Figure 3. This VPP is constructed of a number of modules which interact with DER, markets, and back-end databases to provide the grid service:

1. A **market commitment engine** uses stochastic optimization to maximize expected profit by determining day-ahead VPP bids for the hourly energy and reserve markets. The commitment engine contracts with the ISO/RTO to determine the VPP day-ahead commitments.
2. Once commitments have been established, the **dispatch engine** uses stochastic optimization to find the dispatch schedule for the next short time horizon. This time horizon will be at minimum the amount of time it takes to solve the optimization. The optimizer finds the dispatch schedule for DER assets using a cost minimization objective function with a large penalty for not meeting the commitment. This penalty can be changed to represent the cost of purchasing energy on the spot market or made more severe if this is not an option and the VPP is critical to providing these services.
3. To maintain the desired VPP output magnitude, a **centralized controller** sends DER power setpoints via a communication system. The centralized controller uses multiple

PID controllers and a swing controller (typically for a battery) to ensure the VPP meets the aggregate demand in real-time.

4. A **forecasting** component provides long-term (24-60 hours) forecast of RE anticipated power to the commitment engine, and short term (0-12 hours) forecasts to the dispatch optimization engine.
5. The secure **communications system** monitors and issues commands to each DER through a range of DER gateways which communicate using a variety of open and proprietary communication protocols.

This components and optimization of this VPP will be discussed at length in Section 4.

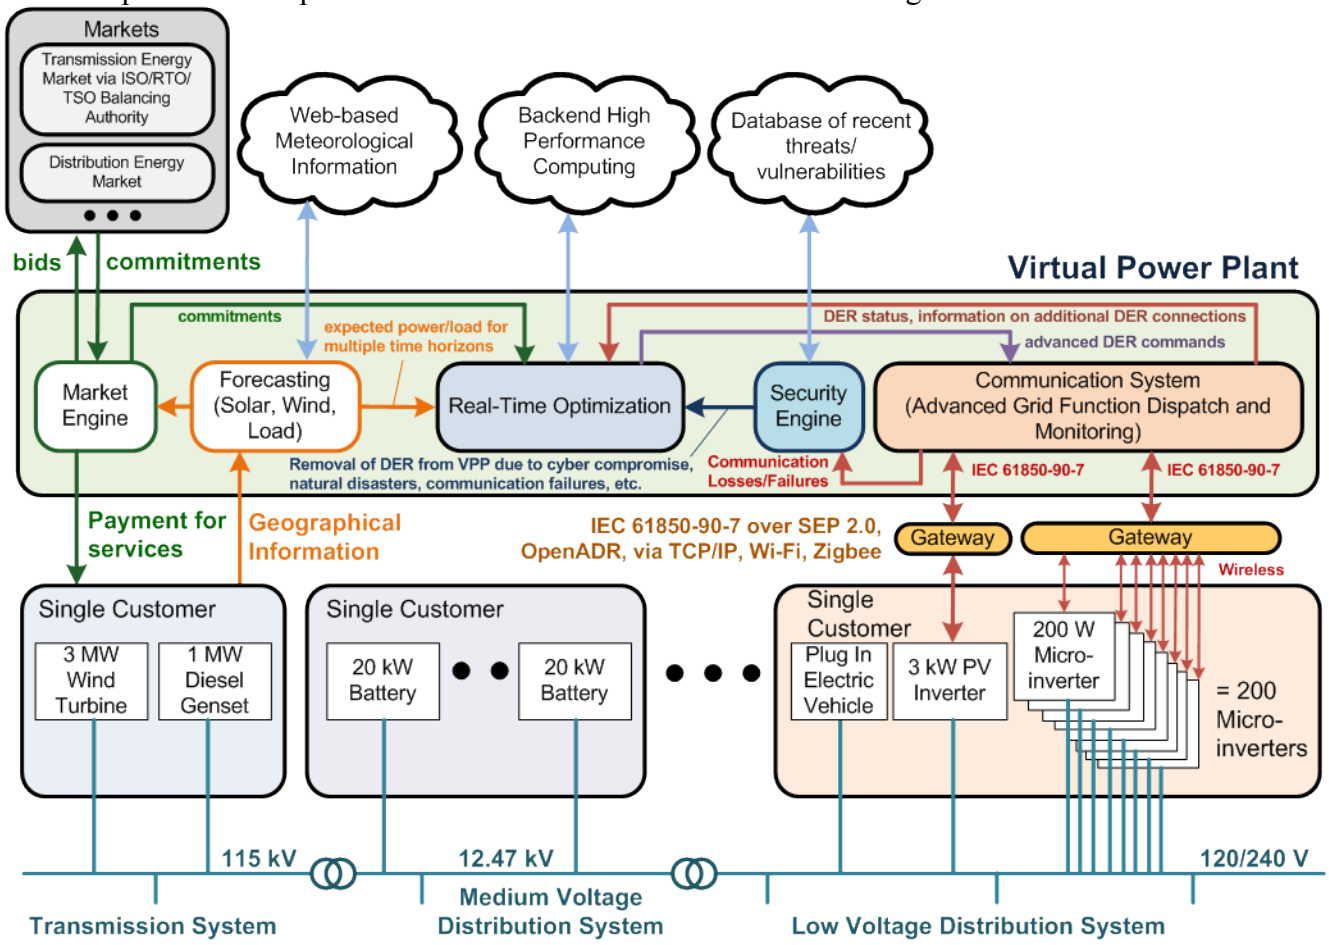

Figure 3. The VPP construct for providing power to the energy and contingency reserve markets.

3.2 VPP Design for Primary Frequency Response Reserves

In this VPP design, we envision a virtual power plant capable of meeting primary frequency response reserve metrics by incorporating a hybrid control scheme consisting of coordinated centralized commercial/utility-scale DERs and decentralized control to dozens of small, residential-scale (<1 kW) DERs using standardized interoperability functions. For primary frequency response reserves, it is not possible to rely solely on centralized control of DERs

because communication latencies will prevent the aggregation from meeting the response and ramp time metrics. Fortunately, the synthetic frequency response reserve is called upon during the initial frequency drop/rise during a fault to simulate spinning inertia and governor response. Individual DERs will be programmed with VPP-optimized advanced interoperability grid functions (e.g., frequency-watt functions) every 15 minutes based on updated renewable and demand response (DR) forecasts. When there is a grid frequency change, the DERs with autonomous functions will exchange active power with the grid to partially stabilize the disturbance. As shown in Figure 4, the VPP is constructed of a number of modules:

1. A stochastic **commitment engine** determines the maximum reserve capacity the VPP can provide toward the grid reserve requirements on a day-ahead, forward basis. The commitment engine contracts directly with the utility or has a cleared offer in the ancillary service markets[61]. The contract/offer is based on long-term forecasts of RE and DR units within the VPP pool, estimated operational and opportunity costs, desired profit margin, and reliability requirements.
2. Once a commitment has been established, the **stochastic optimizer** is enabled to adjust storage states-of-charge and turn on/off the gensets. The optimizer monitors the status of the DERs and, based on short term DER forecasts and DER availability, the optimizer will charge[62] the ESS or start gensets to maintain enough headroom to always meet the reserve requirements. The VPP maintains the ability to meet the reserve with economic stochastic optimization while improving expected costs over alternative options.
3. In order to reach the desired VPP output magnitude, **hybrid control** is employed to generate the appropriate system response. Distributed controls will be enabled in inverter-based DER with SunSpec, IEC 61850, IEEE 2030.5, or IEEE 1815 frequency-watt functions shown in Figure 4 and in engine-generators using droop controls. Since the behavior of the DERs relies on local measurement of grid frequency, the autonomous response of the DERs will be less than 10 cycles (0.16 secs). Unfortunately, the response will rarely meet the reserve requirements alone, so the aggregate autonomous power response must be measured such that the centralized control can then provide the remainder of the reserve magnitude.
4. A **forecasting** component is the same as the previous VPP and provides long-term (24 hour) forecast of RE and DR anticipated power to the commitment engine, and short term (<1-60 minute) forecasts to the dispatch engine.
5. The secure **communications system** monitors and issues commands to each DER through a range of DER gateways which communicate using a variety of open and proprietary communication protocols. This functionality will be established with a Distributed Energy Resource Management System (DERMS).

[61] There is currently no market in the U.S. for primary frequency reserves (i.e., fast frequency response reserves), although it is possible this market will exist in the future.
[62] ESSs are likely to perform other grid-support functions or economic analyses to schedule discharge.

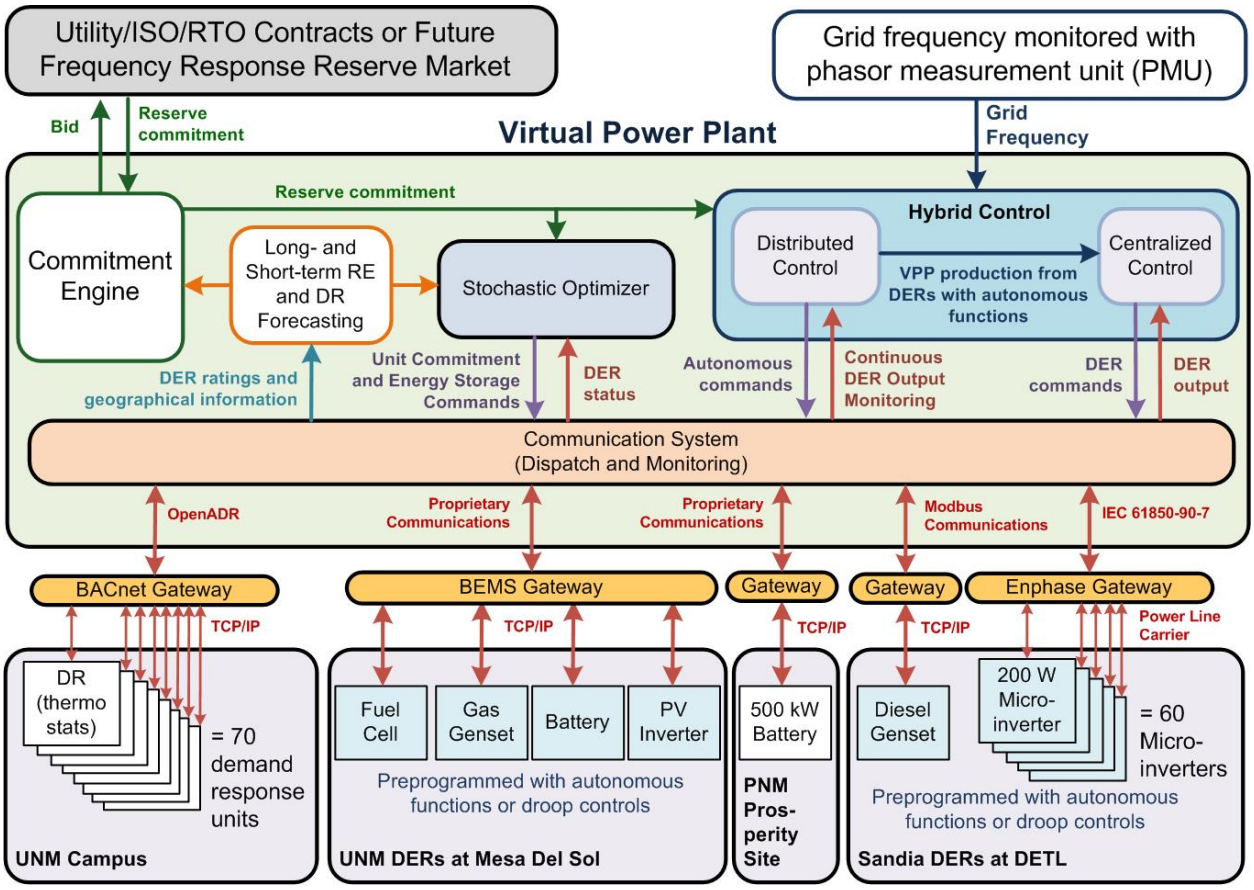

Figure 4. The VPP concept for providing frequency response reserves. Exemplary DERs from four sites in Albuquerque, NM are used to illustrate the diversity of resources, communication protocols, and gateways.

The following sections describe how the VPP for primary frequency reserves would operate.

3.2.1 Forecasting

The long-term forecasting component of the VPP is necessary to calculate the reserve capacity that can be offered to the ISO/RTO over the commitment period. The bid into the day-ahead market essentially promises to the utility/ISO/RTO that the VPP can meet the Reserve Magnitude Target (RMT) for that duration. For example, if a wind turbine was a DER in the VPP and the next 24 hours were forecast to be windy the VPP would be able to provide a higher reserve magnitude than if the next 24 hours were calm. Similarly, if the next day was going to be hot with large demand response-enabled air conditioning loads, the VPP could provide a higher reserve because those devices could be turned off temporarily to provide more 'negawatts' in aggregate.

For long-term forecasting, the forecasting component converts meteorological forecasts to PV power, wind power, and thermostatically or otherwise controllable load forecasts. First, hourly, day-ahead NOAA forecasts of irradiance, wind speed, and temperature are collected. Then PV power forecasts can be computed using a PV model (like the Photovoltaic Array Performance Model[63]), wind production can be computed from wind power models, and demand response

capabilities for thermal units (air conditioning) can be predicted from a 1st order thermal model with outdoor ambient temperature.

The short-term forecasting component supports unit dispatch decisions accounting for renewable energy variability. For example, if PV power is forecasted to decrease in the next 15 minutes, e.g., due to cloudiness, the PV system may not be able to provide energy in the event of a frequency deviation. In that case, the VPP can prepare for that loss in reserve capacity by charging energy storage systems or starting a generator that is off. This ensures that the risk of not providing the committed reserve capacity is minimized. Short-term forecasting can be done for all DR assets using statistical methods fit to recent DR performance data, e.g., PV power production.

3.2.2 Commitment Engine

As stated above, the capabilities of the VPP must be provided to the utility/ISO/RTO to meet FERC requirements[64]. This robust solution provides greater than 99% availability by appropriate reserve magnitude commitments/bids. This could be confirmed with Monte Carlo simulations for the next 24 hours with probabilistically-generated forecasts, DER communication dropouts, and DER failures. The probability densities for each of the losses and the stochastic forecasting process model would be determined with historical information from each of the DERs. Generally, the larger DERs would be more reliable because they have dedicated communication connections maintained by utilities.

3.2.3 Stochastic Optimization

The ESS optimizer uses Sandia's Python-based Pyomo[65] optimization framework in conjunction with a commercial solver, Gurobi Optimizer[66]. While more technically challenging, stochastic inputs make the optimizer robust to residual uncertainties in the forecasts of PV, wind and other RE output and demand loads. In particular, residual uncertainty in the RE/DR forecast is minimized (ideally, less than 5%), but accommodates forecast outliers where there are large deviations in predicted and actual renewable energy production. The goal is to prepare the storage DERs for losses in renewable energy and demand response available power. In practice, the optimization function could either be performed in-house by a VIU or it could be contracted to a VPP aggregator. In either case, the goal is to operate the VPP in such a way to maximize availability and minimize utility/ISO/RTO/VPP costs while maintaining ANSI, IEEE, and NERC standards. VPP expenses are a combination of opportunity costs, operation costs, and maintenance costs associated with added wear on the VPP components (for example, loss of battery life due to cycling). The decision variables in the optimization include charge/discharge

---

[63] D. L. King, W.E. Boyson, J.A. Kratochvil, Photovoltaic Array Performance Model, Sandia Technical Report SAND2004-3535, Dec 2004.
[64] J. F. Ellison, L.S. Tesfatsion, V.W. Loose, R. H. Byrne, "Project Report: A Survey of Operating Reserve Markets in U.S. ISO/RTO-managed Electric Energy Regions," Sandia Technical Report, SAND2012-1000, Sept. 2012.
[65] W.E. Hart, C.D. Laird, J.P. Watson, D.L. Woodruff, (2012) Pyomo: Optimization Modeling in Python. Springer, Berlin.
[66] Gurobi Optimization, the state of the art mathematical programming solver, URL: http://www.gurobi.com/, accessed 7/16/15.

commands to the ESSs using IEC 61850-90-7/SunSpec Alliance or other DER communication protocols.

3.2.4 Hybrid Control

The hybrid controller consists of a portion of DERs configured with autonomous functionality and others configured to receive centralized control commands.

3.2.4.1 Distributed Control

The objective of primary frequency response reserves is to operate within the first few seconds of a frequency change to arrest the system imbalance until secondary control reserves are brought online with automatic generation control (AGC) to replace the frequency response reserves. An example of the time scales that frequency response reserves operate during a grid fault (e.g., the loss of a generator) is shown in Figure 5. The VPP would ideally provide the reserve magnitude within the first 2-3 seconds to reduce the severity of the nadir. To do this, the distributed controller must carefully select the parameters for the inverter-based frequency-watt functions and the slope of the generator governors. Normally a governor slope is selected to be 5% (i.e., 3 Hz deviation would result in the full output of the generator), but for the VPP the distributed controller will likely select a far more aggressive slope (e.g., ~1%) so that the prime movers produce full output at 0.6 Hz grid frequency deviation. Similarly, the frequency-watt settings, shown in Figure 6, would be chosen by the distributed controller to aggressively respond to frequency deviations. Prior research on selecting PV inverter frequency-watt settings to maintain grid frequency during an N-1 contingency on the island grid of Lanai, HI could be leveraged to optimize the distributed control settings for the VPP DERs[67,68,69].

[67] J. Neely, J. Johnson, J. Delhotal, S. Gonzalez and M. Lave, "Evaluation of PV frequency-watt function for fast frequency reserves," 2016 IEEE Applied Power Electronics Conference and Exposition (APEC), Long Beach, CA, 2016, pp. 1926-1933.

[68] J. Johnson, J. Neely, J. Delhotal, M. Lave, "Photovoltaic Frequency-Watt Curve Design for Frequency Regulation and Fast Contingency Reserves," IEEE Journal of Photovoltaics, vol. 6, no. 6, pp. 1611-1618, Nov. 2016.

[69] A. Hoke, A. Nelson, J. Tan, V. Gevorgian, C. Antonio, K. Fong, M. Elkhatib, J. Johnson, R. Mahmud, J. Neely, D. Arakawa, The Frequency-Watt Function: Simulation and Testing for the Hawaiian Electric Companies, Grid Modernization Laboratory Consortium (GMLC) Technical Report, July 2017.

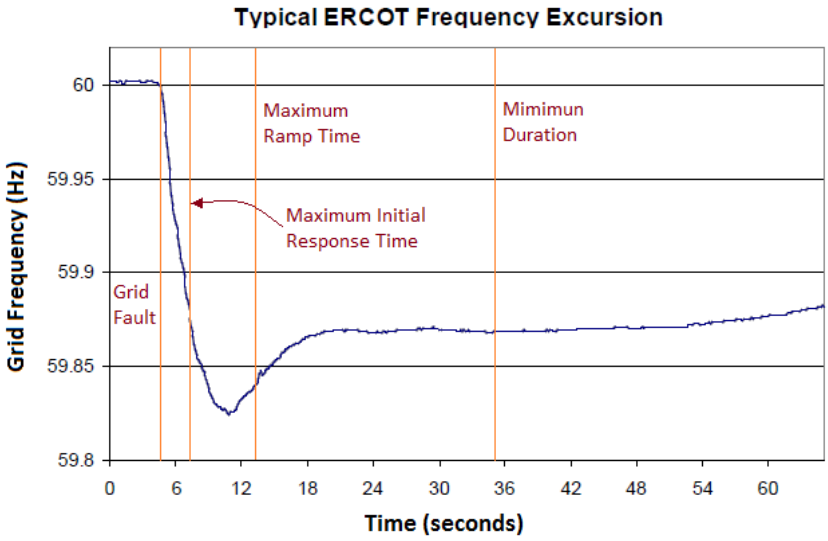

Figure 5: Typical ERCOT Frequency Excursion[70] with the VPP frequency response reserve time metrics.

Distributed droop-like controls will be programmed into converter-based DER to autonomously respond to locally-measured frequency deviations in sub-second time frames. Renewable energy sources can be curtailed to provide headroom during frequency drops; but in doing so, there are significant opportunity costs for the DER owners because they will not be gathering all the available renewable energy revenue. Therefore, the selection of optimal frequency-watt (FW) curve points—illustrated in Figure 6—is necessary to minimize the lost energy revenue while maintaining enough headroom to meet the reserve commitment with a given (e.g., 99.9%) confidence. This trade-space is the basis for the optimal control strategy, which provides the required capabilities for minimal DER opportunity and operation costs. One of the major challenges for this type of VPP is associated with stochastic optimization computation times and communication failures/latencies.

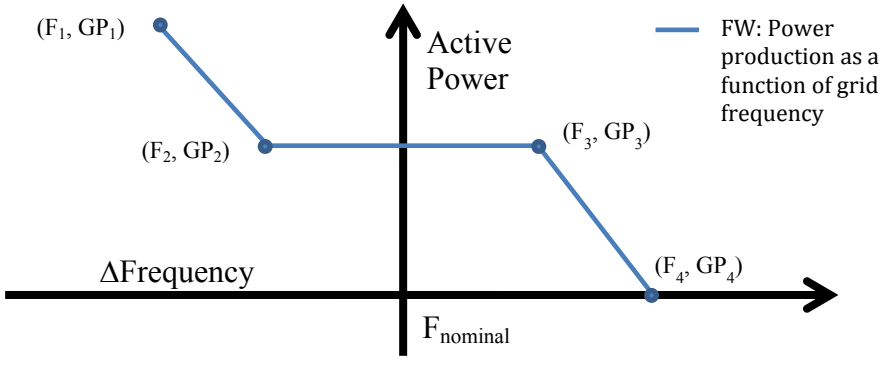

[70] NREC Resources Subcommittee, Balancing and Frequency Control, January 26, 2011.

Figure 6. Autonomous frequency-watt curve which curtails active output power at nominal grid frequency, increases generation at low grid frequency, and decreases generation at high grid frequency.

3.2.4.2 Centralized Control

Once the distributed controllers respond to the frequency deviation, the centralized control is responsible for adjusting the power from the larger DERs to meet the reserve magnitude. The VPP will predict the expected response of the autonomous control based on the grid frequency deviation. The frequency will be monitored continuously with a Phasor Measurement Unit (PMU) so that the centralized controller is aware when it must send DER commands and to estimate the response of the DERs with distributed-controls. The distributed control DERs will be continuously monitored to verify the response of the distributed controls. Based on Sandia experiments at the DETL, the response of power electronic devices with frequency-watt functions is highly accurate, but the DR response will vary depending on the duty cycles of the DERs, ambient temperatures, thermostat settings, etc. As information is received from the distributed DER, the estimate for the distributed response becomes more accurate and the centralized controller updates the DER active power commands.

3.2.5 Communications

The VPP construct necessitates control communications to distribution-level devices at residential and commercial facilities over wired and wireless communication channels. In the hypothetical VPP use case communications to DETL equipment would use SunSpec Modbus RTU and TCP/IP commands[71] to 10 advanced inverters with the ability to curtail the output of the PV system. The remaining generators and demand response units at MdS, Prosperity and UNM would be controlled with a variety of communication protocols:

- Control of the MdS devices will be conducted through a communication link to the Toshiba-developed Building Energy Management System (BEMS). In a previous U.S.-Japan bilateral project, BEMS-connected DERs at MdS were controlled to provide PV smoothing using an OSIsoft PI historian[72].
- The prosperity battery is controlled using gain adjustments and auxiliary inputs in the control algorithm[73], programmed into the Battery Energy Storage System (BESS) by Ecolt. Previously, Sandia led a project to smooth prosperity PV output by adjusting the battery control[74]
- The thermostats in the UNM mechanical engineering building will be control with OpenADR commands via BACnet.

---

[71] J. Johnson, B. Fox, "Automating the Sandia Advanced Interoperability Test Protocols," 40th IEEE PVSC, Denver, CO, 8-13 June, 2014.

[72] J. Johnson, K. Morino, A. Denda, J. Hawkins, B. Arellano, T. Ogata, T. Shinji, M. Tadokoro, A. Ellis, "Experimental Comparison of PV-Smoothing Controllers using Distributed Generators," Sandia Technical Report SAND2014-1546, Feb 2014.

[73] A. Ellis and D. Schoenwald, PV Output Smoothing with Energy Storage, Sandia Technical Report, SAND2012-1772, March 2012.

[74] J. Johnson, A. Ellis, A. Denda, K. Morino, T. Shinji, T. Ogata, M. Tadokoro, "PV Output Smoothing using a Battery and Natural Gas Engine-Generator," 39th IEEE Photovoltaic Specialists Conference, Tampa Bay, Florida, 16-21 Jun, 2013.

Latencies associated with monitoring distributed control DERs and commanding centralized control DERs are especially critical. The precise requirement depends on a tradeoff between the accuracy of the distributed control response estimation and accuracy of the centralized controller response. If the centralized controller can hit the reserve magnitude target directly, the communications requirements can be relaxed; however, if it takes multiple attempts to adjust the DER aggregation to reach the reserve setpoint, the communication time must be short.

3.3 VPP Design for Regulation, Following Reserve, and Contingency Reserves

The broadly-applicable Sandia VPP architecture can provide grid stability with operating reserves. The end goal is a VPP software system which provides utilities, balancing authorities, and third parties a portfolio of services to adjust the operation of a DER fleet and in this case, we designed the system to provide frequency control reserves and non-event reserves (regulation and following reserves). As shown in Figure 7, the VPP is constructed of several modules:

1. A stochastic **market engine** which determines the appropriate supply curve (price vs provided active power) for one or multiple grid-support services (e.g., frequency reserves). The market engine contracts with the utility, ISO, or ancillary service markets based on long-term forecasts of renewable energy and demand response units within the VPP pool, estimated operational and opportunity costs, desired profit margin, and reliability requirements.
2. Once a commitment has been established in the day ahead, hourly, or sub-hourly markets, a **DER optimization** is performed to adjust storage states-of-charge and monitor the status of the DER. Based on the renewable energy and demand short term forecasts, the DER optimization establishes the optimal economic dispatch of DER resources. Note that the optimization will not be able to run in real-time, so there will be some error in the initial response of the VPP based on communication latency and failures, forecasting errors, and inaccuracies in DER statuses.
3. In order to reach the desired output magnitude required for the reserve type, a **real-time feedback control** is employed to dynamically reallocate DER operations. This operation is extremely fast, but not optimized for economic dispatch of the DER resources. Instead, the real-time feedback control relies on heuristic rules to meet the magnitude requirement within the designated response time.
4. A **forecasting** component provides long-term (24 hour) forecasts of demand response and renewable energy components to the market engine, and short term (<1-15 minute) forecasts to the optimization engine.
5. The secure **communications system** monitors and commands each of the DER through a range of DER gateways. These gateways will communicate using a variety of open and proprietary communication protocols. This functionality will be built into a Distributed Energy Resource Management System (DERMS), such as the EPRI DERMS[75].

[75] B. Seal, "Enterprise Integration of Distributed Energy Resources," EPRI IntelliGrid Smart Grid Information Sharing Webcast, 22 Jan 2014.

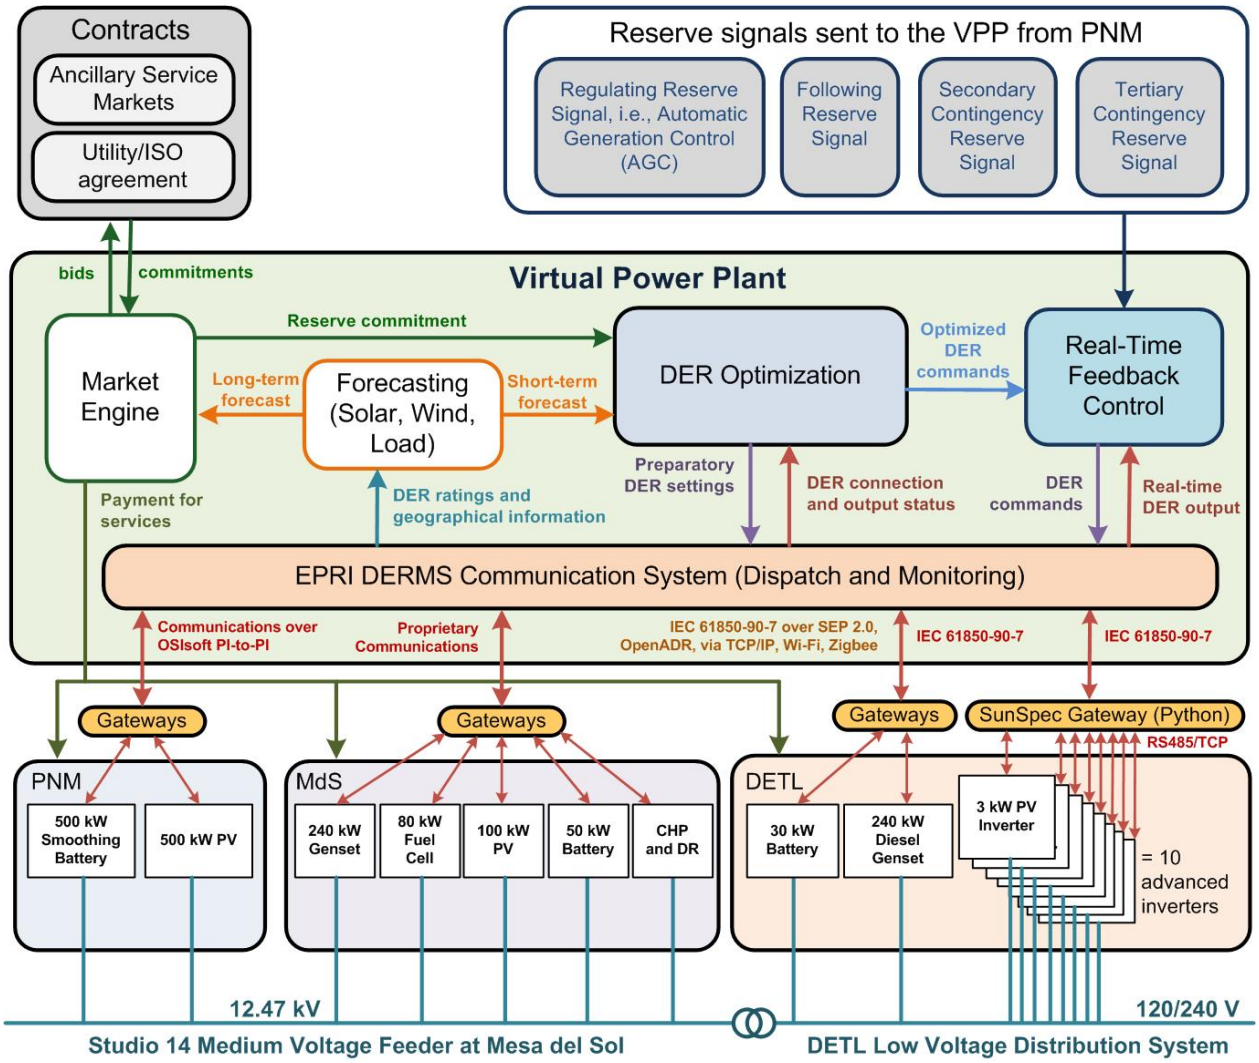

Figure 7: Secure Virtual Power Plant design for multiple operating reserves.

While there are variations in the requirements for each of the reserves depending on the ISO/RTO, a generalized set of requirements for each of the reserves is provided in Table 1. In addition to the applicable NERC standards, these will be the target metrics for each of the reserve functions in the project (based on ISO requirements). The regulating reserve time constraints and tertiary contingency reserve magnitude tolerance requirements are the most demanding for the VPP, so during the optimization and control design phase, those critical metrics will drive much of the design. An example of how the *DER optimization* and *real-time feedback control* operate together to meet these reserve metrics is shown in Figure 8. There are multiple steps to reach the target reserve magnitude: (A) is the response time needed to receive the reserve command signal and send the optimized DER commands through the DERMS system to the DERs; (B) depending on communication latency and DER response characteristics, the VPP output ramps up; multiple feedback iterations occur as the VPP works toward the target magnitude by (C) monitoring the DER and (D) issuing new commands; (E) eventually the VPP reaches a magnitude where the feedback control no longer issues DER commands but the controller monitors the output of the VPP to ensure it remains within the tolerance. If the VPP

magnitude approaches the tolerance limits the controller will adjust DER outputs accordingly. This deviation could result from renewable energy or demand response changes, storage or fuel depletion, or other unexpected events. During this final stage, the VPP can also reallocate DER commitments based on additional economic optimizations.

Table 1: Reserve metrics for the virtual power plant.

Reserve Metric	Regulating Reserve	Following Reserve	Secondary Contingency Reserve	Tertiary Contingency Reserve (based on Spinning Reserve)
Initial Response Time	2 sec	10 min	1 min	5 min
Magnitude Tolerance	±10%	±5%	±5%	±3%
Ramp Time	5 sec	10 min	5 min	10 min
Duration	30 min	60 min	60 min	60 min

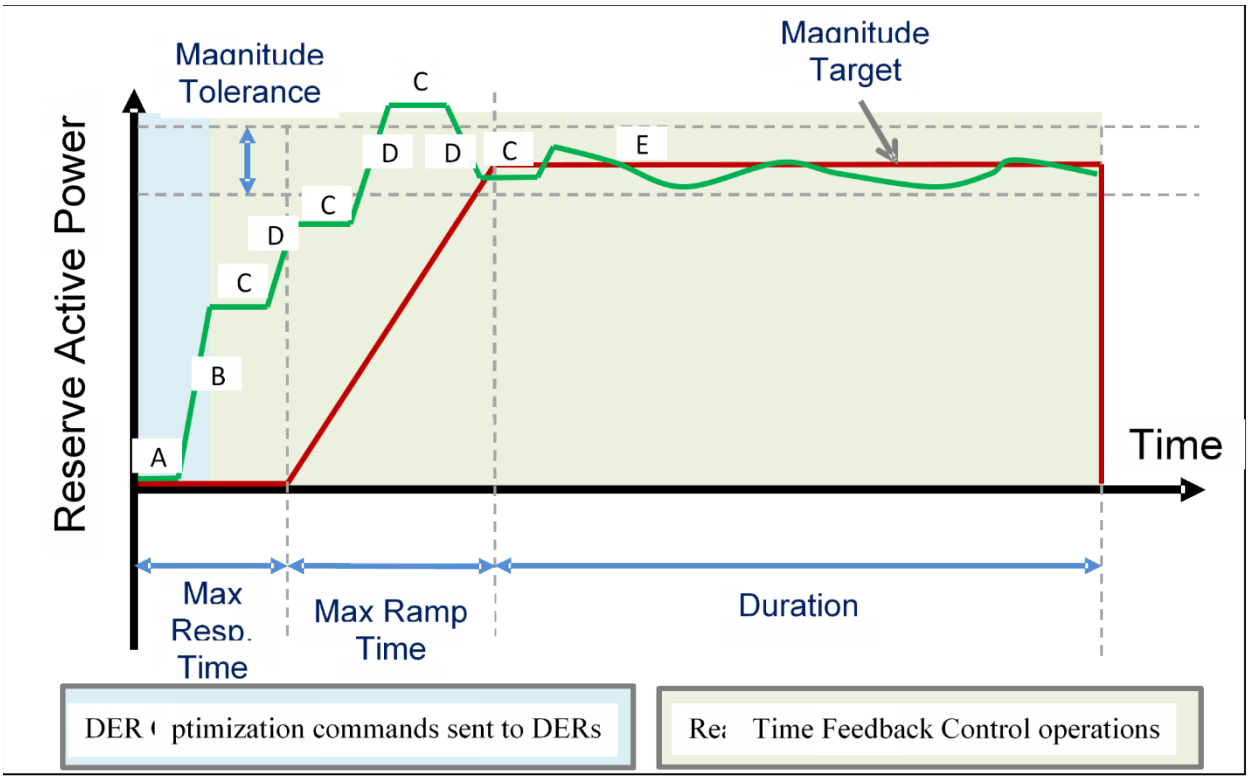

Figure 8: The reserve metrics with example VPP output.

The stochastic optimization will use Sandia's Python-based Pyomo[76] optimization framework in conjunction with a commercial solver. The optimization will be stochastic to make it robust to residual uncertainties in the forecasts of PV output, demand response availability, and other DER capacities. In particular, the residual uncertainty in the PV forecast will be minimized (ideally, less than 3%), but must accommodate forecast outliers where there are large deviations in predicted and actual renewable energy production and demand response. In practice, the optimization function could be performed in-house by a vertically integrated utility or it could be

[76] W.E. Hart, C.D. Laird, J.P. Watson, D.L. Woodruff, (2012) Pyomo: Optimization Modeling in Python. Springer, Berlin.

contracted out to a VPP aggregator. In either case, the goal is to operate the VPP in such a way to maximize profit while maintaining ANSI, IEEE, and NERC standards. The profit generated by the VPP primarily comes from value streams corresponding to the multiple types of reserves listed above (regulating, following, secondary contingency, and tertiary contingency). Subtracted from the value streams are opportunity costs and other costs associated with added wear on the VPP components (for example, loss of battery life due to cycling). The decision variables in the command optimization include the PV inverter settings. Other decision variables are the commitments of the small generators and the charging and discharging rates for energy storage systems. The commands may be structured in a hierarchical fashion corresponding to local, regional and global scales. Experiments will determine the benefit of optimizing at various levels of electrical and geographical aggregation in comparison to a fully disaggregated optimization of all VPP components.

3.4 VPP Design for Distribution Voltage Regulation

Increasing photovoltaic (PV) penetration on distribution circuits can cause voltage swings that could compromise protection systems and damage customer equipment. Sandia National Laboratories and the Georgia Institute of Technology have been collaborating to develop a VPP construct capable of providing distribution voltage regulation. Our design consists of an open-source, safe, cyber-secure, ADMS architecture that dispatches PV inverters and other power electronics-based DER to maintain voltage control, monitor and protect the distribution system and improve grid performance and reliability. This technology will ultimately result in lower capital expenditures for distribution circuits and introduce a breakthrough method for modeling, protecting, and managing these circuits in the future. As an example, our ADMS technology will allow utilities to defer—possibly indefinitely—fielding additional expensive voltage regulation equipment, like LTCs and capacitor banks.

Scalable distribution system planning and real-time operations software dispatches resources that have DER grid-support functions to provide voltage regulation while maintaining distribution protection schemes—thereby optimizing the system operations over the planning horizon. SCADA sensors, other telemetry, and DERs feed data to a Distribution System Distributed Quasi-Dynamic State Estimator (DS-DQSE). A stochastic optimization engine, called Prescient[77], in conjunction with a Quasi Static Time Series (QSTS) model, to optimize DER active and reactive power settings (P/Q) that maximize economic value while also minimizing the risk of exceeding the ANSI C84.1 Range A voltage limits and compromising protection systems. Since inverter grid-support functions (e.g., volt-var (VV), specified power factor (SPF), and active power curtailment) reduce active power and PV owner revenue, the control system will also co-optimize the operations of the DERs to minimize lost production.[78] Co-optimizing DER operations to meet multiple objectives is novel in distribution system management.

[77] Sandia Develops Stochastic Production Cost Model Simulator for Electric Power Systems, press release. 2015.
[78] While there are some incentives (payment structures) in the wholesale markets for reactive power capacity and output, this is not typically the case at the distribution level. In the future, there may be new regulatory models that incentivize reactive power supply on feeders—in which case the optimization problem would be redesigned to account for this additional revenue stream.

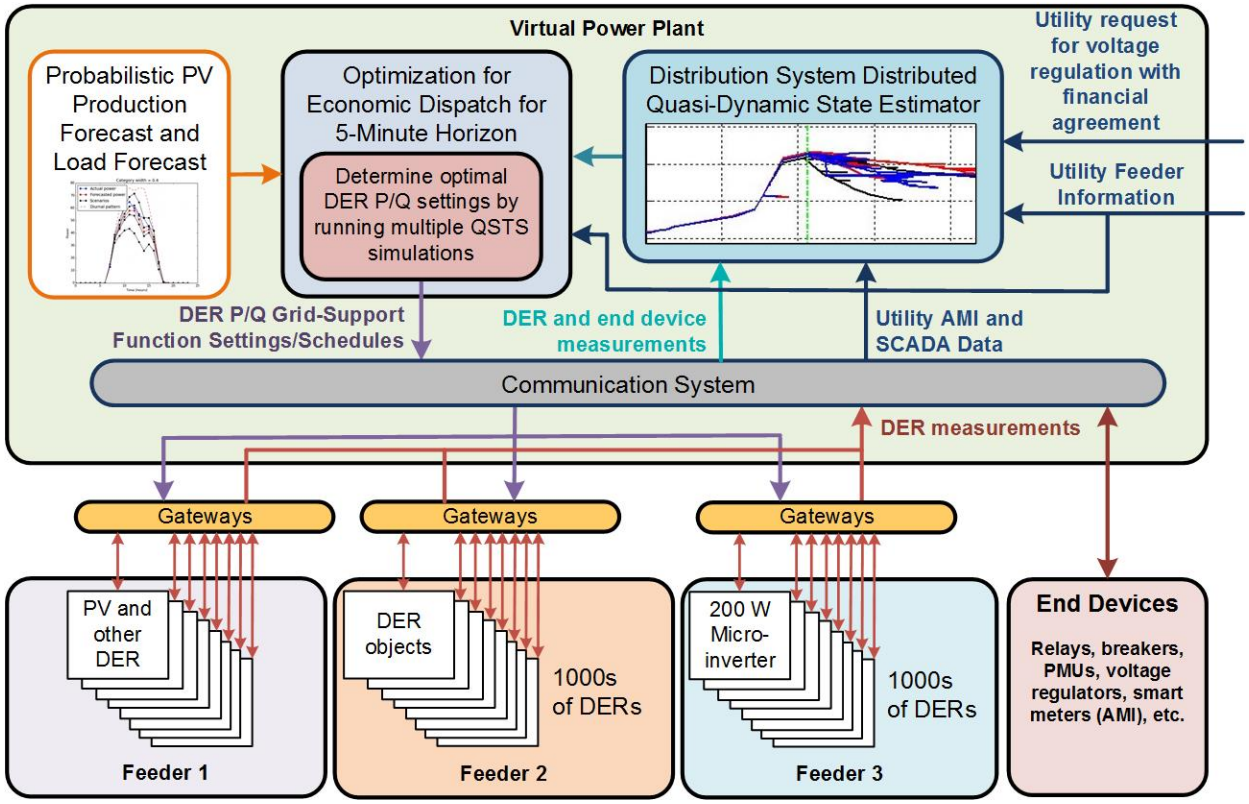

Figure 9: The VPP construct for providing voltage regulation.

The open-source software system will provide real-time protection, voltage regulation and visualization by integrating the following modules, shown in Figure 9:

1. The Distribution System Distributed Quasi-Dynamic **State Estimator** (DS-DQSE)[79] takes feeder telemetry, DER and customer data, and generates the voltage and power flow estimation and validates the real-time model.[80] This information is used to populate the QSTS simulations within the dispatch optimization engine.
2. (Optional) the **Estimation-Based Protection** (EBP) detects faults and protects the system by isolating the faulted section of the distribution circuit by recloser/breaker/switching operations. The EBP solves the issues associated with the present protection schemes for distribution circuits with distributed resources.
3. The **probabilistic forecasting** component provides short-term (e.g., 10 minute) forecasts of PV power output and load using recent system states and statistical irradiance modeling in conjunction with PV performance models. Uncertainty in forecasts is quantified using an empirical technique[81] which refines uncertainty estimates using feeder telemetry data over time.
4. A **dispatch optimization engine** determines the necessary P/Q power settings for groups of DERs to maintain voltage and distribution protection systems for the next time period

[79] R Huang, et al., IEEE Transactions on Smart Grid, Vol PP, Issue 99, pp 1-10, February 2016.
[80] M. H. Nazari, et al., Trans. Power Systems, vol. 29, no. 6, pp. 2934-2942, Nov. 2014.
[81] Y. Feng, I. Rios, S.M. Ryan, et al. "Toward scalable stochastic unit commitment. Part 1: load scenario generation," Energy Systems, Volume 6, Issue 3, pp 309–329 September 2015. doi:10.1007/s12667-015-0146-8

(5 minutes) value. The optimizer evaluates circuit performance using QSTS simulations given the state estimate and the range of power and load forecasts to minimize the risk of voltage or protection violations while also maximizing economic value.

5. The **communications system** uses a DER to communicate via internet channels to DER or aggregators that send signals to inverters. VV/SPF and active power curtailment commands will be issued to the DERs via TCP/IP using SunSpec, IEEE 2030.5, IEC 61850-90-7, or revised IEEE Std. 1547 information models. **Cybersecurity** is paramount to successful interoperable operations and the success of any DERMS system.

PV-induced voltage fluctuations along distribution feeders are problematic but will be managed with the system. Present utility practice of regulating voltage with regulators at the start of the feeder and capacitors along the feeder are inadequate for feeder with high penetration of PVs and the variability of PVs causes wear and tear of voltage regulators, transformer load tap changers (LTCs), and switched capacitors.[82] The number and the frequency of the PV fluctuations determine the impact on the number of tap changes and cap switches. PV systems distributed along a feeder exhibit less variability than the same capacity of PV concentrated at a single location, and therefore less impact on the voltage regulators.[83] To efficiently model the complex interactions between load and PV variability, quasi-static-time-series (QSTS) simulation tools are being developed.[84]

By appropriately employing the advanced inverter functions in coordination with legacy voltage regulation equipment, the distribution system PV hosting capacity can be significantly increased.[85] PV inverter reactive power (e.g., volt-var) functions can control the voltage locally and provide some voltage regulation, reducing the number of voltage regulator tap changes and capacitor switches. The approach also determines the appropriate advanced inverter settings. Prior results show that poor volt-var settings can increase the number of tap changes significantly from the unity power factor case (>2 times the number of taps);[86] whereas, by selecting the correct volt-var curve, the number of tap changes can be reduced 20%. Methods have been proposed to determine site specific inverter settings, but the settings are highly dependent on the specific scenario.[87]

This VPP design determines reactive power settings for groupings of inverters through stochastic optimization. The state estimator produces a real-time model of a feeder from telemetry data

---

[82] H. Ravindra, M. O. Faruque, K. Schoder, M. Steurer, P. McLaren, and R. Meeker, "Dynamic interactions between distribution network voltage regulators for large and distributed PV plants," in IEEE PES Transmission and Distribution Conference and Exposition (T&D), 2012.

[83] S. Dhakal, P. Tripathi, and M. F. Baroughi, "Distributed versus Centralized Photovoltaic Plants in IEEE-34 Node Test Feeder for Reduced Power Fluctuation and Stress on Feeder Components," in Conference on Power and Energy Systems, 2013.

[84] R. Aghatehrani and A. Golnas, "Reactive power control of photovoltaic systems based on the voltage sensitivity analysis," in IEEE Power and Energy Society General Meeting, 2012.

[85] J. Seuss, M. J. Reno, R. J. Broderick, and S. Grijalva, "Improving Distribution Network PV Hosting Capacity via Smart Inverter Reactive Power Support," IEEE PES General Meeting, Denver, CO, 2015.

[86] S. R. Abate, T. E. McDermott, M. Rylander, and J. Smith, "Smart inverter settings for improving distribution feeder performance," in IEEE Power & Energy Society General Meeting, 2015.

[87] M. Rylander, M. J. Reno, J. E. Quiroz, F. Ding, H. Li, R. J. Broderick, et al., "Methods to Determine Recommended Feeder-Wide Advanced Inverter Settings for Improving Distribution System Performance," in IEEE Photovoltaic Specialist Conference (PVSC), 2016.

primarily from interoperable PV inverters and other DERs, but also from AMI, microPMUs, and traditional SCADA equipment as available from the utility partners. Short-term forecasts of PV production and load, along with uncertainty in these forecasts, and the real-time feeder model are fed to the optimization engine. Within the engine, QSTS simulations estimate voltage and power flows to determine the P/Q settings that simultaneously minimize risk of voltage violations while also maximizing the economic value of PV generation. The reactive power will maintain the feeder voltage within a tight tolerance around the nominal grid voltage and the active power control will ensure that operating constraints are satisfied. In turn, this approach guarantees that protection functions will not be triggered due to wide variability of PV output and abnormal voltages and/or loading.

The VPP design for voltage regulation is presented in a separate report.[88]

[88] A.P. Meliopoulos, G. Cokkinides, B. Xie, C. Zhong, J. Johnson, "Full State Feedback Control for Virtual Power Plants," Sandia Technical Report, September 2017.

4 VPP COMPONENTS

Of the multiple VPP designs in Section 3, only the VPP Design for Energy and Reserve Markets was implemented in code. For this VPP, the team created modular components which ran as multi-processing servers in a Python environment. The components of the VPP interacted to exchange pertinent information through a backend process. This section describes each of the VPP components in more detail as well as the experimental design, results, and validation taken for each VPP component prior to integration.

4.1 Forecasting

The forecast component was implemented in two servers: a long-term forecast engine supporting day-ahead unit commitment, and a short-term forecast engine supporting within-day unit dispatch. The long-term and short-term forecast engines differ in data requirements and forecast methods as illustrated in Figure 10. The forecast methods implemented for the VPP are less sophisticated (and less accurate) than state-of-art methods available in the renewable power industry. For the VPP demonstration, the selected methods suffice; the VPP structure permits replacement with alternative forecasts if desired.

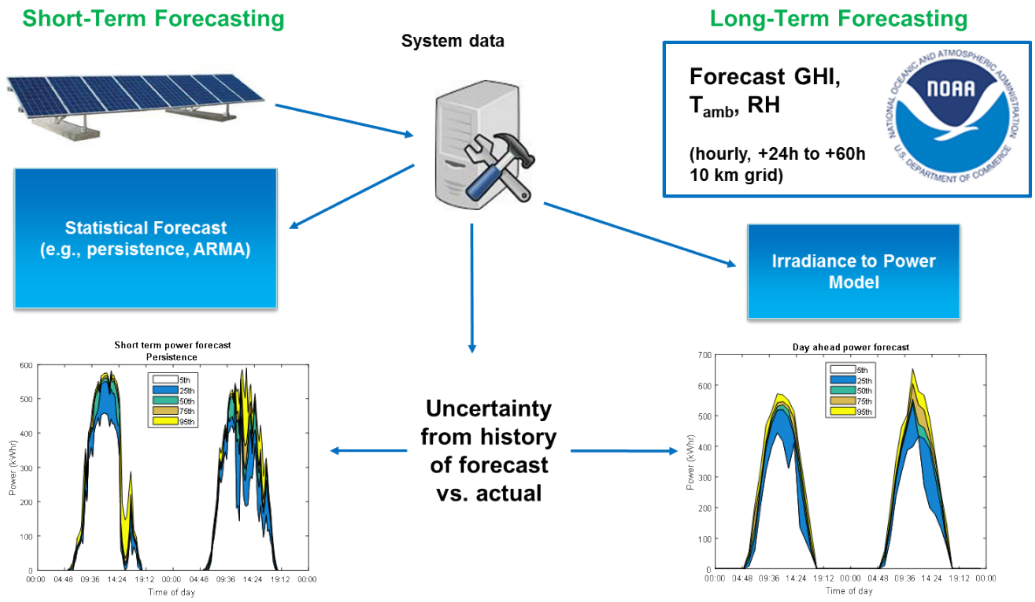

Figure 10: Example forecast

4.1.1 Long-term forecasting

Commitment to day-ahead energy and reserve markets requires forecasts of energy and power production. The VPP forecast engine provides forecasts of hourly PV power levels for the next operating day by 1) obtaining pertinent weather forecasts from NOAA and 2) applying a weather-to-power transfer model. The NOAA North American Mesoscale (NAM) CONUS model provides gridded (approx. 12km x 12km) forecasts of irradiance, surface temperature and surface wind speeds at hourly resolution out to 84 hours, with model initiation at 00:00, 06:00,

12:00 and 18:00 daily. Forecasts are available approximately 3 hours after model initiation. The VPP code downloads the relevant data fields at specified geographic coordinates from the NOAA forecasts (available for most recent 7 days at nomads.ncep.noaa.gov). The download routine runs continuously to accumulate a long-duration archive of forecast values; considerable effort was made to enable the code to ride through forecast delays, network outages, download interruptions and corrupted data transfers. Historical forecast values are accumulated to characterize forecast error for the stochastic optimization that determines day-ahead unit commitments.

Table 2. Day-ahead forecast weather fields obtained from NOAA NAM

Quantity	NOAA NAM field name
Irradiance (horizontal, W/m^2)	dswrfsfc
Air temperature (K)	tmp2m
Surface wind (m/s)	gustsfc

A weather-to-power transfer model is implemented using PV system modeling components from the pvlib-python toolbox[89] which originated at Sandia National Laboratories and which provides reference implementations of commonly-used functions for solar power modeling. PV systems are defined by choosing modules and inverters, specifying geographic location, system topology, mounting and orientation. The pvlib-python components handle translation of forecast irradiance, air temperature and wind speed to PV power.

Figure 11 illustrates long-term forecast results. The challenge inherent in forecasting irradiance at a specific geographic location 24 hours in advance is evident in the Figure. The forecast irradiance (which correlates well with power) may reasonably anticipate the power trend (DOY 91 and 93) or may miss widely (DOY 92 and 95); the wide scatter in the bottom panel of Figure 11 indicates a significant change of a large forecast error in any particular hour. Forecast performance is summarized by a mean bias error of –32 kW (6% of capacity) mean absolute error of 92 kW (18% of capacity) and a root mean square error of 124 kW (25% of capacity) (error statistics exclude nighttime). Our results are similar to those reported elsewhere.[90] Errors for irradiance forecasting using models tailored for solar power can be substantially less[91]; these forecasts can be purchased from a number of vendors. The NOAA product is the only publicly available forecast which includes hourly irradiance.

[89] Pvlib-python available at https://github.com/pvlib/pvlib-python
[90] D. Larson, L. Nonnenmacher, C. Coimbra, Renewable Energy, vol. 91, pp. 11-20, 2016.
[91] R. Perez et al., Solar Energy, vol. 94, pp. 305-326, 2013.

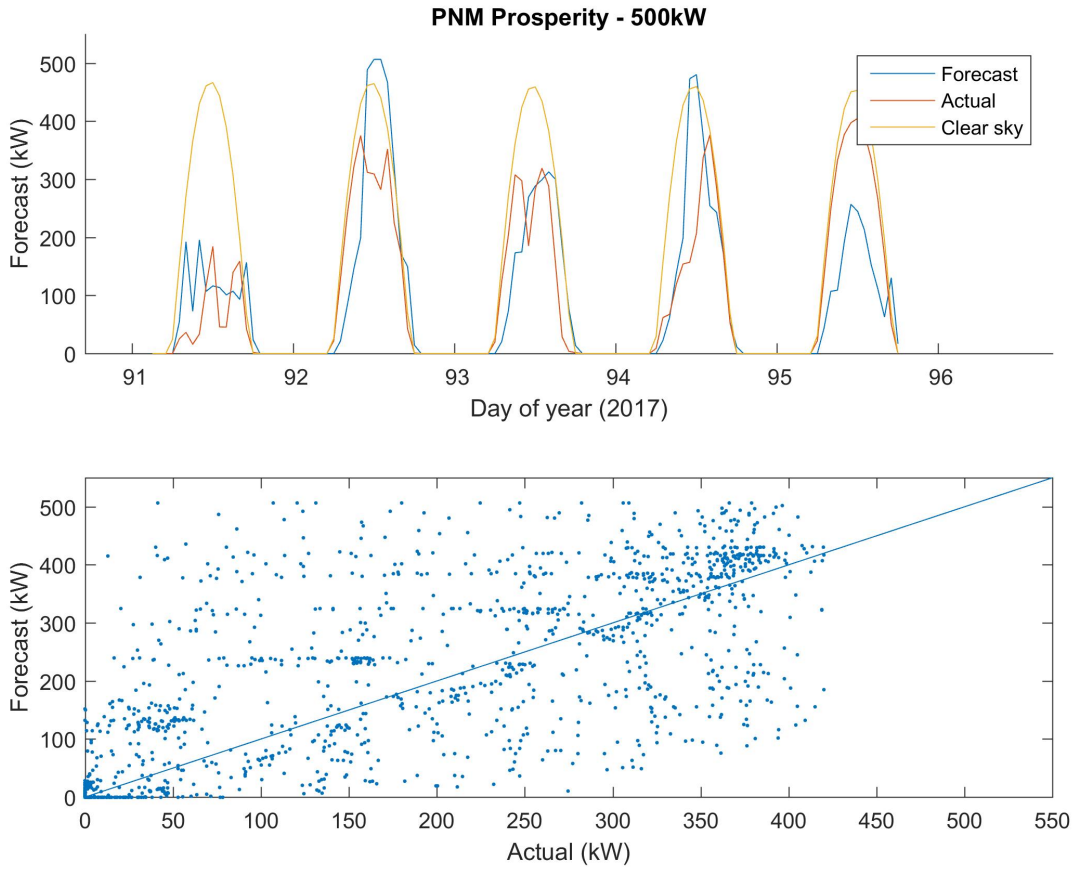

Figure 11. Long-term forecasts of power for the PNM Prosperity 500 kW PV plant: forecast power, measured power and forecast clear-sky power (top), and comparison of forecast and measured power for Feb 2, 2017 through May 17, 2017.

4.1.2 Short-term forecasting

Dispatch of VPP resources to meet energy and reserve commitments must anticipate changes to the VPP power generation levels. Forecasts of PV power at, e.g., 15 minute intervals are made to the end of each operating day using a persistence technique; the VPP code includes capability to replace the persistence technique with autoregressive/moving average (ARMA) models (with or without differencing) fitted on rolling windows, or to download hourly forecasts from NOAA's High Resolution Rapid Refresh (HRRR) model and employ the weather-to-power transfer model used in the long-term forecast engine.

For PV power forecasts, persistence was implemented on the clear-sky power index, i.e., the power of the PV system normalized by the PV system power under clear sky conditions. Historical PV power was obtained from system meters; for the VPP demonstration, data for the PNM Prosperity 500 kW system was available through an OSIsoft® PI historian that collected field measurements from the PNM site. PV power, air temperature and wind speed data were downloaded on a continual basis from the historian. For the short-term persistence forecast at

15-minute resolution, the most recent hour of power data was converted to clear-sky power index, and the forecast was made by assuming the clear-sky power index persists into the future. The clear-sky power index was multiplied by the forecast clear-sky power to obtain the power forecast. The clear-sky power was forecast by applying the weather-to-power transfer model to forecasts of the clear-sky irradiance, air temperature and wind speed; clear-sky irradiance on the PV system plane was calculated using models in the pvlib-python toolbox, air temperature and wind speed were assumed to persist by using the last hour's measured values.

Figure 12 displays example time series of short-term forecasts and compares forecast and measured power from the 500 kW PNM Prosperity PV system. The lag inherent in persistence forecasting is evident when comparing time series of forecast and actual values (e.g., DOY 91): a significant change in irradiance results in a corresponding change in power output, which is repeated by the forecast at a later time. This lag largely would be absent from forecasts produced by fitting an ARIMA model to rolling windows; the error could also be corrected by applying a model output statistics (MOS) technique, in which a statistical model is formed to predict forecast error, and the predicted error is subtracted from the forecast. Despite the lag, forecast performance is reasonably clustered about the 1:1 line in the bottom panel. Forecast accuracy is summarized by: mean bias error of –15 kW, mean absolute error of 39 kW, and root mean square error of 64 kW.

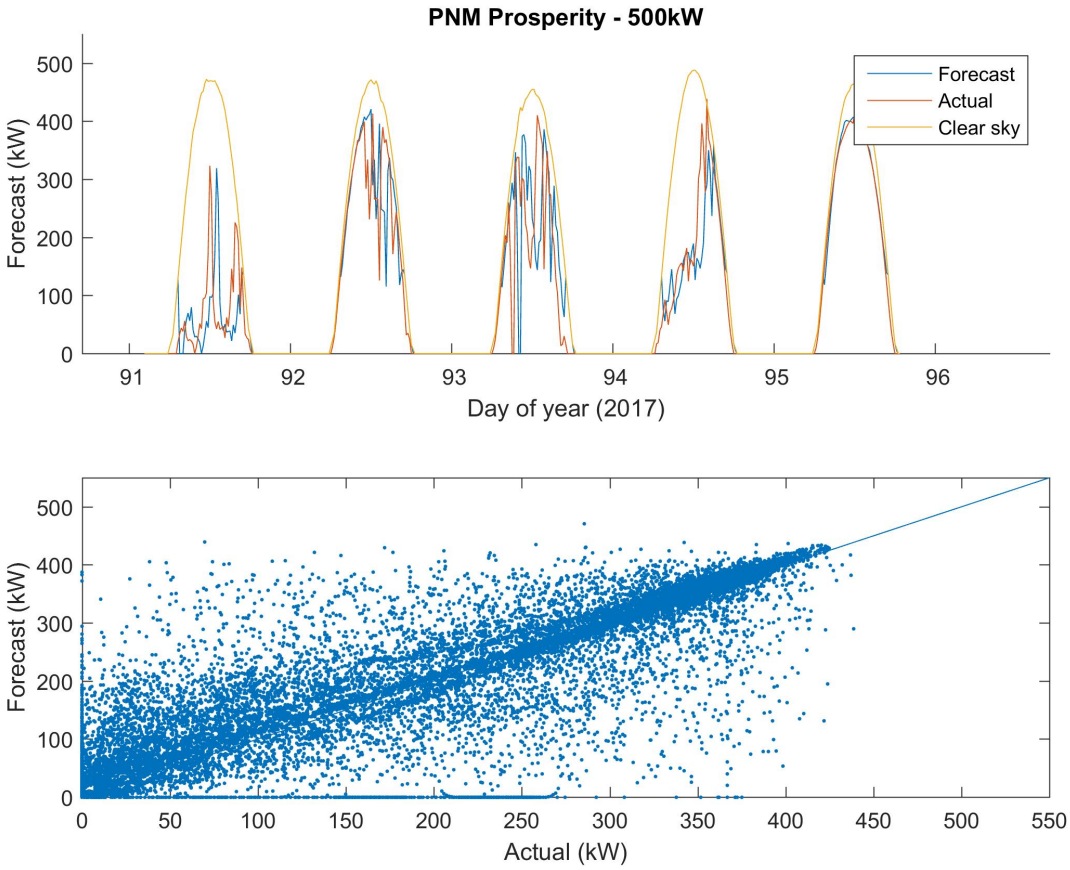

Figure 12. Short-term forecasts of power for the PNM Prosperity 500 kW PV plant: forecast power, measured power and forecast clear-sky power (top), and comparison of forecast and measured power for Feb 2, 2017 through May 17, 2017.

4.2 Optimization

The day-ahead scheduling and real-time operations are two separate optimization routines that are solved on a rolling horizon basis. These optimization problems were developed in Sandia's Python-based Pyomo mathematical programming framework, and then solved with the commercial solver Gurobi. The Prescient software tool[92] is leveraged for both optimization problems to generate scenarios that quantify historical forecast errors in the long-term and short-term forecasting engines.

Both optimization routines can be classified as two-stage stochastic programs. The variables that represent decisions that are made before the uncertainty is revealed are known as first-stage variables, and decisions variables that represent decisions that are made once the uncertainty is revealed are known as second-stage. In a two-stage stochastic program, an optimal policy is

[92] A. Staid, J.-P. Watson, Roger J.-B. Wets, and D. L. Woodruff, "Generating short-term probabilistic wind power scenarios via nonparametric forecast error density estimators." Wind Energy, 2017.

solved for all possible future realizations of the uncertainty parameters. We assume that the VPP is an infra-marginal resource participating in the market, and therefore does not impact the market prices.

4.2.1 Quantifying Uncertainty through Scenario Generation

The VPP must operate under uncertainty in the forecast variability that, when not appropriately handled, may result in an offer strategy that yields sub-optimal outcomes. The sub-optimal outcomes in the long-term effect the economic viability of the VPP. Therefore, to account for forecast variability, we incorporate uncertainty modeling into our stochastic programming. A common practice is to approximate a probability distribution of an input parameter by a set of scenarios with associated probabilities of occurrence. For example, a random parameter $\theta_\omega, \omega \in \Omega$ where ω indices the set of scenarios Ω. Each realization of θ_ω is associated with a probability of occurrence $\Pr(\omega)$, defined as:

$$\Pr(\omega) = Pr\left(\omega | \theta = \theta_\omega\right),$$

where $\sum_{\omega \in \Omega} \Pr(\omega) = 1$.

The software tool Prescient constructs a set of relevant scenarios that capture forecast errors by including low-probability, but not extremely unlikely, events for more reliable and resilient operations. In particular, we construct probabilistic scenarios that capture the range of intermittent generation and market pricing behaviors that have been historically observed. Prescient characterizes forecast error based on nonparametric density estimates of historical forecast errors through epi-spline basis functions. The scenarios are generated in such a way that they capture known relationships present in solar photovoltaic and pricing patterns, such as temporal dependencies, as well as biases and correlation in the forecasts. These scenarios along with their estimated probabilities of occurrence provide a set of possible future outcomes over which the dispatcher can optimize. Therefore, the optimal decisions in the stochastic program is based on data available at the time that the decisions are made, and not on future observations. The basic idea is that the stochastic optimization enables the VPP operator to be prepared for all possible future realizations of the unknown parameters.

4.2.2 Day-Ahead Energy and Reserve Market Stochastic Co-optimization

In the day-ahead, a two-stage stochastic optimization problem is solved to determine optimal energy and reserve schedules for the subsequent operating day that maximizes the expected profit,

$$E[\Pi_\omega] = \sum_{\omega \in \Omega} \Pr(\omega) \Pi_\omega,$$

across all scenarios $\omega \in \Omega$ in the day-ahead market, where the *first-stage variables* include:

(1) The power schedule for the VPP, p_t^s, which can be a withdrawal $\left(p_t^s < 0\right)$, injection $\left(p_t^s > 0\right)$, or no participation $\left(p_t^s = 0\right)$ for each time period $t \in T$ of the subsequent operating

day. Note that the power schedule is not for each VPP subresource but is an aggregated schedule for the VPP to offer;

(2) The reserve capacity for the VPP, $\bar{r}_t^s$, which can be a positive capacity $(\bar{r}_t^s > 0)$, or no participation $(\bar{r}_t^s = 0)$ for each time period $t \in T$ of the subsequent operating day. Note that the reserve capacity is not for each VPP subresource but is an aggregated capacity for the VPP to offer;

(3) The binary commitment variable, $b_{t,g}$, that indicates whether thermal unit $g \in G$ is online for each time period $t \in T$. Note that the binary commitment variables are used for all VPP subresources that have start-up and/or no-load costs.

The *second-stage variables* indicate the power production, $p_{\omega,t,w}$ and the reserve capacity, $\bar{r}^{s}_{\omega,t,u}$, of all VPP subresources $u \in U$ for each time period $t \in T$ in each scenario $\omega \in \Omega$. For storage devices, the energy level, $e_{\omega,t,u}$, is also defined as a second-stage variable. As a result, the scheduled power and reserve for the VPP as a whole are scenario-independent and therefore are feasible for all future realizations of scenarios $\omega \in \Omega$, whereas the power output of each VPP subresource depends on the realization of each scenario $\omega \in \Omega$.

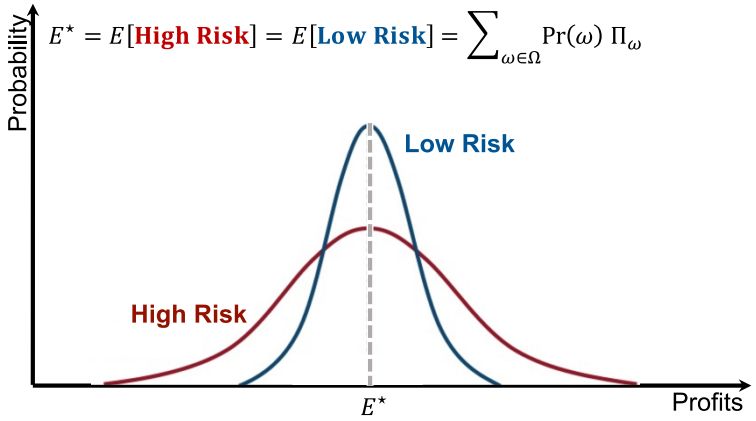

Figure 13. Illustration of profit variability per scenario, where the expected profit is the same for both the high risk and low risk uncertainty sets.

In the proposed stochastic optimization approach, the VPP operator has the option to handle the risk associated with profit variability (see Figure 13) due to the inherent uncertainty in day-ahead energy prices, $\lambda^{DA}_{\omega,t}$, day-ahead reserve prices, $\lambda^{SP}_{\omega,t}$, and intermittent resource generation, $P^{F}_{\omega,t,i}$, for the solar photovoltaic subresources $i \in I$ in the VPP. By incorporating a multi-objective conditional value-at-risk (CVaR) optimization approach, the VPP operator can determine optimal schedules that result in low-risk operating strategies.

We explicitly incorporate this through an α-confidence level CVaR approach. This approach balances the trade-off between profits and risk, enabling the VPP to adaptively schedule and operate as a more risk-averse or more risk-neutral market participant. The $\alpha \in (0,1)$ denotes the confidence level and for an equivalent α used for VaR as shown in Figure 14, the CVaR is the mean of the tail distribution exceeding VaR, i.e., the average loss exceeding the VaR value.

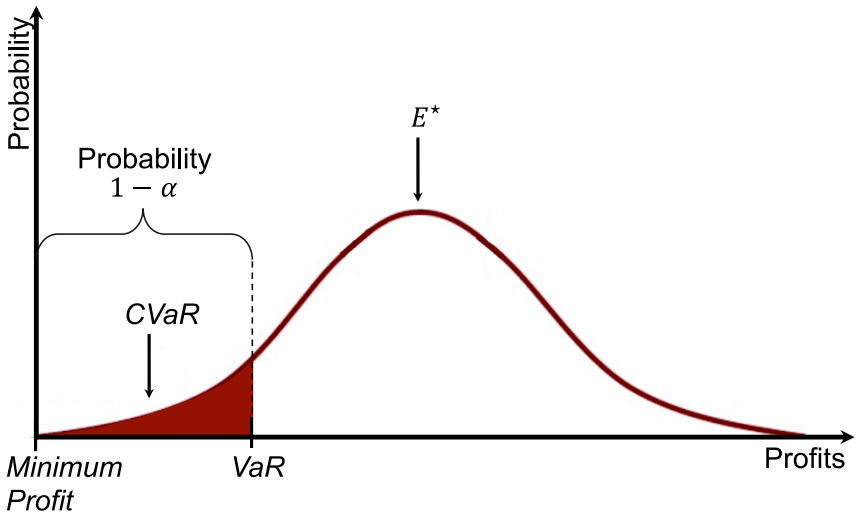

Figure 14. Multi-objective function of the VPP Operator maximizes both the expected profits and the CVaR of the profits.

Therefore the VPP scheduling model is formulated as a multi-objective stochastic optimization problem, with objective function weight $\beta \in [0,1]$ (e.g., $\beta = 0$ is expected profit maximization only), where the operator maximizes the expected profit and the CVaR of the expected profit over an uncertainty set of scenarios $\omega \in \Omega$ in:

$$\max (1-\beta)\sum_{\omega \in \Omega} \Pr(\omega)\Pi_{\omega} + \beta CVaR - \Phi$$

where,

$$\Pi_{\omega} = \sum_{t \in T}\left(\lambda_{\omega,t}^{DA} p_t^s + \lambda_{\omega,t}^{SP} \bar{r}_t^s\right) - \sum_{u \in U} f_u\left(p_{\omega,t,u}, \bar{r}_{\omega,t,u}^{\ s}\right) \qquad \text{VPP Profits per Scenario}$$

$$CVaR = VaR - (1-\alpha)^{-1}\sum_{\omega \in \Omega} \Pr(\omega)\sigma_{\omega}$$

$$\Phi = \sum_{\omega \in \Omega} \Pr(\omega)\sum_{t \in T} \lambda_{\omega,t}^{DA}\left(\epsilon^{-1} p_{\omega,t}^{+} + \epsilon p_{\omega,t}^{-}\right) \qquad \text{Penalty function}$$

subject to

$VaR - \Pi_{\omega} \le \sigma_{\omega}$	$\forall \omega \in \Omega$	CVaR Formulation
$0 \le \sigma_{\omega}$	$\forall \omega \in \Omega$	Nonnegativity
$p_t^s = \sum_{u \in U} p_{\omega,t,u} - p_{\omega,t}^{+} + p_{\omega,t}^{-}$	$\forall \omega \in \Omega, \forall t \in T$	VPP Energy Schedule Offer
$\bar{r}_t^s \le \sum_{u \in U} \bar{r}_{\omega,t,u}^{\ s}$	$\forall \omega \in \Omega, \forall t \in T$	VPP Reserve Capacity Offer

and VPP subresource constraints — See Appendix A

Note that the subresource constraints are specified in Appendix A. The above problem is risk-constrained when the VPP operator specifies $\beta > 0$, and is risk-neutral otherwise. In other words, the weighting parameter β enforces a trade-off between profit and risk such that the higher the value of β, the more risk-averse the VPP operator becomes.

The profit calculation Π_ω per scenario $\omega \in \Omega$ incorporates the forecasted prices $\lambda^{DA}_{\omega,t}$ and $\lambda^{SP}_{\omega,t}$ and the CVaR calculation quantifies the losses associated with the lower tail of the profit distribution, where σ_ω is an auxiliary continuous nonnegative variable defined as the maximum between zero and the difference between the VaR and the profit of a scenario ω. The penalty function Φ is to disincentivize the VPP operator from waiting to sell ($p^{+}_{\omega,t}$) or buy ($p^{-}_{\omega,t}$) in the real-time balancing market by creating a loss of $\lambda^{DA}_{\omega,t}(1-\epsilon^{-1})$ or $\lambda^{DA}_{\omega,t}(\epsilon-1)$, respectively, where $\epsilon > 1$.

4.2.3 VPP DER Dispatch

For real-time operations, a two-stage stochastic optimization problem is solved to determine optimal VPP subresource (DER) setpoints that maximizes the expected profit,

$$E[\Pi_\omega] = \sum_{\omega \in \Omega} \Pr(\omega)\Pi_\omega ,$$

across all scenarios $\omega \in \Omega$ and meets the specified energy schedule, $p^{s\star}_{t}$, and reserve capacity, $\bar{r}^{s\star}_{t}$, which are the accepted offers of the VPP operator from the day-ahead market for all time periods $t \in T$ in the subsequent 24-hour timeframe (where $1 \le t \le |T|$ and $|T| = 24$). This approach may also be used for real-time balancing markets in addition to meeting the day-ahead energy and/or reserve market bids.

The *first-stage variables* correspond to decisions of the VPP operator in the current time period, $\tau \in T$, while the *second-stage variables* correspond to the future time periods $\tilde{T} := \{\tau+1,\ldots,\tau+|T|\}$ where $\tilde{T} \subset T$. Therefore, the *first-stage variables* are non-anticipative across all scenarios $\omega \in \Omega$ and include:

(1) The power setpoint, $p_{\omega,\tau,u}$;
(2) The reserve capacity, $\bar{r}_{\omega,\tau,u}$; and
(3) The reserve output setpoint, $r_{\omega,\tau,u}$

for all VPP subresources $u \in U$ in the current time period τ; in other words, each VPP subresource must maintain the same power and reserve capability in the current time period regardless of the future realization in the uncertainties, as enforced by the "non-anticipative constraint set" below. The second-stage variables indicate the power and reserve capability for each VPP subresource by scenario $\omega \in \Omega$ for future time periods $t \in \tilde{T}$.

In the proposed stochastic optimization approach, the VPP operator has the option to handle the risk associated with profit variability (see Figure 13) due to the inherent uncertainty in real-time energy prices, $\lambda^{DA}_{\omega,t}$, and intermittent resource generation, $\bar{P}^{F}_{\omega,t,i}$, for the solar photovoltaic subresources $i \in I$ in the VPP. By incorporating a multi-objective conditional value-at-risk (CVaR) optimization approach, the VPP operator can determine optimal subresource setpoints for the current time period that result in low-risk operating strategies for future time periods.

We explicitly incorporate this through an α-confidence level CVaR approach. This approach balances the trade-off between profits and risk, enabling the VPP to adaptively schedule and operate as a more risk-averse or more risk-neutral market participant. The $\alpha \in (0,1)$ denotes the

confidence level and for an equivalent α used for VaR as shown in Figure 14, the CVaR is the mean of the tail distribution exceeding VaR, i.e., the average loss exceeding the VaR value.

Therefore the VPP dispatch model is formulated as a multi-objective stochastic optimization problem, with objective function weight $\beta \in [0,1]$ (e.g., $\beta = 0$ is expected profit maximization only), where the operator maximizes the expected profit and the CVaR of the expected profit (see Section 4.2.2) over an uncertainty set of scenarios $\omega \in \Omega$ in:

$$\max (1-\beta) \sum_{\omega \in \Omega} \Pr(\omega) \Pi_{\omega} + \beta CVaR - \Phi$$

where

$$\Pi_{\omega} = \sum_{t \in T, u \in U} \left[\lambda^{RT}_{\omega,t} \left(p_{\omega,t,u} + r_{\omega,t,u} \right) - f^{e}_{u} \left(p_{\omega,t,u} + r_{\omega,t,u} \right) \right] \quad \text{VPP Profits per Scenario}$$

$$CVaR = VaR - (1-\alpha)^{-1} \sum_{\omega \in \Omega} \Pr(\omega) \sigma_{\omega}$$

$$\Phi = \sum_{\omega \in \Omega} \Pr(\omega) \sum_{t \in T} \phi \left(p^{\varepsilon +}_{\omega,t} + r^{\varepsilon}_{\omega,t} \right) \quad \text{Penalty function}$$

subject to

$VaR - \Pi_{\omega} \leq \sigma_{\omega}$	$\forall \omega \in \Omega$	CVaR Formulation
$0 \leq \sigma_{\omega}$	$\forall \omega \in \Omega$	Nonnegativity
$p^{s\star}_{t} = \sum_{u \in U} p_{\omega,t,u} + p^{\varepsilon +}_{\omega,t} - p^{\varepsilon -}_{\omega,t}$	$\forall \omega \in \Omega, \forall t \in T$	VPP Energy Dispatch
$\bar{r}^{s\star}_{t} - r^{\varepsilon}_{\omega,t} \leq \sum_{u \in U} \bar{r}^{s}_{\omega,t,u}$	$\forall \omega \in \Omega, \forall t \in T$	VPP Reserve Capacity
$0 \leq r^{\varepsilon}_{\omega,t} \leq \bar{r}^{s\star}_{t}$	$\forall \omega \in \Omega, \forall t \in T$	Reserve Penalty Bound
$0 \leq r_{\omega,t,u} \leq \bar{r}^{s}_{\omega,t,u}$	$\forall \omega \in \Omega, \forall t \in T, \forall u \in U$	VPP Subresource Reserve Output Limits
$p_{1,\tau,u} = \ldots = p_{\lvert\Omega\rvert,\tau,u}$	$\forall u \in U$	Non-anticipative Constraint Set
$\bar{r}^{s}_{1,\tau,u} = \ldots = \bar{r}^{s}_{\lvert\Omega\rvert,\tau,u}$	$\forall u \in U$	
$r_{1,\tau,u} = \ldots = r_{\lvert\Omega\rvert,\tau,u}$	$\forall u \in U$	
$p^{\varepsilon +}_{1,\tau} = \ldots = p^{\varepsilon +}_{\lvert\Omega\rvert,\tau}$		
$p^{\varepsilon -}_{1,\tau} = \ldots = p^{\varepsilon -}_{\lvert\Omega\rvert,\tau}$		
and VPP subresource constraints		See Appendix A

Note that the subresource constraints are specified in Appendix A. Similar to the stochastic scheduling problem, the stochastic dispatch problem is risk-constrained when the VPP operator specifies $\beta > 0$, and is risk-neutral otherwise. The above problem is solved on a rolling horizon so that each time period in the forward 24-hour horizon becomes the current time period τ. Therefore, model references to $\tau - 1$ refers to actual operations for the VPP subresources in the previous time period, which accounts for deviations of the controller from the setpoint determined by the dispatcher.

4.3 DER Controller

VPP optimization has been the subject of many studies. The EU sponsored projects to create a VPP composed of fuel cell DER[93] and the FENIX project investigated (a) technical VPPs consisting of DER in one geographical region that accounted for the local power network (e.g., voltage regulation) and (b) commercial VPPs designed to bid into wholesale and other markets[94]. Many researchers have studied VPP bidding mechanisms and market interactions. Centralized bidding strategies for VPPs were investigated extensively[95,96,97]; and a detailed optimization formulation to optimize the day-ahead thermal and electrical scheduling of large scale VPPs has been proposed[98]. Once the VPP has contracted power delivery, a control system must issue commands to DER to produce the desired aggregate power. To do this, a few architectures have been proposed, including direct, hierarchical and distributed management architectures for VPPs[99] and decentralized multi-agent based techniques for VPP operations[100,101]. However, there is little emphasis in the literature on the design and implementation of real-time feedback control for VPP operations. The Sandia VPP team investigated three different control techniques: distributed control, hybrid control, and centralized control. Ultimately the team employed a centralized controller for the Energy and Reserve Market VPP demonstration.

4.3.1 Distributed Control

Distributed control of DER devices has been a widely-accepted concept for voltage and frequency regulation using the volt-var and frequency-watt functions. For instance, there has been prior research using FW functions to provide grid-services to Hawaii[102, 103, 104]. Unfortunately, DER autonomous functions cannot provide time-sensitive ancillary services alone because they do not optimize the entire system for specific dispatch power levels. The system

---

[93] A. Dauensteiner, "European virtual fuel cell power plant," Management Summary Report, Feb 2007.
[94] D. Pudjianto, C. Ramsay, G. Strbac, and M. Durstewitz, "The virtual power plant: Enabling integration of distributed generation and demand," FENIX Bulletin 2, Feb 2008.
[95] E. Mashhour and S. Moghaddas-Tafreshi, "Bidding strategy of virtual power plant for participating in energy and spinning reserve markets - Part I: Problem formulation," IEEE Transactions on Power Systems, vol. 26, no. 2, pp. 949–956, May 2011.
[96] E. Mashhour and S. Moghaddas-Tafreshi, "Bidding strategy of virtual power plant for participating in energy and spinning reserve markets - Part II: Numerical analysis," IEEE Transactions on Power Systems, vol. 26, no. 2, pp. 957–964, May 2011.
[97] D. Pudjianto, C. Ramsay, and G. Strbac, "Virtual power plant and system integration of distributed energy resources," Renewable Power Generation, IET, vol. 1, no. 1, pp. 10–16, March 2007.
[98] M. Giuntoli and D. Poli, "Optimized thermal and electrical scheduling of a large scale virtual power plant in the presence of energy storages," IEEE Transactions on Smart Grid, vol. 4, no. 2, pp. 942–955, June 2013.
[99] A. Raab, et al., "Virtual power plant control concepts with electric vehicles," 16th International Conference Intelligent System Application to Power Systems (ISAP), pp. 1–6, Sept 2011.
[100] M. Vasirani, R. Kota, R. Cavalcante, S. Ossowski, and N. Jennings, "An agent-based approach to virtual power plants of wind power generators and electric vehicles," Smart Grid, IEEE Transactions on, vol. 4, no. 3, pp. 1314–1322, Sept 2013.
[101] H. Yang, D. Yi, J. Zhao, and Z. Dong, "Distributed optimal dispatch of virtual power plant via limited communication," Power Systems, IEEE Transactions on, vol. 28, no. 3, pp. 3511–3512, Aug 2013.
[102] M. El-Khatib, J. Neely, and J. Johnson, "Evaluation of Fast-Frequency Response Functions in High Penetration Isolated Power Systems," IEEE PVSC, Washington, DC, 25-30 June, 2017.
[103] J. Neely, J. Johnson, J. Delhotal, S. Gonzalez, M. Lave, Evaluation of PV Frequency-Watt Function for Fast Frequency Reserves, IEEE Applied Power Electronics Conference (APEC), Long Beach, CA, March 20-24, 2016.
[104] J. Johnson, J. Neely, J. Delhotal, M. Lave, "Photovoltaic Frequency-Watt Curve Design for Frequency Regulation and Fast Contingency Reserves," IEEE Journal of Photovoltaics, vol. 6, no. 6, pp. 1611-1618, Nov. 2016. doi:10.1109/JPHOTOV.2016.2598275

response from autonomous DER functions are not deterministic *a priori* without detailed information about the current operating conditions, such as available PV power (e.g., from pyranometers), which will not be available for smaller DER installations. As such, it is recommended to use hybrid control when operating a VPP with variable energy resources.

4.3.2 Hybrid Control

A hybrid control structure that provides extremely fast response to grid disturbances by combining autonomous, decentralized, agent-based control with a centralized controller. When there is a grid disturbance, a portion of the RE/DER units respond to changes in grid frequency with pre-programmed autonomous control. The power change from these devices will be monitored to determine the initial VPP aggregate response. The remainder of the reserve magnitude will be made up through DR units and a centralized feedback controller for larger DER resources that have higher bandwidth communications. Conceptually, two primary considerations factor into the decision to issue the DER centralized or decentralized control commands: communication capabilities and local control capabilities. Larger DER are can generally be controlled directly whereas smaller DER—which have slower communication rates, less availability, and less bandwidth–will not receive the commands in time to act on the signal. For instance, many microinverters have 15-minute control update rates in the United States, so these devices are not suitable for centralized control. However, to use decentralized control of DER, the devices must have autonomous controls, such as droop control, frequency-watt or volt-var functions.

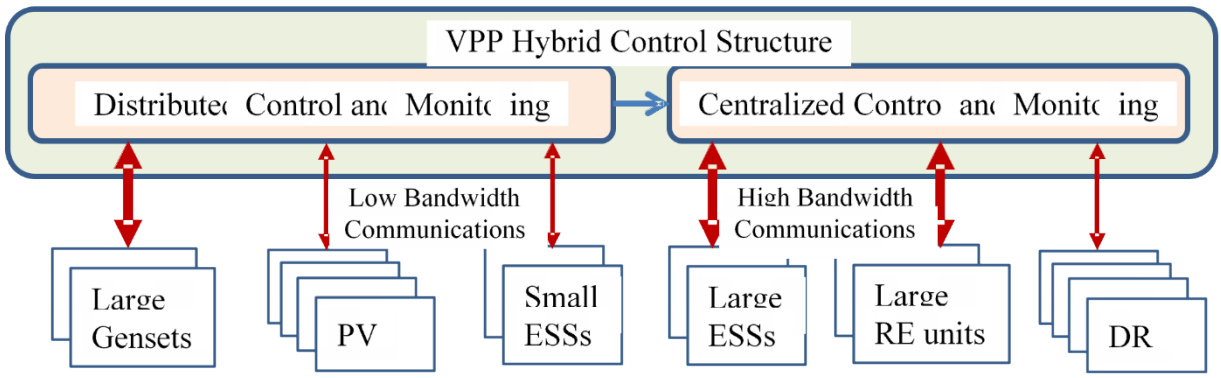

Figure 15. Hybrid VPP controller in which some DER are controlled directly and other DER are controlled using decentralized, autonomous controls.

For example, to provide synthetic inertia, centralized control schemes cannot update residential and commercial systems through Internet channels and receive feedback within the tight 2 second response time, and decentralized control cannot guarantee the aggregate reserve magnitude target or variability tolerance. In that case a novel two-step, hybrid method, should be used where thousands of distributed energy resources (DERs) are programmed with FW curves and centralized control mechanisms fine-tune the VPP response with fast DER. Every 15 minutes a stochastic optimization could be performed based on renewable generation and demand forecasts to set the appropriate distributed autonomous controls, initialize the necessary resources (e.g., charge batteries), and prepare commands for larger, faster DERs in the event of a contingency. Once the reserve is called upon, autonomous functions adjust the inverter-based

DER active power outputs based on the grid frequency and the real-time VPP dispatcher sends the optimized command set to the centrally-controlled, quicker resources (e.g., utility-scale DER) that have larger communication bandwidth. In the second step, the VPP would measure all the DER outputs and adjusts larger DER assets to meet the reserve magnitude by accounting for variations in autonomous DER responses, renewable forecasting errors, and communication failures.[105] PV, wind, energy storage systems and other inverter-based DER are pre-programmed with autonomous FW functionality to meet the initial response time requirement. Demand response could be controlled with price signal.

While there are variations in the requirements for each reserve type depending on the ISO/RTO[106], the ARPA-E NODES metrics for frequency response reserves are provided in Table 3. In order to meet these metrics, the distributed and centralized control operate together as shown by the letters in Figure 16:

(A) After the fault, the DERs programmed with autonomous frequency-watt functions or droop controls ramp up their output power;

(B) Depending on communication latency and DER response characteristics, the VPP will measure or predict the response of the distributed controllers based on the frequency deviation;

(C) The VPP issues commands to the DERs operating under the centralized control and monitors their output;

(D) The VPP ensures it remains within the tolerance. If the VPP magnitude approaches the tolerance limits the centralized controller will adjust DER outputs accordingly. This deviation could result from renewable energy or demand response changes, storage or fuel depletion, or other unexpected events.

Table 3: Frequency Response Reserve metrics for the virtual power plant.

Reserve Metric	Target Value
Initial Response Time	< 2 sec
Reserve Magnitude Target (RMT, % of load)	> 2%
Reserve Magnitude Variability Tolerance (RMVT)	< ±5%
Ramp Time	< 8 sec
Duration	> 30 sec
Availability	> 95%
Cascaded Contingency Support	> 2

[105] Note that in many cases, communications to these DER will be slow and the response may need to be estimated until DER measurement data is received.

[106] J. F. Ellison, L.S. Tesfatsion, V.W. Loose, R. H. Byrne, "Project Report: A Survey of Operating Reserve Markets in U.S. ISO/RTO-managed Electric Energy Regions," Sandia Technical Report, SAND2012-1000, Sept. 2012.

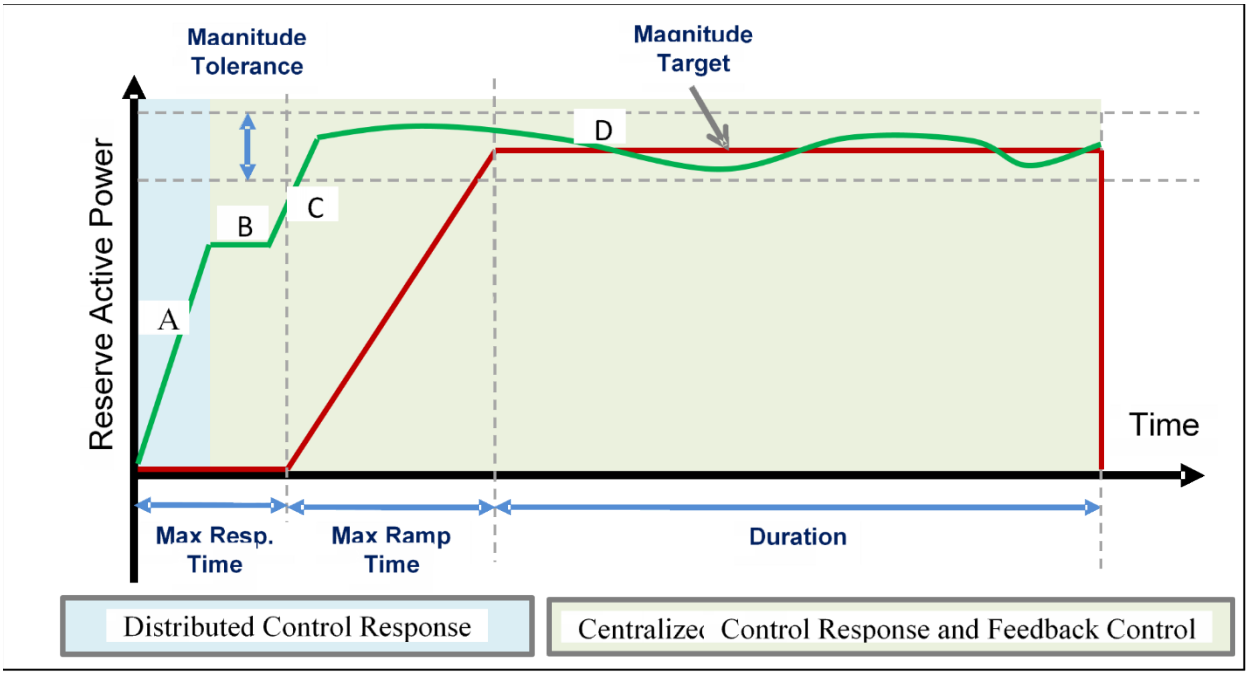

Figure 16. The reserve metrics with example VPP output.

4.3.3 Centralized Control

Ultimately, the VPP in this project used a centralized control algorithm. The main goal of the VPP control system is to ensure that the real-time output of the VPP is maintained within an acceptable error margin. This control task is challenging for several reasons. First, the presence of renewable energy DERs in the VPP cause the VPP output to fluctuate. Second, the VPP controller should compensate in real-time for the loss of any particular DER (communication failures, DER disconnection, etc.) or the inability of any DER to attain its reference power output. Third, due to the geographical diversity of DERs in the VPP, a communication network must connect the VPP controller to the DERs through public internet channels. Unlike many previous VPP implementations, DERs included in this work extend down to the residential level (e.g., rooftop microinverters on homes). The presence of the internet-based communication network introduces additional difficulties to the design of the control system because of communication latencies and data loss. In this section, we present a VPP controller design utilizing PID and proportional controllers to provide fast, reliable aggregate power production.[107]

The VPP controller receives optimal dispatch setpoints for each DER from the optimization routine at a specified interval (e.g. every 15 minutes). From this starting operating condition, the VPP controller is responsible for keeping the total output of the VPP within an acceptable error margin from the VPP reference power defined by:

[107] Also presented in M. El-Khatib, J. Johnson, and D. Schoenwald, "Virtual Power Plant Feedback Control Design for Fast and Reliable Energy Market and Contingency Reserve Dispatch," IEEE PVSC, Washington, DC, 25-30 June, 2017.

$$VPP_{ref} = E(t) + \alpha(t)R(t) \quad (1)$$

where E is the energy market commitment, R is the reserve commitment, and α is a binary variable indicating if the reserve is required at time, t.

The design of VPP controllers is challenging for several reasons. First, VPPs aggregate heterogeneous DERs with wide ranges of ramp rates which makes it hard to tune the controller and ensure stable response. Second, the controller must compensate for small variations in the VPP output due to the variability of RE DER resources and respond to changes in VPP output due to unexpected DER tripping or communication failure. Third, reliance on communication network introduces significant latencies and a probability of data loss which could destabilize the controller.

Figure 17 shows a schematic overview of the proposed VPP controller structure. The optimizer (Optimization Block) resolves for the optimal DER dispatch settings every 15 minutes to account for changes in short-term forecast and other DER status changes—e.g., loss of DER communications. These new setpoints are issued to the VPP controller to re-adjust DER reference powers. The proposed controller consists of the Feedback Controller and the Re-dispatch Processor as detailed below.

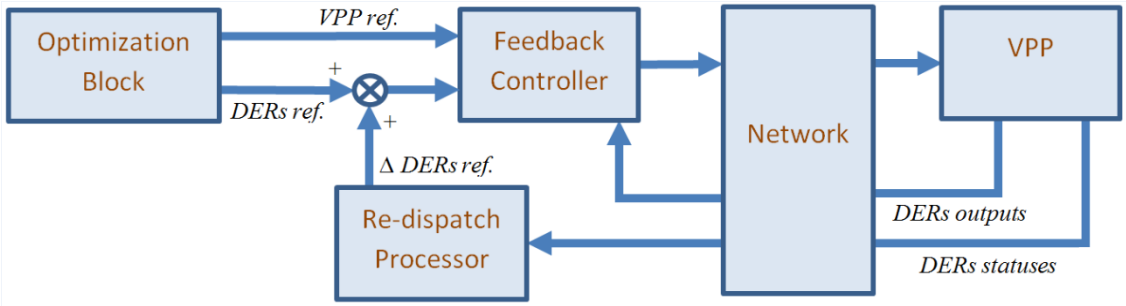

Figure 17. VPP controller consisting of feedback control and re-dispatch processor.

4.3.3.1 Feedback Controller

The feedback controller is responsible for maintaining the VPP output at the target level by compensating for changes in DERs outputs. The proposed controller structure is shown in Figure 18 for a VPP with three DERs—though this architecture can be expanded to any number of devices. The controller uses overall VPP error to derive the output of different DERs. Due to the wide range of DER ramp rates, only one DER is equipped with PID controller and the rest of the DERs are equipped with proportional gain control to avoid output ringing.

The DER equipped with PID controller is designated the swing DER of the VPP and is responsible for smoothing the output of the VPP and eliminating any steady state errors. Typically, a large storage-based DER should be used as a swing DER to ensure adequate controller response because of its fast ramp ability.

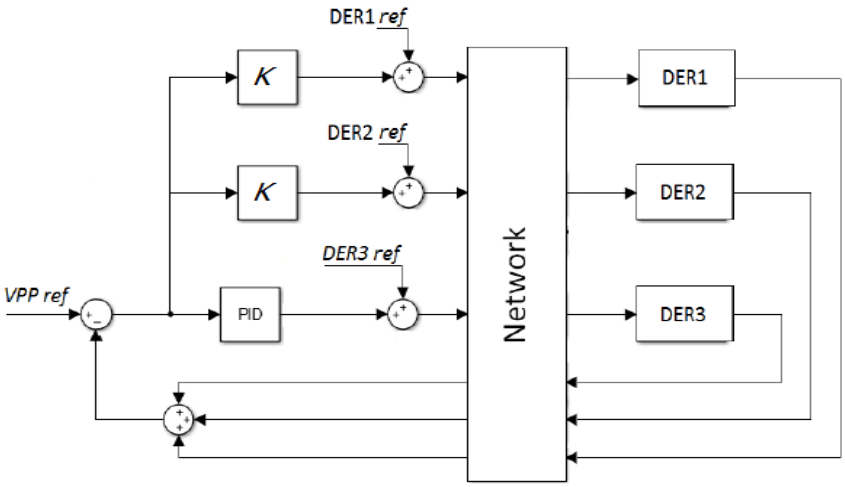

Figure 18. VPP feedback controller structure for 3 DER.

4.3.3.2 Re-dispatch Processor

During real-time operation, for large VPP errors, DERs may drift significantly from their reference powers determined by the VPP optimizer. As a result, the VPP operates in a suboptimal economical state. One possible solution to this problem is to actively re-adjust DERs reference powers in real-time to ensure that the VPP output is restored using the most economical DERs. Due to time constraints of real-time operation, it is hard to formulate and solve a complete optimization problem in the re-dispatch processor. However, the initial DER reference powers from the optimizer represents the most economical solution using DER cost curves to meet the VPP bids. Therefore, we propose to dispatch DERs proportional to their initial reference powers. In other words, if P_{error} is the difference between the VPP reference and actual powers—due to communication failures, renewable energy reductions, or tripping of DER k—then for each available DER i in the VPP, the reference output power will be updated as follows.

$$\Delta P_i = P_{error} \frac{P_{i,initial}}{\sum_{m=1,m\neq k}^{N} P_{m,initial}} \tag{2}$$

$$\Delta P_{i,new} = P_{i,initial} + \Delta P_i \tag{3}$$

where, $P_{i,\ initial}$ is the initial output power of DER i before the contingency $P_{m,\ initial}$ denote the output power of DER m. Note that, once updated reference powers are received from the optimization engine at the beginning of the subsequent optimization period, DERs will follow the new reference powers and the re-dispatch processor will be reset.

4.3.3.3 Control Simulations

In order to study the impact of different factors on the performance of the VPP controller, a simulated collection of DERs was created based on the equipment located at Mesa del Sol (MdS), Public Service Company of New Mexico (PNM) Prosperity Site, and Sandia's Distributed Energy Technologies Laboratory (DETL) in Albuquerque. The equipment at MdS and Prosperity sites was controlled previously for PV smoothing[108,109] so this collection of

devices could form a VPP with the correct control structures. The DER included in the simulations is listed in Table 4 with their size, dispatchable power levels, and swing settings. In order to create a stable VPP controller first the swing controller settings were determined and then the gain was selected for the non-swing DER.

Table 4: DER VPP Parameters.

DER	Size (kW)	Dispatchable Power (kW)	Swing?
Miller Cycle Genset at MdS	240	200	No
Diesel Genset at DETL	250	90	No
Battery at Prosperity Site	500	300	Yes
PV at Prosperity Site	500	500	No
Battery at MdS	163	140	No
Fuel cell at MdS	80	40	No
Rooftop PV at MdS	100	100	No
Eight Inverters at DETL	8 x 3	24	No

4.3.3.4 Controller Tuning

The swing PID controller for the above VPP was tuned using the Ziegler–Nichols method. The VPP scenario in Table 5 was simulated, shown in Figure 19, to illustrate the basic operation of the VPP controller with different controller parameters. For each setpoint command issued to the DERs, a delay and probabatility of packet loss were simulated. The simulation time step was set to 0.01 s but the control setpoints were only recalculated and re-issued every 0.2 seconds to represent the communication delay in sending and receiving power data from the equipment. Fig. 3 shows the performance of the VPP controller which quickly reaches the VPP power reference, but with different overshoot levels and settling times. The final swing control parameters were chosen to be $K_p = 0.7$, $K_i = 1.0$, and $K_d = 0$.

---

[108] J. Johnson, A. Ellis, A. Denda, K. Morino, T. Shinji, T. Ogata, M. Tadokoro, "PV Output Smoothing using a Battery and Natural Gas Engine-Generator," 39th IEEE Photovoltaic Specialists Conference, Tampa Bay, Florida, 16-21 Jun, 2013.

[109] J. Johnson, K. Morino, A. Denda, J. Hawkins, B. Arellano, T. Ogata, T. Shinji, M. Tadokoro, A. Ellis, "Experimental Comparison of PV-Smoothing Controllers using Distributed Generators," Sandia Technical Report SAND2014-1546, Feb 2014.

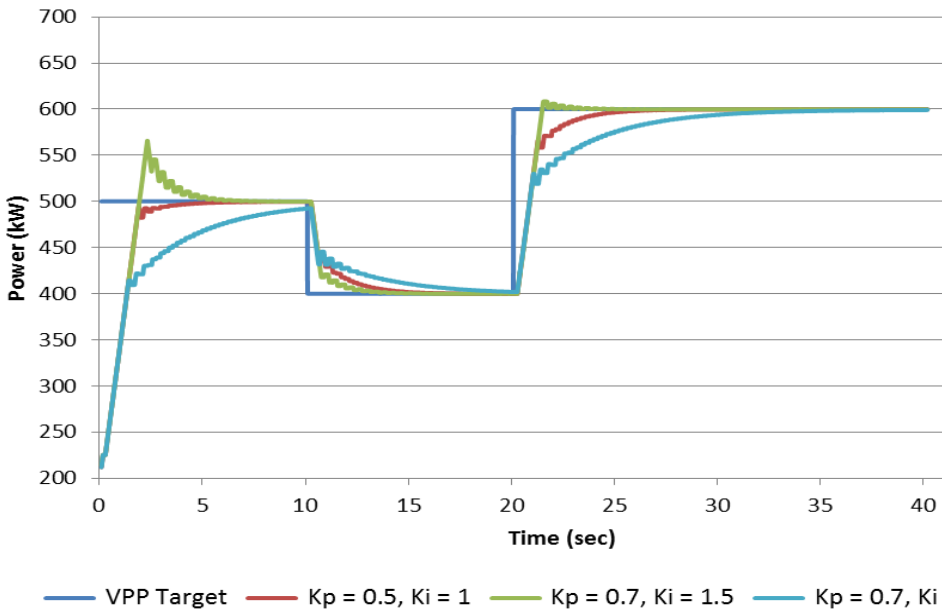

Figure 19. Swing PID parameters influence on the response of the VPP.

Table 5. VPP Operation Scenario

Time (s)	Energy Market Power Commitment (kW)	Reserve Market Power Commitment (kW)	Reserve Request, α
0	500	200	0
10	400	200	0
20	400	200	1

Once the swing controller PID settings were selected, the gain for the non-swing DER was determined. All the PV systems included in Table I were simulated by replaying one of seven 24-hour AC power 1-second datasets recorded and scaled from the 500 kW Prosperity Site PV plant. The effect of K_p gain on the VPP response is shown in Figure 20. The final non-swing DER gain was selected to be 0.1.

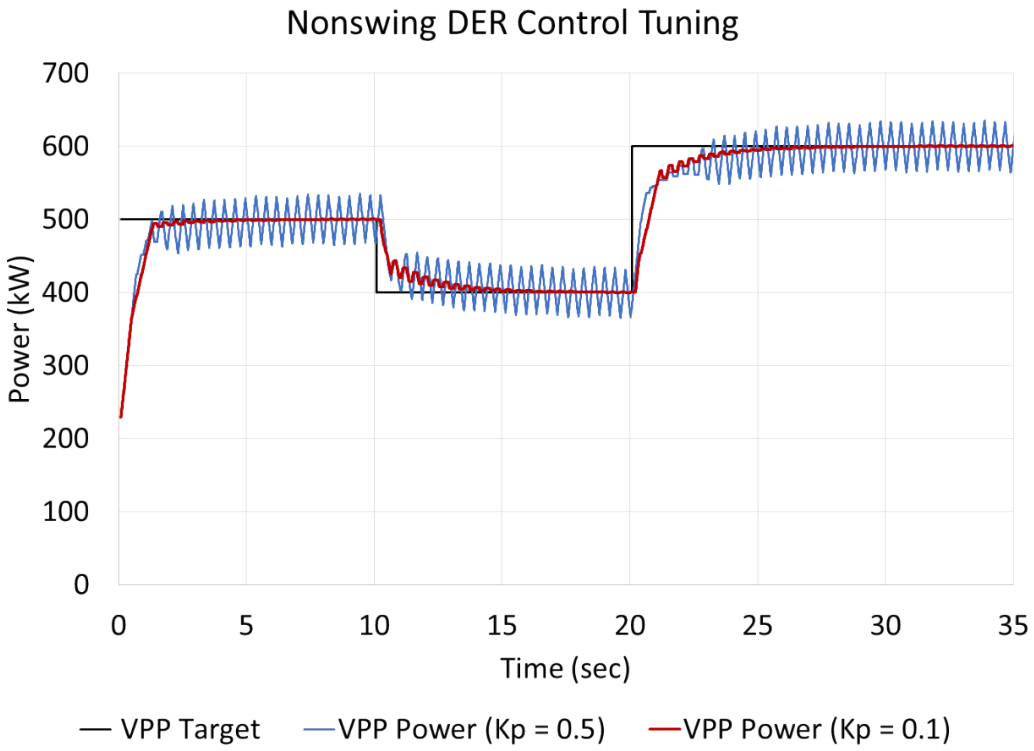

Figure 20: VPP response for two non-swing gains.

A. Impact of Communication Rate and Delay

The scenario from Table 5 was repeated with different communication rates. The control rate is the speed at which new setpoints are issued to the DER and represents the aggregate time to measure the DER outputs and issue new setpoints. Communications to physical DER devices at DETL takes approximately 0.2 sec, which does not significantly influence the stability or effectiveness of the VPP controller, as shown in Figure 21. It is clear from Figure 22 that the slower the controller rate is, the more oscillations will appear in the swing DER response as well at the VPP power.

The impact of different communication delays on the VPP response was studied as well. After control information is issued to the DER, the device does not respond for a period of time while the data packet is routed through the communication network. Depending on the transport media, communication protocol, and network topology this time could be quite short (< 10 ms) or relatively long (seconds). In the past, this was a challenge in the MdS and Prosperity PV smoothing project[110] and was ultimately a challenge for the VPP, as described below. Simulations of 100 and 150 ms delays showed the VPP controller was robust to some network

[110] J. Johnson, A. Ellis, A. Denda, K. Morino, T. Shinji, T. Ogata, M. Tadokoro, "PV Output Smoothing using a Battery and Natural Gas Engine-Generator," 39th IEEE Photovoltaic Specialists Conference, Tampa Bay, Florida, 16-21 Jun, 2013.

latency. The delay in the DER output from network delay is seen when the VPP target changes in Figure 23.

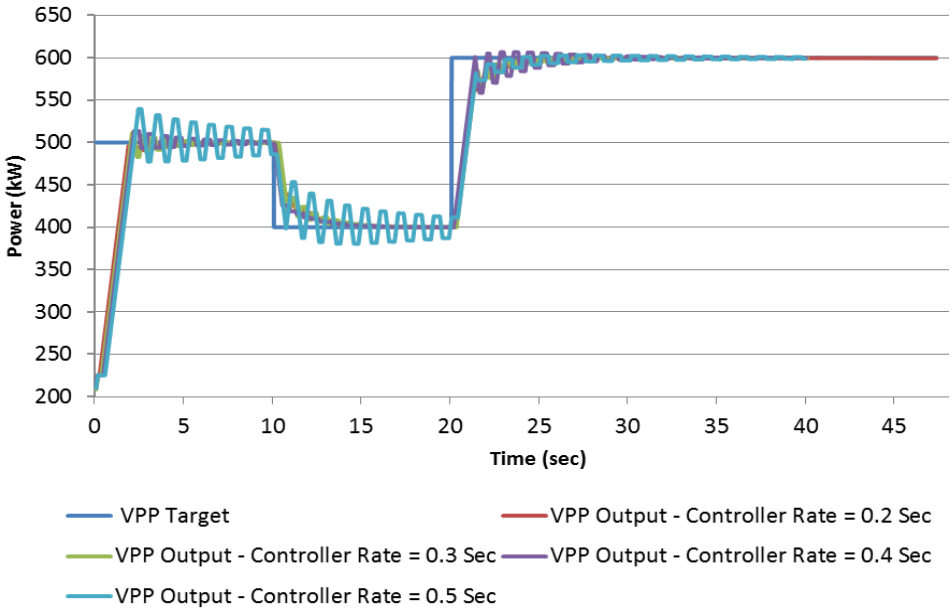

Figure 21. VPP Output under different controller rates.

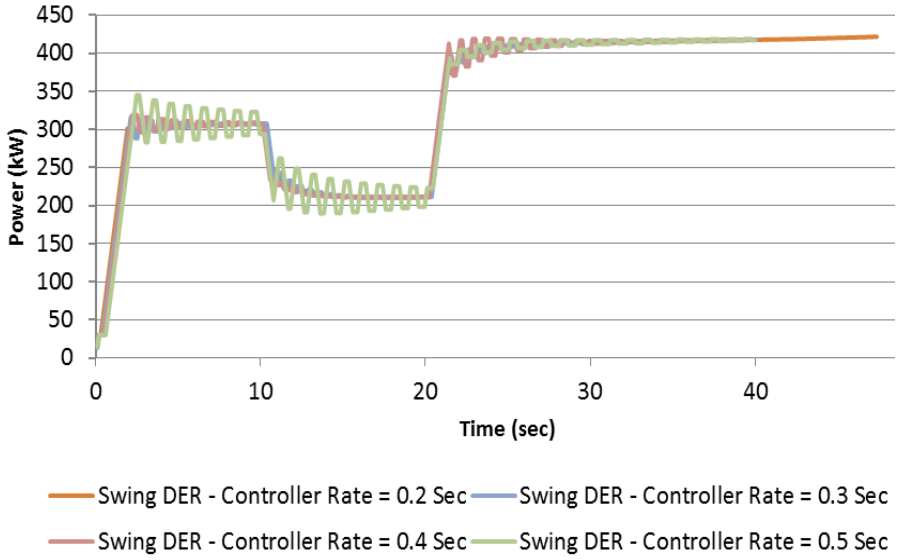

Figure 22. Swing DER output under different controller rates.

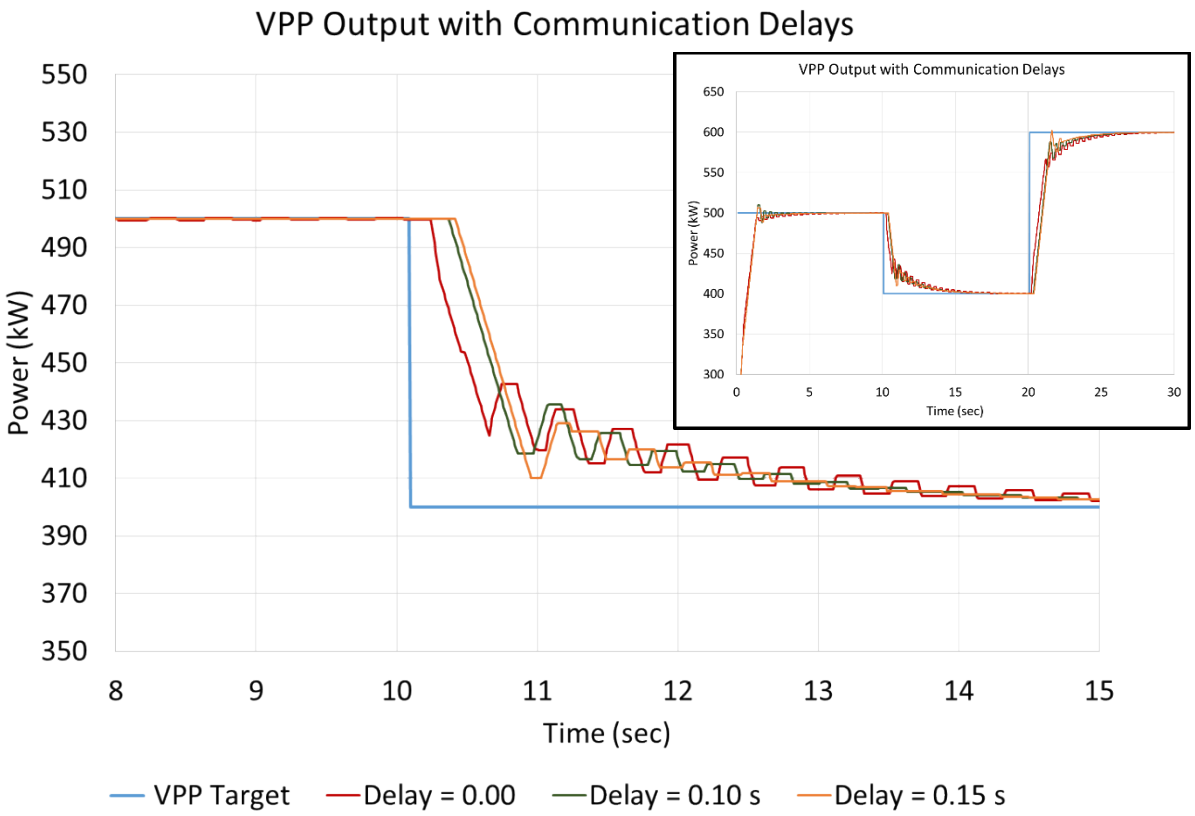

Figure 23. Influence of DER Delay on VPP output.

4.4 VPP Cyber Security

Residential-, commercial-, and utility-scale DERs are configured with VPP-optimized advanced grid functions (AGFs) settings to provide ancillary services via real and reactive power flows to the grid. Although supervisory control and data acquisition (SCADA) communications are generally proprietary stovepipe systems running on dedicated communication channels with a dependence on perimeter defenses, a fundamental challenge of this project is securing VPP communications running over public and home-area networks. The VPP construct necessitates control communications to distribution-level devices at residential and commercial facilities over wired and wireless Internet channels. The inherent unsecure nature of these communications increases the risk of an adversary potentially leveraging a VPP to control large quantities of grid-connected generators, ignore or spoof SCADA signals, manipulate bids on energy exchanges, or instigate grid instabilities. This cyber security aspects of this project established broadly-applicable VPP network architectures and intrusion detections systems for secure dispatch of interoperability commands. Consequently, novel methods for detecting, mitigating, and recovering from cyber attacks must be developed. Techniques of identifying and removing compromised/unauthorized DERs, segmenting DERs into resource pools to minimize damage in the event of successful comprise, and safeguarding the VPP from mass compromise were developed for the VPP framework.

There were three different cyber security concepts investigated for the VPP project. The first was an intrusion detection system (IDS) that correlated DER information and measurements with other power system data sources to determine when there was a high probability of end device manipulation. The second was an IDS and classification technique using SCADA data which could identify equipment whose VPP control manipulated by an adversary. Novel data correlations between network cyber information and VPP power and control data were developed to more quickly and accurately detect and mitigate compromised devices. The third was a segmentation or enclaving technique to minimize the common-mode vulnerabilities of the system by using cyber-secure network architectures. New segmentation strategies specific to VPP functions and devices were designed to cleave the VPP into resource pools that are more resistant to widespread corruption by minimizing common-mode vulnerabilities.

4.4.1 Intrusion Detection System

The IDS system uses data from DER systems and other sources of real-time data (uPMUs, SCADA, AMI, etc.) to determine when there is suspicious data from the DER. This would indicate DER data is being spoofed on the DER is being controlled by some other entity. The cyber security overlay checks only for cyber security indicators (not equipment problems). The IDS would be an auxiliary data-requesting agent on the OPC server since it reaches all DER enclaves. This data can be harvested from Modbus or any other DER protocol using Bro.

Specific plans:

- For simplicity, data interpretation is limited to MODBUS, TCP, IP, and ARP
- The system can sense and parse directionality for SuiteLink and SMB
- The system can also detect messaging that are not MODBUS, ARP, SuiteLink, or SMB
- BroBounds may be used to set up “out of bounds” conditions for MODBUS messages
- Use exiting Archimedes analytics: “moving average” and “statistical outlier”
- May add custom analytics with Python
- Eventually the system may learn what is reasonable, but the first set will be deterministic
- Tests may be solely cyber, if no physical data is involved, or cyber-physical otherwise

Types of cyber security checking:

- Identical: is the data the precisely the same as other data?
- Consistent: is the data, or its implications, consistent with other data (within some pre-programmed error thresholds)?
- Appropriate: is the data within its permissible envelope?
- Reasonable: is the data, or its implications, consistent with the expected system operational character? (Operation references values \& error thresholds are preprogrammed and conditional.)

Potential implementations (first cut):

- Identical
    - Check if MODBUS message (control or status) is the same (or even exists) from to/from OPC and DER

    - Check if ARP requests resolve the same as previous ones
- Consistent
    - check if current DER power from MODBUS agrees with last power based on frequency change and DER FW curve, e.g. unlikely.
    - check if DER voltage times current is the same as its reported apparent power
    - check if multiple DER voltage measurements at a bus are basically identical
- Appropriate
    - check if DER output (or other parameters) is greater than nameplate (plus a margin)
    - check if DER reports capability values different than installed settings
    - check if a non-BESS DER is absorbing power for any significant interval
    - Check if BESS DER average energy is slightly positive, on average
    - check if DER ramp rate is within device capability
    - monitor for location-specific un-allowed protocols on network segments (should be only MODBUS, SuiteLink, or SMB
    - check if MODBUS coils/registers are known
    - monitor for un-allowed network management messaging (routing, spanning tree, terminal connections, etc.)
    - check for MODBUS directionality on status and command messaging
    - after setup time, check for new Ethernet MAC addresses
- Reasonable
    - monitor for un-allowed post-installation changes to settings (like networking, maximum power, or ride-through settings)
    - check if solar output is measurable at midnight
    - check if DER are given conflicting FW curves
    - check if slope of FW curve change exceeds expected rate (or is inverted)
    - check if volume of FW settings changes exceeds expected rate
    - check if timing of DER FW settings changes from established patters
    - check if multiple DER are disconnected within a small time window
    - check if a host is alive, based on communications within a reasonable window

4.4.2 Cyber Attack Identification and Classification

The objective of the cyber classification tool was to leverage high-resolution and high-quality measurement data of the physical power grid to identifying manipulations of DER controls. It was observed that for certain attacks on the VPP controllers (such as disconnectivity and changes of setpoint), the power system voltage and frequency would respond differently. This enabled the creation of an intrusion detection system (IDS) by correlating the measured voltage and frequency with the signature response for each type of attack. We first introduce our developed dynamic case for the development of our IDS module. Second, we will discuss a list of intrusion scenarios and analyze the corresponding effects to the physical power system. Finally, we demonstrate the possibility of identifying the intrusion scenarios using high sampling rate measurements of the physical power grid.

4.4.2.1 Creation of a dynamic simulation case with DER devices

The MdS-Prosperity-DETL VPP was used as the basis for the analysis. Base on the type and capacity information, the Illini 42-bus test case in PowerWorld (PW)[111] was modified to incorporate these DER devices. This was a 345/138 kV transmission system with transient stability simulation capability. Note that for compatibility, the original DER capacities were scaled by a factor of approximately 1000 to match the Illini 42-bus ratings. The DER placement and capacity information is listed in Table 6.

Table 6: DER sizes and locations in the transmission simulation

Bus Number	Capacity	Resource Type	Finalized Capacity	PW Model
Apple345	500 kW	BESS	565 MW	CBEST
Dolphin345	500 kW	Inverter	565 MW	PVD1
Oak345	240 kW	GENSET	271.2 MW	GENROU
Steel138	25 kW	DR	28.25 MW	GENROU
Viking345	100 kW	Inverter	113 MW	PVD1
Viking345	80 kW	Fuel Cell	90.4 MW	PVD1
Owl138	163 kW	BESS	183.19 MW	CBEST
Oak345	Five 3 kW PVs 3 kW PV 24 kW PV Total = 42 kW	Inverter	47.46 MW	PVD1

A Python client for this power simulation was also created to be used as the real-time transmission model for hardware-in-the-loop simulations using PW Dynamic Studio server. With the PW Dynamic Studio connectivity, dynamic simulations in a HIL environment with actual DER devices can be completed.

4.4.2.2 Intrusion actions and the corresponding physical effects

Once an attacker gains access to a single DER or the centralized controller, the adversary may manipulate active or reactive power of the generators to disrupt a normal power system operation. The attacker could possibly perform the following actions on a generator: disconnect/reconnect, change active power setpoint, ramp up/ramp down active power, change voltage setpoint.

To evaluate the physical effects of such attacks, simulations of the transient response of the Illini42 case under different attacks on the distributed energy resources (DERs) were performed and voltage and frequency measurements at all the buses were measured. The following attack scenarios are considered:

- *Attack 1:* Disconnect the battery at bus Apple345
- *Attack 2:* Disconnect the solar farm at bus Viking345
- *Attack 3:* Disconnect the wind farm at bus Prairie345

[111] Illinois Center for a Smarter Electric Grid (ICSEG), "Illini 42 Tornado," http://icseg.iti.illinois.edu/illini-42-tornado/

- *Attack 4:* Ramp up active power of the wind farm at bus Prairie345: the active power is commanded to ramp from 500 MW to 1000 MW in 1 second
- *Attack 5:* Ramp down active power of the wind farm at bus Prairie345: the active power is commanded to ramp from 500 MW down to 100 MW in 1 second
- *Attack 6:* Change active power setpoint: the setpoint at the wind farm at bus Prairie345 from 500 MW to 100 MW
- *Attack 7:* Change reactive setpoint: the reactive setpoint at the wind farm at bus Prairie345 is commanded to change the bus voltage to 96% of the nominal value

For the simplicity of visualization, we show the voltage and frequency profiles in time for the Prairie345 bus in Figure 25 and Figure 26, and discuss the physical effects on the system.

- **Battery attack (*Attack 1*)**: The battery at bus Apple 345 supplies 0 MW and 44.62 MVAr. When it is disconnected under the scenario Attack 1, the system voltage immediately drops, e.g., the voltage profile at Prairie345 bus is illustrated by the dash-red line in Figure 24. Because the battery does not provide active power to the system, once disconnected, the system frequency does not vary significantly (red-dash line in Figure 25).
- **Solar farm attack (*Attack 2*)**: The solar farm at bus Viking345 provides 113 MW and 0 MVAr. When it is disconnected, the physical impacts are mostly on the system frequency, while the voltage is mildly affected, as shown by the dash-blue lines in Figure 24 and Figure 25.
- **Wind farm attacks (*Attacks 3, 4, 5, 6, 7)***: The wind farm at bus Prarie345 provides 500 MW and 250 MVAr.
    - *Attack 3*: Once disconnected, both voltage and frequency varied. The voltage at bus Prairie345 drops immediately after the disconnection and recovers thanks to the control mechanism embedded to the system. However, the voltage level could not recover to the pre-attack level, as shown by the dash-blue line in Figure 24. Similarly, the frequency immediately drops and then settles to a post-attack level, as shown by the dash-blue line in Figure 25.
    - *Attack 4 and Attack 5:* the active power is ramped up/ down. The physical effect is mostly on the system frequency, as shown by red-dash-dot and green-dash-dot lines in Figure 25.
    - *Attack 6:* the active power set point is changed. Like the *Attack 4* and *Attack 5*, the physical effect is mostly on the frequency, as shown by the blue-dash-dot in Figure 25. Compare to the *Attack 5*, the system settles to the same steady state values in both voltage and frequency. However, right after the attack, the rate of change in both voltage and frequency is higher compared to the *Attack 5*.
    - *Attack 7:* the voltage set point is changed. The physical effect of this attack is mostly observed on the voltage measurement, as shown by the red-dot line in Figure 24. The rate of change in voltage is slower than the cases where reactive power is suddenly cut-off such as in the *Attack 1* and *Attack 3*.

In general, after an attack, the voltage and frequency settle to a constant level for all the scenarios. From this point on, an attack could not be identified.

4.4.2.3 Identifying the intrusion actions using high sample rate measurements

Using high sampling rate data acquisition system, by comparing the high-resolution measurements, one could identify and classify an attack incident. Shortly after an attack incident, voltage and frequency evolves differently for different attack scenarios. For example, for the attacks that disconnects reactive power source, voltage drops immediately. If attacker changes the voltage setpoint, voltage varies gradually. In addition, the dynamics of the attacked target also play a role in the system response. Those patterns are illustrated in Figure 24 and Figure 25 and mentioned in the foregoing discussion.

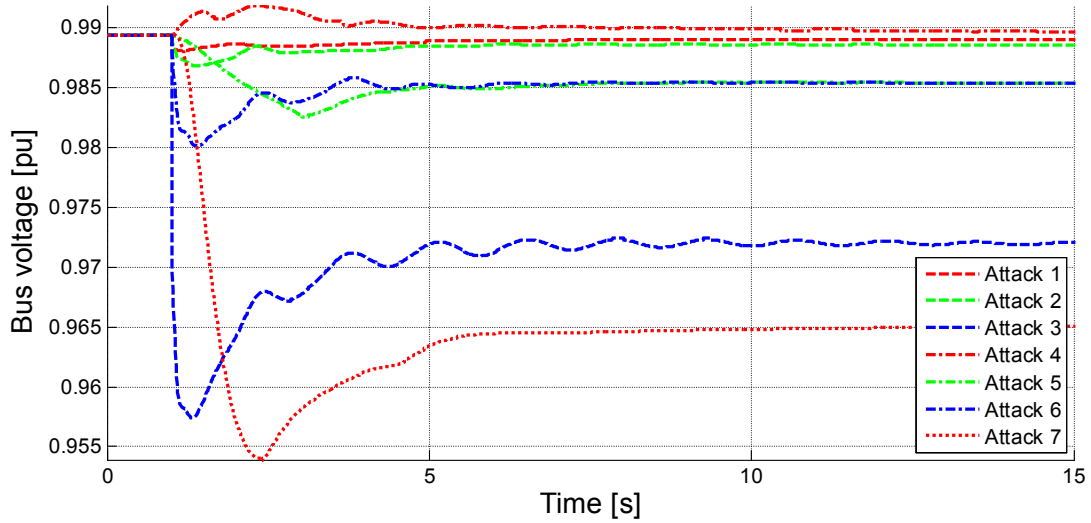

Figure 24: Voltage profile at bus Prairie345 under different attack scenarios.

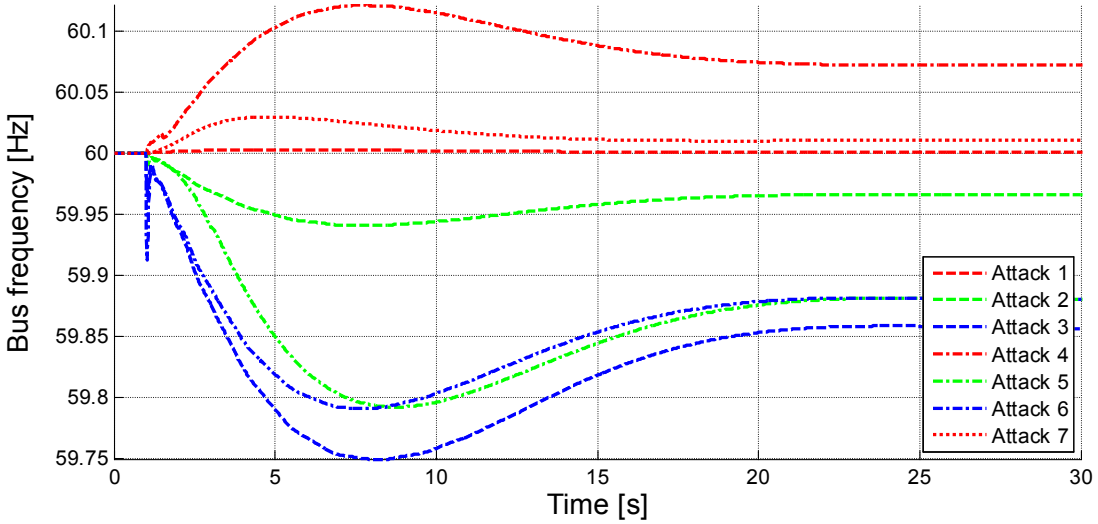

Figure 25: Frequency profile at bus Prairie345 under different attack scenarios.

By comparing the voltage and frequency profiles with the known signatures, the VPP could classify the DER that was attacked. One method of comparison is based on correlation index. The correlation index is calculated from the normalized measurements because the type of attack

is the question of interest. By normalizing the measurement, the type of attack could be identified regardless of its magnitude. Given the two sampled measurements $f[k]$ and $g[k]$, with $k = 1...N$ to be the time index, starting when an attack occurs and ending when the system has come back to a steady-state condition. The normalized measurements are calculated as:

$$\bar{f}[k] = \frac{f[k] - f_{ss}}{max(|f[k] - f_{ss}|)}$$

where f_{ss} is the steady-state value of f, max is the maximum value operator, and $|\cdot|$ is the absolute value operator. The steady-state value could be calculated by average several measurement samples once the system has got into a steady-state condition. The measurement $g[k]$ is normalized using the same method. We use the correlation index for the measurements f and g defined as follow:

$$c_{fg} = \frac{\sum_{i=1}^{N} (\bar{f}[i] - f_0)(\bar{g}[i] - g_0)}{\sqrt{\sum_{i=1}^{N} (\bar{f}[i] - f_0)^2} \sqrt{\sum_{i=1}^{N} (\bar{g}[i] - g_0)^2}}$$

where $f_0 = \frac{1}{N}\sum_{i=1}^{N} \bar{f}[i]$, $g_0 = \frac{1}{N}\sum_{i=1}^{N} \bar{g}[i]$ are the average values of f and g.

The correlation indices for voltage and frequency measurements at bus Prairie345 under the listed attack scenarios are illustrated in Figure 26 and Figure 27. Higher correlation coefficient indicates high similarity between the two measurements. If the value is 1, the two normalized measurements are identical. If the value is -1, they are opposite but could have the same nature, e.g., the attacks of increasing and decreasing the same set point by the same amount. If the value is 0, the two attacks are likely to be different in nature, e.g., an attack that changes voltage setpoint versus the one that changes active power setpoint.

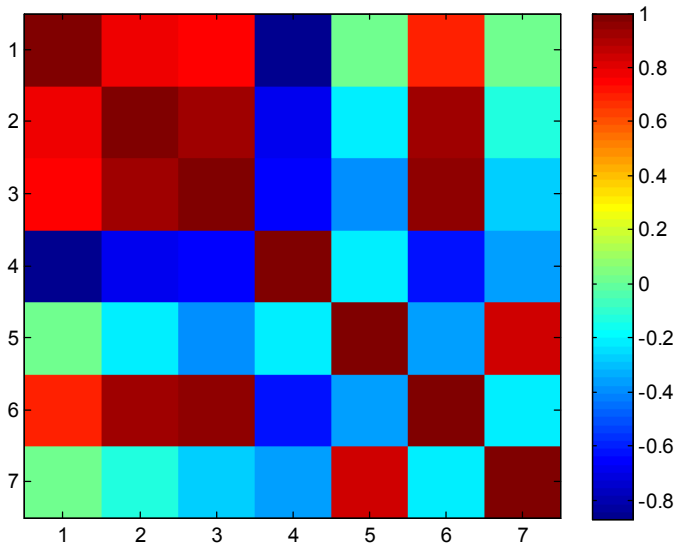

Figure 26: Voltage correlation indices for the measurements at bus Prairie345.

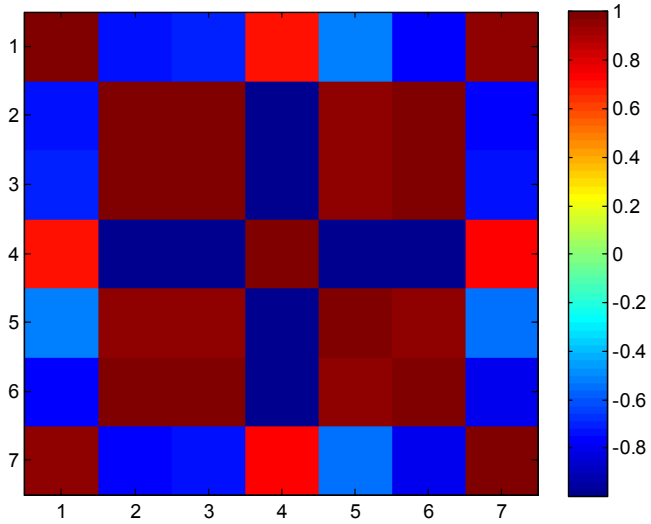

Figure 27: Frequency correlation indices the measurements at bus Prairie 345.

Combining the voltage and frequency correlation indices, the VPP could detect, differentiate, and identify the attack type. Once the typical measurements for attack scenarios are studied in simulation *a priori* and a classification tool is developed (e.g., based on Bayesian networks, neural networks, etc.) the VPP could classify new attack event. For example, from the voltage correlation indices, the Attack 1 could be clearly differentiated from the Attack 6 and Attack 7 because the correlation indices at position (1,6) and (1,7) are small. There could be other scenarios that the voltage correlation indices alone could not clearly differentiate the two attacks. In this case, using additional indices, e.g., frequency correlation indices, could help. For example, the voltage correlation indices between the Attack 1 and Attack 2 is high, which indicates that these two attacks could be similar. However, the corresponding frequency correlation index is low. Which helps predicting that those two are different. Indeed, they are

different because the Attack 1 affect the reactive power while the Attack 2 is on the active power.

4.4.2.4 Outlook

We have shown the capability of simulating the physical effects of a cyber-attack on a power system and the concept of using transient measurements to classify an attack. From the simulation data, different attacks could result in different transient effects on voltage and frequency (and other) measurements. Observing those signatures, the VPP could differentiate several types of attack. The similarity between the two measurements is quantified by a correlation index so that the classification process could be automated. Using this index for different types of measurement (e.g., voltage, frequency, etc.), we have conceptually shown the possibility of differentiating the attacks but an automated classification system is necessary to quickly identify the DER or DERs which are no longer controlled by the VPP or trustworthy; at which point the controller can be updated to compensate for these losses in the VPP DER pool.

For future developments, a focus on developing a more complete and realistic attack scenario, a more versatile similarity measure, and a methodology to use measurement at more than one bus to perform attack classification should be considered. Currently, the attack scenarios only consider a single action. However, an attacker is not limited to only a single action in a real attack. A more versatile similarity measure is necessary to identify the attacks that could happen in various time scale. For example, with the active power ramping attack, the attacker could select the ramp rate to create different attacks. The defined correlation index would identify those attacks to be very different although they are from the same nature. Further, using a moving average or only a single time stamp of data would accelerate the cyber attack classification speed. Lastly, although looking at a single bus (Prairie345 in the example above) gives sufficient information to classify the given attack, more information may be required to classify more complicated attack schemes. A systematic way to pick measurements at multiple buses to identify the attacks would be valuable in this case.

4.4.3 DER Enclaving

To provide resilient grid services disruptions to the DER network must be isolated. Selecting the appropriate cybersecurity architectures is essential to maximizing control network security and minimize the impact of an adversary. Through the use of firewall rules, enclaves can be created which segment the DER assets. If an adversary gains access to one of these enclaves through physical or remote hacking, they are only capable of changing the settings on the devices that are isolated into that enclave. Therefore, by segmenting the VPP assets into a large enough collection of enclaves, the VPP will be robust to cyber security attacks of this nature. However, this feature comes at the cost of increased network complexity, maintenance, and hardware.

For this project, three different network topologies were designed:

- **Flat network** with one enclave
- **Segmented network** with DERs randomly placed in X enclaves.
- **Critically segmented network** with the aggregate nameplate capacity of the DER in the enclaves not exceeding 20% of the total VPP capacity.

Cyber-secure communications and resilient controls were created for the example system described previously that includes DERs at the Distributed Energy Technologies Laboratory (DETL), PNM Prosperity Site, and Mesa Del Sol (MdS) Aperture Center by aggregating renewable and traditional energy DERs. The critical segmented network is shown in Figure 28.

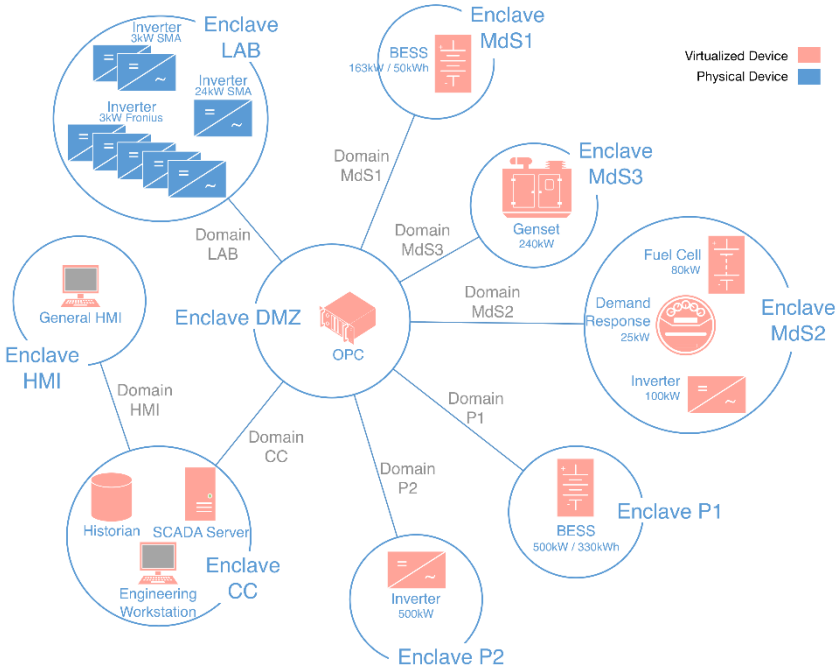

Figure 28: Example Cyber Reference Architecture which enclaves DER devices to minimize common-mode vulnerabilities. In this configuration if an adversary gains access to one of the enclaves, they cannot control the utility/aggregator power devices in the other enclaves—reducing the risk of widespread grid failures.

5 VPP SOFTWARE STRUCTURE, TIMING, AND VISUALIZATION

The VPP system was constructed as a multi-processing Python environment in which each of the VPP components could execute asynchronously. The design and integration of the software components to create a functioning VPP system ultimately accounted for a large portion of the project effort.

5.1 Software Architecture

A master VPP Python script initiated the VPP operations. This script contained a command line interface (CLI) that:

- provided lists for available CLI commands
- displayed the python version, compiler, and other computer information
- launched the VPP by creating a collection of DER objects contained in memory (lists of PV, fuel cell, gensets, batteries and demand response assets) and then initiating the multiprocessing manager. The manager created a number of multiprocessing servers which read/wrote to the DER objects. These servers represented the components in Section 4:
    - Long-Term Forecasting Server: Generates a 24-hour forecast with 1 hour intervals for the VPP resources based on NOAA weather data and saves this information for each PV inverter.
    - Short-Term Forecasting Server: Generates a forecast for the VPP resources based on persistence or autoregressive moving average (ARMA) model and saves this information for each PV inverter.
    - Bid Server: Co-optimizes the VPP into to maximize profit by bidding into energy and ancillary service markets (e.g., tertiary contingency reserve market).
    - Cybersecurity Server: Determines which DER objects cannot be trusted to deliver power based on data correlations between their reported status and other data streams (AMI data, SCADA data, etc.).
    - Dispatch Server: Determines the DER dispatch schedule to meet the energy/ancillary service commitment(s) while minimizing costs to the VPP. These setpoints are used by the control server as the initial power setting for the individual DER PID controllers.
    - Control Server: Quickly moves the VPP aggregate power to the power commitment using PID controllers for the DERs and a swing DER (typically a fast battery) to make up the difference.
    - Report Server: Sampled DER power and plotted the output (redacted with export server).
    - Export Server: Sends DER data to the Elasticsearch database for visualization within Kibana.
- displayed server status (e.g., if the processes are alive or dead)
- changed the target VPP power target manually for debugging purposes
- killed the servers

Each of the servers had the ability to get DER and VPP information from the DER objects through read and write functions. In cases where the server needed to write data to the DER objects, it would lock the DER shared memory objects to prevent other servers from overwriting this information. In general, this method was found to effectively exchange data between the servers so long as the control server which issued commands to the DER devices did so internally and the export server only dumped information back to Elasticsearch every 100+ control loop iterations.

5.1.1 Shared DER Objects

Information about the state of the VPP system was exchanged through shared memory elements representing each of the DER devices. These DER instances contained generic information that all DER assets had, such as:

- DER name
- Initialization time and date
- Simulation or real DER Boolean
- Nominal voltage and frequency
- Nameplate power, reactive power, apparent power
- Ramp rate
- Geographic information (Latitude and Longitude)
- DER communications information
    - Communication type, TCP/IP, Serial/Modbus RTU, etc.
    - IP address, IP port, timeout
    - Communication ID, parity, baud rate, timeout
- Expected communication delay
- Target power
- Control parameters (k_p, k_i, k_d), swing boolean
- Export server queues so that data was not lost during control operations
- Bid optimization parameters (interval, optimization horizon, etc.)
- Dispatch optimization parameters (interval, optimization horizon, etc.)
- Bid optimization schedule for VPP for energy and reserve
- Dispatch schedule for VPP
- Dispatch schedule for DER
- Cyber security warning flags (Booleans)

The DER devices also included methods for:

- Getting VPP target and DER target for the bid (energy and reserve) and dispatch for any time, *t*, or current time from the Pandas dataframe
- Setting control settings
- Setting DER power using absolute values or per unit

There were also specific attributes and methods provided for each of the DER types:

- PV
    - Available power
    - Forecasting inputs
        - Inverter model

        - Module information
        - Number of inverters
        - Site info
        - Historical data
    - Forecasting outputs
        - Long term forecast dataframe
        - Short term forecast dataframe
    - Methods for getting forecast data from the dataframes for given times
    - For simulated PV devices, an 'actual' 1-second PV output profile was generated for 24 hours using scaled power from historical data from a utility PV installation. One of five days of historical data could be selected for the PV profile to represent sunny, cloudy, or partly cloudy conditions.
    - Methods were created to read and write to SunSpec-compliant PV devices
        - DER grid-support functions and settings via pysunspec
- Engine-Generators (Gensets)
    - Cost curve pieces for the optimization
    - Initial operating condition (on/off)
    - Minimum up/down time
    - Start up and shut down ramp limits
    - Fuel costs
- Battery
    - Energy capacity
    - State-of-charge (SOC)
    - Operating SOC limits
    - Cost curve pieces for the optimization
- Fuel Cell
    - Cost curve pieces for the optimization
    - Initial operating condition (on/off)
    - Minimum up/down time
    - Start up and shut down ramp limits
    - Fuel costs
- Demand Response
    - Hot water heater
        - Temperature
        - Mass
        - Thermal conductivity, k
        - Duty cycle
        - Available power
        - Forecasting method (incomplete)
        - 1st order thermal model
        - Set temperature method
        - Read temperature

The DER Object Class also included a number of methods which provided:

- aggregate power of the VPP
- aggregate nameplate power of all the DER

- aggregate forecast power of all the DER
- projected power of the VPP over a given time horizon
- initialization of the DER fleet
- pre-solved dispatch schedule (for debugging)

Each of the servers then interacted with these DER devices to change their attributes and call on the methods to perform their operations.

5.1.2 Long-Term PV Forecasting Server

The long-term PV forecast server completed the following actions:

1. long_server creates an instance of LongFcstObj for each PV system to contain data and methods for forecasting the system's performance.
2. The local archive of NOAA NAM forecasts is updated using a script GetNOAAWeatherForecast_multisite_VPP.py, which requires python 2.7
3. long_server.forecast() loops over the list of LongFcstObj performing the following:
    a. measured data for the system is retrieved from the PI server using methods and functions in PI_access_tools.py
    b. clear-sky power is computed and forecast
    c. missing data are interpolated and forecasts missing from previous periods are generated
    d. forecasts for the next period are generated
    e. forecasts are written to csv files, to each DER's object, and for visualization.
4. After each PV system's forecast is updated, a forecast of aggregate power is compiled and written by summing over the forecasts for each PV system.

5.1.3 Short-Term PV Forecasting Server

The short-term server operated using the following sequence:

1. short_server creates an instance of FcstObj for each PV system to contain data and methods for forecasting the system's performance.
2. short_server.forecast() loops over the list of FcstObj performing the following:
    a. measured data for the system is retrieved from the PI server using methods and functions in PI_access_tools.py
    b. clear-sky power is computed and forecast
    c. missing data are interpolated and forecasts missing from previous periods are generated
    d. forecasts for the next period are generated
    e. forecasts are written to csv files and to each DER's object.
3. After each PV system's forecast is updated, a forecast of aggregate power is compiled and written by summing over the forecasts for each PV system.

5.1.4 Bid Server

The purpose of the Bid Server is to establish an optimal balance between energy production and ancillary tertiary reserve that maximizes profit for the following day. This server is designed to reach a solution for the day-ahead bulk energy and reserve levels for the VPP before 6:00 p.m. every day (or a custom time determined by when the operating authority requires day-ahead

bids). The outcome of this server are energy and reserve levels for the entire VPP for each hour of the following day without regard to the specific schedule of individual assets.

The inputs for the bid server are:

1. projected hourly energy and reserve prices (from NYISO)
2. historical data on reserve events (from PJM)
3. historical power production and forecasts for stochastic assets (wind and solar)
4. day-ahead forecasts for stochastic assets
5. detailed operational information for thermal generation and battery assets

The basic process of the bid server is to:

1. formulate scenarios for stochastic assets using Sandia's Prescient software
2. co-optimize the energy/reserve levels for each scenario and blend using CVaR

5.1.4.1 Scenarios Generation for Stochastic Resources

The incorporation of variable and intermittent resources (i.e., wind and solar) and energy storage is the most challenging aspect of VPP optimization. Variable resources, by their nature, are not baseload power and may provide more or less power than is forecasted. This difference in power production vs. forecasted power can either be stochastic (due to cloud cover) or systematic (due to deficiencies in the forecasting algorithm or the physical asset). Sandia has been investigating the optimization of stochastic resources for years and has developed a software program, Prescient[112], which is designed to develop weighted scenarios for power production of stochastic resources based on historical vs. forecasted power production (Figure 29).

[112] Sandia Press Release, "Sandia Develops Stochastic Production Cost Model Simulator for Electric Power Systems," 13 Nov 2015. http://energy.sandia.gov/sandia-develops-stochastic-production-cost-model-simulator-for-electric-power-systems/

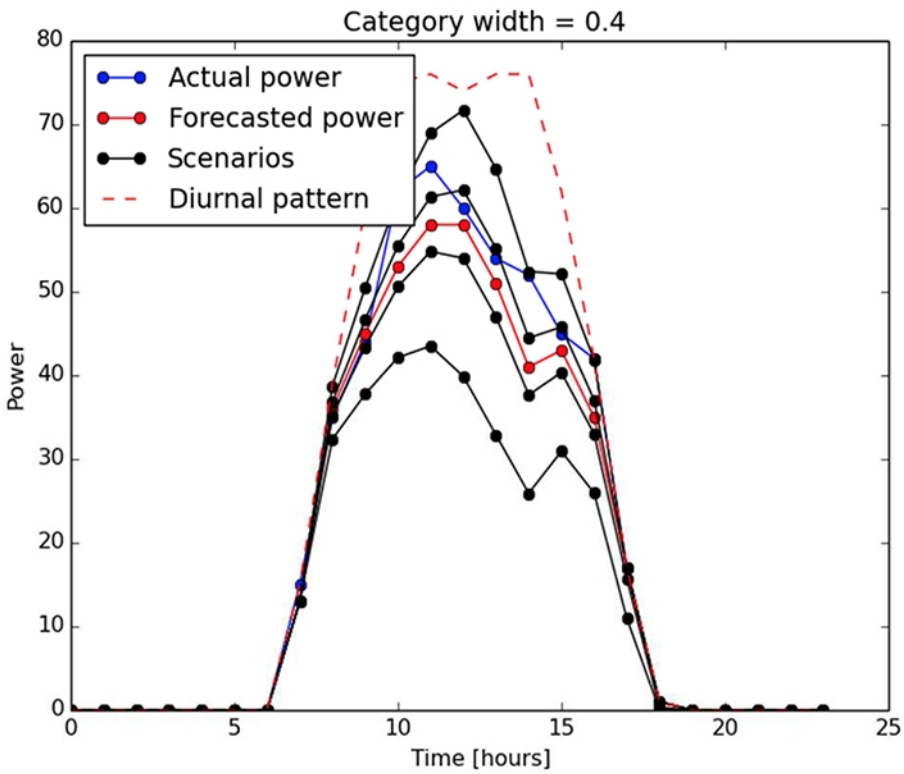

Figure 29: Scenarios for stochastic PV resource (black trace) based on forecasted power (red trace)

To generate scenarios for stochastic resources in the DER resource list (§5.1.1), Prescient requires three basic pieces of information: a history of hourly forecasted power production, a corresponding history of actual power production, and an upper bound of the power production (known as the diurnal pattern). To run Prescient on day N-1 for scenarios on day N, the bid server provides historical power production, forecasted power production, and upper bound power production for days 0 through N-1 and a forecasted power production for day N. The longer the history of forecasted and actual power production, the more accurately Prescient can construct scenarios which accurately reproduce the error distribution.

Prescient then compares the historical hourly error between the forecasted power production and actual power production to determine the error distribution for any given hour. Prescient samples this distribution to develop individual scenarios for each stochastic resource. The greater the historical record of an individual asset, the more accurate the error distribution, especially regarding tail-end probabilities.

The historical forecasted power production from the forecast server (§5.1.2) is collected and appended daily in a csv file. This historical record for each stochastic resource grows, and becomes a more accurate representation of the stochastic error distributions, the longer an asset participates in the VPP. In addition, this csv file contains historical power production for each asset. For the purposes of the LDRD, this historical data was collected daily from the Mesa del Sol PI server and then scaled to the size of each individual asset. Both the historical power production data as well as the forecast server produce time stamps in GMT, while Prescient

requires local time stamps. In the VPP software, the forecast server time is converted to local time before running Prescient.

The scenario generation is bounded by a maximum hourly power production for each asset, known as the diurnal pattern. This upper bound is computed by the forecast server based on the estimated clear-sky power production for a given asset.

With the required information, the bid server creates a bash script for calling Prescient as well as the file locations for the information Prescient provides. Upon executing the bash file, Prescient develops Pyomo-compliant files for the different projected scenarios for each stochastic asset. The total number of scenarios scales quickly, as A^s, where A is the number of asses and s is the number of scenarios per asset. Therefore, computing time/resource can quickly become constrained as the number of stochastic assets or number of scenarios per assets increases.

In order to limit computing resources, the bid server aggregates all VPP stochastic resources into a single asset. In the bid server process this results in relatively little loss in granularity, since the bid server is unconcerned with individual schedules and only concerned with power availability of the bulk VPP. Additionally, this allows for a far greater number of scenarios to be generated than would be possible if the assets were considered individually.

5.1.4.2 Forecasted Energy/Reserve Prices

To determine the profit from energy/reserve, it is necessary to have estimates for prices for tomorrow's energy or reserve. While price forecasting is a research topic unto itself and out of the scope of this LDRD, some ISO's do provide limited estimates of day-ahead energy and reserve prices. In order to get a basic input for tomorrow's prices, we scrape NYISO estimated future prices.

When the bid server is executed, the day ahead market locational marginal price (LMP) for energy and the ancillary service prices are pulled from NYISO's Market Information System. These prices are posted daily at roughly 11AM Eastern Time. Since LMP requires a location, the bid server pulls the information for the NYISO "CENTRL" zone, which encompasses the central New York region.

5.1.4.3 Thermal/Battery Resource Information

Thermal and battery resources are characterized by a variety of parameters (e.g. startup rate, maximum power, SOC) that bound their abilities to provide power. The formalization of these parameters is included in Appendix A. This information is communicated to the optimization engine via a skeleton file. This skeleton file is created automatically from the resources listed in the shared DER objects and noted as able to participate in VPP functionality (e.g. marked as available by the cybersecurity server).

5.1.4.4 Optimization

The bid optimization formalism is fully described in (§4.2.2). Briefly, the optimization of the bid server utilizes the Pyomo stochastic optimization toolbox for python to solve the optimal energy/reserve levels for a given hour over a 48-hour horizon subject to CVaR considerations.

These energy/reserve levels act as VPP targets the dispatch server and control server are required to allot to subresources (DER) in the VPP.

5.1.5 Dispatch Server

The dispatch server is an optimization server that partitions the required energy/reserve levels to targets for all VPP subresources to minimize cost. This optimizer is executed every 15 minutes and hands the optimal energy/reserve setpoints for each resource to the control server. In between 15 minute runs of the dispatch server, the control server may deviate from the optimal setpoint path while maintaining the target power.

The dispatch server has the same basic framework as the bid server. It utilizes Prescient to develop scenarios followed by Pyomo to optimize the solution. The integral difference between the two optimizations is the time-step and horizon. The dispatch server operates in 15-minute time steps (compared to 1 hour for the bid server) and a 24 horizon (as compared to a 48-hour bid server horizon) to prepare devices for the future requirements. Additionally, since the purpose of the dispatch server is to allot energy/reserve setpoint targets to each subresource so that the aggregate energy/reserve power levels for the VPP are met (while minimizing the cost of providing energy/reserve), the dispatch server does not aggregate stochastic (renewable) subresources.

The prescient software automatically utilizes 1 hour time-steps for the creation of scenarios, regardless of the time-step of the input forecasts. Therefore, in order to optimize at 15-minute time-steps, the 1-hour scenarios from Prescient are interpolated into 15-minute increments. Similar time-step scaling in required for the non-stochastic (thermal or battery) resource parameters provided by the DER resource list.

Once the inputs are scaled at 15-minute increments, the Pyomo optimization is carried out to determine the ideal setpoints for each subresource for each 15-minute increment. The setpoint for the following 15-minute increment is sent to the control server for control of each resource.

5.1.6 Control Server

The control server operated rapidly to maintain the VPP near the target power via the following sequence:

1. The optimal DER power setpoints were collected from the DER object shared memory.
2. The PID control parameters were configured for the DER and the swing DER.
3. Start loop:
    a. For all simulated DER devices, the output power was calculated from available PV power in the historical dataset for the give time of day
    b. For each DER, the power setpoint is calculated from the PID control parameters and the VPP error
        i. For physical devices, this value was issued to the DER through the communications network and the DER power was measured
        ii. For simulated devices:
            1. the setpoint is sent to a network simulation tool that delayed the signal and dropped packets with a given probability, and
            2. the ramp rate of the DER was also simulated

3. power limit constraints were also preserved

c. Every X times through the controller, data would be printed to the screen and sent to the DER objects shared memory so it could reach the export server.
d. Every Y times through the controller, the optimal DER power setpoints were gathered from the DER object shared memory.

4. Return to loop start at Step 3.

5.1.7 Export Server

The export server was created to interface with Elastic Stack (Elasticsearch, Logstash, and Kibana) to provide real-time visualization tools to monitor the behavior of the system. DER data that was exported to the Elasticsearch database included:

- DER names
- DER power
- DER nameplate power
- DER available power
- DER energy dispatch setpoint
- DER reserve dispatch setpoint
- DER compromised from cyber security warning flag
- Forecasts for all PV systems
- VPP target power
- VPP actual power
- VPP power error
- VPP nameplate power
- VPP available power
- VPP reserve target
- VPP energy target
- VPP reserve bid
- Meta data including step time, export server size, DER list size

Example visualization screenshots are presented in Section 5.3.

5.2 VPP Timing

The asynchronous execution of the VPP servers occurred at different times and at different rates depicted in Figure 30. The long and short-term forecasting ran regularly (with specified sleep cycles) to update the DER forecasts for use with the bid and dispatch servers. The bid server ran once a day around 4:00 to issue the offer to the ISO/RTO at 5:00 or 6:00 PM as required for the given market. DER optimization (e.g., the dispatch server) ran during the day of service (day N) and the control server ran continuously to meet the bid server (and VPP) market requirements. The cyber server operated continuously to update the trustworthiness metric of the DER assets and inform bid and dispatch server if certain DER equipment should not be included in the optimization.

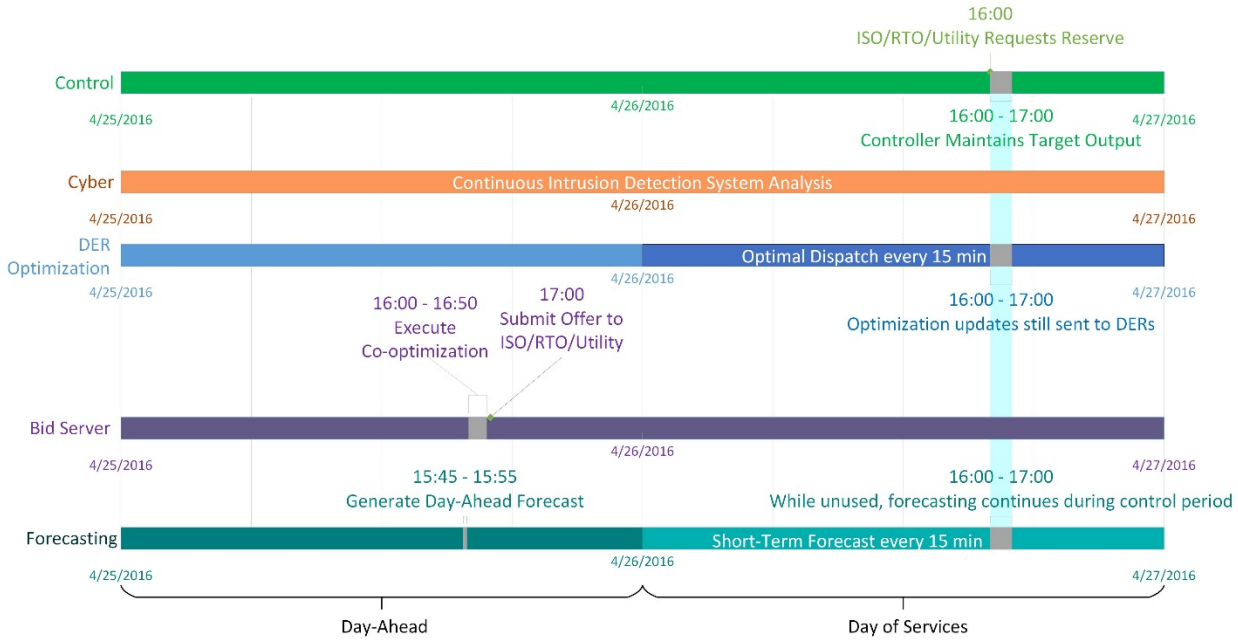

Figure 30: Example two-day operation of the VPP when bidding into the contingency reserve market. In the case of co-optimization, the energy market demand is met through the 24-hour period of the 'Day of Services'.

5.3 Visualization

Real-time visualization of individual DER and total VPP statuses is necessary for debugging purposes and for grid operators to gain confidence in the system. Kibana takes data from Elasticsearch and graphs it through a web interface, allowing easy exploration of the data and exploring past logged data.

The "metric display" shows, for each DER in the system, the nameplate capacity, the current output (smoothed over the last few seconds), and the percentage of nameplate capacity currently being produced, as seen in Figure 31. This is the most precise readout in the system, giving the operators accurate information about the state of the DERs.

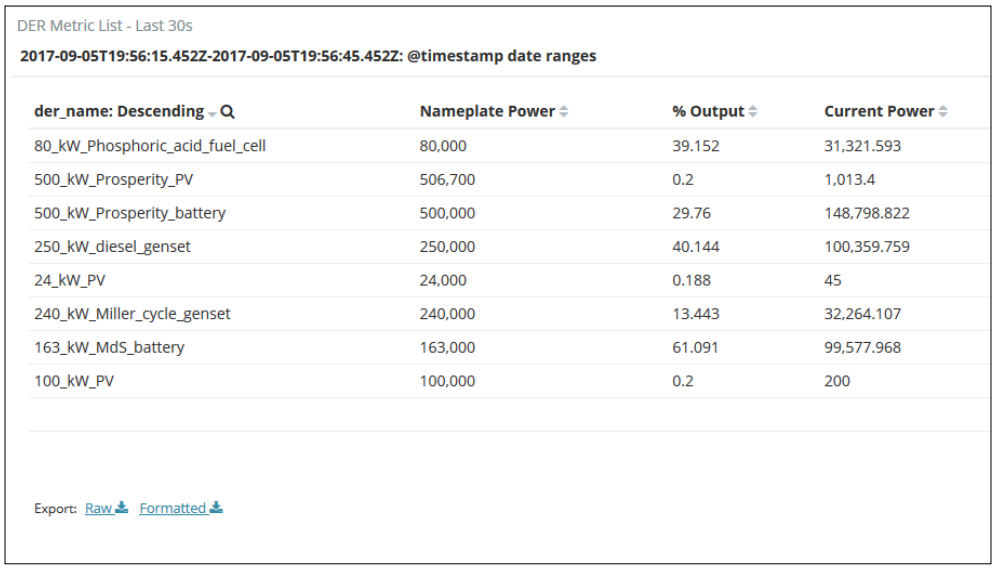
DER Metric List - Last 30s

2017-09-05T19:56:15.452Z-2017-09-05T19:56:45.452Z: @timestamp date ranges

der_name: Descending	Nameplate Power	% Output	Current Power
80_kW_Phosphoric_acid_fuel_cell	80,000	39.152	31,321.593
500_kW_Prosperity_PV	506,700	0.2	1,013.4
500_kW_Prosperity_battery	500,000	29.76	148,798.822
250_kW_diesel_genset	250,000	40.144	100,359.759
24_kW_PV	24,000	0.188	45
240_kW_Miller_cycle_genset	240,000	13.443	32,264.107
163_kW_MdS_battery	163,000	61.091	99,577.968
100_kW_PV	100,000	0.2	200

Export: Raw Formatted

Figure 31: Kibana metric display screenshot.

Another useful display is the stacked power graph shown in Figure 32. This graph stacks the power from the DERs on top of each other, allowing a quick visual display of how much of the current output power each DER is producing (as well as the total system power). This view works acceptably if all DERs are exporting power, but the stacked graph concept (as implemented by Kibana) is not well suited to displaying data from battery systems, as they can be charging as well as discharging. This graph type does not work well for this type of data.

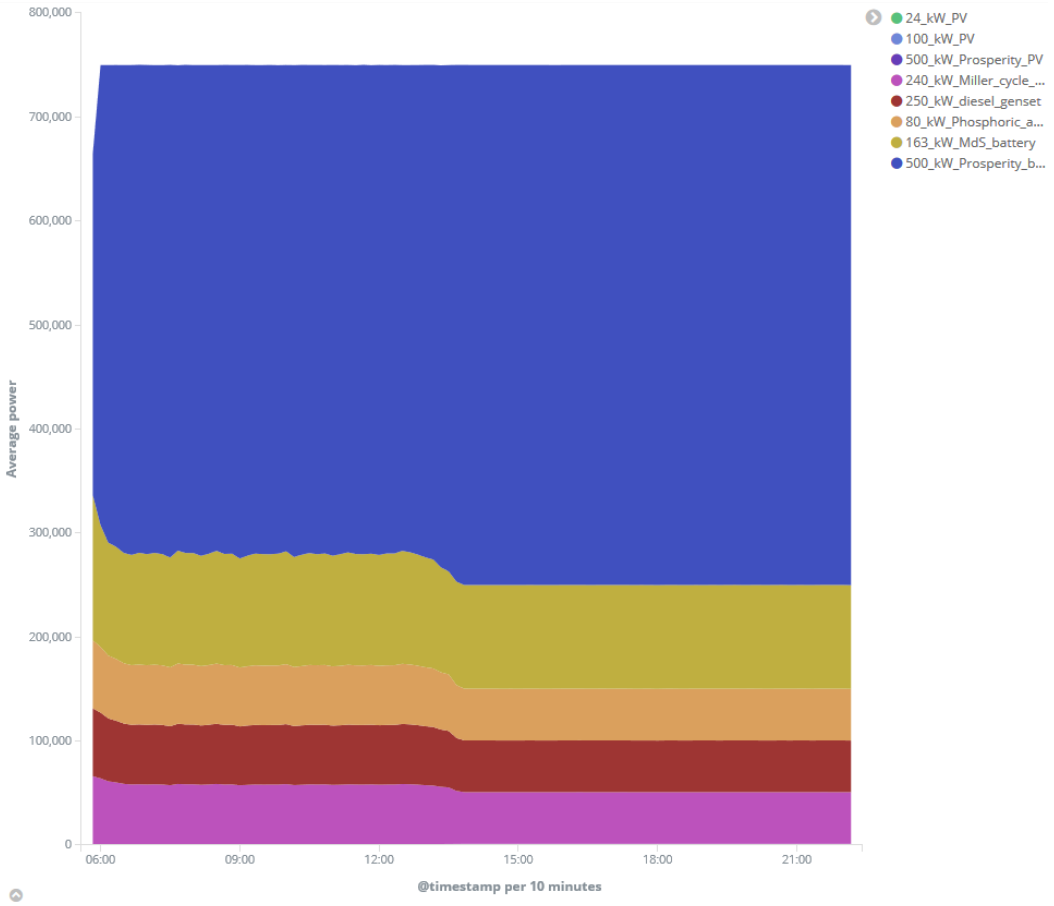

Figure 32: Kibana stacked power screenshot.

To improve the performance of data export (and reduce any impact on simulation performance), a separate process was created for data export. For each step of the control loop, data was exported into the export process (including the state of all the DERs). The export process caches this data for several seconds, and uses the Elasticsearch Bulk Import API to rapidly push all data into the datastore. This is significantly faster than exporting each data point individually, and can be done without slowing the simulation execution.

Other charts showed the total VPP output power as compared to the target power (Figure 33), DER power production (Figure 34), and the percent error in the output power. Unfortunately, the team discovered partway through the visualization system design that Kibana is not well suited to displaying this sort of data. Kibana is designed to aggregate data over automatically determined periods of time, and cannot be coerced into displaying every data point across a reasonably large time scale. Showing every data point is useful for analysis work (looking for short duration oscillations and similar misbehavior), and the aggregated data can hide this type of behavior. Kibana is a great tool, but it is likely not the right visualization front end for a project of this nature. Elasticsearch, however, works well as a data back end, and can easily scale to a cluster size for large data collection and long term storage/analysis.

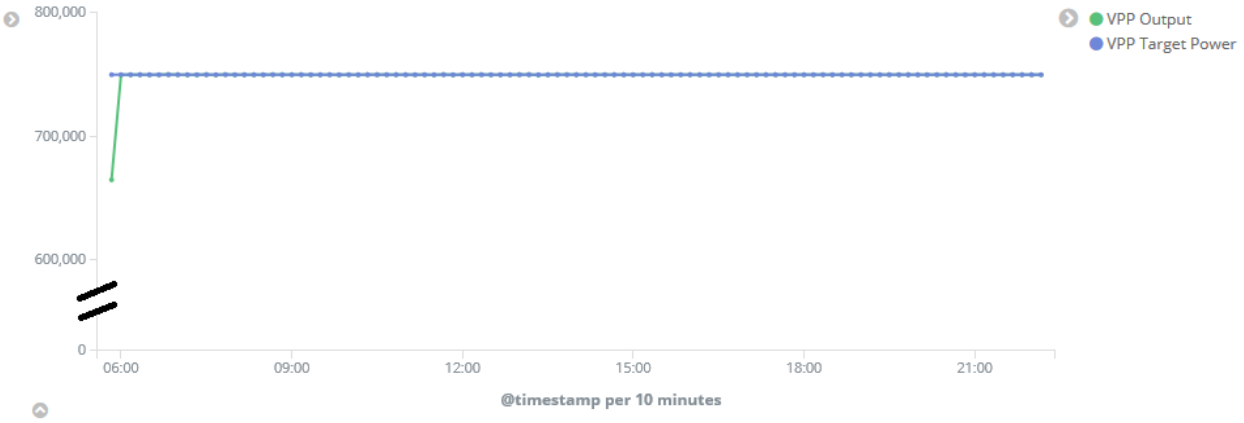

Figure 33: VPP target vs power output.

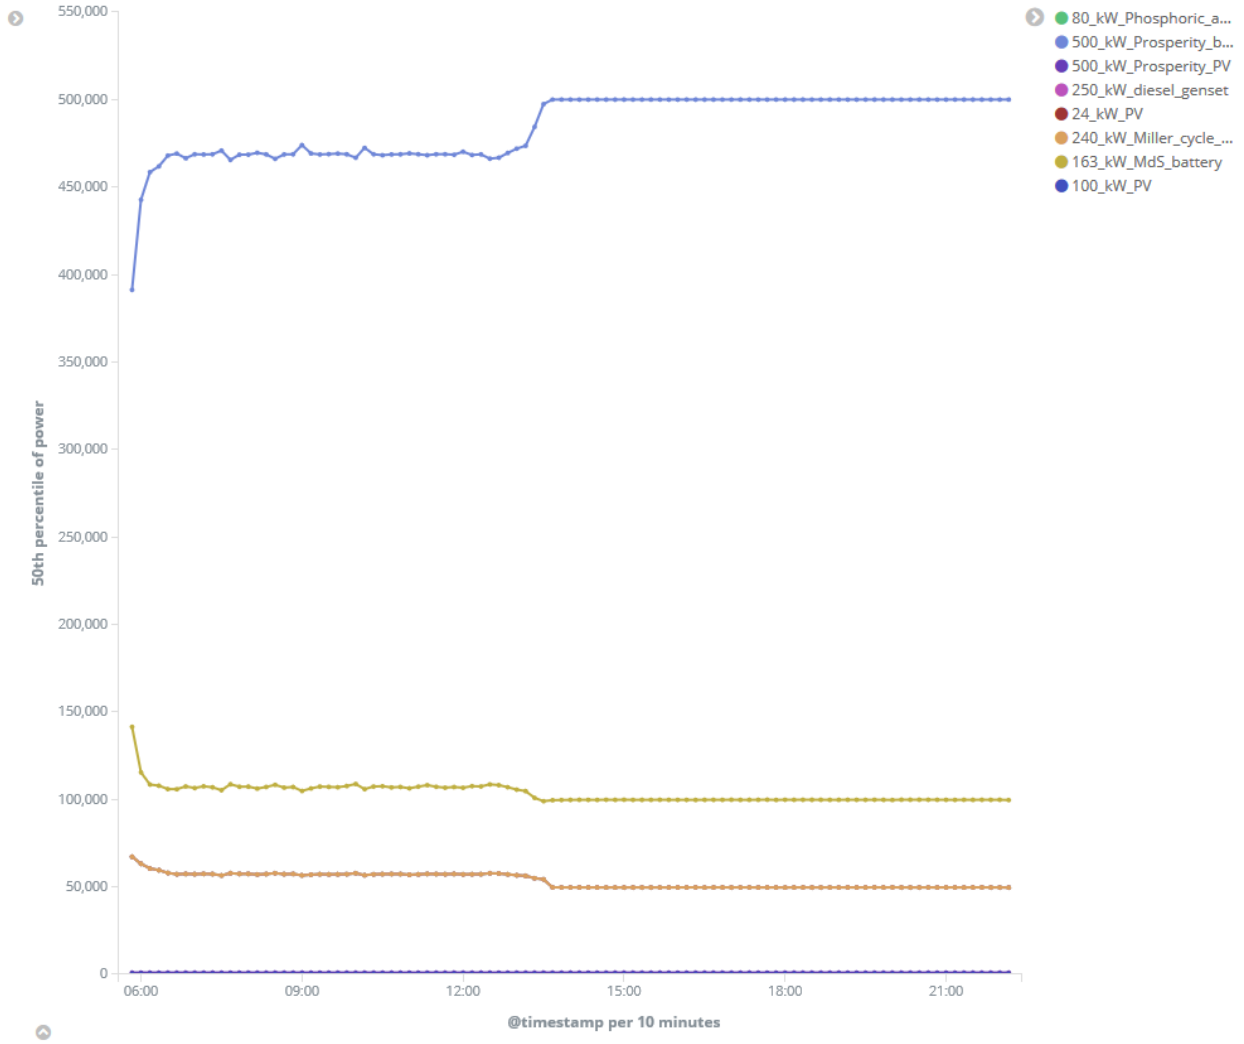

Figure 34: Average DER power output for each time step.

6 VPP SYSTEM OPERATIONAL RESULTS

Only limited operational runs were completed with the entire Energy and Reserve Market VPP system due to funding limitations. When the entire VPP operated, the long and short-term forecasts were gathered; the day-ahead optimization was completed and successfully stored; and the dispatch and control server executed simultaneously while feeding data to the export server.

Physical DER devices were also connected as part of the VPP aggregation, but communication latencies prevented the VPP from reaching the VPP setpoint because the PID control was not tuned for a system with those delays. It is believed a communication server that could issue commanded to multiple DER simultaneously could remediate those issues.

6.1 VPP Control with Simulated DERs

The VPP was run with the tuned control settings, 0.2 second communication rate and no network delay for the scenario shown in Table 5 on page 54. The output of the VPP and the DER devices is shown in Figure 35.

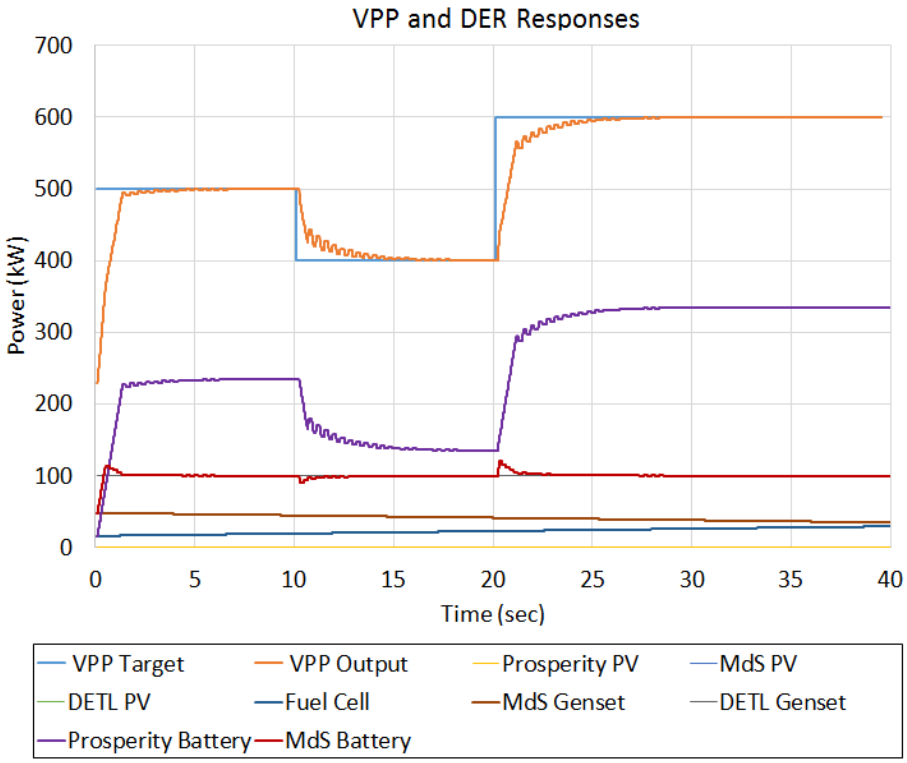

Figure 35. VPP and DER outputs for a commitment scenario.

The operation of the VPP ran for, at most, one hour before programming or integration problems caused one of the servers to crash. In the case of the VPP simulations, the control server succeeded in gathering three different optimal DER dispatch setpoints for the simulated DER. For that run, the weather was fair, so the VPP targets were within 20 W of each other and there was limited change in the output of the DER devices. Next, the commitment and optimization

engines were run to determine the energy and reserve bids and DER setpoints for a day in June 2017 based on live forecasts of the DER assets. Data from the controller was captured for 40 seconds with the reserve called at t = 20 s. The response of the VPP is shown in Figure 36.

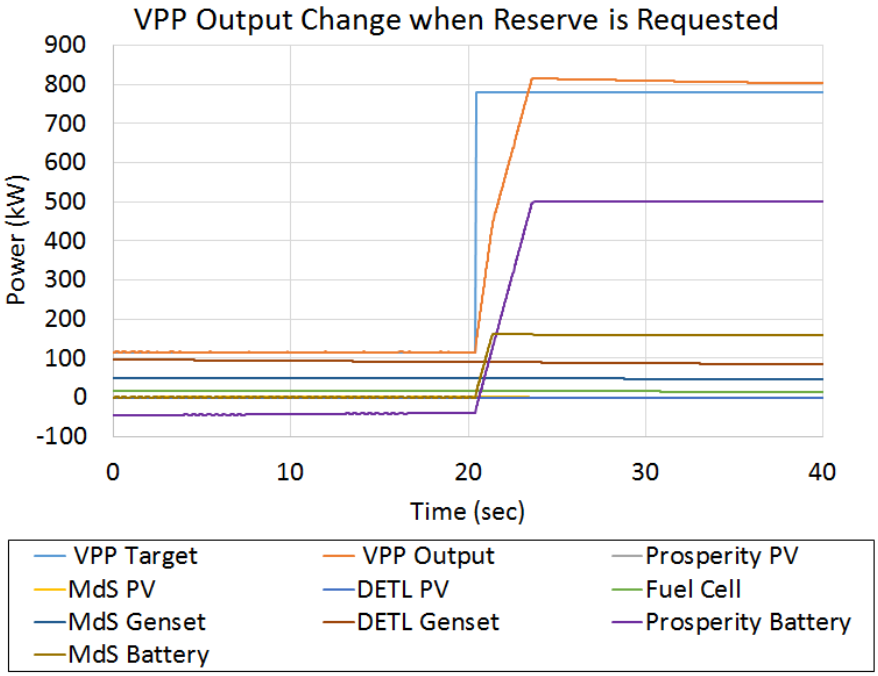

Figure 36. VPP and DER outputs based on commitment and optimization targets at a time when the reserve is requested.

6.2 VPP Control with Real DERs

To validate the VPP control with a real communication network, three PV inverters in DETL were issued curtailment commands from the VPP dispatch controller via SunSpec Alliance Modbus TCP commands. The DER output power was sequentially read and the level of active power curtailment of the DER equipment was adjusted. An example of the PV controls reaching a specified power level is shown in Figure 37. In cases where there was insufficient PV power available, the active power level was not meet, as shown in Figure 38.

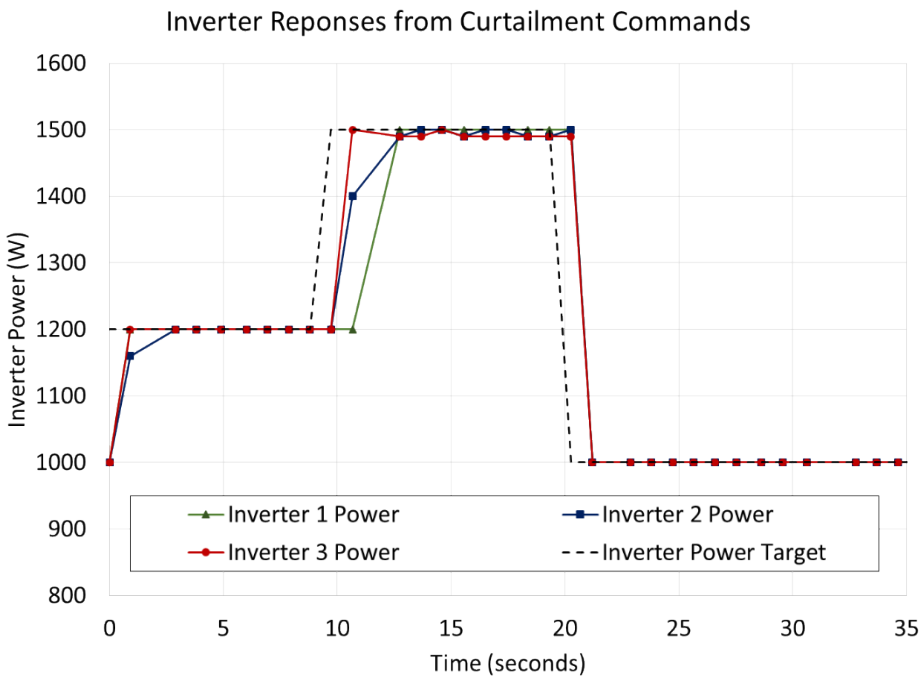

Figure 37. Response of three inverters to a target power signal.

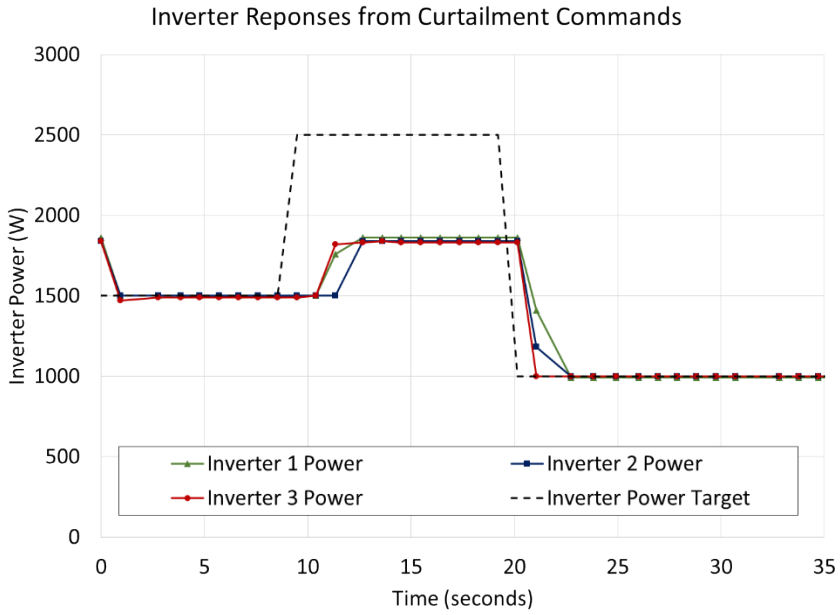

Figure 38. Response of three inverters to a target power signal, where the power level is above the available power of the renewable source.

Using simulated DER, the control loop was configured to execute in 0.01 seconds, but when adding the physical devices the loop time increased and the duration became variable. As shown in Figure 39, the read times for the DER was consistently ~200 ms for the inverters, but the write times varied between ~50 and ~1200 ms and the tuned VPP controls were no longer effective. To have stable control, the loop time must be consistent, so the variability forced VPP operator to execute the control loop at the largest duration, i.e., 2 seconds. This control speed produced poor

VPP system performance. One option to improve the VPP response would be to issue set points via parallelized communications, as opposed to sequentially. This would also allow the VPP to scale as more DER resources are added to the pool.

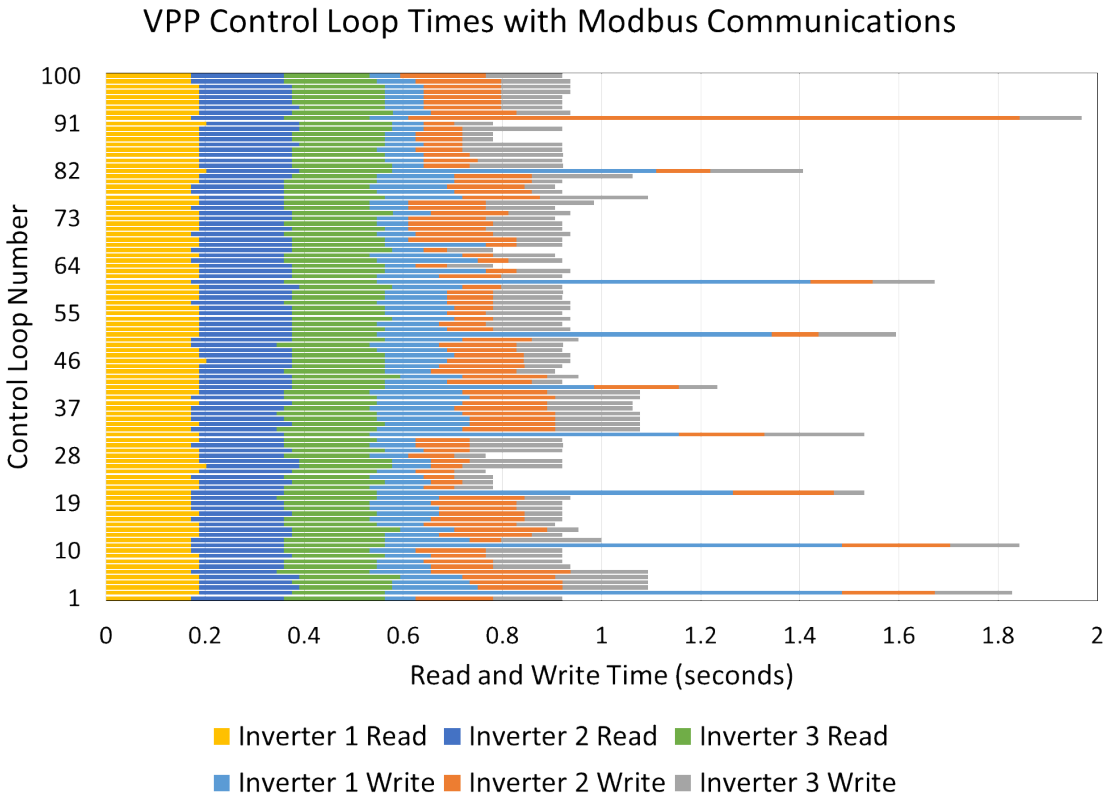

Figure 39. Inverter read and write rates for three physical DERs.

In summary, the centralized feedback control architecture for virtual power plants was operated effectively with simulated devices but, in the case of using physical devices to provide grid services, significant communication latencies prevented real-time operations. In the future, it is recommended to use communication dispatchers, multi-threading or multicast communications to control VPPs to avoid communications-related scaling issues.

7 CONCLUSIONS

In the future, ubiquitous distributed energy resources will be controlled in aggregate to provide grid services that have been traditionally reserved for large thermal generators or other dedicated equipment. These virtual power plants will be cost competitive with thermal generators and, as the nation moves toward a 100% renewable energy future, mandatory to maintain grid operations and resilience.

Multiple virtual power plants were designed to provide several grid services. The VPPs all consisted of DER power forecasting, optimization, and control through a communication network. The Energy and Reserve Market VPP was programmed as a multiprocessing software program at the Distributed Energy Technologies Laboratory at Sandia National Labs and operated under several scenarios. A stochastic optimization generated market bids based on renewable power forecasts; the VPP then met grid operator demands with an optimal DER dispatch mechanism, centralized controller, and communication network. Ultimately, limited VPP simulations were conducted with the entire system, but initial results were promising. The VPP was capable of quickly and accurately reaching target power levels with a diverse and nonhomogeneous collection of DER devices. Additional research should be conducted with each of the VPP components but, more importantly, in the integration of these elements into a fault-tolerant, monolithic software platform that quickly communicates with DER equipment.

The principal science and engineering contributions were in the areas of (a) creating a stochastic co-optimization methodology which enabled the VPP to participate in multiple day-ahead markets with high confidence of meeting the offers; (b) creating a cascading centralized controller with a swing DER to meet the real-time objective of the VPP, (c) providing voltage regulation using a distribution state estimation system and optimal power flow solver, and (d) establishing new cyber security intrusion detection systems based on power system state.

Finally, VPPs face many technical and regulatory challenges which may prevent market adoption. This project showed how many of the technical hurdles were surmountable with the right science and engineering approaches. Building this system provided evidence that VPPs are fully plausible in the near-term and a review of regulations that prevent VPPs from participating in ISO/RTO markets is warranted.

APPENDIX A: VPP SUBRESOURCE MODELS

A.1 Storage Model

The operating constraints for a given storage technology is defined as:

$$e_{\omega,t,s} = \mu_s e_{\omega,t-1,s} + \eta_s^c p_{\omega,t,s}^{\ c} - p_{\omega,t,s}^{\ d}/\eta_s^d \quad \text{Energy Storage Level} \quad (1)$$

$$\underline{E}_s \le e_{\omega,t,s} \le \overline{E}_s \quad \text{Energy Storage Capacity Limits} \quad (2)$$

$$p_{\omega,t,s} = p_{\omega,t,s}^{\ d} - p_{\omega,t,s}^{\ c} \quad \text{Net Power Injection/Withdrawal} \quad (3)$$

$$0 \le p_{\omega,t,s}^{\ c} \le \overline{P}_s^C b_{\omega,t,s}^{\ c} \quad \text{Charge Power Rating} \quad (4)$$

$$0 \le p_{\omega,t,s}^{\ d} \le \overline{P}_s^D (1 - b_{\omega,t,s}^{\ c}) \quad \text{Discharge Power Rating} \quad (5)$$

$$p_{\omega,t,s} + \bar{r}_{\omega,t,s}^{\ s} - p_{\omega,t-1,s} \le \overline{RR}_s \quad \text{Ramp Rate of Charge Limit} \quad (6)$$

$$-p_{\omega,t,s} + p_{\omega,t-1,s} \le \underline{RR}_s \quad \text{Ramp Rate of Discharge Limit} \quad (7)$$

$$\bar{r}_{\omega,t,s}^{\ s} = p_{\omega,t,s}^{\ c} + \left[\overline{P}_s^D\left(1 - b_{\omega,t,s}^{\ c}\right) - p_{\omega,t,s}^{\ d}\right] \quad \text{Reserve Capacity Availability} \quad (8)$$

$$f_s\left(p_{\omega,t,s}, \bar{r}_{\omega,t,s}^{\ s}\right) = f_s^e(p_{\omega,t,s}) + (\lambda_{\omega,t}^{DA} - C_s^1)\bar{r}_{\omega,t,s}^{\ s} \quad \text{Energy + Reserve Opportunity Cost} \quad (9)$$

$$f_s^e\left(p_{\omega,t,u} + r_{\omega,t,s}\right) = C_s^1(p_{\omega,t,s} + r_{\omega,t,s}) \quad \text{Energy Cost} \quad (10)$$

for all $\omega \in \Omega, t \in T, s \in S$, where the storage devices are a subset of the VPP subresources (DERs), i.e., $S \subseteq U$.

Equation (1) defines the total energy stored in the storage unit s for time t in scenario ω and incorporates a self-discharge rate $\mu_s \in (0,1]$ on the storage level as well as charging and discharging efficiencies, $\eta_s^c \in (0,1]$ and $\eta_s^d \in (0,1]$, respectively. In (1) when $t = 1$, then $e_{\omega,t-1,s}$ is equal to the energy storage level in the previous operating period. Constraint (2) limits the energy storage capacity based on a minimum and maximum energy level, $\underline{E}_s$ and $\overline{E}_s$, that can be specified with state-of-charge characteristics. Constraint (3) calculate the net power from the storage device, and then constraints (4) and (5) characterize the charge and discharge power ratings where $b_{\omega,t,s}^{\ c}$ is a binary variable that specifies whether the storage device is in charge ($b_{\omega,t,s}^{\ c} = 1$) or discharge ($b_{\omega,t,s}^{\ c} = 0$) mode. Constraint (6) represents the limit on the rate to charging in addition to providing reserve without physical violations, and constraint (7) represents the limit on the rate to discharging. Constraint (8) specifies that the unit's reserve is equal to the amount of energy that is being charged plus any unused capacity to discharge; note that only one of these terms can be nonzero. Equation (9) represents the cost function used in the day-ahead VPP scheduling optimization, and equation (10) represents the cost function that is used in the real-time VPP dispatch optimization.

A.2 Solar Photovoltaic Model

The operating constraints for a solar photovoltaic resource is defined as:

$$p_{\omega,t,i} = \underline{P}_i b_{\omega,t,i} + p_{\omega,t,i}^{\ \Delta} \quad \text{Power Output} \quad (11)$$

$$p_{\omega,t,i}^{\ \Delta} + \bar{r}_{\omega,t,i}^{\ s} \le (\overline{P}_{\omega,t,i}^{\ F} - \underline{P}_i) b_{\omega,t,i} \quad \text{Power + Reserve Limit} \quad (12)$$

$$f_i\left(p_{\omega,t,i}, \bar{r}_{\omega,t,i}^{\ s}\right) = f_i^e(p_{\omega,t,i}) + (\lambda_{\omega,t}^{DA} - C_i^1)\bar{r}_{\omega,t,i}^{\ s} \quad \text{Energy + Reserve Opportunity Cost} \quad (13)$$

$$f_i^e\left(p_{\omega,t,i} + r_{\omega,t,i}\right) = C_i^1(p_{\omega,t,i} + r_{\omega,t,i}) \quad \text{Energy Cost} \quad (14)$$

for all $\omega \in \Omega, t \in T, i \in I$, where the solar pv devices are a subset of the VPP subresources, i.e., $I \subseteq U$.

Equation (11) represents the power output as a function of its minimum operating level (MOL), $\underline{P}_i$, where $b_{\omega,t,i}$ is a binary variable that specifies whether the solar PV device is at or above the MOL to generate a usable power output. Then constraint (12) limits the generation over the MOL and the reserves, where $\overline{P}^{F}_{\omega,t,i}$ is unit i's forecasted output rating for the given scenario ω in time t. A positive reserve capacity, $\bar{r}^{s}_{\omega,t,i}$, denotes that the solar PV device may incur spillage during operations. Equation (13) represents the cost function used in the day-ahead VPP scheduling optimization, and equation (14) represents the cost function that is used in the real-time VPP dispatch optimization.

A.3 Thermal Generation Model

The operating constraints for the thermal generation model[113] is defined as:

$$b_{t,g} - b_{t-1,g} = b^{su}_{t,g} - b^{sd}_{t,g} \quad \text{Start-up/Shut-down State} \quad (15)$$

$$p_{\omega,t,g} = \underline{P}_g b_{t,g} + p^{\Delta}_{\omega,t,g} \quad \text{Power Output} \quad (16)$$

$$p^{\Delta}_{\omega,t,g} + \bar{r}^{s}_{\omega,t,g} \le \left(\overline{P}_g - \underline{P}_g\right) b_{t,g} - \left(\overline{P}_g - \overline{C}^{U}_g\right) b^{su}_{t,g} - \left(\overline{P}_g - \overline{C}^{D}_g\right) b_{t+1,g} \quad \text{Power + Reserve Limit} \quad (17)$$

$$p^{\Delta}_{\omega,t,g} + \bar{r}^{s}_{\omega,t,g} - p^{\Delta}_{\omega,t-1,g} \le \overline{RR}_g \quad \text{Ramp-up Limit} \quad (18)$$

$$-p^{\Delta}_{\omega,t,g} + p^{\Delta}_{\omega,t-1,g} \le \underline{RR}_g \quad \text{Ramp-down Limit} \quad (19)$$

$$\sum_{f \in F_g} K^{U}_{g,f} b^{0}_{t,g,f} \le c^{su}_{t,g} \quad \text{Start-up Cost} \quad (20)$$

$$b^{0}_{t,g,f} \le \sum_{t'=T^{L}_{g,f}}^{T^{L}_{g,f+1}-1} b^{sd}_{t-t',g} \quad \text{Start-up Type} \quad (21)$$

$$\sum_{f \in F_g} b^{0}_{t,g,f} = b^{su}_{t,g} \quad \text{Start-up Limit} \quad (22)$$

$$\sum_{t'=t-T^{U}_{g}+1}^{t} b^{su}_{t',g} \le b_{t,g} \quad (23)$$

$$\sum_{t'=t-T^{D}_{g}+1}^{t} b^{sd}_{t',g} \le 1 - b_{t,g} \quad (24)$$

$$b_{t-1,g} = B^{0}_{g} \quad (25)$$

$$f_g\left(p_{\omega,t,g}, \bar{r}^{s}_{\omega,t,g}\right) = \sum_{t \in T} c^{su}_{t,g} + f^{e}_{g}\left(p_{\omega,t,g}\right) + \left(\lambda^{DA}_{\omega,t} - C^{1}_{g}\right)\bar{r}^{s}_{\omega,t,g} - C^{2}_{g}\left(\bar{r}^{s}_{\omega,t,g}\right)^2 \quad \text{Energy + Reserve Opportunity Cost} \quad (26)$$

$$f^{e}_{g}\left(p_{\omega,t,g} + r_{\omega,t,g}\right) = C^{2}_{g}\left(p_{\omega,t,g} + r_{\omega,t,g}\right)^2 + C^{1}_{i}\left(p_{\omega,t,g} + r_{\omega,t,g}\right) + C^{0}_{g} b_{t,g} \quad \text{Energy Cost} \quad (27)$$

[113] G. Morales-Espana, J.M. Latorre, and A. Ramos. "Tight and compact MILP formulation for the thermal unit commitment problem." IEEE Transactions on Power Systems, (28)4, 4897-4908, 2013.

for all $t \in T, g \in G$ in (15), (20), and (22); for all $\omega \in \Omega, t \in T, g \in G$ in (16)-(19), (26), and (27); for all $t \in \left\{T_{t'+1,g}^{L},\ldots,T\right\}, g \in G, f \in F_g$ in (21); for all $t \in \left\{T_g^U,\ldots,T\right\}, g \in G$ in (23); for all $t \in \left\{T_g^D,\ldots,T\right\}, g \in G$ in (24); and for all $t \in \left\{1,\ldots,T_g^{D0} + T_g^{U0}\right\}, g \in G$ in (25) where $T_g^{D0} = max\{0,(T_g^D - T_g^{D'})(1 - B_g^0)\}$ and $T_g^{U0} = max\{0,(T_g^U - T_g^{U'})B_g^0\}$ for the minimum downtime, T_g^D, the initial periods offline prior to the scheduling horizon, $T_g^{D'}$, the minimum uptime, T_g^U, and the initial periods online prior to the scheduling horizon, $T_g^{U'}$. The thermal generators are a subset of the VPP subresources, i.e., $G \subseteq U$. The binary constraint in (15) enforces that the binary variables $b_{t,g}^{su}$ and $b_{t,g}^{sd}$ take on the appropriate values when a unit starts up or shuts down, respectively.

The active power dispatch operations for a thermal generating unit is characterized in constraints (16)-(19). Constraint (16) represents the total unit production, $p_{\omega,t,g}$, as the aggregate of the minimum operating level (MOL), $\underline{P}_g$, and the generation over that minimum, $p_{\omega,t,g}^{\Delta}$, where $b_{t,g}$ is a binary variable that specifies whether the unit g is on or off in time period t. Constraint (17) is the upper bound limit on the generation over the MOL along with the reserves, $\bar{r}_{\omega,t,g}^{s}$. The ramp-up constraint in (18) ensures that the unit can provide reserve without violating the upwards ramp limit, and the ramp-down constraint is specified in (19).

The scheduling limits for a thermal generating unit is characterized in constraints (20)-(25). The start-up cost is calculated in (20). The binary variable $b_{t,g,f}^{0}$ associates the corresponding start-up cost $K_{g,f}^{U}$ for a given start-up type, denoted by increasing start-up timeframes in $f \in F_g$, to the amount of time elapsed since the last time period that unit was online in (21) and (22). The minimum uptime and downtime constraints are denoted in (23) and (24), respectively. The initial up/down time constraint is enforced in (25).

Equation (26) represents the cost function used in the day-ahead VPP scheduling optimization, and equation (27) represents the cost function that is used in the real-time VPP dispatch optimization.

8 DISTRIBUTION

2 U.S. Department of Energy
Solar Energy Technologies Office
Attn: Guohui Yuan
M. Kemal Celik
950 L'Enfant Plaza
Washington, DC 20585

1 U.S. Department of Energy
Attn: Dan T. Ton
1000 Independence Ave. SW
EE-2A, FORS
Washington, DC 20585

1 Electric Power Research Institute
Attn: Brian Seal
942 Corridor Park Blvd.
Knoxville, TN 37932

1 National Renewable Energy Laboratory
Attn: Murali Baggu
1617 Cole Blvd.
Golden, CO 80401-3305

1	MS1033	Jay Johnson	8812
1	MS1084	Jack Flicker	5267
1	MS1033	Cliff Hansen	8812
1	MS1027	Anya Castillo	5853
1	MS1140	David Schoenwald	8813
1	MS1188	Mark A. Smith	8833
1	MS0757	Russell Graves	6612
1	MS0671	Jordan Henry	5828
1	MS0671	Trevor Hutchins	5828
1	MS1033	Abraham Ellis	8812
1	MS1033	Ross Guttromson	8812
1	MS1104	Charlie Hanley	8810
1	MS0899	Technical Library	9536 (electronic copy)
1	MS0359	D. Chavez, LDRD Office	1911


Sandia National Laboratories