



Inverter-based Resource Control for Grid Support

Advanced Solar Photovoltaic Plus Energy Storage System Demonstration and Technology Assessment

2021 TECHNICAL REPORT

Inverter-based Resource Control for Grid Support

*Advanced Solar Photovoltaic Plus Energy
Storage System Demonstration and
Technology Assessment*

EPRI Project Manager
A. Huque



3420 Hillview Avenue
Palo Alto, CA 94304-1338
USA

PO Box 10412
Palo Alto, CA 94303-0813
USA

800.313.3774
650.855.2121

askepri@epri.com

www.epri.com

3002023056

Technical Report, October 2021

DISCLAIMER OF WARRANTIES AND LIMITATION OF LIABILITIES

THIS DOCUMENT WAS PREPARED BY THE ORGANIZATION(S) NAMED BELOW AS AN ACCOUNT OF WORK SPONSORED OR COSPONSORED BY THE ELECTRIC POWER RESEARCH INSTITUTE, INC. (EPRI). NEITHER EPRI, ANY MEMBER OF EPRI, ANY COSPONSOR, THE ORGANIZATION(S) BELOW, NOR ANY PERSON ACTING ON BEHALF OF ANY OF THEM:

(A) MAKES ANY WARRANTY OR REPRESENTATION WHATSOEVER, EXPRESS OR IMPLIED, (I) WITH RESPECT TO THE USE OF ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT, INCLUDING MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, OR (II) THAT SUCH USE DOES NOT INFRINGE ON OR INTERFERE WITH PRIVATELY OWNED RIGHTS, INCLUDING ANY PARTY'S INTELLECTUAL PROPERTY, OR (III) THAT THIS DOCUMENT IS SUITABLE TO ANY PARTICULAR USER'S CIRCUMSTANCE; OR

(B) ASSUMES RESPONSIBILITY FOR ANY DAMAGES OR OTHER LIABILITY WHATSOEVER (INCLUDING ANY CONSEQUENTIAL DAMAGES, EVEN IF EPRI OR ANY EPRI REPRESENTATIVE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES) RESULTING FROM YOUR SELECTION OR USE OF THIS DOCUMENT OR ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT.

REFERENCE HEREIN TO ANY SPECIFIC COMMERCIAL PRODUCT, PROCESS, OR SERVICE BY ITS TRADE NAME, TRADEMARK, MANUFACTURER, OR OTHERWISE, DOES NOT NECESSARILY CONSTITUTE OR IMPLY ITS ENDORSEMENT, RECOMMENDATION, OR FAVORING BY EPRI.

THE ELECTRIC POWER RESEARCH INSTITUTE (EPRI) PREPARED THIS REPORT.

NOTE

For further information about EPRI, call the EPRI Customer Assistance Center at 800.313.3774 or e-mail askepri@epri.com.

Electric Power Research Institute, EPRI, and TOGETHER...SHAPING THE FUTURE OF ELECTRICITY are registered service marks of the Electric Power Research Institute, Inc.

Copyright © 2020 Electric Power Research Institute, Inc. All rights reserved.



Acknowledgments

The Electric Power Research Institute (EPRI) prepared this report.

Principal Investigators

A. Magerko

A. Huque

M. Ortega-Vazquez

Other Contributors

A. Tuohy

D. Larson

D. Freestate

Y. Ma

This report describes research sponsored by EPRI.

EPRI would also like to acknowledge the support of the following organizations that contributed to the research and report:

Direct Involvement:

- Los Angeles Department of Water and Power (LADWP)
- 8minute Solar Energy
- Doosan GridTech, Inc.
- Clean Power Research

Indirect Involvement:

- Reuniwatt
- Power Electronics

This publication is a corporate document
that should be cited in the literature in
the following manner:

*Inverter-based Resource Control for Grid
Support: Advanced Solar Photovoltaic
Plus Energy Storage System
Demonstration and Technology
Assessment.*

EPRI, Palo Alto, CA: 2021. 3002023056.



Abstract

A field demonstration of a utility-scale solar photovoltaic (PV) system paired with energy storage (ES) aimed to evaluate the functional capabilities of smart inverters, PV plus ES hybrid systems, and solar forecasting. Research was organized into four areas: 1) Solar forecasting applications and accuracy analyses 2) Performance evaluation of PV plus storage grid support functions 3) Feasibility investigation of PV ramp rate control without energy storage and 4) Evaluation of 24-7 voltage regulation support from PV smart inverters.

Long-term solar forecasts with high accuracy proved to be critical for an ES control function. Advanced PV plus ES control functions generally performed as defined in the MESA standard with some exceptions between real world operation and an idealized simulation environment. Short-term solar forecasts showed the need for further improvement in accuracy and reliability before being useful in solar ramp rate control type applications. Impact of increased solar forecast accuracy on dynamic reserve needs was assessed in the context of a risk-based resource planning approach. PV smart inverters with volt-var functionality demonstrated transmission-level voltage regulation support both in daytime and nighttime hours. Active power demand of the PV plant to provide reactive power support during nighttime has also been calibrated and reported.

Keywords

Solar forecast accuracy
MESA control functions
PV ramp rate control without storage
Smart inverter
Nighttime reactive power support
Dynamic reserve capacity

Deliverable Number: 3002023056

Product Type: Technical Report

Product Title: Inverter-based Resource Control for Grid Support: Advanced Solar Photovoltaic Plus Energy Storage System Demonstration and Technology Assessment

PRIMARY AUDIENCE: Los Angeles Department of Water and Power (LADWP)

SECONDARY AUDIENCE: Utilities with large-scale inverter based solar and energy storage resources; plant designers and operators

KEY RESEARCH QUESTION

Research aimed to assess the capabilities of inverter-based resources which are replacing traditional synchronous machine-based generation sources. Investigation of solar forecasts, smart inverters, and PV plus storage system controls would reveal to what extent variability and the lack of full dispatchability could be mitigated, while other PV and PV plus storage functions attempted to illustrate what additional value inverter-based resources could bring to the grid.

RESEARCH OVERVIEW

Where possible, research consisted of demonstrating, monitoring, and analyzing the performance of various inverter capabilities. Test plans were developed in advance and approved by participating parties. Tests were then executed at the utility-scale field demonstration site (that was already being designed and built at the inception of this research) and measurement data was fed back to researchers for analysis. Finally, analyses were performed to calculate how well the functions/capabilities performed and what potential benefit these capabilities had. In the case of solar forecast accuracy analysis, forecasts were gathered from multiple vendors and assessed by comparing it to measured values obtained from the field site.

KEY FINDINGS

- Improved long-term solar forecasts demonstrated median error of 3.5% for the Springbok 3 plant, down from a median error of 12.5% in the forecasts that LADWP currently uses (for comparable forecast sets). The blending of multiple numerical weather prediction inputs and satellite-based imagery into the forecast generation algorithm as well as a more refined plant model are likely the sources of improvement. Short-term forecasts from two different vendors (Clean Power Research & Reuniwatt) showed accuracies within 0.5% of one another.
- Nighttime reactive power support capability of a solar PV plant was demonstrated and active power demands to absorb and inject different magnitudes of reactive power were calibrated. The transitions into and out of its nighttime control mode were not seamless at the time of testing.
- Reduced voltage variations were observed at the point of interconnection during the test periods when a volt-var function was enabled; however, hundreds of kW of active power were drawn from the grid to keep the smart inverters functioning during night times.
- Field tests of advanced control modes showed a coordinated PV and battery energy storage system (BESS) system actively smooth PV ramps, respond to grid frequency conditions, and shift PV energy to different times of day. The component functions continued to perform as expected when combined into various mode stacks, achieving multiple mode objectives simultaneously with no significant

degradation in individual mode performance. However, an interconnection agreement at the Springbok 3 site that restricted charging the battery from the grid limited capabilities of some functions at night, and plant-specific response time limitations negatively impacted the capabilities of rapid response functions like frequency-Watt and automatic generation control (AGC).

- Conceptually, short-term solar forecasts can be used to limit solar PV ramp rates without energy storage. However, this project showed that, at present day, regulating PV-induced ramp rates at the individual plant level may still necessitate energy storage because solar forecast accuracy is insufficient during periods of high solar variability. Further research needs to be carried out to identify the right scale (i.e., at the individual plant level or across multiple plants across a wider geographical region) that is more cost efficient for managing PV ramp rates.

WHY THIS MATTERS

Assessing the capabilities and shortcomings of inverter-based resources helps technology vendors improve their products as well as helps utilities understand the state of the industry so that they can make informed choices, come up with stop-gap solutions where necessary, and better define interconnection requirements. Through repeated real-world demonstration efforts such as this one, technologies mature, familiarity and confidence in new systems builds, and ultimately, renewable generation is enabled to displace larger and larger amounts of fossil-fuel-based generation.

HOW TO APPLY RESULTS

Field demonstration of smart inverters and PV plus ES systems and corresponding analysis included in this report should be considered as example results to better understand the potential grid support capabilities of these technologies. This research also identifies gaps in the technologies that need to be addressed before the technologies can be relied upon for long-term planning and day-to-day operation of the power grid.

LEARNING AND ENGAGEMENT OPPORTUNITIES

- Work with EPRI to design and demonstrate similar pilot projects with longer duration testing and validation
- Perform techno-economic analyses of inverter-based resources relative to traditional solutions that provide specific grid services
- Participate in future design specification meetings with the [Modular Energy System Architecture \(MESA\) Standards Alliance](#)

EPRI CONTACTS: Aminul Huque, Principal Project Manager, mhuque@epri.com and Alex Magerko, Engineer/Scientist II, amagerko@epri.com

PROGRAM: DER Integration (Program 174)

Together...Shaping the Future of Electricity®

Electric Power Research Institute

3420 Hillview Avenue, Palo Alto, California 94304-1338 • PO Box 10412, Palo Alto, California 94303-0813 USA

800.313.3774 • 650.855.2121 • askepri@epri.com • www.epri.com

© 2020 Electric Power Research Institute (EPRI), Inc. All rights reserved. Electric Power Research Institute, EPRI, and TOGETHER...SHAPING THE FUTURE OF ELECTRICITY are registered service marks of the Electric Power Research Institute, Inc.

Table of Contents

Abstract	V
Executive Summary.....	VII
Chapter 1: Motivation, Overview, and Demonstration	
Site Details 1-1	
Demonstration Site Details.....	1-1
Full Plant Specifications.....	1-1
Power Block #5 Specifications.....	1-1
Springbok3 Plant Components and Equipment	1-2
Site Location.....	1-4
Control Architecture	1-4
Project Partners.....	1-7
Chapter 2: Solar Forecasting Accuracy Assessment	2-1
Introduction, Objectives, & Chapter Overview	2-1
Solar Forecast Accuracy Analyses & Comparisons	2-2
Background.....	2-2
Long-term Forecasts.....	2-6
Short-term Forecasts.....	2-9
Forecast Improvements Over Time	2-12
Investigation of Improved Forecast Accuracy on	
Dynamic Reserve Requirements.....	2-14
Method to Generate Improved Forecasts	2-14
Descriptive Statistics of Error	2-14
Impact on the Reserve Requirements	2-16
Probabilistic Forecasts.....	2-21
Hybridized Forecasts	2-25
Forecast Horizons and Use Cases	2-25
Qualitative Logistics, Costs, and Benefits of	
Blending Forecasts	2-25

Chapter 3: MESA-ESS/MESA-DER PV plus BESS Plant

Control Modes.....	3-1
Objective.....	3-1
Background and Test Plans.....	3-1
Test Plan Development.....	3-1
Demonstration of Isolated MESA-DER Control Modes ...	3-3
Coordinated Charge/Discharge (CCD)	3-3
PV Active Power Smoothing	3-14
Frequency Response (Frequency-Watt).....	3-18
Automatic Generation Control	3-27
Demonstration of Stacked Control Modes	3-31
Impact to CCD Accuracy.....	3-31
Impact to APS Accuracy.....	3-35
Impact to FW Accuracy	3-37
Summary of Results.....	3-39

Chapter 4: Ramp Rate Limiting with Solar Forecasts but No Energy Storage

Objective & Usefulness	4-1
How the Function Works	4-2
In More Detail.....	4-2
Performance	4-2
Emulated Scenario	4-3
Up-Ramp Control.....	4-3
Down Ramp Control	4-4
Performance Challenges	4-4
PV Forecasting Accuracy Limitation	4-4
Response Time and Sub-Control Interval PV Change	4-6
Conclusion	4-7

Chapter 5: Full-plant Voltage Regulation with Volt- var Control Function

Objective & Significance	5-1
Active Power Consumption Calibration Test	5-3
Test Procedure	5-3
Results.....	5-4
Extended Volt-var Testing.....	5-7
Test Procedure	5-7
Performance	5-8
Continued Challenges.....	5-12
Review of Existing Transmission-level Voltage Regulation Technologies	5-12
Voltage Regulation Solutions and Their Operating Principles.....	5-12
Comparison of Voltage Regulation Technologies .	5-14

Representative Commercial Voltage Regulation Products.....	5-18
Chapter 6: Conclusion.....	6-1
Findings Summary	6-1
PV Forecasting	6-1
Demonstration of MESA-DER Functions for PV+BESS.....	6-1
Ramp Rate Limiting Without Energy Storage	6-2
Full-Plant Voltage Regulation with Reactive Power	6-3
Lessons Learned.....	6-3
Potential Future Work	6-5
Appendix A: Dynamic Reserve Requirements to Accommodate VRES Vagaries.....	A-1
Acronyms.....	A-1
Nomenclature.....	A-1
Introduction.....	A-2
Dynamic Reserve Requirements.....	A-3
Data Requirements and Preparation	A-4
Historical Data.....	A-5
Anticipated Data	A-6
Typical Problems with Historical Data	A-6
Descriptive Statistics of Historical Data	A-6
Operating Reserves Calculation	A-11
LADWP Reserve Policy	A-11
Exact Reserves for LADWP	A-12
Comparison of Reserve Policies.....	A-13
Risk as a Reserve Performance Metric	A-13
Reserve Procurement Comparison	A-14
Effect of Improved Forecasts	A-16
Method to Generate Improved Forecasts	A-16
Descriptive Statistics of Error	A-17
Impact on the Reserve Requirements	A-18
Conclusions	A-23
References	A-23

List of Figures

Figure 1-1 Power converter topologies used at different power blocks within the Springbok3 power plant	1-2
Figure 1-2 Diagram of communications architecture at power block #5	1-5
Figure 1-3 Diagram of solar forecast data flows implemented in this demonstration	1-6
Figure 1-4 Photograph of Reuniwatt sky camera installed as part of short-term forecast service	1-6
Figure 1-5 Logos of primary project partners who participated in this pilot	1-7
Figure 2-1 Visualization of selective sampling from CPR forecasts used to generate the equivalent “day-ahead” forecasts to be compared to those from Meteologica	2-4
Figure 2-2 Visualization of selective sampling from CPR (left) and Reuniwatt (right) forecasts used to analyze equivalent short-term forecasts for accuracy.....	2-4
Figure 2-3 Absolute error by forecast horizon for GHI (left) and power (right); CPR short-term forecast used in illustration.....	2-5
Figure 2-4 Flowchart of possible (not comprehensive) sources of error between a GHI forecast and a plant power forecast	2-6
Figure 2-5 Median absolute error by month for equitable long-term forecasts from Meteologica and CPR, filtered down to the time periods when all component power blocks were reporting and operational †11 am forecast set only	2-7
Figure 2-6 One week sample of forecasts and actuals from two weeks in different seasons	2-7
Figure 2-7 Median bias error by month for equitable long-term forecasts from Meteologica and CPR, filtered to when all component power blocks were reporting and operational †11 am forecast set only	2-7

Figure 2-8 Bias error by forecast horizon for CPR day-ahead portion of long-term forecast.....	2-8
Figure 2-9 Visualization of where raw forecasts (before considering plant export limits) would collapse down to a known export limit, reducing error	2-8
Figure 2-10 Median absolute error by month for equitable short-term forecasts from CPR and Reuniwatt, filtered to when all component power blocks were reporting and operational † First 31 minutes only ‡ XX:00 or XX:30 forecasts only	2-9
Figure 2-11 Box and whisker plot indicating distribution of absolute error by solar variability classification for the CPR short-term forecasts.....	2-9
Figure 2-12 Median bias error by month for equitable short-term forecasts from CPR and Reuniwatt, filtered to when all component power blocks were reporting and operational † First 31 minutes only ‡ XX:00 or XX:30 forecasts only	2-10
Figure 2-13 Standard deviation of absolute error by month for equitable short-term forecasts from CPR and Reuniwatt, filtered to when all component power blocks were reporting and operational † First 31 minutes only ‡ XX:00 or XX:30 forecasts only	2-10
Figure 2-14 Cumulative distribution function (CDF) indicating the percentage of Reuniwatt InstaCast forecasts (within viable periods) with a given absolute error value or less.....	2-10
Figure 2-15 CDF indicating the percentage of CPR short-term SolarAnywhere forecasts (within viable periods) with a given absolute error value or less.	2-11
Figure 2-16 Measured/realized PV powers and corresponding forecasted value (Reuniwatt InstaCast) from 10 minutes prior, demonstrating inconsistent prediction of rapid PV changes.....	2-11
Figure 2-17 Side by side comparison of short-term forecasts from two different vendors with realized (measured) values behind each for reference	2-12
Figure 2-18 Short-term forecasted and measured power profiles (from before the correction) highlighting the error previously induced by the missing export limit feature in CPR's PV plant model ...	2-12
Figure 2-19 Global horizontal irradiance forecast absolute error and standard deviation of that error by month	2-13

Figure 2-20 Summary of the method to produce improved forecasts	2-14
Figure 2-21 System deviations for a) the original forecast case and b) the improved solar forecast case; full year left, sample day right	2-15
Figure 2-22 Isolated solar deviations for a) original forecast and b) improved forecasts; full year left, sample day right.....	2-16
Figure 2-23 LADWP reserve determination policy for a) the original solar forecast and b) the improved solar forecast. Color bands represent percentile ranges of calculated reserve needs over the course of a year at each given time of day.	2-17
Figure 2-24 Exact up and down reserve requirements for the LADWP system for both the original forecast and the improved forecast. Color bands represent percentile ranges of realized reserve requirements over the course of a year at each given time of day.	2-18
Figure 2-25 Up and down reserve requirements for the test day in LADWP using both the original forecast and the improved forecast for solar generation.....	2-19
Figure 2-26 Up and down-reserve requirements and associated short and long risk for both, the original solar forecast and the improved solar forecast	2-20
Figure 2-27 Average reserves (top) and average risk (bottom) assessment for the LADWP reserve policy and a collection of outputs from EPRI's reserve determination tool with different confidence intervals. Each reserve policy has one set of results for the original forecast, 'o', and one set for the improved/enhanced forecast, 'e'	2-21
Figure 2-28 Brier scores for the 10 th and 90 th percentiles at different forecast horizons, lower is better (data from two months of long-term CPR forecasts)	2-23
Figure 2-29 Quantile scores for the 10 th and 90 th percentiles at different forecast horizons, lower is better (data from two months of long-term CPR forecasts)	2-23
Figure 2-30 Reliability scores for the 10 th and 90 th percentiles at different forecast horizons (data from two months of long-term CPR forecasts).....	2-24

Figure 3-1 Profile of simulated SoC values and targets (top) with corresponding simulated power outputs from the PV, BESS/ESS, and net power at point of interconnection (bottom)	3-4
Figure 3-2 Profile of reported SoC estimates in response to SoC target values using “greedy” CCD.....	3-5
Figure 3-3 Profile of reported SoC estimates in response to SoC target values using predictive CCD (CCD-P).....	3-6
Figure 3-4 “Greedy” CCD and CCD-P algorithms performing identical SoC target profile with primary difference in behavior called out	3-7
Figure 3-5 PV, BESS/ESS, and net (power block 5 point of interconnection) power profiles resulting from consecutive days of greedy CCD (left) or CCD-P (right) testing. Top: field results, Bottom: simulated results.....	3-8
Figure 3-6 Dot plot indicating distribution of SoC errors for the collection of SoC targets from Figure 3-2.....	3-10
Figure 3-7 Dot plot indicating distribution of SoC errors for the collection of SoC targets from Figure 3-3.....	3-10
Figure 3-8 Distribution of reported SoC error relative to the desired target value at each target for a variety of test days	3-11
Figure 3-9 Distribution of reported SoC error relative to the desired target value, categorized by whether target occurred at night (left) or during daylight hours (right)	3-11
Figure 3-10 Power profiles (simulated on top, measured on bottom) in response to a simple CCD-P profile with daytime charging and nighttime discharging	3-12
Figure 3-11 Sample period from ramp-rate-only test day .	3-15
Figure 3-12 Sample period from full APS function test day	3-15
Figure 3-13 Distribution of observed ramp rates for PV and at POI (=PV+BESS) for the ramp-rate-only test day (left) and full APS function test day (right)	3-16
Figure 3-14 Histogram of BESS utilization while performing APS both with and without gradient filtering	3-17
Figure 3-15 Scatter plot indicating the frequency-Watt setpoint and measured response to the desired frequency-Watt curve for the daytime and nighttime test	3-18

Figure 3-16 Brief, time series sample of power responses (setpoint in red and measured in blue) relative to the calculated ideal response value in green.	3-19
Figure 3-17 Zoomed in version of Figure 3-15 with cluster of points deviating from desired curve circled in yellow.....	3-20
Figure 3-18 Extended time series plot indicating measured frequency (top plot), various response types (middle plot), and the available battery SoC (bottom plot) for the nighttime test performed on 1/14-15/2021	3-21
Figure 3-19 Extended time series plot indicating measured frequency (top plot), various response types (middle plot), and the available battery SoC (bottom plot) for the daytime test performed on 1/15/2021	3-22
Figure 3-20 Histograms of the percent error relative to the ideal additive power response (normalized by rated power in that direction) for daytime values	3-23
Figure 3-21 Histograms of the percent error relative to the ideal additive power response (normalized by rated power in that direction) for nighttime values	3-24
Figure 3-22 Time series sample of nighttime frequency-Watt responses exemplifying overshoot only in positive direction	3-24
Figure 3-23 Cumulative distribution function of frequency-Watt response error (relative to ideal additive power response)	3-25
Figure 3-24 Histograms of the percent error relative to commanded ECC AGC signals (normalized by max committed capacity of +/-200 kW and excluding nighttime periods with <100 kW requested)	3-27
Figure 3-25 Extended time series plot indicating measured additive power response relative to the external AGC command (top plot), plus the available battery SoC (bottom plot) for the AGC test performed 10/27/2020	3-28
Figure 3-26 Sample of ESS response to the AGC command relative to measured and target responses at the meter in front of the system (POI); demonstrates power offset absorbed by system	3-29
Figure 3-27 Cumulative distribution function of response error (relative to ideal additive power response)	3-29

Figure 3-28 Zoomed-in view of the setpoint response, instantaneous measurement, and time delayed measurements relative to the commanded AGC signal value	3-29
Figure 3-29 Dot plot of SoC target errors associated with various test days; dots with a white 'X' represent days impacted by a hardware limitation	3-32
Figure 3-30 Time series plot illustrating the frequency conditions and corresponding plant response that caused larger-than usual error for the 18:00 SoC target on this CCD-P+FW test day.....	3-33
Figure 3-31 Time series plot illustrating the frequency conditions and corresponding plant response that caused larger-than usual error for the 18:00 SoC target on this CCD-P+APS+FW test day	3-34
Figure 3-32 Alternate SoC accuracy metrics for various CCD-P tests	3-35
Figure 3-33 Time series plots of three APS test days without (top) and with (middle and bottom) additional stacked functions active	3-35
Figure 3-34 Bar plot depicting the mean ramp rate at various measurement points for three different tests involving the APS function.....	3-36
Figure 3-35 Distributions of observed PV, basepoint, and POI ramp rates for various tests involving APS (zoomed in to better see distributions; values not shown are predominantly ramp rates very close to 0).....	3-37
Figure 3-36 Intended (i.e., setpoint) and measured responses to observed frequency values with ideal response curve (black) for reference and deviations from this curve highlighted.....	3-38
Figure 3-37 Setpoint and measured response error for various tests involving FW	3-38
Figure 4-1 Simulated desired plant output (orange curve) for hypothetical, highly variable PV generation scenario (blue) and amount of curtailment (in gray) required to obtain that desired output	4-1
Figure 4-2 Sample of emulated plant response (gray line) to rapid changes in available PV power (blue line).....	4-3

Figure 4-3 Sample field measurement data from neighboring power blocks (with PB5 under a RRP test) demonstrating three instances of proper up-ramp performance	4-3
Figure 4-4 Sample field measurement data from neighboring power blocks (with PB5 under a RRP test) demonstrating one superfluous and one successful instance of proper down-ramp performance	4-4
Figure 4-5 Sample day showing forecasted PV power profile (consisting of only the first forecast from every half hour to avoid visual clutter from overlapping forecasts) relative to measured power profile	4-5
Figure 4-6 Forecasted PV power estimates at different forecast horizons plotted for successive forecast sets (unmodified)	4-6
Figure 4-7 Forecasted PV power estimates at different forecast horizons plotted for successive forecast sets after shifting (modifying) purported forecast horizons	4-6
Figure 5-1 Quoted active and reactive power capability of the Springbok 3 plant (for VPOI = 1.00 pu) with allowed range outside of 0.95 PF highlighted	5-2
Figure 5-2 Example reactive power setpoint profile (and closely matching measured response) from a fixed var active power consumption calibration test....	5-4
Figure 5-3 Scatter plot of raw active power and reactive power measurements before necessary modifications were made to isolate true active power consumption calibration curve	5-5
Figure 5-4 Active power consumption calibration curve from nighttime reactive power testing	5-6
Figure 5-5 Volt-var curve used for voltage regulation tests (black) overlaid onto historic voltages and reactive powers (blue) from periods of pre-testing operation at the Springbok3 site	5-7
Figure 5-6 Inherent voltage response to reactive power injection/ absorption at the Springbok 3 point of interconnection	5-8

Figure 5-7 Reactive power response (bottom) to measured voltage (top) with two highlighted periods where output did not align with expected behavior at that time (right highlight occurred around sunset). The estimated “no Volt-var” voltage utilized the slope found in Figure 5-6 to back calculate what the voltage might have been in the absence of any reactive power flow.....	5-9
Figure 5-8 Observed Volt-var response from second string of test days (10-days long) as measured by power plant controller (PPC).	5-10
Figure 5-9 Distribution of measured voltage at plant point of interconnection with and without Volt-Var function enabled.....	5-1
Figure 5-10 CDF indicating percentage of time that reactive power contribution was a certain value or less	5-1
Figure 5-11 Alternative technologies for transmission-level voltage regulation	5-13
Figure 5-12 Illustrative TCR voltage and current profile to demonstrate its basic operating principle.....	5-14
Figure 5-13 Illustrative voltage and current profile of a TSC to demonstrate its basic operating principle	5-14
Figure 5-14 Illustrative voltage and current profile of a STATCOM to demonstrate its basic operating principle	5-14
Figure 5-15 Illustrative comparison of reactive current (top) and reactive power (bottom) capability between SVC and STATCOM.....	5-15
Figure 5-16 Illustrative comparison of operational active power consumption for SVC and STATCOM.....	5-16
Figure 5-17 Cost comparison between SVC and STATCOM recreated from a research paper published in 2004 ⁶⁴	5-17

List of Tables

Table 1-1 Specifications of hardware equipment included at Springbok 3 plant	1-3
Table 1-2 Specifications and justifications of hardware equipment selected for powerblock #5 at Springbok 3 plant	1-4
Table 2-1 Summary of forecast sources available in this project	2-3
Table 3-1 Various performance indices as applied to "greedy" CCD and CCD-P test days.....	3-12
Table 3-2 Table of parameters used in subsequent tests of APS and the ramp-rate-only portion of APS.	3-14
Table 3-3 Performance metrics indicating percentage of observed responses within 5% and 1% tolerance bound of ideal additive power response value broken down by portion of day and response type....	3-25
Table 3-4 Median SoC accuracy by test.....	3-31
Table 3-5 Qualitative change in CCD performance metrics between various stacked mode test days (involving CCD-P) and a greedy CCD test day	3-32
Table 5-1 Power plant controller parameters used for Volt-var voltage regulation tests	5-7
Table 5-2 Long-term Volt-var test periods.....	5-8
Table 5-3 Comparison of voltage regulation technologies where the number of '+' symbols qualitatively indicates the strength of that technology in that metric (more is better)	5-18
Table 5-4 A sample of representative manufacturers and products	5-18
Table 6-1 Sample of historical forecasts	A-5
Table 6-2 Sample historical actuals.....	A-5

Chapter 1: Motivation, Overview, and Demonstration Site Details

The Los Angeles Department of Water and Power (LADWP) has in place initiatives that strive to continuously achieve higher Renewable Portfolio Standards (RPS). When setting out on this project, LADWP was targeting 65% RPS, but it is now targeting 100% RPS by 2035.

Renewable generation has advantages over conventional generation, but it is not without challenges, too. Renewable generation using variable energy resources (VER) such as solar and wind is heavily weather dependent and varies over time, often making it difficult to schedule and count on with a high level of certainty. New technologies, including energy storage, advanced inverter functions, improved forecasting systems, and enhanced monitoring & controls potentially have the capability to mitigate some of the challenges of VER integration and make them a suitable alternative to conventional generation. Implementing these potential technologies in a field demonstration and investigating their current potential benefits was the aim of this project.

In particular, the following four topics constituted the foci of this project and outline the chapter structure of this report:

1. Quantitative forecast accuracy analyses and qualitative use case analysis of improved day-ahead and real-time solar PV forecasts
2. Demonstration and performance analysis of various advanced DER control functions involving both PV and storage
3. Analysis of PV power ramp rate limiting capabilities *without* energy storage (using short-term PV forecasts)
4. Demonstration and performance analysis of full-plant voltage regulation capabilities utilizing inverter reactive power capability

Demonstration Site Details

Research questions within the scope of this project were tested on a utility-scale PV plant that was already under development by 8minute and LADWP when discussions with EPRI began. This subsection outlines the plant specifications, hardware components, plant location, and reasons for why the plant and its constituent power block #5 were designed the way they were.

Full Plant Specifications

- Power ratings:
 - 120.09 MW_{DC}
 - 105.094 MW_{AC} |
 - 90 MW power injection limit, per interconnection agreement
- 32 constituent Power Blocks (PBs)
 - 30 Power Electronics 3225 kVA inverters
 - 1 Power Electronics 3225 kVA *bidirectional* inverter with dc-coupled energy storage behind it
 - 1 PB consisting of 33 Huawei Sun-2000-100KTL inverters
- Power Plant Controller (PPC): GreenPowerMonitor (GPM) integrates plant and provides main plant controller (“PPC”)

Power Block #5 Specifications

- Battery Size: 1.5 MW, 1.5 MWh (dc)
- Vendors: LG Chem JP2 batteries, Power-Electronics (PE) inverter, and dc/dc converters (3 x 500 kW)
- Integrator: NEC Energy Solutions (“NEC”) integrates system and provides main storage controller (“AEROS”)
- Dispatcher: Doosan GridTech Intelligent Controller (“DG-IC”) is layered over the

integrator's plant controller and provides MESA-ESS controls

Springbok3 Plant Components and Equipment

The 105 MW_{ac} plant used in this project consisted of 32 power blocks. 30 of those power blocks had unidirectional Power Electronics (PE) inverters (left side of Figure 1-1). Meanwhile, Powerblock 3 used Huawei inverters and a variety of PV module manufacturers as part of a technology testing ground (not part of this research project), and Powerblock 5 used a bi-directional Power Electronics inverter as part of its dc-coupled energy storage solution (right side of Figure 1-1).

The full plant was designed to be solar-only with a small battery energy storage system (BESS) added to one power block for testing purposes. The high cost of battery energy storage and the

fact that this was a pilot project would have made deployment of energy storage to all power blocks impractical. As a result, only Powerblock 5 was equipped with a BESS alongside its 3 MW_{ac} of PV. All the active power control mode tests delineated in Chapter 3 were performed using this NEC-supplied energy storage system.

8minute chose a dc-coupled BESS architecture for the pilot block because it wanted proof that dc coupled PV plus storage (PV+S) could capture what would otherwise be clipped PV energy (on an ac coupled PV+S). The dc/dc converter was placed in between the BESS and the dc bus connecting the PV to the inverter as depicted on the right half of Figure 1-1. 8minute also wanted to demonstrate that the technology could operate reliably over many months.




Figure 1-1
Power converter topologies used at different power blocks within the Springbok3 power plant

The 1.5 MW battery was selected because the solar inverter blocks at the Springbok3 project are 3 MW_{ac} each, and many large-scale PV+S projects are sized with the BESS power at 50% of the PV power. 8minute decided that having the BESS sized to 50% of the solar generation would provide representative operation that could be translated to larger projects that were in planning stages at the time. The battery was selected to have a 1-hour

capacity (1.5 MWh) because the pilot project budget was limited and could not support a 2- or 4-hour system.

Power Electronics (PE) was selected for the pilot block (PB#5) because (1) PE was already selected as the inverter provider for all but one of the other power blocks on site, (2) PE offered a dc/dc converter that was designed from the same OEM to specifically pair with this inverter, and (3) PE

reported at the time that another developer would have commissioned a similar project with the same technology approximately 1 year prior to Springbok3, which would give PE time to fix any

initial technical challenges before building Springbok3. 8minute decided that these three factors made Power Electronics the most likely solution to be reliable for this pilot project.

Table 1-1

Specifications of hardware equipment included at Springbok 3 plant

Equipment	Vendor	Quantity	% of Site
Modules	First Solar Series 6 (Variety in Tech Block #1)	283,662 1,460	99.5% of MW _{dc} 0.5% of MW _{dc}
Inverters	Power Electronics central (single direction) Power Electronics (bi-directional with DC/DC converters) Huawei string inverters (single direction) (Tech Block #1)	30 1 (3 DC/DC converters) 33 string	93.7% of MW _{ac} 3.1% of MW _{ac} 3.2% of MW _{ac}
Battery System	NEC integrated container (LG chem batteries)	1.5MW x 1hr	0.1% of total MW
Battery Control System	NEC Energy Solutions ("NEC") main storage controller ("AEROS")	1	N/A
Software Dispatcher	Doosan GridTech Intelligent Controller ("DG-IC")	1	N/A
Power Plant Controller	GreenPowerMonitor	1	N/A

Table 1-2

Specifications and justifications of hardware equipment selected for powerblock #5 at Springbok 3 plant

Li-ion Battery Module		DC-to-DC Converter		Inverter	
Vendor					
 LG Chem		 POWER ELECTRONICS®		 POWER ELECTRONICS®	
Product					
Li-ion Battery Energy Storage System		FD0500 Freemaq DC-to-DC Converter		FP3225M Bi-Directional Inverter	
8minute's Reason for Selection					
- Stack & folding technology results in stable, high energy density structure with superior lifespan - Battery Management System (BMS) technology creates highly reliable and integrated system		- Ability to perform multiple grid support functions (e.g., load leveling, renewable integration, utility curtailment recovery) - Modular design allows for flexibility in amount of dispatched storage power		- Similar in functionality and specifications to Power Electronics' FS3225M inverters - Ability to flow power in both directions to facilitate charging of ESS from grid, if desired	
Specifications					
Cell Technology	Li-ion Nickel Manganese Cobalt	Voltage Range	700 Vdc ~ 1,310 Vdc	Rated Power	3550kVA @ 25°C / 3225kVA @ 50°C
Cell Size	100mm x 353mm x 16mm	Rated Power	500 kWdc	Rated Voltage	34.5kV
Cell Nominal	3.6 V / 72.5 Ah	Rated Current	722 Adc	Rated Frequency	50Hz / 60Hz
Cells / Module	28	Voltage Range	800 Vdc ~ 1,310 Vdc	Rated Power Factor	0.5 leading / lagging
Cells in Series	238	DC Voltage Ripple	<3%	Maximum Voltage	1,500 Vdc
Module Nominal	51.4 Vdc	PCS Maximum Efficiency	99%	MPPT Operation Range	913 Vdc ~ 1,310 Vdc
DC Standby Consumption	5.2 kW	Full Power Temperature Range	-35 ~ 60°C (-31 ~ 140°F)	Maximum Efficiency	98%
ESS Rated Capacity	1,500 kWh			CEC Efficiency	N/A
ESS Dimensions	40ft x 8.5ft x 16ft			Full Power Temperature Range	-35 ~ 60°C (-31 ~ 140°F)

Site Location

Springbok 3 is located in Cantil, CA (Kern County), about 100 miles north-northeast of Los Angeles. The site was selected for the purposes of maximizing solar yield as well as other interconnection and land considerations. This project, along with its sister projects Springbok1 and Springbok2, connect to the Beacon substation.

Control Architecture

LADWP interfaces and implements MESA-ESS operating modes through the Doosan GridTech Intelligent Controller (DG-IC) via DNP3. The DG-IC communicates via MODBUS TCP/IP with the NEC AEROS controller to set a PV+S set point accordingly. The NEC AEROS system is continuously providing feedback to the DG-IC on the real time data of the system including PV generation, SoC, communication status, and other PV and BES status points. The NEC Aeros system communicates via MODBUS TCP/IP to the PV

inverter, each dc/dc converter, and each battery rack. It collects information from each and likewise distributes setpoints from the DG-IC. Additionally,

the DG-IC receives solar forecasts from an 8minute server via SFTP (Figure 1-3).




Figure 1-2
Diagram of communications architecture at power block #5




Figure 1-3
Diagram of solar forecast data flows implemented in this demonstration

Solar Forecasting at Springbok3

Various forecast services were assessed by 8minute prior to EPRI's involvement in the research. From their findings, 8minute selected the following two vendors based on performance, cost, and other considerations: Clean Power Research (CPR) and Reuniwatt.

Forecasts service horizons actively used by control algorithms in this project fell into two categories:

1. Long-term forecasts
 - T+1 hour to T+7 days at 1 hour granularity provided by CPR; refreshed every hour
2. Short-term forecasts
 - T+1 minute to T+30 minute at 1 minute granularity generated by Reuniwatt using an on-site sky camera (see photograph in Figure 1-4); refreshed every minute
 - T+1 minute to T+30 minute at 1 minute granularity generated by a satellite-based forecast from CPR; refreshed every 30 minutes



Figure 1-4
Photograph of Reuniwatt sky camera installed as part of short-term forecast service

The long-term forecasts were used by the DG-IC controller for the Predictive Co-ordinated Charge-Discharge (CCD-P) algorithm (discussed in the MESA Functions chapter, Chapter #3) and the short-term forecasts were used for the Ramp-Rate Product algorithm (managing PV ramp-rate by curtailing without the use of the BESS as discussed in Chapter #4). Accuracy of these forecasts would also be compared against one another and, where applicable, against day-ahead forecasts that LADWP already used (accuracy analysis in Chapter #2).

Project Partners

This pilot project was a collaborative effort between LADWP, EPRI, 8minute and many subcontractors who participated in discussions and development either directly or indirectly as orchestrated by 8minute.



Figure 1-5
Logos of primary project partners who participated
in this pilot

Chapter 2: Solar Forecasting Accuracy Assessment

Introduction, Objectives, & Chapter Overview

The intermittency of variable renewable energy sources (VRES) such as wind and solar adds significant uncertainty to resource scheduling and comes with reduced reliability. While technological alleviations like advanced PV plus BESS control functions exist (some of which will be addressed in subsequent chapters of this report), sufficiently accurate and reliable forecasts could go a long way toward mitigating the generation uncertainty of VRES. The following are a few examples of how forecasts (or improvements in those forecasts) could be applied:

- Reduce the amount of regulation reserve from traditional generation necessary to confidently be able to respond to real-time conditions with a certain degree of risk (due to increased confidence in forecasted day-ahead powers being realized)
- Provide partial PV power firming over the course of a charging window (and provide other benefits to the BESS) by using forecasts to charge more when PV exceeds the forecasted average over the charging window and charge less when it is below (Performance would improve as forecasted PV better aligns with realized production.)
- Potentially limit power ramp rates observed at the plant point of interconnection (POI) without the use of energy storage (to prevent sharp curtailments during momentary cloud cover that would otherwise result in conventional generation ramping up and consequently over-generating when PV returned)

The objectives of this project were to first, assess the accuracy of the forecast services provided (and any potential benefit that blending multiple forecasts might have), and second, determine if the above applications could be achieved with the use of those high-accuracy forecasts. In cases where

applications of the forecasts yielded positive results, quantifications of the potential benefits would be provided.

As an overview of this chapter, the accuracies of alternative and blended forecasts were assessed relative to the forecasts LADWP currently uses as well as against each other (where applicable). The first section will detail accuracy analyses from the past year of forecasts and measured plant data. The next section analyzes the potential for these improved forecasts to benefit dynamic reserve requirements with a thorough investigation of the existing method and an alternative risk-based method to which a higher accuracy forecast is applied for comparison. The chapter concludes with a brief investigation into the idea proposed at the project outset that a fully blended forecast might provide utility benefits.

Solar Forecast Accuracy Analyses & Comparisons

Background

One of the first tasks of the project was to compare the accuracy of LADWP's existing PV forecast supplier with those supplied by vendor 8minute. Unfortunately, apples-to-apples comparisons were complicated by differences in forecast durations, resolutions, and refresh rates. The following subsections will detail which forecast services were available from different vendors and how subsets of those services were compared.

Available Forecasts

Up through April of 2021, LADWP used numerical weather prediction or NWP-based forecasts from Meteologica for use in resource planning; however, these forecasts were only day-ahead forecasts provided at 11 am on the day prior to the forecasted day. No short-term forecasts were being used by LADWP at the time of this demonstration.

For reasons discussed in the Hybridized Forecasts section of this chapter, 8minute supplied LADWP with multiple different forecast types from two different vendors. Accuracy of the long-term CPR SolarAnywhere forecast, which blends multiple numerical weather predictions (NWPs) with satellite imagery, was compared against the NWP-based, Meteologica, day-ahead forecasts that LADWP currently receives. Meanwhile, short-term forecasts from CPR's satellite-based SolarAnywhere forecasts series and Reuniwatt's Instacast forecasts (based off inputs from a ground-mounted sky imaging camera) were compared to determine which had higher accuracy and would thus be the better choice for additional applications such as advanced PV plus BESS functions. The full collection of forecasts available for analysis consisted of two short-term forecasts, one medium-term forecast, and two long-term forecasts (one of which was the baseline Meteologica forecast that LADWP was already using). Table 2-1 summarizes the forecast services made available in this project, their sources, and key attributes of each.

Table 2-1
Summary of forecast sources available in this project

Horizon Classification (Project Specific)	Precise Horizon	Resolution	Update Frequency	Forecast Name from Commercial Supplier			Algorithm Input Source
				Meteologica	Clean Power Research	Reuniwatt	
Long-term	0-7 days ahead	1 hour	Every 30 min		<i>Solar-Anywhere</i>		Blend of satellite & multiple NWP
Day-ahead	Next day (13-36 hours ahead)	1 hour	Once per day at 11 am local	Solar Generation	(Subset of <i>Solar-Anywhere</i> forecast) ¹		NWP
Medium-term	0-6 hours ahead	15 min	15 min			<i>HourCast</i>	Satellite
Short-term	0-1 hour ahead	1 min	Every 30 min		<i>Solar-Anywhere</i>		Blend of satellite & multiple NWP
	0-31 min ahead	1 min	Every 1 min			<i>InstaCast</i>	Sky imaging camera

Forecast Assimilation Process

Solar forecast accuracy evaluation would typically be performed across all available horizons and for every update received. If any accuracy comparison is to be made between vendor products, however, only equitable sets of forecasts can be used for the comparison to be fair. To clarify, all forecasts are evaluated against realized measurement in the field to determine their accuracy, but by considering only equitable forecast horizons in the analysis, the resulting accuracy metrics for each can then be fairly compared.

In the case of long-term forecasts, CPR does not explicitly provide a day-ahead forecast like what LADWP receives from Meteologica; however, CPR provides forecasts sufficient in horizon, resolution, and update rate from which an equivalent forecast can be generated via sampling. Figure 2-1 illustrates how an equitable “day-ahead” forecast can be found by selecting the segment of CPR’s forecast series generated at 11 am each day. Considering only the portion highlighted in yellow in the analysis allows the accuracy metrics from CPR’s long-term forecast and Meteologica forecasts to be fairly compared.

¹ Not an official forecast offering




Figure 2-1

Visualization of selective sampling from CPR forecasts used to generate the equivalent “day-ahead” forecasts to be compared to those from Meteologica

A similar process was used to be able to compare the accuracy metrics between short term forecasts from CPR and ReuniWatt, except in this case, the intersection of available forecast horizons and update rates required sampling of both forecasts. The left side of Figure 2-2 shows how CPR supplies a 1-hour ahead forecast whereas Reuniwatt only

supplies 31-minute ahead forecasts, so only the first 31 minutes could be compared. Reuniwatt also updates their forecast every 60 seconds whereas CPR only updates theirs every 30-minutes, so the right side of Figure 2-2 shows how only Reuniwatt forecasts generated on the hour or half hour were considered.




Figure 2-2

Visualization of selective sampling from CPR (left) and Reuniwatt (right) forecasts used to analyze equivalent short-term forecasts for accuracy

Irradiance Forecasts vs. Power Forecasts

The term “solar forecasts” can mean a few different things. In the context of PV, solar forecast usually refers to a prediction of either irradiance expected to hit the site (can be that which hits a horizontal plane or that in the plane of the PV arrays) or the power/energy generation expected from the PV plant during that time. In almost all cases, the irradiance forecasts will generally be more accurate than the power forecasts since power forecasts are

usually based on the irradiance forecasts and then add in many other models and assumptions, each that can contribute to increased error. Using error over different forecast horizons as an example, Figure 2-3 demonstrates that the median error for global horizontal irradiance (GHI) is about 2% lower across all time horizons than the median error for predicted plant power output. Figure 2-4 provides more than a dozen potential sources for the error along the multistep process of transforming irradiance forecasts into power forecasts.




Figure 2-3
Absolute error by forecast horizon for GHI (left) and power (right); CPR short-term forecast used in illustration




Figure 2-4

Flowchart of possible (not comprehensive) sources of error between a GHI forecast and a plant power forecast

While investigating the accuracy of irradiance forecasts can be informative for forecast suppliers to tune their models, this project predominantly focused on power forecasts. This is because, whether the power forecast is supplied by the vendor or calculated from irradiance forecasts by the utility themselves, PV forecasts ultimately need to be in terms of power if they are to be useful to utility planning and operations.

Long-term Forecasts

In the category of long-term forecasts, equivalent day-ahead forecasts from Meteologica and CPR were compared against plant power measurements to see whether the blended forecasts (from CPR) provided an accuracy improvement over forecasts based on a single NWP source. Accuracy metrics for each source are then presented side by side for comparison.

Full Plant Availability Filtering

The focus of the forecast analysis was to calculate the accuracy of different power forecasts. However, the pilot project experienced many periods,

especially early on, where one or more power blocks within the PV plant were not operational or were not reporting for various reasons. Low, full-plant availability was primarily a consequence of downtime to address issues with inverter switching devices early in the project. These issues were resolved by May 2021.

The research team decided that the intent of the research in this project was to assess the inherent accuracy of the PV power forecasts rather than the ability of the forecasts to account for plant outages. Therefore, forecast accuracy analyses only considered those time periods in which *every* power block in the plant was fully functional (and thus presumably capable of operating at full power). Since no such status point had been defined or communicated that would indicate “functionality”, a proxy was used instead. Each power block would be considered functional if its measured PV production was greater than zero and non-NUL during that period. Then, if even one power block was not reporting power generation, the period was excluded. The availability plots in Figures 2-5, 2-7, 2-10, 2-12, and 2-13 represent this strict approach.

Accuracy Analysis

The first metric looked at median absolute error² across all horizons and days in each month as depicted by the bars in Figure 2-5. As noted, however, only the subset of forecasts pertaining to periods of full plant operation were considered, and not every month had the same plant availability. The dot plot in Figure 2-5 represents the percentage of hours in that month included in the plant availability filter. The higher that number, the more reliable the result is for that month. Of the 11 months with data (after filtering), blended forecasts from CPR had a lower median error than the forecasts (from Meteologica) in 10 of the 11 months. Taking the average across months, median absolute error was 3.5% for CPR and 12.5% for Meteologica. This difference was especially pronounced in late spring through early fall time periods.




Figure 2-5
Median **absolute** error by month for equitable long-term forecasts from Meteologica and CPR, filtered down to the time periods when all component power blocks were reporting and operational †11 am forecast set only

The primary reason for the seasonal difference in forecasts can be explained by the timeseries plots in Figure 2-6 and median bias error plots in Figure 2-7. The Meteologica forecasts appeared to be more

consistent year-round in their forecasts, particularly in the noon-peaking shape and peak magnitude of around 75 MW (See Figure 2-6).³ This underprediction, seen as negative values in the bias plot (Figure 2-7), led the Meteologica forecast to be less accurate during the summer months when production regularly reached 90 MW but more in line with realized power output in the months with less sunlight and cloudier conditions.




Figure 2-6
One week sample of forecasts and actuals from two weeks in different seasons




Figure 2-7
Median bias error by month for equitable long-term forecasts from Meteologica and CPR, filtered to when all component power blocks were reporting and operational †11 am forecast set only

² Absolute error is the magnitude of error in either the higher or lower direction

³ It almost appeared as if the plant model for Meteologica did not take into account the single axis tracking that was present at Springbok 3, but this was not confirmed.

Digging one level deeper than a single number per month, the bias error by time of day was analyzed for the more accurate, CPR forecasts (Figure 2-8). These long-term forecasts, which tended to

overpredict overall, displayed a regular pattern of greater overprediction in the mornings, especially, but also late afternoons (see the median value in red).




Figure 2-8

Bias error by forecast horizon for CPR day-ahead portion of long-term forecast

The reason for the lower error during the middle of the day was mostly due to the plant controller limiting power to the 90 MW interconnection limit, even though higher production was sometimes possible. Figure 2-9 illustrates how a theoretical

forecasted profile before application of the power limitation can well exceed the limit. Most any error present in this raw production forecast gets collapsed down to nearly zero since 90 MW is a known maximum during this high-production period.




Figure 2-9

Visualization of where raw forecasts (before considering plant export limits) would collapse down to a known export limit, reducing error

Short-term Forecasts

In the category of short-term forecasts, equivalent half-hour-ahead forecasts from CPR (based on satellite imagery) and from Reuniwatt (based on a ground mounted camera) were compared against plant power measurements. Accuracy metrics for each source are then presented side by side for comparison. As in the long-term forecast analyses, only periods assumed to have full plant availability were considered.

Accuracy Analysis

Considering all months, median absolute error for the two forecasts was similar, with a mean separation between forecasts of about 0.4% and a maximum separation of 4.4% in October of 2020. The performance of both forecast sources appeared to vary seasonally, as shown in Figure 2-10; error was generally elevated in Fall and Winter and depressed in Spring and Summer. Peak error for the two forecast sources appeared offset by about two months, with error for the satellite-based CPR forecasts peaking in early Fall whereas error for the Reuniwatt, ground-mounted, sky imaging camera forecast peaked in late Fall / early Winter.




Figure 2-10
Median **absolute** error by month for equitable short-term forecasts from CPR and Reuniwatt, filtered to when all component power blocks were reporting and operational
† First 31 minutes only
‡ XX:00 or XX:30 forecasts only

A major contributing factor to short-term forecast accuracy, especially for short-term forecasts, was variability of the solar resource. Figure 2-11 illustrates how the bulk of error values (median is the red line and interquartile range is the blue box) increase with the variability classification of the day. The more variable the day, the higher the typical error observed.




Figure 2-11
Box and whisker plot indicating distribution of absolute error by solar variability classification for the CPR short-term forecasts

The correlation with variability also explains why absolute error was typically lower in the summertime for both forecasts. Springbok 3 resides in a place where summer months experience a preponderance of clear days, which have the lowest error whereas Fall and Winter months experience much more cloud cover that induces variability, and results in less accurate forecasts.

Most of the error for the CPR forecast came from its propensity to overestimate output, especially in the fall and winter months, as observed in Figure 2-12. Reuniwatt forecasts were less consistent in this regard, overpredicting some months and underpredicting others.




Figure 2-12

Median **bias** error by month for equitable short-term forecasts from CPR and Reuniwatt, filtered to when all component power blocks were reporting and operational

† First 31 minutes only

‡ XX:00 or XX:30 forecasts only

As well as being about 0.4% more accurate throughout the year, Reuniwatt short-term forecasts usually had a smaller standard deviation of error values (smaller by anywhere from 0% to 6% of rated power) relative to CPR short-term forecasts (Figure 2-13). Despite this, forecasts errors occasionally experienced large misses. Error for Reuniwatt's InstaCast forecasts exceeded 50% of the plants power rating 1% of the time (Figure 2-14) and CPR's short-term forecasts exceeded the same threshold 2.4% of the time (Figure 2-15). In summary, while both satellite-based short-term forecasts from CPR and sky-camera-based short-term forecasts from Reuniwatt were accurate overall, especially for clear days, individual forecast errors could still be quite large.




Figure 2-13

Standard deviation of absolute error by month for equitable short-term forecasts from CPR and Reuniwatt, filtered to when all component power blocks were reporting and operational

† First 31 minutes only

‡ XX:00 or XX:30 forecasts only





Figure 2-14

Cumulative distribution function (CDF) indicating the percentage of Reuniwatt InstaCast forecasts (within viable periods) with a given absolute error value or less.



SolarAnywhere forecasts (within viable periods) with a given absolute error value or less.

The potential for large error can be seen in the sample day shown in Figure 2-16. In this figure, the measured PV output (yellow) is compared against the Reuniwatt forecast from 10 minutes prior (green).⁴ In some cases, the forecast tracks the actual changes. Other times, sharp changes/“features” are completely missed or offset in time by a few minutes, which has significant implications for functions that need accurate short-term forecast inputs such as the Ramp Rate Product discussed in the next chapter.

Figure 2-15
CDF indicating the percentage of CPR short-term




Figure 2-16
Measured/realized PV powers and corresponding forecasted value (Reuniwatt InstaCast) from 10 minutes prior, demonstrating inconsistent prediction of rapid PV changes

The issue is not unique to Reuniwatt, either. While a constant horizon forecast was not available for comparison, Figure 2-17 shows successive 30-minutes ahead forecasts from both short-term forecast providers. Short term forecasts from both vendors exhibit periods where their forecasts capture rapid event characteristics and periods where they seem oblivious to them.

⁴ While Reuniwatt was used as an example, short-term forecasts from CPR exhibit similar behavior.




Figure 2-17
Side by side comparison of short-term forecasts from two different vendors with realized (measured) values behind each for reference

Forecast Improvements Over Time

At the project outset, the research team hoped to observe improvements in PV forecast accuracies over the one-year test period arising from the combination of measurement feedback and machine learning helping to fine-tune algorithms. Unfortunately, excluding one manual update, numerical improvements over time were not observed due to the lack of cleansed feedback data available to forecast vendors. While data flows were put in place to transfer telemetry data to forecast vendors, defining and implementing data cleansing criteria were necessary actions but outside the scope of this project.

A quick aside on the manual correction of a plant model by one of the forecast suppliers: In observing time series plots like that in Figure 2-18, researchers noticed that the plant model used by CPR to generate forecasts originally did not consider an interconnection injection limit of 90 MW. As a result, forecasts overpredicted the measured power that saturated at this limit (blue line in figure). The correction consisted of updating the model to saturate output at the 90 MW interconnection limit.




Figure 2-18
Short-term forecasted and measured power profiles (from before the correction) highlighting the error previously induced by the missing export limit feature in CPR's PV plant model

Excluding this one manual change, however, numerical improvements over the past year were not discernable, including in global horizontal irradiance (GHI) forecasts that typically feed into plant models to estimate power. Forecast accuracy metrics that compared irradiance forecasts to measured irradiance on site can be seen in Figure 2-19. Neither the median absolute error nor the standard deviation of the error was noticeably reduced over time. If improvements were present, seasonal changes and non-linearities (such as the generally greater power magnitudes observed during Summer) accounted for large enough differences that algorithmic improvements were in the noise.

returned to forecast suppliers. If supplied, forecasting companies could then filter the measured values by this point when working to improve their algorithms and plant models. Forecast vendor CPR informed EPRI that they were unable to use the telemetry data being sent back to them to improve their model because doing so with all the periods of plant downtime would have degraded the performance.

Getting ‘real-time’ telemetry data, including the most recent irradiance and power measurements, could also be a beneficial form of feedback for forecast vendors. This, however, proved elusive for short-term forecasts in this pilot due to data latency issues.




Figure 2-19
Global horizontal irradiance forecast absolute error and standard deviation of that error by month

Forecast vendors are continuously working to improve their models, however. CPR, for example, reported that they plan to implement updates soon related to satellite hardware calibration and aerosol optical depth & clear sky modelling, both of which expect to improve the weather-based portions of their forecasts. It is just that these further improvements may be difficult to discern without being able to keep other factors fixed or at least compensate for or normalize them. (i.e., isolate improvements from other factors that can obscure those improvements).

One learning from this project is that data such as a ‘full plant availability’ bit, which would indicate periods of outages (both planned and unplanned), might be a valuable addition to plant data being

Investigation of Improved Forecast Accuracy on Dynamic Reserve Requirements

This subsection assesses the impact of improved solar forecasting on flexible reserve requirements. That is, given an improved system-wide PV forecast, by how much could the resulting reserve requirements be reduced, since the expected system uncertainty would effectively be reduced. Content in this subsection is only a portion of a fuller investigation performed during this project that compares LADWP's current reserve methods with a dynamic, risk-based reserve calculation method developed by EPRI. For the in-depth comparison and context surrounding the assumptions that preceded the following analysis, see Appendix A.

Method to Generate Improved Forecasts

The method used to generate improved forecasts simulates a desired percentage improvement in uncertainty applied uniformly to all forecasted values. First, the error between historical forecasts and realized measurements is computed. This error is then reduced by the desired percentage and added to the existing historical actual. A visual depiction of the process applied to a sample day is shown in Figure 2-20. In it, the resulting output (purple) lies a fixed fraction of the way between the original forecast (green) and the actual output (blue).




Figure 2-20
Summary of the method to produce improved forecasts

For this analysis, an improvement of 50% was chosen. This relatively large forecast improvement was selected in order to extoll the potential advantages of improved forecasts in general; it does not suggest that such improvements are easily attainable, nor does it provide suggestions as to how they would be attained. Likewise, the uniform improvement across all forecast values was done for simplicity; real, incremental forecast improvements would likely benefit select conditions more than others. Finally, please note that, at present, system-wide PV generation forecasts are generally not derived by aggregating production forecasts from individual PV plants like those analyzed elsewhere in this section. Unless this fact changes, improvements in individual plant forecasts may not influence bulk system forecasts, though independent improvements in bulk PV forecasts are most certainly still possible.

Descriptive Statistics of Error

Figure 2-21 shows the magnitude and direction of realized deviations from the DER forecast for both the original DER forecast (top) and the DER forecast that contained the improved PV forecast (bottom). It can be seen that if the solar forecast is improved, the aggregated system error magnitude reduces significantly during hours of solar production.




Figure 2-21

System deviations for a) the original forecast case and b) the improved solar forecast case; full year left, sample day right

To illustrate the reduction in error more clearly, the isolated deviations for just solar are shown in Figure 2-22.




Figure 2-22

Isolated solar deviations for a) original forecast and b) improved forecasts; full year left, sample day right

In Figure 2-22, the effect of the improved forecasts on the magnitude of the deviations is quite evident. High-quality forecasts can reduce the magnitude of these deviations, and therefore to reduce the overall needs for reserve. The theoretical limit would be to have forecasts that are 100% accurate, and in such cases, there would be no need to allocate reserves for uncertainty.

Impact on the Reserve Requirements

The most valuable feature of the improved forecasts is that it allows drawing more precise operating plans for the system, and in particular, reduces the need for additional recourse to deviations (i.e., reserves). The figure below (Figure 2-23) shows the impact on the LADWP policy for year 2019.




Figure 2-23

LADWP reserve determination policy for a) the original solar forecast and b) the improved solar forecast. Color bands represent percentile ranges of calculated reserve needs over the course of a year at each given time of day.

Figure 2-23 shows that the amounts of reserve that the LADWP estimate with or without the improved solar forecasts are virtually the same. This is because the LADWP reserve policy focuses on the magnitude of VRES production, as opposed to the quality of the forecast of the VRES production. That means that even if it were possible to predict the VRES power with complete accuracy, the LADWP policy would still produce roughly the same reserve requirements.

EPRI's dynamic reserve determination methods compute the reserve requirements based on the difference between historical forecast and historical actuals (see Appendix A for the complete description of the reserve determination method inside EPRI's DynADOR tool). That means that if the quality of the forecasts is improved, the magnitude of the deviations reduces, and this is translated into lower reserve requirements. The exact reserve requirements for the LADWP system using both, the original forecasts and the improved forecasts are shown in Figure 2-24.




Figure 2-24

Exact up and down reserve requirements for the LADWP system for both the original forecast and the improved forecast. Color bands represent percentile ranges of realized reserve requirements over the course of a year at each given time of day.

Figure 2-24 shows that as the quality of the forecasts improves, the exact needs for reserves in the system reduces. This is a desirable feature since it will translate into a more efficient utilization of the recourse resources while maintaining the same level of system reliability. The actual reserve requirements for the LADWP system are shown in Figure 2-25.

In Figure 2-25, it is evident that the reserve requirements reduce as the quality of the forecast improves, even maintaining the same confidence interval (CI). For instance, for a CI = 99% (which

means that 99% of the expected deviations would be captured by the reserve procured) the amounts of up-reserve for the original forecast are above those obtained with the LADWP policy. However, considering an improved forecast, the same CI of 99% results in a reserve profile that is below the LADWP policy during the high solar power production period. The amounts of up-reserve for the night and early morning are the same using EPRI's dynamic reserve determination method since these reserves are driven by factors other than PV and were assumed to be the same for both cases in this analysis.





Figure 2-25

Up and down reserve requirements for the test day in LADWP using both the original forecast and the improved forecast for solar generation


To assess the expected performance of the reserve, it becomes necessary to quantify the expected number of potential violations (i.e., materializations above the reserve requirements) that this reserve profiles would not hedge against. Risk is a good metric since it considers the likelihood of the potential violation, as well as the potential extent/magnitude. Figure 2-26 shows the

reserve profiles and their associated risk for the LADWP policy (left pair of figures, a & c) and the dynamic reserves obtained with EPRI's clustering method for a CI of 90% (right pair of figures, b & d). The top row of Figure 2-26 shows the reserves procured for the original forecast, and the bottom row shows the same but for the enhanced forecast.


a) Reserve and risk for the LADWP reserve policy, original forecast



b) Dynamic reserves and risk for CI = 90%, original forecast



c) Reserve and risk for the LADWP reserve policy, improved forecast



d) Dynamic reserves and risk for CI = 90%, improved forecast




Figure 2-26

Up and down-reserve requirements and associated short and long risk for both, the original solar forecast and the improved solar forecast

In this prior figure, reserve profiles under the LADWP policy are virtually the same in both cases, with the only exception being a reduction in the already low, mid-day risk of being short. This risk is slightly lower in the case of the improved forecast, but the reduction is incidental since the amount of reserve is the same in both cases. The reason for this reduction is because the quality of the forecast is better, meaning that the deviations are lower in magnitude. The risk for this policy in the periods in which there is no solar production remains large and the same for both cases.

When using the dynamic reserves from EPRI's reserve determination tool, the profiles are similar

in shape for both cases, but the overall amount is lower for the case of the improved forecast. Additionally, the risk profile is also lower for the improved forecast. This would translate in actual system operation as economic savings (lower reserves), while maintaining or potentially enhancing the system reliability.

The above reductions can be difficult to see Figure 2-26, but can be better appreciated in Fig. 24, which shows the average amounts of up and down reserves over the whole period (top) as well as the average short and long risks over the whole period (bottom).




Figure 2-27

Average reserves (top) and average risk (bottom) assessment for the LADWP reserve policy and a collection of outputs from EPRI's reserve determination tool with different confidence intervals. Each reserve policy has one set of results for the original forecast, 'o', and one set for the improved/enhanced forecast, 'e'.

Reiterating, the reserves provided using the LADWP policy are the same in both the original ('o') and enhanced ('e') cases, and the short risk is slightly lower. As for the dynamic reserve requirements obtained using EPRI's reserve determination tool, the average reserves for any confidence interval (CI) are lower when using the enhanced ('e') forecasts compared to the original forecast ('o'), and depending upon the confidence interval required, can be lower than what LADWP currently uses for their reserve policy. Additionally, the average risks are also lower considering the improved forecasts. These results point to the conclusion that if used in actual system operation, the result of improved forecasts would be economic savings while attaining the same or even greater reliability levels.

For those interested in a refresher on dynamic reserve requirements, data requirements and preparation needs to compute those requirements, and a comparison of the EPRI vs. existing LADWP approach used to calculate the operating reserves, Appendix A provides the full reserve requirements analysis (including the above content in context).

Probabilistic Forecasts

In addition to forecasting the most probable irradiance or PV power output for a certain time (also known as the deterministic forecast value), probabilistic forecasts provide other values that help indicate the range of possible values that could occur at a given time and the chance of those values occurring. Since the range of possible values is continuous, discrete probabilities or quantiles are provided as part of a probabilistic forecast. For

example, a P-10 forecast would give an approximate value of power that (for similar forecast conditions) is expected to be exceeded 90% of the time (i.e., the realized values should be below the P-10 forecast in about 10% of trials). As another example, a P-90 forecast would give an approximate value of power that (for similar forecast conditions) is expected to be exceeded only 10% of the time (i.e., realized values should be below the P-90 forecast in about 90% of trials). Alongside deterministic forecasts, probabilistic forecasts could provide value to planners looking to reduce risk in their plans by giving them a better idea of the range of possible values to expect at a given time.

In this project, probabilistic long-term PV forecasts from CPR were analyzed as a proof of concept using the Solar Forecast Arbiter, an open-source forecast evaluation platform.^{5,6} Many metrics could be used to assess the performance of probabilistic forecasts, but the following three were selected:

- **Brier Score** – Paraphrased meaning: Judging in an all-or-nothing fashion at each observation, how often the probabilistic forecast correctly captured events & dynamics of the plant; BS of 0 = perfect

Equation:

$$BS = \frac{1}{n} \sum_{i=1}^n (f_i - o_i)^2$$

forecast probability
(0=0%, 1=100%) observed outcome
(0=didn't occur, 1=did occur)

- **Quantile Score** - Paraphrased meaning: How close the probabilistic forecast value was to the quantile value it was meant to represent with a penalty for how confident you were in that answer; QS of 0 MW = perfect

Equation:

$$QS = \frac{1}{n} \sum_{i=1}^n (\text{error [MW]}) \cdot (\text{confidence [\%]})$$

error [MW] confidence [%]

- **Reliability** - Paraphrased meaning: Of the forecasts associated with a given forecast probability, how close the actual frequency of occurrence was to the quoted probability; REL of 0 = perfect reliability

Equation:

$$REL = \frac{1}{n} \sum_{i=1}^I N_i (f_i - \bar{o}_i)^2$$

The Brier Score (BS) can be a useful tool to determine whether a forecast falls on the desired side of a certain probability threshold. Figure 2-28 shows the BS for forecasts being less than the 10th and 90th percentile thresholds at various forecast horizons. CPR's long-term probabilistic forecasts were better at forecasting when power would be less than its 10th percentile forecast than it was at predicting when power would be less than its 90th percentile forecast, likely due to increased power variability at higher powers. One shortcoming of the BS is that it is an all or nothing assessment, so even if the probabilistic forecast was only slightly greater than the value of the precise percentile threshold, it would count against the performance.⁷

⁵ <https://solarforecastarbiter.org/>

⁶ This tool was developed by the University of Arizona in partnership with EPRI, Sandia National Lab, and Sharply Focused as part of a Department of Energy Solar Energy Technologies Office project.

⁷ Going forward, this metric might make more sense on the upper end if the reverse convention is used [Prob(f >= x)] since the probability of actual powers exceeding an upper forecasted probability threshold could be equally as undesirable as the probability of actuals falling below a lower forecasted probability threshold.




Figure 2-28

Brier scores for the 10th and 90th percentiles at different forecast horizons, lower is better (data from two months of long-term CPR forecasts)

The Quantile Score (QS) aims to address the “all-or-nothing” shortcoming of the Brier Score, by considering how close the actual power value was to the power value of the target percentile (as opposed to just the percentile that power value belonged to and whether it was greater than the threshold). Figure 2-29 reveals that even though the 90th percentile forecasts were less accurate than the 10th percentile forecasts in terms of closeness of

percentile (Brier Score), the 90th percentile forecasted power values were generally closer to the true 90th percentile power value than the 10th percentile power forecasts were to the true 10th percentile power value.

As an aside, both the BS and QS were largely unaffected by forecast horizon, though the 10th and 90th percentile scores did slightly trend toward the same score for longer horizons.




Figure 2-29

Quantile scores for the 10th and 90th percentiles at different forecast horizons, lower is better (data from two months of long-term CPR forecasts)

The third and final metric considered in this project was the Reliability Score (REL), where a lower score means that the probability of an event (percentile of forecast being less than the realized percentile) occurs as often as it says it will. By this metric, the

longer-horizon probabilistic forecasts (e.g., 36 hours ahead) better captured the true probability as opposed to the shorter horizons (Figure 2-30), possibly due to differences in models used for shorter vs longer horizons.⁸




Figure 2-30
Reliability scores for the 10th and 90th percentiles at different forecast horizons (data from two months of long-term CPR forecasts)

Probabilistic forecasts and assessment thereof may prove valuable as utilities advance their probability- and risk-based methods of planning and reserves. As a reminder, though, the probabilistic forecasts presented in this report are for a single PV plant. Equivalent probabilistic forecasts for the PV at the bulk level will likely be necessary before such forecasts can be employed for reserves or other planning calculations.

Additional applications of probabilistic forecasts on the plant level could be probabilistic ramp rate forecasts and, for internal battery SOC control, statistical scenarios. Investigation of both was outside the scope of this project.

⁸ CPR uses a blend of a satellite-based method for shorter horizons (horizons in minutes) and numerical weather prediction (NWP) for longer horizons. The generally lower reliability values for 12, 24 and 36 hours ahead (compared to

1, 3 and 6 hours) might be due to difference in source data (NWP vs satellite).

Hybridized Forecasts

A variety of multiple input sources and numerous proprietary algorithms have historically been employed to generate the most accurate PV generation forecasts possible. However, algorithms are usually based on a single input type such as meteorological weather data, satellite imagery, or ground-based cameras. When designing this project, it was suspected that combining the best features from one or more data input sources might yield an improved PV forecast. However, during the course of the project, it was discovered that the various inputs 1) were usually associated with distinct forecast time horizons, 2) targeted different use cases, and 3) were often employed by distinct vendors with their own proprietary algorithms. For all these reasons, further blending of these forecasts (beyond what was provided by individual forecast suppliers) was not believed to be worth the effort, nor necessary in this pilot.

Forecast Horizons and Use Cases

When assessing the uses of solar forecasts in this pilot project, two main categories emerged:

1. Longer term (e.g., day-ahead or longer), lower time resolution forecasts for day-ahead planning and as input to the Coordinated Charge/Discharge with Prediction (CCD-P) function detailed in the MESA-DER Functions chapter
2. Short-term (sub hourly), higher time resolution forecasts for use in the Ramp Rate Product function detailed in the Ramp Rate Limiting with Solar Forecasts chapter

To best address each need, different technologies (and their associated algorithms) are used. For example, the ground-mounted sky imaging camera that Reuniwatt used to generate its short-term forecasts proved valuable in generating forecasts with slightly higher accuracy than those obtained from CPR's short-term satellite imagery; however, Reuniwatt does not offer forecasting of next-day generation with this technology. Likewise, satellite imagery that CPR uses in the nearer-term part of its long-term forecasts provided an accuracy bump in day-ahead forecasts over those from Meteorologica that are informed by numerical weather predictions (NWP); however, satellite-based forecasts become

less capable of providing accurate forecasts beyond about 24 hours. It is for this reason that CPR already combines the satellite-based forecasts with NWP to generate estimates 2-7 days ahead. Thus, since the use cases are already relatively discrete, and certain technologies are best suited to different time horizons, combining them does not inherently improve the value that each can provide separately.

Qualitative Logistics, Costs, and Benefits of Blending Forecasts

The only possible time horizon left to blend would have been the short-term satellite-based forecasts with the sky imaging camera. However, unlike longer-term forecasts where blending can potentially cancel out errors to converge toward reality, the key feature of short-term forecast will be their sharpness in predicting ramps and rapid changes. Simply taking a weighted average will smooth out these features and make the forecast less useful. Thus, any further improvement would have had to come from the algorithmic level rather than simply merging after the fact. Doing so presented a logistical challenge in this project since the two forecast types came from competing vendors. If there are advantages to blending ground and satellite-based imagery, investigations into the logistics and potential accuracy advantages would likely be done by the individual vendors.

The other theoretical benefit of blending forecasts is that utilities could potentially use a single source for all their PV forecasting needs. This potential convenience in set-up/configuration did not seem to warrant the merging of forecasts with two different time resolutions, however. The forecast would still need to be parsed for the relevant portion belonging to each use case, and the mixing of forecast time resolutions would either result in heterogeneous time resolutions or very large file sizes (if homogenized to the finest granularity timestep).

In the end, it was concluded that keeping short-term and long-term forecasts separate was the most pragmatic approach for this project.

Chapter 3: MESA-ESS/MESA-DER PV plus BESS Plant Control Modes

Objective

Energy storage is playing an increasingly large role in the integration of renewable energy sources, as it provides flexibility and dispatchability in the face of intermittent renewable generation. In the ideal case, PV combined with storage would be at least as dispatchable as thermal generators are today. However, energy storage is only useful if combined with intelligent algorithms that determine when and how much to charge or discharge the finite-capacity battery. In part, this project aimed to demonstrate and evaluate a sample of such algorithms, both to determine their technical readiness and to better understand what features and limitations they have.

Background and Test Plans

In this project, one power block of the 32 at the Springbok 3 site, Power Block #5, was equipped with 1.5 MW, 1.5 MWh battery energy storage system (BESS). The battery is dc-coupled with 3.77 MW of PV behind a 3.225 MW/3.225 MVA inverter as shown in Figure 1-1. By coupling energy storage with the PV array, the combined system can perform advanced inverter functions defined in the MESA-DER specifications (a superset of the former MESA-ESS standard).⁹ The specific functions and their performance are detailed in subsequent subsections, first as individual/standalone functions, then as combined or “stacked” sets of functions.

8minute worked with engineers from LADWP’s Energy Control Center (ECC) to implement multiple interfaces with the combined PV plus BESS system.

In addition to the human-machine interface (HMI) available over remote desktop protocol (RDP), connectivity was successfully demonstrated in August 2020 between the onsite controller and LADWP ECC SCADA via DNP3 communications. Integration with SCADA allows ECC to configure settings and parameters on the remote unit according to the instruction manual provided by 8minute¹⁰ without needing to log into the HMI.

Test Plan Development

Testing of the MESA functions consisted of multiple parts that were organized into a single test plan. EPRI worked in conjunction with LADWP and 8minute to define appropriate test parameters, schedule the tests, and, once all parties approved, coordinate execution of the tests with LADWP’s ECC.

The test plan itself stepped through at least one test day for each control mode to demonstrate that the system could perform that function and to set a baseline accuracy. The latter portion of the test plan then stepped through combinations of the individually tested function/control modes with the same parameter arguments and reassessed the accuracy of each component function in the “stack”. The accuracy of the function could then be compared between the standalone test and the stacked test to determine if mode stacking resulted in decreased performance for any one capability.

Settings for all tests were primarily selected to demonstrate functionality of each mode; they should not be considered as recommended settings for this or other power blocks. In some instances,

⁹ The MESA Standards Alliance works with utilities and technology suppliers to help inform communications and data exchange specifications to aid in the standardization of integration with systems like this PV plus storage unit at Springbok 3, PB#5. More information about their organization

and links to their specifications can be found at <http://mesastandards.org/>.

¹⁰ Doosan GridTech, Inc. (Sep 2020). Doosan GridTech Intelligent Controller Operating Manual (User Manual v3, Rev 1.03).

surface-level assessments were made to determine settings that would induce action (such as the recent range of frequency variances), but longer-term and more in-depth system studies would be needed to determine the best possible settings for each mode. Some recommended considerations have been provided in subsequent sections.

Demonstration of Isolated MESA-DER Control Modes

The first component of demonstrating MESA-DER control modes consisted of testing individual functions with hypothetical test settings developed by EPRI in agreement with 8minute. These functions could be classified as either basepoint or additive functions as outlined in the Basepoint vs. Additive sidebar.

Coordinated Charge/Discharge (CCD)

The Co-ordinated Charge/Discharge (CCD) control mode is a basepoint mode used to reach a specified target State of Charge (SoC) by a specified target time. This function can enable energy arbitrage with strategic selection of SoC target values and times.

Consider a hypothetical scenario in which mid-day energy is assumed to be cheap due to an abundance of solar energy whereas energy is assumed to be expensive in the evening due to peak loads. Figure 3-1 illustrates how a low SoC target in the morning followed by high SoC target in the afternoon (targets represented by blue stars in top figure) encourages charging with solar energy during the middle of the day (see bottom figure). Similarly, high SoC targets in the afternoon followed by low SoC targets at night (top figure) encourages discharging of the battery in the evening (bottom figure). While the algorithm is not factoring in energy price values or time of use rates at this time, Figure 3-1 demonstrates how low and high SoC targets can be used to control the charge/discharge behavior of the battery in response to typical energy price patterns.

Basepoint vs Additive:

Basepoint: Basepoint control modes are planned active power control modes that can be “stacked” on top of other control modes. Basepoint control modes can be combined using an assigned priority in such a way that the functional goals of multiple control modes can be satisfied.


Each active power **basepoint** control mode outputs a minimum, maximum and ideal BESS power setpoint. The min and max setpoints of an active power control mode represents the minimum and maximum BESS power a control mode can accommodate while still being able to satisfy its defined requirements. The ideal setpoint represents what the control mode would ideally like to do. The combination of all the basepoint control modes together contribute to a final “basepoint” (as opposed to a setpoint).

Additive: The additive power control modes result in a single output and do not follow the min, max, ideal paradigm. The additive power control mode setpoints are simply added on top of the basepoint.

The modes in **blue** result in BESS active power setpoints that may contribute to the “basepoint”. The modes shown in **green** are “additive” modes that add positive or negative power to the basepoint.

- Active Power Limit control mode
- Active Power Response control mode
- Charge / Discharge storage mode
- Co-ordinated Charge / Discharge (CCD) control mode
- Active Power Smoothing control mode
- Automatic Generation Control mode
- Frequency-Watt control mode

The figure below illustrates an example of stacked control modes coalescing into one resultant command.






Figure 3-1
Profile of simulated SoC values and targets (top)
with corresponding simulated power outputs from
the PV, BESS/ESS, and net power at point of
interconnection (bottom)

There are two flavors of the CCD algorithm that can perform this energy shifting, the base algorithm informally called “greedy” CCD and a proprietary implementation that is called Predictive CCD (CCD-P). Each will be analyzed in turn and then their

qualitative impact on net power profiles, battery availability, and other effects will be compared.

“Greedy” CCD

“Greedy” CCD attempts to reach the desired SoC charging targets (those targets above the current SoC) as rapidly as energy and power availability limits allow. It will then attempt to maintain the SoC at that target level until the target time. All versions of the CCD algorithm (not just “greedy” CCD) attempt to discharge at a constant rate as determined by the following equation any time the next SoC target is lower than the current SoC:

$$\text{Discharge Rate (kW)} = (\text{Target SoC} - \text{Initial SoC}) * \text{battery capacity (kWh)} / \text{Time to Target (h)}$$

Both charging and discharging SoC strategies can be observed in Figure 3-2, which provides the reported (“measured”) SoC profile (in light red) from a sample greedy CCD test day. Blue stars represent the target SoC values and times that the algorithm is aiming for. The most notable feature of the profile is that SoC jumps from about 10% to 75% in about 90 minutes and then remains there for about six more hours until the target time. Explanations of the teal-colored callouts in Figure 3-2 can be found at the end of the CCD analysis in a subsection called Project-Specific Performance Issues.




Figure 3-2
Profile of reported SoC estimates in response to SoC target values using “greedy” CCD.

Predictive CCD

8minute’s vendor-specific implementation of the control mode called CCD-Predictive (CCD-P) can leverage a PV generation forecast to optimize charging/discharging based on the forecasted power available during each portion of a solar shift charge window. For instance, CCD-P aims to hit SOC targets at their respective time targets instead of as soon as possible, leaving more power and SOC

headroom available throughout the period for additional use cases (frequency response, power smoothing, AGC, etc.) Hitting charging targets as late as possible¹¹ within an allotted charging window lowers average SOC and reduces battery degradation.¹² Figure 3-3 provides the reported SoC profile (light red) and SoC targets (blue stars) from an example CCD-P test day. Explanations of the teal-colored callouts in Figure 3-3 can be found at the end of the CCD analysis in a subsection called Project-Specific Hardware Performance Issues.

¹¹ Clarification: Reaching the target as late as possible is not the same as waiting to charge for as long as possible. CCD-P may start charging right away, but only completes charging shortly before the target time.

¹² The CCD-P charging strategy may also result in more PV energy being captured that would otherwise be clipped, but this was not investigated in this project.




Figure 3-3
Profile of reported SoC estimates in response to SoC target values using predictive CCD (CCD-P).

Further, implementation of the proprietary CCD-P algorithm is designed to produce a firmer plant output profile at the POI as charging will occur roughly in proportion with the forecasted amount of available PV, so at a higher charge rate during times with more PV and at a lower charge rate during times with less PV. This benefit will be most visible on days where a significant portion of the day (e.g., a few hours in the morning) is overcast while the rest of the day is mostly clear.

“Greedy” CCD vs. Predictive CCD

On the surface, “greedy” CCD and Predictive CCD (CCD-P) both appear to accomplish the same goal—reaching certain SoC targets by certain times, but under normal circumstances¹³, CCD-P is the

preferred choice for a couple reasons. The main purpose behind the development of CCD-P was that it would provide various qualitative benefits over greedy CCD without sacrificing SOC accuracy. These qualitative benefits can be captured with the use of certain metrics. What follows is an investigation of some of those metrics and the qualitative effects they represent.

Battery Degradation

Figure 3-4 shows how, relative to greedy CCD, CCD-P spends less time at a high state of charge for an identical SoC target profile, which in general means reduced battery degradation. For this specific test scenario, the average SoC was about 11% higher for the greedy CCD than it was for CCD-P.

¹³ Greedy CCD may be preferable in unplanned or emergency situations where assurance of sufficient energy as soon as possible is paramount.




Figure 3-4
 “Greedy” CCD and CCD-P algorithms performing identical SoC target profile with primary difference in behavior called out

Grid Impact

Additionally, CCD-P typically charges at a more even rate than greedy CCD, as seen by comparing the shape of the green curves in the top left and right windows of Figure 3-5 (or the simulated red curves in the bottom left and right windows of Figure 3-5). The more uniform charge rate means smaller magnitude changes in net power output at the POI (blue line) when the battery power (green line) starts and stops charging. The PV+BESS (=POI) output produced by CCD-P should also be smoother than greedy CCD since it attempts to charge at a faster rate when more PV is forecasted and at a slower rate when less PV is forecasted. These two benefits can be measured using the following two metrics:

1. **Maximum Ramp Rate** - This is the maximum of the instantaneous differences between the PV+BESS power at time t and $t+60$ seconds for the entire test data set
2. **Average Ramp Rate** - This is simply the mean of the instantaneous differences between the PV+BESS (i.e., net) power at time t and $t+60$ seconds for the entire test data set

The maximum ramp rate, which typically occurs when the battery rapidly stops charging, was 38% lower for CCD-P than greedy CCD, and the average ramp rate was about 26% lower.




Figure 3-5
PV, BESS/ESS, and net (power block 5 point of interconnection) power profiles resulting from consecutive days of greedy CCD (left) or CCD-P (right) testing. Top: field results, Bottom: simulated results

As a more qualitative metric, CCD-P can better reduce peak solar generation (when desired) compared to greedy CCD by charging throughout the peak solar period. CCD-P will attempt to “smooth shift” by charging the BESS over the full charge period for later use in proportion to the amount of available energy expected.¹⁴ For example, if it were predicted to be cloudy in the morning, CCD-P could reduce charge rates in the morning and increase them in the afternoon so that the net power would be flatter, while still achieving its target SoC. In contrast, greedy CCD would exacerbate the issue by further reducing the net plant output immediately upon charging and then

likely not charging at all in the afternoon when PV was abundant.

Future Use Case

The idea behind CCD-P’s “firming” capabilities (where PV forecasts are used to inform charging levels based on the amount of forecasted PV energy) could potentially be extended to a new use case involving load. In such a scenario, generation could adjust in response to load profile forecasts from the utility to better align with expected demand. Combining the two, future algorithms could hypothetically include forecasted load and PV profiles to better align charging with both forecasted PV availability and expected demands.¹⁵

¹⁴ While CCD-P accomplishes some degree of “smooth shift”, stacking the active power smoothing or ramp rate product control mode together with CCD-P can help mitigate the rapid fluctuations that the long-term PV forecasts would not be able to foretell.

¹⁵ Utilities may choose to take other approaches to this problem such as using markets to incentivize generation at the right time or using a centralized controller to schedule generation from different plants, among other potential solutions.

Availability for Other Stacked Modes

The more even, optimized charge rate of the CCD-P profile (evidenced by the reduced peak magnitudes of the green line in Figure 3-5) also allows for increased flexibility to simultaneously perform other functions. In essence, CCD-P reserves space in the interconnection power limit (headroom) and in the BESS energy capacity so other functions can be performed along with CCD.

In the case of “greedy” CCD, the battery is likely to charge up at full rated power (assuming sufficient PV is available) leaving no room for additive active power commands from functions like frequency-Watt or active power smoothing during this time.

The increased ability to accommodate other stacked functions can be captured using the following three metrics:

1. Average Available Discharge Capacity (headroom) – This is the mean difference between the maximum allowable power capacity of the PV+BESS system and the current PV+BESS power
2. BESS Energy Availability (Avg % SoC below target) – Captures the mean SoC capacity in the charging direction (usable energy between current SoC and target SoC) that is available during the day for other control modes to perform their desired operation
3. BESS “Stackability” (% of time) – Captures the percentage of time during the primary charging period that the BESS has energy and power capacity available for other control modes to perform their desired operation (without exceeding the upcoming SoC target)

During the primary charging period from 6 am to 3 pm, CCD-P made available (to potential stacked functions) an average of 120 kW (8% rated) more discharge capacity and 31% more energy take capacity (i.e., energy flexibility in the charging direction that could be used by other functions) relative to its greedy CCD counterpart. During the

same period, the BESS was also available for charging 92% of the time¹⁶ under CCD-P vs. 30% for greedy CCD.

SoC Target Accuracy

While the multiple auxiliary benefits of CCD-P can be useful, it is imperative that the function still achieve its foundational goal of accurately reaching predefined SoC target at the desired time(s). As a summary, both “greedy” CCD and CCD-P reach targets with >98% accuracy, but what follows is a detailed breakdown as to where the error comes from.

“Greedy” CCD:

SoC error at each target time was typically¹⁷ less than 1% for greedy CCD (Figure 3-6). Mostly, this is due to the conservative approach of the control mode. The 15:00 target (see Figure 3-2), for example, receives an additional accuracy benefit from the CCD algorithm because CCD periodically recalculates the SoC before reaching the target to account for changes due to other functions. Thus, since the algorithm had sufficient time after reaching that target, it attempted to adjust the SoC back to the target value after deviations from the SoC recalculation.

¹⁶ Periods where charging would cause the battery to exceed its upcoming target were excluded from charging availability totals.

¹⁷ 7 out of 9 instances/day for this particular SoC target profile



Figure 3-6
Dot plot indicating distribution of SoC errors for the collection of SoC targets from Figure 3-2.

Missed targets were primarily caused by an SoC recalculation issue and the inability to charge from the grid at night, which is not allowed as per the interconnection agreement. Take for example the case of the target at 23:00 in Figure 3-2; greedy CCD did not correct the error. The SoC recalculation issue resulted in SoC estimates dropping from the 22:00 target value that was supposed to be maintained, and the system was unable to correct the SoC because doing so would have required charging the BESS from the grid (sunset was at about 18:00 that day), which was not permitted due to restrictions on grid charging.

See the “Project-Specific Performance Issues” section at the end of the Coordinated Charge/Discharge section for more detail on why certain targets were regularly missed.

CCD-P:

In general, CCD-P performs similarly to “greedy” CCD regarding SoC target accuracy. For example, the 23:00 target was still missed (as is pointed out in Figure 3-3) by upwards of about 4.2% due to the SoC recalculation issue in conjunction with not being able to charge at night. Additionally, putting the system into maximum power point tracking (MPPT) mode overnight resulted in a slight SoC increase by morning. (Again, see the “Project-Specific Hardware Issues” for a detailed breakdown of these issues)

One target time (of the specific target profiles selected) saw a noticeable difference between CCD-P and greedy CCD. CCD-P incurred additional error at the 16:00 target relative to what greedy CCD would have likely seen. This is because 8minute put the system into MPPT mode at 15:00 to maintain the SoC, not allowing the algorithm to correctively discharge in response to the SoC recalculation. In greedy CCD, error from this recalculation occurred early in the day, leaving it time to correct. Since CCD-P typically reaches targets just before the target time, though, the SoC recalculation occurred after the target time while the system was in MPPT mode and could not correct for it, negatively impacting the accuracy of the next SoC target.

Despite the various sources of error, the majority of target values were reached within 2% for the example day shown in Figure 3-3. The complete distribution of errors for this day is presented in Figure 3-7.




Figure 3-7
Dot plot indicating distribution of SoC errors for the collection of SoC targets from Figure 3-3

Overall SoC Accuracy:

Regardless of whether greedy CCD or CCD-P is used, average SoC error was found to be below the required 2% accuracy target for the various charge profiles tested in this project. The mean average error for all SoC targets was 1.38% (across the two test days already shown in Figure 3-2 and Figure 3-3 as well as five more CCD/CCD-P test days with varying profiles). Figure 3-8 shows the distribution of error values broken down by test day.



Figure 3-8
Distribution of reported SoC error relative to the

desired target value at each target for a variety of test days

As pointed out in the individual test days, most of the error above about 2% comes from recalculation of the battery SoC combined with not being able to charge from the grid or being put into MPPT mode (where BESS does not charge/discharge). Since these periods mostly occur at night, SoC error for nighttime targets was generally higher than error for daytime targets (Figure 3-9), averaging 1.60% for nighttime and 0.91% for daytime.




Figure 3-9
Distribution of reported SoC error relative to the desired target value, categorized by whether target occurred at night (left) or during daylight hours (right)

Qualitative Metrics Summary

Table 3-1 summarizes the previously discussed metrics as applied to the greedy CCD and CCD-P test days performed for this project. In every metric, CCD-P shows a quantitative advantage over greedy

CCD apart from SoC accuracy, where, at present, there is an almost-negligible decrease in accuracy.¹⁸

¹⁸ The slight shortcoming in SoC accuracy will likely disappear once the SoC recalculation issue is resolved.

Table 3-1

Various performance indices as applied to “greedy” CCD and CCD-P test days.

Metric	Greedy CCD	CCD-P
Average Ramp Rate [%/minute]	0.58%	0.43%
Maximum Ramp-Rate [% nameplate]	70.1%	32.0%
Average SOC	48.45%	37.39%
Avg Available Discharge Capacity (6am-3pm only) [(% nameplate)]	0.92 MW (61.5%)	1.04 MW (69.2%)
BESS “Stackability” (6am-3pm only) [% of time]	29.8%	92.0%
BESS Energy Availability [Avg % SoC below trgt.]	8.2%	39.6%
SoC Accuracy	98.97%	98.44%

Project-Specific Performance Issues

Two major issues hindered accuracy of the MESA-ESS operating modes during the testing period. These behaviors are specific to the Springbok 3 plant and include the DC-coupled hardware configuration installed at power block 5 and the inability to grid charge due to power purchase agreement (PPA) restrictions. While these characteristics impacted the tests being performed, neither are necessarily required or necessarily indicative of the industry standard.

SoC Recalculation Issue

While not specific to CCD algorithms, the SoC recalculation issue was uncovered when analyzing CCD results and is thus being addressed here. The issue is visible in Figure 3-2, where after significant amounts of charging or discharging, the SoC profiles appears to rise after charge targets and fall further after hitting discharge targets. This is because the battery management system (BMS) is miscalculating the SoC while charging or discharging. When the charging/discharging stops, the chemistry settles and the BMS recalculates the SoC estimate, which causes the SoC reading to move without the battery actually charging or discharging. At present, this issue causes

unintentional changes of up to about 5% to appear in the SoC estimation after significant (intentional) changes to the SoC.

In addition, the miscalculated SoC during charging/discharging distorts the intended BESS & POI power profiles, somewhat degrading the auxiliary benefits of CCD-P. Figure 3-10 illustrates the difference in BESS and POI charging profiles between a simulated (top) and measured (bottom) CCD-P response (to the same charging profile). The BESS profile in the simulated case (top, red) a) is much more uniform throughout the day, b) charges slightly faster in the morning when more PV was forecasted, and c) completes charging before the 16:00 target. The BESS profile in the measured case (bottom, green) a) begins later (separate issue), b) delays charging toward the end of the period, and c) continues charging right up until the target time. This “explosive” behavior of the ESS profile around 16:00 is a result of the algorithm trying to catch up to inaccurate SoC values that fail to keep pace with its expectation of what SoC should be given the energy input being received.




Figure 3-10
Power profiles (simulated on top, measured on bottom) in response to a simple CCD-P profile with daytime charging and nighttime discharging

Overnight Charging Issue

Since the BESS was only allowed to charge from PV in this project (per terms of the PPA for this plant), its SoC can only decrease throughout the night. Unless in a special SoC preservation mode during this time, this particular dc-coupled system will continuously dissipate around 10 kW of power just to idle. This power draw from the battery primarily supplies the dc-dc converters (to keep them “awake”), but also supplies controls and other auxiliary loads. While a consequence of the design for this particular dc-coupled system, any dc-coupled system will need to address the issue of where power can be drawn from.

To avoid the significant discharges, especially overnight, 8minute would put the batteries in MPPT mode¹⁹ whenever it wanted to hold a given SoC (i.e., when the next SoC target was equal to the prior one). The mode was not intended to be used this way, though; it was simply a workaround to the SoC recalculation issue. Unexpectedly a small amount of energy seems to be charged into the BESS at night while in this mode because of unknown causes in the PE/NEC hardware. Especially over long periods such as between 01:00 and 07:00 in Figure 3-2, this charging resulted in a small SoC increase when none should have been present.

Potential Daytime Discharging Limitation

Note that due to the dc-coupled topology in combination with the inverter sizing at this site, there is the potential for a bottleneck if the BESS wants to discharge during high PV generation periods. If power output exceeds the inverter limit, PV power may be curtailed.

¹⁹ MPPT mode is usually used during the day to maximize power output from PV by regulating the dc voltage, but it has

the side effect of minimizing power flow into or out of the BESS.

PV Active Power Smoothing

The active power smoothing (APS) function was designed to reduce rapid changes in PV by strategically injecting and absorbing energy from the battery energy storage system (BESS). The desired result is for the net power profile to be smoother than the raw PV power output and net power ramp rates to be reduced to within target limits.

The APS function is composed of two parts, one that focuses on picking a basepoint value such that the net power output follows a filtered version of the previous power outputs, and a last-minute additive component that works to mitigate any rapid changes beyond an allowed ramp rate. Before demonstrating the full APS function consisting of both parts, a field test was conducted to demonstrate the performance of the second, ramp rate fail-safe component by itself.

Ramp Rate Only Test

In this test, the gradient parameter of the APS function was set to 0.00 or 0% as seen in Table 3-2. This meant that the function would attempt to compensate for 0% of the difference between the moving average and the planned (basepoint) output, effectively disabling the filtered target component. Only the ramp rate limiting component of the algorithm attempted to regulate output using the real-time, “additive” power command from the battery. Its objective was to always keep net output power between the previous setpoint plus or minus the max allowable ramp rate. Such a reactionary response results in the output frequently being pinned up against the ramp rate limits during rapid PV power swings as observed in Figure 3-11. Here, the power at the POI (red) is often at the maximum or minimum of the allowed ramp rate range (light red band) due to the significant changes in PV power output (green) during that time.

Table 3-2

Table of parameters used in subsequent tests of APS and the ramp-rate-only portion of APS.

Parameter	Argument		Unit
	Ramp Rate Only Test	Full APS Test	
Point of Reference	PB#5 POI [Meter ID #0]		Meter at which measured active power is to be referenced
Lower and Upper Deadband Limit	0	0	Watts away from moving average inside of which no smoothing will be applied
Gradient	0.00 (0%)	0.80 (80%)²⁰	Ratio of change in additionally applied smoothing wattage to delta wattage (difference between moving average and reference power)
Filter Time	N/A (any is fine)	900 (=15 min)	seconds
Discharge Ramp Up & Ramp Down Rate	0.08333% ²¹	0.08333%	Percentage of maximum generation ramp rate per second
Charge Ramp Up & Ramp Down Rate	0.08333%	0.08333%	Percentage of maximum charging ramp rate per second

²⁰ Setting the gradient to 80% per recommendations in the Control Mode Specification document not to exceed this value when APS is second priority or lower.

²¹ Equals (5%/min) / (60 s/min)




Figure 3-11
Sample period from ramp-rate-only test day

Nevertheless, the average ramp rate on this day was reduced from 1.74%/min for the PV to 0.63%/min for the PV+BESS (=POI) using just the ramp rate limit capability.

“Full” Active Power Smoothing (APS) Function Test (Filtering Enabled)

Returning the gradient parameter to a recommended value such as 0.80 or 80% yielded the results found in Figure 3-12. Here, the algorithm is keeping a running (but non-linear) average of the past 15 minutes of PV output,²² and using that value as the target toward which it should trend the POI value (at a rate equal to the gradient parameter). In other words, it will target the value 80% of the way between the current output and the filtered value at each time step (for a gradient setting of 0.80). We observe that the output is much less “jagged” in most cases, and that the actual output at the POI (dark red) is smoother, remaining distanced from the absolute ramp rate bounds (light red region edges) for the vast majority of the time.




Figure 3-12
Sample period from full APS function test day

Numerically speaking, the full APS algorithm performed better than the ramp rate only portion when it came to average ramp rate. Full APS reduced the average ramp rate from 1.71%/min for PV to 0.43%/min at the POI (compared to the 0.63%/min for ramp rate only). It should be noted, however, that the two test days, while both containing high-variability periods, were not identical, and exact performance comparisons would require longer-term testing of both variants.

Comparison of APS with Ramp Rate Only vs. Full APS

For a better understanding of the performance benefits of APS, it helps to investigate the distribution of observed ramp rates provided in Figure 3-13 (zoomed). While the ramp-rate-only test day (windows on left) started out with a different distribution of ramp rates than the full APS test day (windows on right), the characteristic behavior of the functions is still evident when comparing the PV ramps observed (top windows) relative to the observed ramps at the POI (bottom windows).

The ramp-rate limiting portion of the function brought most of the PV ramps exceeding the desired boundary (5%/min in these cases) within the desired boundary. However, because the ramp limiting part of the mode only responds once the

²² 15 minutes was an arbitrary filtering time setting chosen for these tests.

ramp threshold has been exceeded, system response times for this plant caused the limit to be exceeded 3.2% of the time. In contrast, the full APS algorithm effectively targets a slow-changing (i.e., near-zero) ramp rate all the time, not just when excessive ramp rates are observed. As a result, the distribution of

ramps is tightened significantly to the point that virtually all observed ramps at the POI (99.93% of them) were within the requested 5%/min boundary.




Figure 3-13

Distribution of observed ramp rates for PV and at POI (=PV+BESS) for the ramp-rate-only test day (left) and full APS function test day (right)

The cost of this increased smoothing and better regulated ramps is increased battery cycling. Even after normalizing by the mean observed PV ramp rate (since the results in Figure 3-13 came from two different test days), the full APS function cycled the battery more than applying only the ramp rate portion of the APS function. For the specific filtering time and gradient configuration selected, the ramp-

rate-only function cycled a total of about 0.63 MWh in and back out of the BESS over the course of that day while the full APS function cycled an adjusted 0.94 MWh in and back out on its test day. Figure 3-14 shows the distributions of power flows into and out of the BESS for these two particular test days, keeping in mind that power and energy throughput will be a function of the intermittency observed.




Figure 3-14
Histogram of BESS utilization while performing APS both with and without gradient filtering

Frequency Response (Frequency-Watt)

To fulfill the frequency response functionality requirement, a frequency-Watt algorithm was developed. The so-called Frequency Watt mode is an additive power control mode used to alter the system's active power output to the grid in response to the measured deviation from a specified nominal frequency.²³

The fixed frequency-Watt curve defining the response used for testing is shown as a black piecewise line in Figure 3-15. It had a deadband of 0.01 Hz to either side of the nominal 60 Hz and saturates another 0.3 Hz away from the deadband.

Currently, it is asymmetric. The saturation value in the reduction/charging direction is the rating of the BESS (1.5 MW) whereas the saturation value in the increase/discharging direction is the rating of the inverter (3.225 MW).

In addition to the desired curve, the plot contains scattered points corresponding to two series. Each series indicates a form of response by the PV + BESS system to the measured frequency—red for measured additive powers, teal for additive power setpoints.²⁴ Performance of the system can be judged by the proximity of response points to the desired curve, where closer is better. (For a zoomed-in version of Figure 3-15, see Figure 3-17.)




Figure 3-15

Scatter plot indicating the frequency-Watt setpoint and measured response to the desired frequency-Watt curve for the daytime and nighttime test

²³ This active power mode implementation of Frequency Watt uses the FERC guidance as published in [“NERC – Reliability Guideline, BPS-Connected Inverter- Based Resource Performance, September 2018”](#)

²⁴ Frequency was measured at the point of interconnection (POI) of the power block under test (PB#5).

Types of System Responses Analyzed

The system response to deviations in frequency was designed to be additive, or in addition to baseline behaviors such as planned charging powers. As a result, there is no direct measurement that represents the additive power only. Instead, we estimate it to be the measured power at the power block POI minus both the power block's measured PV output as well as any basepoint power setpoint.²⁵ In equation form, that is the following:

$$Additive_{est} = POI_{meas} - PV_{meas} - BP_{planned}$$

The other series represents the additive power setpoint. Since only the frequency-Watt algorithm was running for this test, it is assumed that any additive power setpoint represents the controller's intended output. This is useful because it reveals the setpoint sent by the controller, even if hardware response to these commands at this site often takes several (5-20) seconds.

Figure 3-16 shows a small sample of both the intended, additive responses (red) as well as the measured responses (navy blue) relative to the ideal additive target value (lime green) for reference. The target value series was calculated by feeding the measured frequency (assumed to be 100% accurate) at each point in time into the desired, piecewise frequency-Watt curve to calculate the ideal additive power response for that time.




Figure 3-16

Brief, time series sample of power responses (setpoint in red and measured in blue) relative to the calculated ideal response value in green.

The additive setpoint response values (red) almost always track the ideal series (green) very closely with the small offset most likely attributable to rounding error in the reported frequency measurement relative to those used in the control calculations/setpoints. “Measured” values (blue) also track closely to the ideal response but change more slowly due to control filtering delays, resulting in a smoother curve.

Other Notes on the Response Type Accuracy

The setpoint responses are usually going to be more accurate than the measurement estimate, because the setpoints are nearly instantaneous calculated responses to the measured frequency while the measured values are delayed and realized values. The increased accuracy is easier to see in Figure 3-17, which is a zoomed in version of Figure 3-15. “Measured” values include offsets in time due to control signal delay, inverter response time, ramp limits, etc. and also suffer increased error from measurement equipment, among other sources. Both the “measured” response and intended additive response suffer from low resolution frequency data relative to the aggressive curve implemented. Frequencies were only measured to

²⁵ To use this estimate, it is necessary to assume that no error is introduced between the stated basepoint setpoint and the portion of power assumed to be basepoint power being

measured at the POI. In reality, the assumption is not fully true, but no better alternative was found.

0.01 Hz which is why the responses appear to occur at discrete levels in Figure 3-17.²⁶




Figure 3-17
Zoomed in version of Figure 3-15 with cluster of points deviating from desired curve circled in yellow

Day vs. Night

With the zoomed in Figure 3-17, it is clearer to see that the scattered points are generally attempting to follow the frequency-Watt curve, except for responses to high frequency values in the nighttime plot (circled in yellow). The periods of deviation can be seen in the time domain plot of the nighttime test in Figure 3-18. Regardless of how high the

frequency went, output during those times remained approximately zero. This is because the appropriate response during these times would be to increase power consumption from the grid, which was not allowed due to the prohibition on grid charging for this site.²⁷ This is similar to the issue seen for CCD control modes.

²⁶ Some frequencies appear to report with 0.005 Hz resolution, but this is just a consequence of two half-second values being averaged together to form 1-second resolution data.

²⁷ One way to partially get around this limitation would be to set a base discharge rate from which the battery could back off during the night, but the power capability would be limited to roughly 100 kW due to the need to spread available energy capacity of this particular battery throughout the full night.




Figure 3-18

Extended time series plot indicating measured frequency (top plot), various response types (middle plot), and the available battery SoC (bottom plot) for the nighttime test performed on 1/14-15/2021

More generally, system operators should take note that a BESS cannot charge for high frequency response when at 100% State of Charge (SoC) and cannot discharge for low frequency response when at 0% SoC. Power limitations could also limit frequency regulation activity when the BESS is already charging or discharging at or near 100% power for other modes. System operators may want to set SoC targets between 5% and 95% SoC to leave space for the system to always be available for frequency response and use modes like CCD-P to allow for increased power headroom for other functions. Battery capacity, frequency-Watt curve settings, and desired availability for other functions may dictate that more or less SoC buffer is necessary than the 5% and 95% boundaries.

If a utility would like to operate frequency-Watt 24 hours per day, the following options could be considered:

- Preferred option: Charging could be allowed at night, potentially just for additive power functions (with some hardware/software development and reworked interconnection agreements)
- Imperfect workaround: SoC could be charged to a sufficiently high level each afternoon to provide for a fixed power discharge throughout the night (plus losses) large enough to cover a response to probable high-frequency deviations.

If the latter option is selected, energy storage capacity, frequency-Watt settings, other active power functions, and risk tolerance for a saturated

response to extreme frequency deviations will all need to be weighed in determining what SoC value would be sufficient. For example, approximately 22% of battery SoC was discharged during a 4-hour period of the daytime test (see teal SoC curve in bottom window of Figure 3-19) to respond to the extended low-frequency event (see top window in

Figure 3-19).²⁸ This 22% SoC consumption would be reduced to 5.25% on a four-hour BESS,²⁹ but even so, discharging every night just to be able to provide one-way FW service is likely an impractical solution for utilities long-term, especially when this energy may be desired for other purposes.




Figure 3-19
Extended time series plot indicating measured frequency (top plot), various response types (middle plot), and the available battery SoC (bottom plot) for the daytime test performed on 1/15/2021

Accuracy Metrics

Error between the observed responses and the ideal target values was minimal for daytime values, indicative of proper frequency-Watt behavior.

Figure 3-20 depicts this tight distribution around zero for the various response types ranging from the instantaneous target and measurement responses to two time-delayed measurements. Measured responses have broader distributions for the

²⁸ Reminder: The frequency-Watt curve used was more aggressive than would typically be used in regular commercial use.

²⁹ Such as the system planned for LADWP’s Eland project

reasons mentioned above but were still within 1% of the target between 71% and 85% of the time.³⁰




Figure 3-20

Histograms of the percent error relative to the ideal additive power response (normalized by rated power in that direction) for daytime values

Nighttime error distribution is shown in Figure 3-21. Firstly, note that nighttime high-frequency events have been excluded from consideration as the system was precluded from responding to them due to the grid charging prohibition and lack of discharge offset. Error for the remaining values has

a positive bias as a result of the data filtering, but also because of the control algorithm used. The control only overshoots in the positive direction at night, as exemplified by the measured response (blue) in Figure 3-22 relative to the target value (red/green).

³⁰ Noticeable response deviations at discrete powers are typically times when the target value was absent from reporting, but these instances did not impact the actual

responses (i.e., the point likely existed in the controller but just not its reported output)




Figure 3-21

Histograms of the percent error relative to the ideal additive power response (normalized by rated power in that direction) for nighttime values




Figure 3-22

Time series sample of nighttime frequency-Watt responses exemplifying overshoot only in positive direction

Presented differently, the error histograms in Figure 3-20 and Figure 3-21 can be transformed (and combined) into the cumulative distribution function seen in Figure 3-23. Between the two, error is less than 4% nearly 100% of the time. Also notable is that error is reduced (curves are further left) when considering measurements later in time (i.e., allowing for settling time). Specifically, error decreases from the instantaneous error series to the 2-second-delayed series, but then picks up again for the 4-second series. From this, we can conclude that the settling time is around 2-3 seconds for the observed frequency dynamics that were simultaneously at play.




Figure 3-23
Cumulative distribution function of frequency-Watt response error (relative to ideal additive power response)

Table 3-3 lists the percentage of responses, broken down by time of day, for which the absolute error

Table 3-3

Performance metrics indicating percentage of observed responses within 5% and 1% tolerance bound of ideal additive power response value broken down by portion of day and response type

Portion of Day	Response Type	Percentage of Values Within X% Tolerance	
		5%	1%
Full Day ³¹	Additive Setpoint - 0 s	99.5%	90.7%
	Measured - 0 s	98.9%	68.3%
	Measured - 2 s	99.7%	77.5%
	Measured - 4 s	99.9%	79.3%
Daytime	Additive Setpoint - 0 s	99.4%	91.8%
	Measured - 0 s	98.5%	67.4%
	Measured - 2 s	99.5%	74.9%
	Measured - 4 s	99.9%	79.2%
Nighttime ³¹	Additive Setpoint - 0 s	99.6%	87.5%
	Measured - 0 s	100.0%	71.0%
	Measured - 2 s	100.0%	85.4%
	Measured - 4 s	100.0%	79.7%

Overall, the frequency-Watt mode conformed to its prescribed curve with high accuracy when

exceeded 5% and 1% thresholds. From these data, it is evident that the frequency-Watt algorithm can maintain 5% error or less 99% of the time or more when it is fully capable of charging and discharging (e.g., daytime results). When high-frequency nighttime periods are removed from consideration, performance is roughly equal whether operating during the day or at night.

Achieving the more stringent 1% tolerance yielded much lower performance metrics across the response types. This could be because of rapidly changing frequency measurements (on the order of the response time or faster) in combination with the coarse frequency measurements (0.01 Hz resolution) and aggressive frequency-Watt curve tested. Allowing a few seconds for the value to settle did improve the fraction of values achieving the target by up to 11% for the full-day case.

power/energy capacity was available. These tests have shown that with proper planning and setting

³¹ Excludes nighttime high frequency periods since system not set up to achieve these

selection, the PV+BESS system can successfully contribute to frequency regulation.

Automatic Generation Control

The automatic generation control (AGC) mode is an additive control mode that takes in a commanded setpoint from a DNP3 register (or alternatively from a pre-loaded file) and adds or subtracts the requested amount of power to its planned output. The objective of the function is to respond as closely as possible to the commanded target value.

In this single-day test of AGC, the DNP3 point was populated with a new point by LADWP's Energy Control Center (ECC) about every four seconds and acted upon by the Springbok 3 power block (PB#5) in real time. The power block was committed to responding to any power value requested between 200 kW of additional absorption/reduction and 200 kW of additional injection/increase.

Performance

The AGC mode performed very similarly to the frequency-Watt mode both in its accuracy and in its shortcomings. Plotting the error between the commanded setpoint and various response types in Figure 3-24, the following observations can be made:

- The additive power internal control signal matched the incoming request virtually 100% of the time
- During normal operation, such as during the day, measurement error was tightly clustered around 0% (i.e., the target value) with a mean absolute error (MAE) of about 4.1%
- The distribution of measurement responses two seconds after receipt of the target command was even tighter with a MAE of 3.4%




Figure 3-24

Histograms of the percent error relative to commanded ECC AGC signals (normalized by max committed capacity of +/-200 kW and excluding nighttime periods with <100 kW requested)

Atypical operation occurred before sunrise and after sunset whenever charging was required and is attributable to the contractual nighttime charging

prohibition spelled out in the PPA for this plant. Much like FW, AGC relies upon the battery to absorb energy at night when reducing PV output is

not an option (already zero). Charging from the grid was not permitted at this site, preventing responses to negative commands at night.

The battery SoC is also noteworthy in Figure 3-25. Because of the grid-charging limitation and a nighttime power bias to power controls, the system was barely able to perform the discharge commands through midnight before the battery SoC would have been too low to continue operating. Figure

3-26 illustrates the roughly 30 kW dispatch pulled from the BESS SoC (on top of the requested nighttime values seen by the utility at the POI) necessary to keep controls powered and the system available. Allowing grid charging, at minimum for grid services like responding to AGC signals, could mitigate this issue. Alternatively, more constrained AGC capacity limits in the discharge direction could exchange decreased response capacity for operation extending further into the night.




Figure 3-25
Extended time series plot indicating measured additive power response relative to the external AGC command (top plot), plus the available battery SoC (bottom plot) for the AGC test performed 10/27/2020




Figure 3-26
Sample of ESS response to the AGC command relative to measured and target responses at the meter in front of the system (POI); demonstrates power offset absorbed by system

Removing the nighttime charge periods from consideration (but keeping all nighttime discharge periods) the measurement response error is plotted in a cumulative distribution function in Figure 3-27. Results are broken down by response type:

- As seen in Figure 3-24, the additive target internal command (red) has 0% error for virtually all samples (98.8%). The only samples with non-zero error come from the few minutes of transition time in the morning and evening when the inverters are transitioning from nighttime to daytime operation and back.
- Instantaneous measurements (in green) had the greatest error with only 79% of readings having 5% error or less³², mostly because of control delays.
- Allowing for a 2-second (blue) or 4-second (purple) settling time increased the percentage of points within that same 5% error threshold to 86% and 87% respectively. Figure 3-28 helps explain this improvement, as it depicts how the response measured 2, and especially 4 seconds

later, better aligns to the change in commanded value than the instantaneous³³.

The exact optimal time delay appears to vary between about 2 to 4 seconds depending upon the severity of the change and control delay at that time. No benefit was observed beyond this point, likely because new AGC commands were received after four seconds.




Figure 3-27
Cumulative distribution function of response error (relative to ideal additive power response)




Figure 3-28
Zoomed-in view of the setpoint response, instantaneous measurement, and time delayed measurements relative to the commanded AGC signal value

Settling time for the full 200 kW change is also on display in Figure 3-28 and takes between about 10

³² Figure is calculated with nighttime charge periods excluded.

³³ Shifting measurements occurring a few seconds later to the present hints at what the desired response of the system would be in the hypothetical case that response time was minimized.

and 20 seconds, depending upon the tolerance band used. The relatively long ramp, overshoot, and settling time are a limitation of how the dc-coupled topology was configured in this particular system; coordination between the converters (dc-dc and inverter) relied upon stable outputs rather than coordinated algorithms.³⁴ It may be possible to mitigate these delays by utilizing a setup where the battery connects directly to the inverter and PV connects to the DC bus via a dc-dc converter.³⁵

The general takeaway for the automatic generation control (AGC) mode is similar to that for frequency-Watt: During periods of proper operation with sufficient headroom allowed for battery charging and discharging, the algorithm demonstrated a reliable response to the control signals with response times on the order of the control signals it received or longer. However, the BESS is capacity limited, and PV is not always available, so 24-hour operation will take a combination of adequate battery sizing and sufficient battery charge at the beginning of each night to cover expected needs (for this and other functions) throughout that night.³⁶

³⁴ MW-scale PV+BESS systems could have faster response times, but this claim was not verified in this project.

³⁵ *Emerging DER Technologies and DERMS*. EPRI, Palo Alto, CA: 2019. 3002017556.

³⁶ For the BESS size in this system, if the system operator wants to have the BESS available to respond to AGC

commands with both charging and discharging, then SoC targets could be set to between 5% and 95% (or similar) in addition to permitting nighttime charging. This should allow the BESS to respond to frequency regulation signals through AGC or real time 5-minute settlement adjustments, such as those used in the CAISO Energy Imbalance Market (EIM).

Demonstration of Stacked Control Modes

One of the objectives of this project has been to demonstrate proper operation of the PV+BESS system in a so-called “hybrid” mode. The mode attempts to blend the multiple MESA-specified functions for optimal performance by simultaneously running each algorithm and coordinating the outputs. It does this by finding maximum allowed, minimum allowed, and preferred outputs of each component algorithm, then combining the resulting outputs based on pre-defined priority levels.

The agreed-upon test plan outlined tests intended to assess the performance of each function in various stacks. Execution of these tests demonstrated that the system is capable of performing multiple functions simultaneously with minimal impact to individual function objectives. This section details the holistic benefits of stacked modes and compares individual mode performance when stacked to when run by themselves.

The three autonomous functions outlined in this chapter, [Coordinated Charge/Discharge - Predictive (CCD-P), Active Power Smoothing (APS), and Frequency-Watt (FW)] were combined into the following stacked/hybrid modes:

- CCD-P + APS
- CCD-P + FW
- CCD-P + APS + FW

In each test, CCD-P was set to the highest priority over APS (where relevant). In the event that the

preferred outputs of these functions conflict or, in extreme cases, cannot both be satisfied, the system will strive to accomplish the objective of the function with the highest priority before that of one with a lower priority. FW, where relevant, was an additive control mode, meaning that its modification to the basepoint would always get added to the resultant output of the basepoint mode(s) (i.e., last).

Impact to CCD Accuracy

The simplified SoC profile from the top half of Figure 3-1 was used in the comparison between CCD-P in isolation and CCD-P with other functions. In all cases, CCD-P continued to hit SoC targets with minimal impact to accuracy, even while running simultaneously with other functions (see Table 3-4). However, as will be discussed in detail, CCD-P+FW sometimes suffered SoC accuracy issues for specific targets due to the test profile and hardware/software constraints of the Springbok 3 plant.

Table 3-4
Median SoC accuracy by test

CCD-P Only	CCD-P + APS	CCD-P + FW	CCD-P + APS + FW
99.00%	99.30%	98.10%	97.10%

Additionally, CCD-P continued to provide benefits to battery life and mitigate grid impacts relative to greedy CCD.³⁷ Evidence of these benefits is provided in Table 3-5.

³⁷ Average ramp rate increased for the tests with FW involved, but this was to be expected given the highly variable nature of frequency from one minute to the next, the aggressive test

curve used, and the additive nature of the function (appended to output after all other calculations, including ramp rate)

Table 3-5
Qualitative change in CCD performance metrics between various stacked mode test days (involving CCD-P) and a greedy CCD test day

	Test Compared Against Greedy CCD Algorithm			
Metric	CCD-P Only	CCD-P + APS	CCD-P + FW	CCD-P + APS + FW
Average Ramp Rate at POI	Reduced	Reduced	Increased due to nature of FW	Increased due to nature of FW
Maximum Ramp-Rate	Reduced	Reduced	Reduced	Reduced
Mean SOC	Reduced	Reduced	Reduced	Reduced
BESS Availability [% of time] (6am-3pm only)	Increased	Increased	Increased	Increased
BESS Energy Availability [Avg % SoC below target]	Increased	Increased	Increased	Increased

Impact of Project-Specific Hardware Performance Issues on SoC Accuracy Metric

In the case of SoC accuracy, the reported SoC at each target time was assessed relative to the target at that time (just as in the CCD section). If taken at face value, then the SoC accuracy at each of the five targets is represented by a dot in the Figure 3-29 dot plot with colors corresponding to the different stacked function scenarios.

Right away, however, many of the dots have been flagged as having been influenced, sometimes heavily, by limitations outside of the algorithm’s control. Primarily, the target SoC profile was designed for a summer day, and when the test was conducted, multiple of the SoC targets that had previously been during the day now occurred before sunrise or after sunset. Since the battery at Springbok 3 is prohibited from charging from the grid (as was discussed in the section on CCD), this limitation prevented corrective actions that the algorithm would normally have taken if permitted to do so, increasing SoC target error.



Figure 3-29
Dot plot of SoC target errors associated with various test days; dots with a white ‘X’ represent days impacted by a hardware limitation

The two highest error values were a result of frequency-Watt further exacerbating the Springbok 3 nighttime charging limitation by calling for increased output immediately before sunset. Figure 3-30 and Figure 3-31 provide examples from both test nights where CCD-P and FW were operating simultaneously (in Figure 3-31, APS was also active). In both figures, the orange dotted outline highlights where the frequency was low from just after the 15:00 SoC target is reached until about sundown at 16:40. These low-frequency periods induced an additive discharge which reduced the BESS SoC. Under normal circumstances, the CCD-

P algorithm would have become increasingly more aggressive as the target drew near to achieve the desired SoC. However, since the sun had set, this plant could no longer charge and thus missed its

18:00 target by over 15%. So long as operators do not schedule SoC targets that require charging after sunset, these errors should be significantly reduced, if not eliminated.




Figure 3-30
Time series plot illustrating the frequency conditions and corresponding plant response that caused larger-than usual error for the 18:00 SoC target on this CCD-P+FW test day




Figure 3-31

Time series plot illustrating the frequency conditions and corresponding plant response that caused larger-than usual error for the 18:00 SoC target on this CCD-P+APS+FW test day

Given the significant impact that the current nighttime charging limitation has on the SoC mean absolute error (MAE), (an accuracy metric used for CCD-P), a selection of alternative SoC accuracy metrics is presented in Figure 3-32. Using the median error (dark gray bar) better represents algorithm errors while excluding outliers primarily caused by the hardware issues mentioned. The two additional bars represent SoC error for the primary charge target at the end of the afternoon (orange bar) and primary discharge target late at night (blue bar). In both cases, the system should not be hardware or software constrained, so the trustworthiness of the metrics at those times is greater and a better representation of algorithm accuracy.




Figure 3-32
Alternate SoC accuracy metrics for various CCD-P tests

CCD-P Accuracy Impact Summary

CCD-P still provided many of the same benefits over greedy CCD when it performed alongside other functions as it did alone. However, those operating the system should note that CCD-P can be sensitive to the state of charge (SoC) profile selected, and they should adjust SoC target times seasonally so that all SOC targets that require charging to gain or maintain SOC have enough time to reach their targets during daylight hours.³⁸ Even so, CCD-P demonstrated its ability to reach desired SoC with greater than 97% median accuracy for all combinations of stacked function tests.

Impact to APS Accuracy

Figure 3-33 presents the three test days being compared to assess the impact to accuracy that stacking has on the active power smoothing (APS) function. The top window shows the test day in which only the APS function was running. The middle and bottom windows represent stacked function test days in which the APS function was active alongside CCD-P and CCD-P+FW respectively. The relative improvements of PV ramp rates to ramp rates observed at the point of

interconnection (POI) were then assessed and compared for each test day. The APS-only test day is treated as the baseline for all test cases.




Figure 3-33
Time series plots of three APS test days without (top) and with (middle and bottom) additional stacked functions active

Glancing at the raw numbers in Figure 3-34, it appears as if APS reduced the inherent PV ramp rate (teal bar) at the POI (purple bar) by more than 1.1% in the APS and CCD-P+APS case but *increased* it in the CCD-P+APS+FW case. This is because APS only aims to smooth variability coming from PV and planned (basepoint) energy storage behavior; frequency-Watt mode is an additive response,

³⁸ Mostly relevant to the Springbok 3 plant or others which are not allowed to charge from grid.

meaning its output gets added to the smoothed APS output afterward. The pre-FW effect of APS can be approximated by the measured PV plus the planned BESS basepoint (red/salmon-colored bar). Indeed, mean ramp for this point is reduced from the PV ramp in all three tests. Net output from the CCD-P+APS+FW test is only more variable because of the addition of the dynamic frequency response powers.




Figure 3-34
Bar plot depicting the mean ramp rate at various measurement points for three different tests involving the APS function

Distributions of observed ramp rate are shown in Figure 3-35, and they verify that APS successfully reduced high-magnitude incoming PV ramp rates from the perspective of the POI, even with the addition of simultaneous functions. Observe how PV ramp rates (green) in excess of the desired 5% threshold are largely absent from the ramp rates observed at the POI (orange). Only because of the additive frequency response powers did ramps noticeably exceed the desired 5% threshold in the CCD-P+APS+FW case (Figure 3-35) and make the purple POI power profile in the bottom of Figure 3-33 spikier than the incoming PV one (teal).




Figure 3-35

Distributions of observed PV, basepoint, and POI ramp rates for various tests involving APS (zoomed in to better see distributions; values not shown are predominantly ramp rates very close to 0)

Impact to FW Accuracy

Much like CCD-P and APS, frequency-Watt (FW) performance was assessed by running FW in parallel with other functions and comparing it to FW running by itself. Figure 3-36 demonstrates that, in all cases, the response followed the prescribed curve when possible. All test cases saw deviations in the high frequency region due to the

lack of ability to charge at night, while the two stacked functions days also saw deviations in the low frequency region as their depleted state of charge early in the morning prevented the necessary dispatch. A time series example of this missed morning charge can be seen just after 06:00 in Figure 3-30, where the additive target power (in red) intends to dispatch power, but the measured response (in blue) remains at zero due to the depleted SoC (bottom plot, teal).




Figure 3-36
Intended (i.e., setpoint) and measured responses to observed frequency values with ideal response curve (black) for reference and deviations from this curve highlighted

Mean absolute error from the ideal response is summarized for all three test cases in Figure 3-37.³⁹ Differences in response error between test types are nearly negligible. Measured FW responses were between 0%-0.26% higher in the stacked tests compared to the FW-only test, but all mean response errors were below 1.0 %. Measured responses after a 2-4 second response period showed the lowest error in all cases but did not significantly impact the relative accuracy between tests.




Figure 3-37
Setpoint and measured response error for various tests involving FW

Site-specific Recommendations

The results in Figure 3-37 only represent periods where the function was fully capable of performing. At Springbok 3, however, battery SoC and nighttime charging limitations will increase these error figures. To mitigate this, the FW mode can be stacked atop a strategic CCD profile since it has

³⁹ Only considers periods where battery capacity was within operational range and excludes nighttime periods with high frequency (since the appropriate response might require

battery charging from grid, which was not allowed at Springbok 3)

been shown that this does not significantly impact FW performance. The CCD profile would include daytime charging and a planned, after-sunset discharge (decrease in SoC target). This discharge would afford the FW function “headroom” to decrease power outputs in response to high frequencies during this period. An example of this happening can be seen in Figure 3-30 where the measured response for this CCD-P+FW test continued in both directions after sunset (about 18:00).

Summary of Results

MESA-DER/MESA-ESS functions were put to the test to demonstrate proper performance both individually and then in combination.

- Coordinated Charge/Discharge (aka SOC target schedules or CCD) typically achieved SoC targets within 2% under both “greedy charging” and CCD-P, which used PV generation forecasts to provide additional benefits such as smoother charging profiles. Accuracy was most impeded by grid charging limitations for this plant (predominantly at night) that required some workarounds.
- Active Power Smoothing (APS) tests demonstrated how the full APS algorithm (that targets a filtered moving average) improved ramps to within the desired ramp rate bounds, whereas only correcting for excess ramp rates mitigated their intensity but could occasionally exceed the limit due to the response time of this particular system.
- Frequency-Watt (FW) mode followed its respective FW curve settings with less than 5% error more than 99% of the time (when the BESS was permitted and had sufficient SoC to do so). Scheduling a nighttime discharge also allowed responses to high frequency to continue into the night.
- Automatic Generation Control (AGC) likewise responded to setpoints accurately with mean average error of about 4% between the instantaneous response and the intended setpoint (when the BESS was permitted and had sufficient SoC to do so).
- Response times that negatively impacted functions, especially FW and AGC, were due to

stability challenges at this plant that arose from the BESS being dc-coupled to the solar inverter. Further tuning of controls may be able to mitigate some of these impacts.

- Given the grid charging prohibition at the Springbok 3 site, SoC targets that require charging to maintain or increase SOC should be scheduled during daylight hours. Then, scheduled discharges at night can be used to extend the periods during which other MESA-DER functions can be effective.
- The component functions continued to perform as expected when combined into various mode stacks with no noticeable degradation in performance beyond what could be explained by mode prioritization and basepoint/additive behavior. Even then, impact to performance was minimal.

In summary, these tests demonstrated full functionality of the various PV+BESS plant controls with a hybrid of basepoint and additive modes, both individually and in concert with one another.

Chapter 4: Ramp Rate Limiting with Solar Forecasts but No Energy Storage

Objective & Usefulness

PV production's inherently stochastic nature can result in instances of rapid power increases or decreases at a PV plant's point of interconnection (POI). Controlling the magnitude of these sudden changes can smooth plant output and help make PV plant production less dynamic, especially on days with highly variable solar conditions.

As was demonstrated in Chapter 3 of this report, energy storage systems can be added to PV systems and respond quickly enough in both directions to smooth excessive ramps to desired levels. While the addition of energy storage is a fully viable solution at present day, the additional hardware and cycling of the battery could potentially lead to increased costs. An experimental operating mode called Ramp Rate Product (RRP) aims to control PV ramp rate in both up and down directions without using energy storage. Instead, it attempts to use strategic PV

curtailment informed by short-term solar forecasts to limit ramping to a specified rate.

Figure 4-1 illustrates what simulated behavior of the RRP operating mode would look like for highly variable conditions. Note that without some form of energy storage, curtailment of PV (sometimes proactively) is the only source of flexibility that the plant has, both for upward and downward changes in generation. As a result, this function will frequently sacrifice generation to provide the smoother output. Therefore, the orange 'plant output' line in Figure 4-1 is always at or below the 'max potential PV' line in blue. While outside the scope of this project, further research will be required to quantify the extent of this loss for different variability conditions as well as find the breakeven point (if there is one) at which the addition of an energy storage system becomes more economical.




Figure 4-1

Simulated desired plant output (orange curve) for hypothetical, highly variable PV generation scenario (blue) and amount of curtailment (in gray) required to obtain that desired output

The value of constraining ramp up/down events at an individual plant level will also require further research as it is not yet well documented, especially for transmission-connected systems. While EPRI has investigated the potential economic effects of PV variability on distribution equipment such as tap changing transformers,⁴⁰ very large PV plants (such as the Springbok 3 site), and transmission-connected systems in general, benefit significantly from the smoothing effects of geographic diversity. A study by Lawrence Berkeley National Laboratory quantifies these effects for a few case studies involving the size of a PV plant as well as different numbers of PV plants separated by varying distances.⁴¹ While more investigation into this research area could be used,⁴² the general takeaway is that the collective ramp rate of larger and multiple PV systems is lower than individual plants, potentially mitigating the need to regulate ramps at any one site. Regardless of the *practicality*, the scope of this research project encompassed investigating the *feasibility* of ramp rate control without the use of energy storage, and these results are presented below.

How the Function Works

At a high level, the RRP mode controls ramp rates by sending curtailment commands to the PV inverter that limit output to a specified power value. The mode uses different methods to limit ramps in the up and down directions. To limit ramp rate in the up direction, the controller sends curtailment commands anytime a PV up-ramp exceeds a specified ramp rate. Controlling ramp rate in the down direction without the use of a BESS is more complicated. The controller looks at short-term solar forecasts and pre-emptively begins curtailing ahead of forecasted down-ramps that exceed a target ramp rate.

In More Detail

Up ramps

The mode continuously reads PV output and, when solar output is increasing rapidly, sends curtailment commands that limit up-ramps to the operator-specified ramp rate limit. Successive curtailment commands are limited to the allowed ramp per control interval above both the measured PV output and the previous curtailment command. Anytime PV production exceeds the specified ramp rate in the up direction, excess generation is curtailed.

Down Ramps

To control down-ramps, the mode must look to the future and leverage high-accuracy short-term PV forecasts. At the Springbok pilot, the RRP mode uses 30-minute forecasts at 1-minute granularity generated by an infrared spectrum sky camera installed at the site. The mode calculates a gradient defined by the maximum allowed ramp rate and extends it forward from the current time and measured PV value to zero (MW). If any forecasted PV value falls below this gradient, it means that PV is forecasted to exceed the allowed down-ramp rate. When these events are identified, the mode immediately begins pre-emptively ramping down production at the prescribed ramp rate.

The control mode uses the latest forecast from the sky camera (generated once every minute) to refresh its outlook for the production dip. If the forecast changes, it will adjust permissible production levels in response to the new forecasts as soon as they become available.

Performance

Due to hardware response time limitations and forecast accuracy being insufficient (among other challenges) *consistent* demonstration of down-ramp control could not be demonstrated during the testing phase of the pilot project. However, up-ramp

⁴⁰ *The Economics of PV Plant Variability on Distribution Equipment*. EPRI, Palo Alto, CA: 2015. 3002005782.

⁴¹ Hoff, T., et al. "Understanding Variability and Uncertainty of Photovoltaics for Integration with the Electric Power System." *no. LBNL-2855E* (2009): 18.

⁴² Such as to understand what transmission infrastructure is necessary to enable the smoothing or to navigate intricacies associated with smoothing across balancing authority boundaries

control and brief periods of proper down-ramp control were observed during field testing. A couple examples have been compiled as evidence of a successful proof of concept. Further, the core of the algorithm was shown to work through an emulation, helping to isolate what performance challenges (see next subsection) remain to be addressed.

Emulated Scenario

The emulated test case shown in Figure 4-2 uses the same plant dispatcher software and RRP operating mode logic deployed in the field at Springbok 3, but the dispatcher reads a “perfect” forecast and is connected to simulated PV and PCS devices to demonstrate RRP performance without some of the limitations of the Springbok 3 hardware configuration and forecast accuracy.




Figure 4-2
Sample of emulated plant response (gray line) to rapid changes in available PV power (blue line)

PV production in this test case is simulated using actual 1-second Springbok 3 PV production data from a cloudy day with variable production (1/7/2021). The emulation uses a “perfect” forecast generated in the same format and granularity as the real Springbok 3 short-term forecasts by averaging PV generation over each 1-minute interval. This process creates a forecast that allows the RRP mode to accurately predict the timing and magnitude of upcoming drops in PV production and eliminates discrepancy in forecasted values between forecast fetch times. This emulated test setup includes some simulated communications latency and non-instantaneous PCS response time, but the system

responds to and settles at RRP mode curtailment targets on a sub-second timescale. This is 5-10 seconds faster than measured response to curtailment commands using the Springbok 3 hardware configuration. This test setup is therefore highly idealized in terms of forecast accuracy and device response to curtailment commands, but it provides illustration of potential real-world RRP performance if improvements are made in these areas.

Up-Ramp Control

RRP was successfully able to limit ramp rate in the up direction by issuing curtailment commands to the inverter when ramps exceeded the specified rate. The example in Figure 4-3 shows a portion of a cloudy day (3/10/2021) with variable PV generation. The blue line shows the power output of the power block under test while the orange line shows power output from a nearby power block not being tested as an approximate representation of what the blue line would have looked like had it not been running RRP. On this day, RRP was set with an up-ramp limit of 10% per minute (5.4 kW/second for this 3.225 MW system). PV up-ramps that exceeded 5.4 kW/second were curtailed to steady, linear increases in production at the prescribed rate, as exemplified by the three examples that occurred for the blue line after sudden dips in PV (see orange line) in Figure 4-3.




Figure 4-3
Sample field measurement data from neighboring power blocks (with PB5 under a RRP test) demonstrating three instances of proper up-ramp performance

Down Ramp Control

Predictive down-ramp control at the prescribed ramp rate was attempted but only occasionally successful. This was because, to work properly, all the following conditions had to be met:

- Solar forecasts had to be able to predict the exact (or near exact) timing and magnitude of a coming drop in production
- Dips in PV power had to be sufficiently large such that requisite preparation periods exceeded at least one minute (resolution of the PV forecasts received) plus any latency in receiving the forecast as well as an additional 5-6 seconds for the plant response time
- The RRP algorithm had to work as expected, including the dynamic adjustments to the PV curtailment value, which the plant was not originally designed to perform

Figure 4-4 shows two instances of down-ramps being executed by the power block under test (PB5) in blue. Comparing the output to surrounding power blocks not running RRP, it can reasonably be assumed that the first down ramp was a false alarm (decrease in available PV did not occur as forecasted) while the second was an example of proper performance. In this case, the output of the power block running RRP preemptively started ramping down at the prescribed rate (10%/minute) ahead of the rapid drop in available PV power such that by the time the bottom of the dip appeared on surrounding power blocks, PB5 was already curtailed at/below this power level.




Figure 4-4

Sample field measurement data from neighboring power blocks (with PB5 under a RRP test) demonstrating one superfluous and one successful instance of proper down-ramp performance

Performance Challenges

There are two key reasons for the discrepancy in RRP performance between up-ramps and down-ramps:

1. The clearest difference is that RRP's down-ramp logic relies on high-accuracy short-term solar forecasting, while RRP's up-ramp logic does not require any forecasts for the algorithm to work.
2. The other difference is rooted in how up ramp and down ramp control algorithms handle rapid changes in PV within a control interval (PV power changes in the period between successive curtailment commands sent to the power conversion system).

Each of these challenges are addressed in more detail in the subsequent subsections.

PV Forecasting Accuracy Limitation

The RRP mode relies on extremely accurate, high-resolution solar forecasts that need to capture both the exact timing and magnitude of rapid changes to be most beneficial. If actual PV drops out even just a few seconds before forecasted or if the magnitude of the drop is larger than forecasted, the mode cannot perform correctly and a significant ramp violation at the POI meter could result. Additionally, the function would perform best if

forecasts do not vary significantly between fetch times. One bad forecast set has the potential to negatively affect performance, even if all other forecasts are of high accuracy. This combination of consistency and high accuracy has yet to be regularly realized at Springbok 3 and is the most significant hurdle facing proper RRP performance demonstration.

For a visual, Figure 4-5 depicts forecasted values in orange (only the first forecast set of every half hour shown for visual clarity) relative to measured values in blue on a highly variable day. While the forecasts capture the variable nature of the resource over the course of the day, the precise peaks and valleys rarely align with forecasted estimates.




Figure 4-5
Sample day showing forecasted PV power profile (consisting of only the first forecast from every half hour to avoid visual clutter from overlapping forecasts) relative to measured power profile

The sky camera forecasts at Springbok 3 have high granularity and accuracy relative to other forecast options, but to demonstrate consistent down-ramp performance at this scale, sub-minute granularity and hyper-localized accuracy (able to forecast exact timing of clouds arriving not just at a site, but at a specific power block) may be necessary.

Tuning Camera Forecast Algorithm for RRP

On understanding the significant reliance on accurate short-term forecasting to perform ramp-rate control, various initiatives were taken to tune

the camera-based forecasting algorithm for this use-case.

Originally, the forecast service provided a “blended” forecast that weighted outputs from the following two components to generate the forecast:

1. A persistence model that used actual production from the last few minutes
2. A cloud-vector-motion algorithm which is a function of sky images and weather data

This blend was their standard product, and the weights were adjusted and tuned to minimize the Mean Absolute Error (MAE)⁴³ of the forecast stream.

While conventional metrics like MAE were indeed lower for the blended forecast, such metrics did not focus intently on the key attributes of PV ramps: dips and surges in PV production. Instead, such metrics attempted to minimize error across all observations, including many clear days in which the value at any given time was much more likely to resemble those from a few minutes’ prior.

Through experimentation we learned that a forecast service which heavily weights the sky-image-based cloud vector motion forecast relative to the historical production component was much better at predicting changes in production due to cloud cover, both in terms of timing as well as magnitude (even though it would score lower on average when looking at clear and overcast periods alongside highly variable ones). We subsequently adopted this adjusted forecast for the project.

Coordination of Forecast Timestamps

Before any of the above improvements can be made, though, correct implementation of forecast timestamps and data structure should be confirmed. It is suspected that forecasts fed into the algorithm may have experienced erroneous offsets from the time they purported to forecast that increased with horizon. As received, successive

⁴³ Average across all timesteps of the magnitude difference between the forecasted value and the corresponding realized (measured) value

forecast horizons looked like those in Figure 4-6, where forecasted events often appeared to be forecasted two minutes sooner with every new forecast received (one minute passed). Correcting for this suspected offset at each horizon, successive forecasts were more consistently in agreement as to when significant events in PV output would occur, even if the exact output at that time changed with successive forecasts (as seen in Figure 4-7).




Figure 4-6
Forecasted PV power estimates at different forecast horizons plotted for successive forecast sets (unmodified)




Figure 4-7
Forecasted PV power estimates at different forecast horizons plotted for successive forecast

sets after shifting (modifying) purported forecast horizons

Response Time and Sub-Control Interval PV Change

The ramp up and ramp down algorithms in the RRP mode must apply ramp limits differently due to the physical limitations of PV production. As a result, the ramp down algorithm is less equipped to handle the challenges of asynchronous control.

Up Ramps

In both ramp up control logic and ramp down control logic, the RRP mode calculates the allowed ramp for each control interval (in kW) as follows:

$$\text{Allowed Ramp (kW)} = \text{Ramp Rate Limit (kW/sec)} * \text{Control Interval (sec)}$$

To calculate the actual curtailment limit sent to the PCS, the ramp up algorithm uses the following equation:

$$\text{Curtailment Limit (kW)} = \text{Current PV Production (kW)} + \text{Allowed Ramp (kW)}$$

Crucially, however, the ramp up algorithm also limits the curtailment limit within each dispatcher control interval (the time between successive commands sent by the DG-IC to AEROS) using the following constraint:

$$\text{Curtailment Limit (kW)} < \text{Previous Command (kW)} + \text{Allowed Ramp (kW)}$$

Using only the first equation without the second limit, successive curtailment commands could be much further apart than the allowed ramp in cases where PV rises significantly within a control interval. For example, if output rose 2.5 times the allowed ramp limit in one interval and then held steady at that level, the first step would be limited to 1x the ramp limit, but then unconstrained PV production would be used as the new baseline, and the next timestep would pass through the remaining 1.5x ramp limit difference since PV production was not technically changing in this timestep. This is where the second equation is needed to cap the change at another 1x the limit and prevent the first equation from blindly using the max current PV

production available. The constraint remains the active limiting factor until the slower-changing ramp-limited output catches up with the current unconstrained PV output.

Down Ramps

To calculate the actual curtailment limit sent when the mode identifies the need to ramp down, the down-ramp algorithm uses the following equation:

$$\text{Curtailment Limit (kW)} = \text{Current PV Production (kW)} - \text{Allowed Ramp (kW)}$$

Unlike the up-ramp algorithm, however, the down-ramp algorithm cannot apply a secondary limit based on the previous command sent because the correct starting point for the ramp occurs a variable number of steps prior to the critical point, depending upon the cumulative change in PV during that time, and if missed, there is no place from which to add power to the now-lower PV output.

Hypothetically, say the down ramp algorithm applied the following limit:

$$\text{Curtailment Limit (kW)} > \text{Previous Command (kW)} - \text{Allowed Ramp (kW)}$$

and PV dropped faster or more severely than forecasted in the first timestep. In this case, the PV output could swing uncontrolled between the commanded curtailment limit on the upper end (that it has already violated) and zero on the lower end while the curtailment commands slowly descended at the allowed ramp rate to the actual PV production available.

The consequence of not having a viable constraint for within a dispatcher control interval is that changes in PV production within that interval (whether due to changing irradiance or PCS/MPPT control algorithms), change the baseline PV production value from which the allowed ramp must now be applied. If measured PV production does not respond to (or does not fully respond to) a dispatcher down ramp setpoint before the next setpoint is calculated and sent to the inverters, the subsequent setpoint (calculated as *Current PV Production (kW) – Allowed Ramp (kW)*) may not represent the desired ramp rate below the previous

setpoint. Basically, if *Current PV Production* fluctuates and does not settle on the requested setpoint by the subsequent command, it is difficult to enforce a steady decrease at the prescribed rate. This issue is exacerbated when *Allowed Ramp kW* values per control interval are very small and within the realm of typical PV fluctuations.

The challenges raised by the differences between setpoints and current outputs can be partially mitigated by setting a long enough control interval to allow the PCS to settle at the previous curtailment command; however, PV output will still be changing during this time, resulting in inaccuracies and control challenges. To complicate matters further, these field tests were performed using a hardware setup not intended for rapid and accurate response to PV curtailment commands, which compounds the impacts of these RRP mode asynchronous control performance challenges. In fact, in the original installation, PV output was not even metered, but was instead calculated as the settled output of the inverter minus the output of the dc-dc converter (a slow process). Designing the metering and controls for rapid PV responses (faster than natural variability) is the preferred solution to many of these stability challenges.

Conclusion

Despite emulations that demonstrate a properly functioning ramp rate product (RRP) algorithm, ramp rate control without using an energy storage system was intermittently successful at best when field tested. Up-ramp management performed as desired on almost every instance, but there were some challenges faced with continuous management of down-ramps, primarily when it came to PV forecast accuracy and control speed stability. While there are a couple examples where the system curtailed as expected during down-ramps, the vast majority were not or only partially managed. The following steps can be taken to improve operation of down-ramp control with this control mode.

1. Identify a short-term forecast tuned to accurately predict the timing and magnitude of ramps in PV production
2. Select a PV power conversion system capable of and tuned for fast response to changing curtailment commands

As an early demonstration of this type of control, the field tests were significant in their attempt to prove general viability of the RRP function as a proof of concept. However, further advancements in short-term forecast accuracy and refinement of controller tuning practices were identified as necessary improvements before the RRP operating mode can be viewed as a viable alternative to ramp rate control using energy storage.

Finally, the following four questions likely need further investigation before broadly deploying this technological solution to PV variability at any utility:

1. Is controlling ramp rate at the individual plant level necessary (i.e., what problem is it causing?) or does spatial averaging across multiple PV plants sufficiently mitigate the magnitude of instantaneous ramps enough to permit stable grid operation?
 - a. It is important to quantify the extent of ramping at a system level to assess if/how much of a problem the “do-nothing” case is.
2. If it is necessary to mitigate ramps induced by PV, what is the financial benefit of regulating PV ramps at the plant level relative to alternative solutions?
 - a. i.e., it may be more economical to mitigate ramps centrally rather than at the individual plant level.
3. If regulation at the plant level is desirable, what is the cost benefit tradeoff (and potential breakeven point) of the RRP control mode relative to traditional methods using energy storage?
 - a. Factors to consider may include total PV curtailment and forecast subscription costs vs. energy storage hardware costs, energy storage cycling losses, other use cases of energy storage, etc.
4. If ramp rate control is necessary/useful and RRP is the preferred method, what advancements need to be made to short-term forecasts and/or the algorithm to increase performance? Can probabilistic PV ramp forecasts play a role?

Chapter 5: Full-plant Voltage Regulation with Volt-var Control Function

Objective & Significance

The primary objective of this part of the project was to demonstrate voltage regulation support by means of inverter reactive power injection/absorption; however, it is the project details and auxiliary tests that added significance to this increasingly commonplace inverter capability.

- **“Vars at Night”** – Voltage regulation support by PV inverters usually ceases when the sun goes down and the inverter loses dc power, but this pilot aimed to demonstrate var support from PV inverters at night, allowing inverters to provide grid services 24-7.
- **Watt-hour Cost of Vars at night** – Given that nighttime reactive power from a PV inverter is relatively novel, this project provided an opportunity to characterize not only its nighttime performance, but also its active power consumption characteristics.
- **Utility-scale Reactive Power Functions** – Fixed-var and Volt-var capabilities have been demonstrated numerous times on individual inverters, but this project aimed to demonstrate a similar capability at utility scale from a 90 MW PV plant that consists of 32 discrete power blocks.
- **Operation Outside Normal Power Factor Range** – Interconnection agreements between LADWP and large PV plants (like the Springbok 3 plant under test) usually stipulate that the plant operate within ± 0.95 power factor (Up to 29.6 Mvar if Springbok 3 plant is operating at full 90 MW) bounds (unless asked to do otherwise), but this demonstration project afforded the opportunity to operate outside of this range when performing voltage regulation services (region highlighted in yellow in Figure 5-1). Reactive powers up to ± 30 Mvar were attempted, even during times of low or zero active power.

- **Visible Impact on Transmission-Level Voltage** – The sizable reactive power capability of the Springbok 3 plant meant that it could have measurable impacts on transmission level voltage. Through field testing, the project aimed to show improvements in voltage when voltage regulation was enabled relative to when it was absent/disabled.

The following subsections detail the two testing phases performed. The first phase focused on calibrating active power consumption resulting from reactive power activity. The second phase details the voltage regulation performance of the PV plant over multiple, multi-day test periods. During this phase the plant was responding to grid voltage variations using an autonomous Volt-var curve. The chapter concludes with a comparison of inverter-based solutions to more traditional voltage regulation technologies.

Types of Voltage Regulation

Regulating voltage can mean several different things including, but not limited to, the following:

- Acting as an available resource upon which a utility can call for more, less, or a fixed quantity of reactive (or potentially active) power
- Injecting or absorbing a set amount of reactive (or active) power based on the measured voltage according to a predefined response curve (Volt-var or Volt-Watt)
- Injecting or absorbing an increasing/decreasing amount of reactive power in response to a measured voltage until either the voltage is within desired limits or the plant output is maxed out

In this project, a Volt-var function was implemented to demonstrate voltage regulation, although the active power consumption calibration test also proved the capability of LADWP’s Energy Control Center (ECC) to issue reactive power setpoint commands.




Figure 5-1

Quoted active and reactive power capability of the Springbok 3 plant (for VPOI = 1.00 pu) with allowed range outside of 0.95 PF highlighted

Active Power Consumption Calibration Test

Using reactive power from inverters comes at an active power cost. The injection/ absorption of reactive power incurs additional active power consumption (sometimes viewed as losses) from the inverter's switching components and control circuitry, line losses (between the inverter and point of interconnection), and other secondary sources. In addition, inverters operating near full capacity could, in extreme cases, be forced to curtail active power generation in favor of reactive power when inverter current capacity is insufficient for both, though this was not a contributing factor in this project.⁴⁴ If it were, however, this loss of generation would also be considered an active power cost of voltage regulation using reactive power.

Testing in this project included active power consumption calibration tests for three purposes:

1. Improved quantitative understanding of the active power consumption that can be expected from inverter-based reactive power sources/sinks at this scale (for later comparison alternative technologies like synchronous condensers)
2. Ability to estimate incurred active power consumption so that testing could be concluded

within a budgeted 300 MWh of total consumption

3. Validation of PV plant ability to respond to routine reactive power commands via DNP3 from LADWP's Energy Control Center (ECC)

Test Procedure

Tests consisted of a profile of fixed var setpoint commands being issued to the plant controller at night, when smaller, known amounts of active power would typically have been consumed by the inverters.⁴⁵ Figure 5-2 is an example of one such profile. In all, five such profiles were run, one profile per night. 8minute conducted two nights of testing for active power consumption calibration purposes, and LADWP's Energy Control Center (ECC) ran another three nights to demonstrate that they could command reactive power from the plant via SCADA (using DNP3 protocol).

Each reactive power setpoint profile held fixed reactive power values for a period of 10-30 minutes before moving on to the next setpoint value as exemplified by Figure 5-2. This approach allowed for sufficient settling time between setpoints and increased chances to observe any offset inaccuracies.⁴⁶

⁴⁴ The large generator interconnection agreement (LGIA) for this plant required that it be capable of full reactive capability at the same time as full real power production.

⁴⁵ Under normal operation, plant inverters would be turned off and put in standby mode at night

⁴⁶ The staircase profile in Figure 5-2 was used initially to minimize the magnitude of subsequent reactive power changes, which could have sudden and significant effects on transmission voltages. It also afforded opportunities to terminate the test prematurely if impact on voltage became too large.




Figure 5-2
Example reactive power setpoint profile (and closely matching measured response) from a fixed var active power consumption calibration test

After executing the tests, the varying values of measured reactive power were compared against the associated measured active power to inform the calibration curve.

Results

Data from the two active power consumption calibration tests is shown in raw form in Figure 5-3. The curves look upside-down relative to the final consumption calibration curve because the active

power measurements were in terms of generation, not consumption. Additionally, active power consumption on a baseline (non-test) night is non-zero, so the raw data had to be shifted vertically to account for this baseline behavior. This offset value was determined to be about 317 kW and was calculated by taking the median of observed, nighttime, idle power consumption from a period in which no nighttime testing was being run (March 1st – May 18th, 2020).⁴⁷ Finally, data from test night 1 were excluded due to atypical behavior on multiple power blocks potentially contaminating the data.

⁴⁷ The median also excludes the 22.8 kW that was typically consumed by Power Block #5 since this power block was

disconnected for active power consumption calibration testing due to the presence of the BESS.




Figure 5-3

Scatter plot of raw active power and reactive power measurements before necessary modifications were made to isolate true active power consumption calibration curve

The primary result from the consumption calibration test was the active power consumption curve in Figure 5-4 generated by plotting presumed active power consumption⁴⁸ against the reactive power measurements from the same point in time. There are a few things to note about the result.

- Idling in nighttime mode, even when injecting zero kvar of reactive power, consumed an *additional* 575 kW to 725 kW of active power on top of ~317 kW of baseline consumption

- The relationship between active power consumption and reactive power injection/absorption is roughly a quadratic one—active power consumption increases with the square of reactive power commanded⁴⁹

Active power consumption was not quite symmetrical between reactive power injection (orange in Figure 5-4) and absorption (blue)⁵⁰

⁴⁸ Average baseline power consumption from multiple weeks' worth of normal nighttime plant behavior (no reactive power or other nighttime operation) were subtracted from the measured value to leave only presumed additional active power consumption from the reactive power.

⁴⁹ The quadratic nature is likely a result of I^2R losses.

⁵⁰ This asymmetry is expected as it is inherent to differences in inverter control behaviors (specifically necessary voltage levels) when injecting vs. absorbing reactive power; however, the asymmetry did impact the approach taken for calibration.




Figure 5-4
Active power consumption calibration curve from nighttime reactive power testing

Given the asymmetry, two quadratic functions were fitted to the positive and negative halves of the active power consumption curve. The curve was then used to estimate future consumption by measuring reactive power values of the plant when under test and converting each into its corresponding active power consumption estimate.

Additional findings from the fixed var tests are best visualized back in Figure 5-2. Most notably, measured output tracked the desired setpoint with a mean average error of between 0.32% and 1.22%, depending upon the test sequence being run. The plant took several seconds for its reactive power setpoint to settle to the desired value as control signals were sent to the numerous power blocks. Median settling time was between about 10 and 24 seconds depending upon the test night.

Controls Power Source Transition

Power for inverter controls is often supplied from the dc side of the inverter. There are multiple advantages to this choice such as automatic shut-down (and thus energy savings) at night (when PV is unavailable) and the ability to stay powered even when isolated from the grid. However, operating the inverter at night requires power, which must be either supplied by another source (e.g., storage) or drawn from the grid.

At the Springbok 3 site, the requirement per NEC 2014 was to transition the controls from dc PV power to ac grid power daily—just before sunset—and vice versa after sunrise. Originally, inverters conforming to NEC 2014 codes needed to interrupt power to make the transition, since contactors to PV modules were not designed to open under load. This resulted in 1- to 2-minute shut-down periods for each transition. As the project progressed, most inverters were upgraded to the NEC 2017 specifications, which requires series fuses (or a disconnect switch) but permits opening contactors without interrupting load.

Future projects should likewise be able to perform this seamless transition and may be able to avoid the complications associated with the disruption to controls supply power.

Extended Volt-var Testing

Once active power consumption calibration testing was complete, extended testing of a voltage regulation function took place. In the past, ECC had requested various fixed reactive power setpoints depending upon grid conditions, but for this test, an autonomous Volt-var function was to be used as the default mechanism for voltage regulation.

Test Procedure

All extended testing of the Volt-var function occurred with the use of the single curve definition shown in Figure 5-5 and enumerated in Table 5-1.

Table 5-1
Power plant controller parameters used for Volt-var voltage regulation tests

PPC Parameter	Argument
Voltage setpoint	240500 (V)
Low voltage max var	30000 (kvar)
Low voltage deadband	0.208 (%)
Low voltage droop	1.46 (%)
High voltage max var	30000 (kvar)
High voltage deadband	0.208 (%)
High voltage droop	1.46 (%)




Figure 5-5
Volt-var curve used for voltage regulation tests (black) overlaid onto historic voltages and reactive powers (blue) from periods of pre-testing operation at the Springbok3 site

The curve was designed to emulate, or even exceed, past responses to voltage while maintaining stable plant operation. Preliminary Volt-var tests (prior to extended testing) had revealed the relationship

between changes in reactive power output and changes in observed local grid voltage as shown in Figure 5-6, so the slope of the Volt-var curve for extended testing was set to be shallower than the

30Mvar/2.5kV slope that the plant inherently induces.⁵¹

injection/ absorption at the Springbok 3 point of interconnection




Figure 5-6
Inherent voltage response to reactive power
Table 5-2
Long-term Volt-var test periods

	Period 1	Period 2	Period 3
Dates Run	9/9 – 9/16	10/1 – 10/10	10/19 – 10/24@22:15*
Duration in Days	8	10	~6
Consumption During Period	60.1 MWh**	92.4 MWh**	54.2 MWh**
Cumulative Consumption	153.4 MWh (93.3 MWh from prior tests)	245.8 MWh	300 MWh

*Stopped due to hitting active power consumption limit
**For daytime operation, considered consumption only for VARs in excesses of 0.95 PF

Performance

In all, about 24 days’ worth of Volt-var tests were performed. Figure 5-7 portrays a single test day that illustrates the behavior typically seen during testing. At most times of day, reactive power responses were in line with expectations, injecting when voltages were low and absorbing when

voltages were high. The red boxes indicate the two periods this day where the Volt-var response did *not* align with expectations. Plausible explanations for these atypical periods fall into one of two buckets:

- 1. Periods around sunrise and sunset could potentially have been triggered by earlier NEC 2014 versions of the HEM model inverters not being able to support a seamless transition between day and night mode.⁵² Around this

⁵¹ A steeper slope could cause oscillatory behavior in response to voltage deviations due to the induced change being larger than the initial voltage deviation

⁵² The non-seamless transition was not due to any written NEC sections, but rather the nature of the DC contactors used on the inputs of each unit. The inverters in one part of the

transition period, reactive powers would sometimes deviate from their desired setpoints for extended periods (beyond the 1-2 minutes typical of the transition).

2. Other periods that ranged in length from just a minute or two up to around two hours are presumed to have been a result of loss of communications⁵³ and/or software changes⁵⁴.




Figure 5-7

Reactive power response (bottom) to measured voltage (top) with two highlighted periods where output did not align with expected behavior at that time (right highlight occurred around sunset). The estimated “no Volt-var” voltage utilized the slope found in Figure 5-6 to back calculate what the voltage might have been in the absence of any reactive power flow.

To better observe how accurately the Springbok 3 plant followed the desired Volt-var curve, the reactive power response for all the test days in

which Volt-var was active is plotted against measured voltage in Figure 5-8.⁵⁵

plant have since been updated with DC contactors conforming to NEC 2017 specifications, which allows them to transition seamlessly.

⁵³ When the MV switchgears would open on a fault, the SCADA system installed on that unit would run on power from an uninterruptible power supply (UPS). If the inverter that faulted was on a primary fiber node, other inverters connected on this fiber node would lose communication, and as such, would not be able to send read/write commands to support

the Volt-Var testing. The UPS batteries only lasted couple of hours, and depending on when they faulted, there would be a random loss of communications.

⁵⁴ During the Volt-Var test period, the inverters were going through software changes on the equipment to better support issues with the switchgear opening on a fault

⁵⁵ Testing occurred in three sets with the strings lasting 8, 10, and 6 days respectively.

Around 87% of measured response points (orange) lie near the desired Volt-var curve (black) indicating the reliability observed during this ~3.5 weeks of testing. The remaining points reside mostly at zero Mvar, despite being outside the deadband. These points indicate periods where the function failed to perform accurately.⁵⁶ The ‘trails’ of points going to and returning from zero Mvar usually indicate periods where the output was slowly transitioning to or from non-functioning behavior. Most of these

deviations occurred on the voltage absorption (elevated voltage) side of the curve.

The offset to the right is a result of how voltage is averaged across the three phases and not anything to do with the algorithm; however, calibration of source meters and specification of where voltage and reactive power are to be measured are key considerations.




Figure 5-8

Observed Volt-var response from second string of test days (10-days long) as measured by power plant controller (PPC).

Despite the periods where the plant response would deviate from the desired curve definition, the algorithm still had a noticeably beneficial effect on the voltages observed at the site, as supported by Figure 5-9. Measured voltages on the final 18 test days in which the Volt-var function was enabled (green) had a significantly tighter distribution compared to the 22 interspersed baseline days in which it was not enabled (blue). In fact, when putting aside a few anomalous readings (from when

the function was not running but should have been), the range of observed voltages was shrunk from 236 kV to 238 kV on the lower end and 247 kV down to about 244 kV on the upper end.

To make sure that these effects were most likely attributable to the actions of Springbok 3, nearby PV plants connected to the same 230 kV bus were run in a fixed power factor mode the whole time, and capacitor switching at Springbok 3 was also restricted during testing.⁵⁷

⁵⁶ The line of points in the 2-4 Mvar range were also associated with points where Volt-var failed to perform, but for an

unknown reason the plant would sometimes default to a non-zero output

⁵⁷ The same was true for testing of fixed var commands as well.




Figure 5-9
Distribution of measured voltage at plant point of interconnection with and without Volt-Var function enabled

The ‘cost’ of this voltage regulation was regular reactive power participation by the inverter which resulted in active power consumption. While realized participation will vary depending upon grid conditions, the point of interconnection, etc., the cumulative distribution function (CDF) in Figure 5-10 shows the percentage of time a given reactive power (or less) was required during the 24 days of Volt-VAR testing.

Combining the reactive power usage with the active power consumption estimation curve in Figure 5-4, an estimated 9.5 MWh of additional energy per day was consumed to perform the voltage regulation observed.⁵⁸ The final section of this chapter compares inverter-based voltage regulation to alternative voltage-regulation technologies in an attempt to provide better context of the strengths and weaknesses of each solution.




Figure 5-10
CDF indicating percentage of time that reactive power contribution was a certain value or less

⁵⁸ The active power consumption calculation includes other considerations for daytime operation such as excluding

reactive powers within 0.95 PF since these could reasonably be considered part of normal operation requirements.

Continued Challenges

Voltage regulation using a Volt-Var function on PV inverters was proven as feasible but not yet highly reliable at this particular plant. The first hurdle remaining is consistency of response when in Volt-var mode. Departures from the desired response curve in this test can likely all be resolved with different instrumentation and upgraded hardware; however, it will be important to demonstrate consistent operation before PV plants can be comfortably relied upon to perform this service.⁵⁹

A second area of improvement for inverter-based reactive power voltage support will be active power consumption coefficients, particularly those observed during standby operation. Though just one aspect to consider when comparing voltage regulation technologies, operating costs for inverter-based resources typically exceed those of alternative solutions such as static var compensators (SVCs) due to the relatively high active power consumption incurred. More detail on this is provided in the next section.

Review of Existing Transmission-level Voltage Regulation Technologies

In the previous subsections, the capability and performance of using a PV plant to provide reactive power for voltage regulation has been demonstrated. This subsection reviews other transmission level voltage regulation solutions in the industry. The operating principles are presented first, and then compared qualitatively in terms of control speed, granularity, cost, and other relevant aspects. Representative commercial products for the identified technologies also are summarized.

Voltage Regulation Solutions and Their Operating Principles

Three major voltage regulation solutions (besides IBRs) have been identified in this report:

1. Synchronous condensers
2. Static var compensators (SVC)
3. Static synchronous compensator (STATCOM)

Schematic-style depictions of each technology and its primary, constituent, electrical components are shown in Figure 5-11.

⁵⁹ At the Springbok facility, the PE inverters with the HEM model are configured with the earlier NEC 2014 versions of the HEM which cannot support the seamless transition from day to night and vice versa—not due to any written NEC specifications, but rather the nature of the DC contactors used on the inputs of each unit. PE has since updated the DC

contactors with those conforming to the NEC 2017 specs that would allow the inverters to transition between day and night modes seamlessly. PB5 has this spec installed and smooth transitions have been observed.




Figure 5-11
Alternative technologies for transmission-level voltage regulation

Synchronous Condenser

A synchronous condenser is a synchronous generator without a prime mover on its shaft. Although it cannot generate active power, it has the same reactive power capability and performance of a synchronous generator with the same rating. If controlled to be over-excited or under-excited, the synchronous condenser can inject or absorb reactive power for voltage regulation. The amount of the reactive power can be dynamically adjusted both quickly and accurately by the exciter control of the synchronous condenser.

In addition, since the rotational speed of the synchronous condenser is directly coupled with the grid frequency, the kinetic energy stored in the rotor and the optional flywheel connected to the rotor can be exchanged for electrical power during a grid frequency event, contributing to grid inertia for system frequency stabilization. Like all other rotating machines, the synchronous condenser can also contribute to the fault current during a grid short circuit event, which may have positive or negative impact depending on the system needs.

The cost of a synchronous condenser is relatively high compared to an SVC or STATCOM with the

same rating. But a lot of times, the synchronous condenser is converted from a retired synchronous generator, so that the installation cost can be significantly reduced.

Static Var Compensator (SVC)

SVCs employ thyristors to operate the shunt-connected capacitors and reactors to support voltage regulation. An SVC may include any combination of one or more thyristor-controlled reactor (TCR), thyristor switched reactor (TSR), thyristor switched capacitor (TSC), mechanically switched capacitor (MSC), or other compensator.

Figure 5-12 shows an illustrative TCR voltage and current profile to demonstrate its basic operating principle. When the thyristor is turned on, the reactor connects to the system and absorbs reactive power. Since the firing angle of the thyristor can be dynamically adjusted in fine steps, the reactive current absorbed from the reactor can change accordingly. The TCR can be regarded as an adjustable inductor to the system. As seen from Figure 5-12, the thyristor control may produce low order harmonics to the system, and thus may require additional filtering equipment. The active power consumption of TCR also increases as its reactive power contribution increases.




Figure 5-12
Illustrative TCR voltage and current profile to demonstrate its basic operating principle

An illustrative voltage and current profile of a TSC is shown in Figure 5-13. Different from TCR, TSC can only be operated fully-on or fully-off. As a result, there is generally no harmonic distortion generated from a TSC. Usually, there are several TSCs in parallel to achieve some level of granularity, and the active power consumption depends on the number of TSCs connected to the system.




Figure 5-13
Illustrative voltage and current profile of a TSC to demonstrate its basic operating principle

Other possible components in an SVC can include a TSR, MSC, or harmonic filter. Altogether, the SVC may exhibit the characteristic of a variable impedance to the grid, injecting or absorbing reactive power for voltage regulation. Because of the solid-state thyristor control, SVC components such as TCR and TSC can operate in tens of milliseconds,

which is much faster than the traditional solutions such as MSC for system voltage regulation.

Static Synchronous Compensator (STATCOM)

STATCOM, also known as Static Synchronous Compensator, employs an inverter to inject or absorb reactive power. Similar to the analogy between synchronous condenser with synchronous generator, a STATCOM may be regarded as an inverter-based resource (IBR) without its active power source. For reactive power generation or absorption, it shares the same operating principle, capability, and dynamic characteristics of an IBR with the same rating. Figure 5-14 shows the illustrative voltage and current profile of a STATCOM. The inverter can modulate its fast-switching power electronic switching devices to approximate a sinusoidal reference voltage and control its output reactive current to inject/absorb reactive power for voltage regulation. With inverter technology, the STATCOM is able to respond to system changes within one line cycle, which is faster than the SVC.




Figure 5-14
Illustrative voltage and current profile of a STATCOM to demonstrate its basic operating principle

Comparison of Voltage Regulation Technologies

This subsection compares the identified voltage regulation technologies in terms of the following:

- Reactive power capability at low voltage
- Efficiency

- Cost

A comparison chart to summarize the discussions is provided at the end of this subsection.

Reactive Power Capability at Low Voltage

The maximum current of a STATCOM is determined by its thermal design and can usually be maintained even when the voltage is low. However, the current from an SVC is determined by its terminal voltage and its equivalent impedance. When the voltage drops low, the current drops proportionally with the voltage, resulting in lower reactive power capability for the SVC than the STATCOM at low voltage. The illustrative comparison is shown in Figure 5-15.




Figure 5-15
Illustrative comparison of reactive current (top) and reactive power (bottom) capability between SVC and STATCOM

Efficiency

As described in the previous subsections, power consumption in an SVC depends on the number of thyristor switched capacitors/reactors (TSCs/TSRs) connected to the system or how much thyristor-controlled reactors participate in reactive power absorption, dependent upon the configuration. The illustrative example of system active power consumption with respect to the reactive power output is shown in Figure 5-16. For SVC1, only one TCR and one fixed capacitor is in the SVC system. The consumption is lowest when the TCR does not operate and increases when the TCR starts to absorb reactive power. For SVC2, there are multiple TSCs available, and the active power consumption would have a step change when a TSC is switched in to inject reactive power.

On the other hand, the operation principle of STATCOM is the same as an IBR. Similar to the power consumption analysis of the PV plant in the previous subsection, the STATCOM will have lower active power consumption at zero output, and the consumption increases smoothly with more reactive output, as shown in Figure 5-16.

SVC and STATCOM have different efficiencies (i.e., active power consumption characteristics) at different reactive power values. According to the comparison conducted at GE, overall weighted consumption of an SVC is lower than that of a STATCOM with the same rating.⁶⁰

The active power consumption of a synchronous condenser is generally about 1-3% of its rated power.⁶¹ The total consumption includes mechanical friction and windage losses as well as electrical Ohmic and core losses. With more injection or absorption of reactive power, mechanical and core losses tend keep relatively constant, while Ohmic loss on the windings of the synchronous condenser increases. Thus, the total active power consumption of a synchronous

⁶⁰ Neil Kirby and Anurag Guha, "Making the Right Choice: A Comparative Analysis of SVC and STATCOM". Source: <https://resources.gegridolutions.com/facts/making-the-right-choice-a-comparative-analysis-of-svc-and-statcom>

⁶¹ M. Nambiar and Z. Konstantinovic, "Impact of using synchronous condensers for power system stability and improvement of short-circuit power in mining projects," Mining Engineering, January 2015, Vol. 67, No. 1

condenser has a similar shape as that of a STATCOM but is greater than a STATCOM with the same rating.⁶²




Figure 5-16
Illustrative comparison of operational active power consumption for SVC and STATCOM

Cost

In general, an SVC is less expensive than a synchronous condenser that is built new but more expensive than a repurposed synchronous condenser from a retired synchronous generator with the same rating.⁶³

Figure 5-17 recreates a cost comparison between an SVC and a STATCOM from a 2004 research paper.⁶⁴ Although the costs for both technologies have been reduced to less than \$100 per kvar in a more recent

⁶² S. Kynev, G. Pilz, and H. Schmitt, "Comparison of Modern STATCOM and Synchronous Condenser for Power Transmission Systems," in Proc. IEEE Electrical Power and Energy Conference (EPEC). 2016

⁶³ "Synchronous Condenser Market by Cooling Type (Hydrogen, Air, Water), Type (New & Refurbished), Starting Method (Static Frequency Converter, Pony Motor), End User

(Electrical Utilities & Industries), Reactive Power Rating, and Region-Global Forecast to 2025." Source: <https://www.marketsandmarkets.com/Market-Reports/synchronous-condenser-market-189197147.html>

⁶⁴ Klaus Habur and Donal O'Leary: "FACTS – Flexible Alternating Current Transmission Systems For Cost Effective and Reliable Transmission of Electrical Energy"

2016 research, STATCOMs generally remain more expensive than SVCs.⁶⁵




Figure 5-17
Cost comparison between SVC and STATCOM
recreated from a research paper published in
2004⁶⁴

Comparison Chart Summary

Table 5-3 provides a summary of the comparisons made between inverter-based resources (IBRs), SVCs, and STACOMs in regard to providing voltage regulation services.

⁶⁵ FO Igbinovia, G Fandi, Z Müller, “Cost Implication and Reactive Power Generating Potential of the Synchronous

Condenser,” in *Proc. International Conference on Intelligent Green Building and Smart Grid (IGBSG)*, 2016

Table 5-3

Comparison of voltage regulation technologies where the number of '+' symbols qualitatively indicates the strength of that technology in that metric (more is better)

	Inverter-based Resources	STATCOM	SVC		Synchronous Condenser
			TCR/TSC	MSC	
Control Speed	+++ ⁶⁶	+++	++	+	++
'Efficiency' (Active power consumption for reactive power)	+	+	++	+++	+
Cost (Capital and O&M)	\$ (incremental cost for additional voltage reg. functionality)	\$\$\$	\$\$	\$	\$\$\$ (new) \$ (from retired synchronous generator)
Granularity	+++	+++	+++ / +	+	+++
Reactive Power Capability at Low Voltage	+++	+++	+	+	+++
Other Considerations	Inverter hardware & installation likely to be driven/funded for other purposes already	May provide inertia if short term storage is included ⁶⁷	TCR produces low order harmonics		Provides inertia, High short circuit current contribution

Representative Commercial Voltage Regulation Products

In this subsection, a sample of representative manufacturers and their voltage regulation

products are compiled as examples of commercially available products. The technology types, product names, available information on their sizes and websites are summarized in Table 5-4.

Table 5-4

A sample of representative manufacturers and products

⁶⁶ Response times of IBRs in general can be very fast, though results may be slower if aggregating many inverters behind a plant controller. Plants intended to perform voltage regulation services should be designed for faster response times from the start (& thus specified in RFP requirements).

⁶⁷ Siemens provides SVC PLUS (STATCOM) Frequency Stabilizer technology for inertia response. Source: <https://www.siemens-energy.com/global/en/offerings/power-transmission/facts/portfolio/svcplus-frequency-stabilizer.html>

Company	Tech. Type	Product Name	Key Parameters	Product Website
Siemens	SVC	Static Var Compensator (SVC)	Example project includes a hybrid plant with 4 SVC PLUS, 2 TSC, and 1 TSR; -400/+600 Mvar at 380 kV	https://www.siemens-energy.com/global/en/offerings/power-transmission/facts.html
	STATCOM	SVC PLUS® (STATCOM)		
	Sync. Condenser	Synchronous Condenser	Example project includes a 1100Mvar plant	
Hitachi (originally ABB)	SVC	Static Var compensation (SVC)	Not found online	https://www.hitachiabb-powergrids.com/offering/product-and-system/facts/static-var-compensation-svc
	STATCOM	PCS 6000 STATCOM	20 Mvar to 100 Mvar	https://www.hitachiabb-powergrids.com/offering/product-and-system/facts/statcom
		SVC Light® STATCOM	Example project includes a +/- 400 Mvar plant at 400 kV	
GE	SVC	Static Var Compensator (SVC)	69 - 800 kV 40 - 1200 Mvar	https://www.gegridsolutions.com/power/catalog/static_var_comp.htm
	STATCOM	Static Synchronous Compensator (STATCOM)	Range to 300 Mvar	https://www.gegridsolutions.com/power/catalog/statcom.htm
	Sync. Condenser	Synchronous Condenser	10 - 300 Mvar+ per machine	https://www.gegridsolutions.com/power/catalog/synch_cond.htm
NR electric	SVC	PCS-9580 Static Var Compensator	Not found online	https://www.nrec.com/en/index.php/product/productInfo/297.html
	STATCOM	PCS-9583 Static Synchronous Compensator	Not found online	https://www.nrec.com/en/index.php/product/productInfo/300.html
Mitsubishi	SVC	SVC	Not found online	https://www.mitsubishielectric.com/ei/g/energysystems/products/transmission/pssfactsc/index_03.html
	STATCOM	STATCOM	Example project includes a +/- 450 Mvar plant	
Ingeteam	STATCOM	STATCOM	±100kvar to ±40Mvar	https://www.ingetteam.com/en-us/power-electronics/electrical-storage-facts-sssc/pc28_18_236/ingegrid-statcom.aspx

Chapter 6: Conclusion

This large-scale demonstration of inverter-based resources consisted of the following discrete areas of investigation:

1. Improvements in PV forecasting
2. Advanced MESA-DER-specified control capabilities with the combined PV+BESS system
3. PV ramp rate limiting (bidirectional) without the use of energy storage
4. Full-plant voltage regulation with the use of reactive power control

Results that are summarized below by topical area.

Findings Summary

PV Forecasting

Long-term forecasts provided in this project (from vendor Clean Power Research) demonstrated median forecast error of 3.5% for the Springbok 3 plant, down from a median forecast error of 12.5% in the forecasts that LADWP currently uses (for comparable forecast sets). The addition of satellite-based imagery in the forecast generation algorithm as well as a more refined plant model are likely the most significant sources of improvement.

Short-term forecasts based on inputs from a ground mounted, infrared, sky imaging camera (from vendor Reuniwatt) were roughly as accurate in terms of median absolute error to short-term forecasts informed by satellite imagery (from vendor CPR) with a mean error across months of 3.8% for Reuniwatt and 4.2% for CPR; however, the distribution of those errors was the same or tighter for Reuniwatt than CPR in the months observed.

Forecast accuracy was typically highest for clear days whereas error for variable days was substantially higher, hinting at the challenges associated with forecasting rapid changes in PV output.

Long- and short-term forecasts from CPR and Reuniwatt both varied seasonally with lower error in Spring and Summer (when the plant was more frequently saturated at its 90 MW injection limit) and higher error in the Autumn and Winter.

Similarly, forecasted PV power at this plant typically experienced the highest amounts of error in the mornings and afternoons due to physical properties of the site being exacerbated by solar tracking and oversizing of the PV & inverter capacity relative to the interconnection limit.

Many sources of error can arise between forecasts of global horizontal irradiance (GHI) and the forecasts of expected plant power generation, which are often estimated using forecasted GHI and a plant model. A flowchart in Chapter 2 (Figure 2-4) highlighted some of the possible sources.

Forecast services were not able to use telemetric feedback to train forecast models due to issues determining when inverters were out of service. It is expected that further accuracy gains may be possible if such feedback can be established.

Real-time feedback of telemetry (true irradiance and power measurements) to forecast vendors was complicated by latency from data historians and transfer mechanisms.

Analysis of a hypothetical improvement in area-wide PV forecast accuracy showed that together with the implementation of a risk-based dynamic reserve allocation method, reserve requirements could be reduced while maintaining the same degree of risk.

Demonstration of MESA-DER Functions for PV+BESS

Coordinated Charge/Discharge (aka SOC target schedules or CCD) typically achieved SoC targets within 2% under both “greedy charging” and CCD-P, which used PV generation predictions to provide additional benefits such as firmer plant output. In

addition, SoC targets were reached later in the charging/discharging window under CCD-P, which created more headroom/"footroom" for other services (and for more hours) and reduced average SOC and battery degradation. Nighttime grid charging limitations for this plant (as per the interconnection agreement) impeded the system's ability to hit some nighttime SOC targets requiring charging to maintain or increase SOC after sunset.

Active Power Smoothing (APS) tests demonstrated how the full APS algorithm (that targets a filtered moving average) improved ramps to be fully within the desired ramp rate bounds (+/-5% per minute), whereas only correcting for excess ramp rates mitigated their intensity but could occasionally exceed the limit due to the relatively slow response time of this particular system. On the flip side, the more active gradient smoothing (full APS) came at the cost of increased battery throughput relative to only correcting for excess ramps.

Frequency-Watt (FW) mode followed its respective FW curve definition with less than 5% error more than 99% of the time (when the BESS was permitted and had sufficient SoC to do so). Scheduling a nighttime discharge also allowed responses to high frequency to continue into the night.

Automatic Generation Control (AGC) likewise responded to setpoints accurately with mean average error of about 4% between the instantaneous response and the intended setpoint (when the BESS was permitted and had sufficient SoC to do so).

Response times that negatively impacted functions, especially FW and AGC, were due to stability challenges specific to this plant's dc-coupled PV+S hardware components and configuration. Further tuning of device controls and hardware may be able to mitigate some of these impacts.

Given the grid charging prohibition at the Springbok 3 site, SoC targets that require charging to maintain or increase SOC should be scheduled during daylight hours. Then, scheduled discharges at night can be used to extend the periods during which other MESA-DER functions can be effective.

While any number of control modes can be stacked, function prioritization is only applicable to basepoint functions. This is because basepoint functions are planned outputs whereas additions like frequency-Watt and AGC responses will get added on top of the base output whenever there is capacity to do so.

The component functions continued to perform as expected when combined into various mode stacks, achieving multiple mode objectives simultaneously with no significant degradation in individual mode performance.

Ramp Rate Limiting Without Energy Storage

The ramp rate product (RRP) control mode with PV forecasts is an experimental feature that is not yet ready to take the place of more traditional solutions involving energy storage.

Upward ramps were appropriately curtailed, but downward ramps that depended upon accurate short-term solar forecasts frequently exceeded the desired ramp rate limit. The downward ramp algorithm was shown to work properly in simulation, but real-world implementation was typically ineffectual due to insufficient accuracy of the near-term solar forecasts (both in timing and magnitude of events) and limitations on the inverter's curtailment limit response time, accuracy, and stability.

Since ramp rate limiting without energy storage depends solely upon curtailment, constraining ramps will come at the cost of decreased energy production, which may be exacerbated on highly variable days.

Additional research is needed to prove that regulating ramps at individual, transmission-connected plants is necessary and economically viable. If it is, cost benefit comparisons with energy storage or other plant design alternatives should be performed to determine the most economical choice.

Full-Plant Voltage Regulation with Reactive Power

Reactive power, specifically responses from a Volt-var curve at the Springbok 3 plant, narrowed the distribution of observed voltages closer to the desired voltage range.

Reactive power support from PV inverters has an active power cost associated with it, which is especially apparent at night. This project estimated that the cost varied quadratically with the magnitude of reactive power being injected/absorbed and ranged from roughly 600 kW to 1.1 MW while providing between 0 and +/-30 Mvar during a calibration test night.

To keep the PV plant in a “ready” state, in which it could provide reactive power support during nighttime instead of going into sleep mode, it absorbed between 575 kW and 725 kW in addition to its typical nighttime power consumption (baseline). Inverters in traditional PV-only plants are configured to go into sleep mode during nighttime, and would thus have lower nighttime baseline usage than the PV+S plants where inverters are expected to always remain in a ready state.

Multiple, more traditional voltage regulation technologies were presented alongside inverter-based resources; the relative strengths, weaknesses, and characteristics of each technology type were identified then summarized in Table 5-3.

Lessons Learned

Active and Reactive Power Controls

- The current structuring of the investment tax credit (ITC) encourages systems with batteries to be set up to charge only from PV generation and not from the grid. Case and point, the power purchase agreement (PPA) and interconnection terms for the battery at this site prevented grid charging entirely. As a result, MESA-DER functions that performed as expected under daylight conditions required modifications and preparation to demonstrate nighttime operation. Even then, capabilities would have to be scaled back further in some cases to operate 24 hours per day. Allowing inverter-based resources limited power consumption from the

grid to perform grid services, especially during nighttime when PV generation is not available, could be beneficial for power system reliability.

- Operating this plant at night (to enable nighttime inverter functions) required 10s of kW of power. Since this energy could not come from PV like it did during the day and charging from the grid was not permitted at this plant, the configuration at this site required that the BESS discharge a non-trivial amount of capacity to keep the dc-dc converter, inverter, and controls powered. Defining power going into the battery as distinct from consumption by electronics and controls may allow for the necessary flexibility to reduce unnecessary overnight discharge by the battery.
- PV+BESS setpoint response times (when on the order of seconds or longer) can impact performance of fast acting functions such as frequency-Watt or AGC, whereas functions like CCD, which work to achieve SoC targets over longer time horizons, are largely unaffected.
- Calculated state of charge values are only estimations and may drift over time as the battery chemistry settles and/or because of changing conditions such as temperature fluctuations, rate of charging/discharging, etc. Except where necessary, SoC estimates should only be used as a guide. Conservative kWh requirements may sometimes be a more appropriate metric by which to quantitatively judge energy-centric function performance.
- Direct monitoring of PV power is useful in decoupled systems. In Power Block 5 of the Springbok 3 plant, PV output was originally estimated by subtracting off measured ESS output (via a point exposed by the dc-dc converter) from the inverter output (measured by a commercial power meter). This approach led to some discrepancies and ambiguity due primarily to different sampling rates and time asynchronization between sources, especially as the ESS or PV rapidly changed values, but also due to dc/ac conversions and differences in

measurement accuracy of different sources.⁶⁸ The ideal solution would be to have a dc meter on the PV side as well as the BESS side, all with similar sampling rates to the POI meter.

- While performing Volt-var tests, visibility and control was lost to some inverters due to network fiber breaks (after the UPS ran out of battery). This led to times where the Volt-var function did not perform as expected because multiple inverters would stop participating. When designing the fiber ring, it would be prudent to ensure that the failure of one inverter not bring down the entire ring.

Ramp Rate Control Without Energy Storage

- Obvious as it may seem in retrospect, field testing of functions like ramp rate product and active power smoothing depends upon there being highly variable weather conditions, which, depending upon your climate may not frequently be the case. (This project was intentionally sited in a location with many clear days.) Budget ample time for testing to account for clear days that will provide little to no useful data.
- Conventional forecast performance metrics such as mean absolute error, mean bias error, etc. may not be the best indicators of performance for this control mode (nor the best form of feedback for model training) as they focus on overall accuracy rather than forecasts' ability to properly capture high ramp periods (in both time and magnitude).
- Since ramp rate control depends upon fast-acting and often successive commands based on previous actions, systems should be able to respond quickly to curtailment commands (ideally on the order of fractions of seconds) and be able to hold the curtailment command over a period of time.

Forecasting

- Long-term forecasts are useful for day-ahead resource planning or more near-term energy planning whereas short-term solar forecasts are most applicable to ramp rate functions. Use cases in this pilot used either one forecast type or the other.
- Merging of high-resolution short-term forecasts with long-term forecasts would have yielded logistical challenges as either timestamp resolutions would be heterogeneous and the file would have to be parsed for use anyway, or data would have to be stored at the highest resolution, greatly increasing bandwidth & storage requirements.
- Hybridizing short-term forecasts from multiple vendors post production (such as short-term forecasts from Reuniwatt and CPR) yields logistical challenges and no apparent expected benefits. If benefits from hybridization of short-term forecast sources like sky imaging cameras and satellite imagery are to exist, forecast vendors will need to integrate these sources at an algorithmic level, like Clean Power Research did with its long-term SolarAnywhere forecasts, which merge data from numerical weather predictions and satellite imagery.
- Plant availability (partial or full) should be considered and/or communicated while feeding back data to the forecast vendor for learning and tuning purposes.
- It is important to confirm plant models with PV forecast suppliers upon completion of plant installation, especially if PV generation forecasts have higher error than expected. Significant amounts of error can be attributed to outdated or incorrectly implemented plant model assumptions or parameters.

Other

- All intended use cases should be spelled out and implications of each considered during the

⁶⁸ Discrepancies could also come from differences in meter calibration/accuracy and line/inverter losses.

design stages of a project. Retrofits inherently incur additional challenges.

Potential Future Work

This project has uncovered new questions and potential research opportunities outside the current scope. The following is a list of topics potentially deserving of further consideration, investigation, and research.

- Clarity on what data needs to be fed back to forecast vendors for tuning / machine learning based improvements considering the influences of plant availability, followed by a reevaluation of PV forecast improvements once these new pathways and procedures are established
- Potential application of different kinds of probabilistic forecasts to multiple areas including 1) bulk system operations 2) individual plant ramp rate control 3) internal battery SoC control using statistical scenarios
- Dynamic reserve calculations currently consider only deterministic PV forecasts (the most probable power value for a given period). With the availability of probabilistic forecasts, further research could be done on how to implement this data into usable algorithms and quantifying what benefits it might have.
- The predictive coordinated charge/discharge (CCD-P) algorithm takes into consideration forecasted PV. Future research could investigate the value of algorithms that also include load forecasts or market price forecasts
- This project demonstrated that, when proper operation of all components aligned, regulating PV ramp rates at a single plant without energy storage is at least feasible. The next steps will be to determine the value of doing so and performing a cost benefit analysis with alternative technologies (e.g., What is the cost of curtailed energy from RRP relative to cost of energy storage system)
- Power required to have the plant ready and available to provide reactive power (even when not providing any reactive power) was in the hundreds of kW. A more thorough investigation of where these losses occur and what could be done to minimize them could help increase the value of reactive power services from inverters.
- Active power consumption may be different for different inverter functions or different conditions. Calibration of consumption under different functions and operating conditions will better define the true cost/value of these functions.
- During Volt-var testing, the plant experienced intermittent deviations from target voltage regulation value. Further testing at other sites could confirm whether this was site specific or a broader reliability challenge of plant-level functions.
- Likewise, further testing at other sites could help inform whether response delays in dc-coupled plants are inherent to the topology or just a consequence of the specific technologies used in this project. If the latter, further research could shed light on what design considerations could mitigate or eliminate factors that caused slowed response times (e.g., the lack of a designated PV meter, placement of the dc-dc converter in front of the storage rather than the PV, etc.).
- This project compiled a qualitative comparison of traditional transmission level voltage regulation and inverter-based resources. The next step will be to perform *quantitative* cost benefit analyses between the various technologies, likely informed or supported by inputs from additional field tests.

More real-world research projects like this one, preferably with longer periods of performance evaluation, are needed to demonstrate the continued advancements and assess maturity of inverter-based resources and their ability to provide accurate and reliable grid support and services. At this time, examples of relying on these technologies for maintaining grid power quality and stability are limited. Experiences with real time control via standard protocols, originating from utility control centers, will help build power system operators confidence in these relatively newer inverter-based resources.

Appendix A: Dynamic Reserve Requirements to Accommodate VRES Vagaries

Acronyms

ACE	Area Control Error
BA	Balancing Authority
CDF	Cumulative Distribution Function
CI	Confidence Interval
DA	Day-Ahead
DynADOR	Dynamic Assessment and Determination of Operating Reserves
EPRI	Electric Power Research Institute
ISO	Independent System Operator
LADWP	Los Angeles Department of Water and Power
PDF	Probability Density Function
PI	Prediction Interval
RT	Real-Time
RTO	Regional Transmission Organization
US	United States
VRES	Variable Renewable Energy Source(s)

Nomenclature

ε	Deviation from a forecasted quantity
ξ	Extent of a deviation
μ	Mean value
$\sim$	Median
Δ	Rate of change of a variable
ρ	Risk of undergoing a deviation
σ	Standard deviation
μ_3	Skewness
μ_4	Kurtosis
ρ_{long}	Risk of not having sufficient down-reserve

p_{short}	Risk of not having sufficient up-reserve
$F(\varepsilon)$	Cumulative density function of the deviations
$f(\cdot)$	function of its argument
$f(\varepsilon)$	Probability density function of the deviations
ℓ	System-wide load
n	System-wide net-demand
P	Probability of undergoing a deviation
rdn	Down-reserve requirement
rup	Up-reserve requirement
s	System-wide solar generation
w	System-wide wind generation

Introduction

As the penetrations of variable renewable energy sources (VRES) deepens in power systems, the residual demand (i.e., demand minus VRES) features greater degrees of variability and uncertainty. This variability and uncertainty are typically accommodated by re-dispatching or event re-scheduling conventional generation and energy storing resources. These actions might force system operators to either resort to real-time markets or resources which, while agile, are generally expensive, or to reduce the system reliability by system's operating margins.

An effective way to accommodate the expected variability and uncertainty is by means of allocating sufficient recourse in the form of operating reserves when drawing the operating plans of the system. Currently, Los Angeles Department of Water and Power (LADWP) uses a quasi-dynamic method to determine the operating reserve requirements. This method holds in reserve an amount proportional to the wind and solar generation forecasted up to a predefined maximum level.⁶⁹ While simple and practical, such method overlooks the actual needs for operating reserves based on past-observed deviations or quality of the forecasts. As such, LADWP is interested in exploring alternative methods to dimension the operating reserve requirements to ensure an efficient operation of the system while maintaining or potentially improving the levels of reliability in their system.

This subsection describes work done for LADWP to effectively dimension the operating reserve requirements using EPRI's tool DynADOR (Dynamic Assessment and Determination of Operating Reserves), [1], [2], [3], [4]. The results are compared against LADWP's existing method in terms of amounts of expected system risk, [5], [6]. Additionally, the benefits of enhanced forecasts are explored in terms of reserve reduction while maintaining a given system's risk.

The appendix is organized as follows:

- Refresher on dynamic reserves requirements

⁶⁹ LADWP currently uses 1/2 of the forecasted wind and solar generation up to 400 MW, but such numbers are subject to change.

- Data requirements as well as data preparation needed to compute reserve requirements
- Calculated operating reserves and comparison against the existing method at LADWP
- Discussion on the impact of improved forecasts on the reserve requirements
- Conclusion

Dynamic Reserve Requirements

In order to understand the methods behind the dynamic reserve determination tool (DynADOR), it is useful to define the central needs for reserve, which are shown in Fig. 1.




Fig. 1 Central needs for reserve

Fig. 1 shows the central needs from reserve derived from a typical scheduling process. In this process the resolution is of 5 minutes over a horizon on 20 minutes. The process of drawing the operating plans for a power system begins with a forecast of the amount of the net demand that would need to be served by dispatchable resources (blue staircase line). However, forecasts are always wrong (red staircase line). A comparison of the interval average amounts against the interval forecasted reveals the first need for reserve: “uncertainty” (area in red). From these profiles, we can also identify the need to provide sufficient ramping capabilities to be able to make transitions between operating points, and this need is referred here as “inter-interval variability” (area in cyan). Finally, the real-time quantities vary even within scheduling intervals, pointing to the need to carry reserves for “intra-interval variability” (area in light blue). The uncertainty and intra-interval variability needs are combined through a rigorous statistical process to determine the aggregated capacity needs for reserve. The inter-interval needs complement by allocating sufficient ramping capabilities between periods. Additionally, the system operator can adjust these needs using a tolerance in order to capture only a desired percentage of these deviations.




Fig. 2 Dynamic reserve determination method

The dynamic reserve determination method built into DynADOR is composed of four main steps, each denoted in Fig. 2 by a blue rectangle from left to right. These steps are described below:

1. **Reserve Characteristics:** The first step consists in defining the operating characteristics of the balancing area (BA). That is, it needs to define the process in which the reserve that is to be computed is held, as well as the process in which the reserve is intended to be released. This information is important in determining the reserve needs. Additionally, information on the reserve product is needed. This includes the direction of the reserve (i.e., upward, downward or both), what will the reserve be used to hedge against (e.g., variability, uncertainty, both, etc.), the desired confidence interval (or risk tolerance), and additional reserve margin (if any), among others.
2. **Historical Assessment:** The second step consists of assessing the historical reserve needs. That is, given historical values for forecasts and actuals over a historical time period (e.g., 1 year), the method computes the exact level of reserves that would have been needed to meet the deviations. These computed reserves would be such that cover the historical deviations of the actuals from the forecasts exactly. These exact reserve needs are the foundation for the rest of the reserve determination method.
3. **Categorize type and source:** The third step of the process consists in clustering the historical exact reserve needs according to the different needs (i.e., intra-interval variability, inter-interval variability, and uncertainty –blue icons in upper center). The correlation of these needs to the different explanatory variables from the sources (i.e., load, wind, solar) are matched through a correlation matrix (center of the box). Note that different explanatory variables for each source could be used (e.g., production level, hour of the day, rate of change forward, etc.)
4. **Assemble using best explanatory variables:** The final step of the process consists in identifying the dominant explanatory variables for each of the reserve needs, and to create a functional representation so, that given a new forecast (e.g., a new period in the future), the reserve requirements can be computed in a systematic manner. In DynADOR this is referred to as “EPRI clustering method.”

Data Requirements and Preparation

The data needed to operate DynADOR is divided into two groups, historical data (past) and anticipated data (future), Fig. 3.




Fig. 3 Data requirements for DynADOR

Historical Data

The historical data is composed by two main groups: i) forecasts and ii) actuals. Historical forecasts are the recorded past predictions made to operate the system, and should contain information for load, wind and solar. Forecasts typically have a lower resolution than actuals. For instance, a typical set of historical forecasts could contain data in the following format shown in Table 6-1.

Table 6-1 Sample of historical forecasts

timestamp	load	wind	solar
1/1/2019 00:00	2489.318	94.66	0
1/1/2019 01:00	2381.463	91.05	0
1/1/2019 02:00	2385.677	89.74	0
:	:	:	:
12/31/2019 23:00	2172.951	89.24	0

Error! Reference source not found. shows a sample of historical forecasts for year 2019 with a resolution of 1 h. Load, wind and solar are in MW, and note that solar is zero in all cases because of the recorded hour. The data for historical actuals is similar, however this information is the amount of load, wind and solar after the variability and uncertainty was revealed in real time. These quantities are measured at a higher rate than the forecast. For instance, Table 6-2 shows the corresponding actuals for this example with a recorded resolution of 5 minutes.

Table 6-2 Sample historical actuals

timestamp	load	wind	solar
1/1/2019 00:00	2521.544	80.88539	0
1/1/2019 00:05	2493.541	87.08394	0

1/1/2019 00:10	2486.621	94.46409	0
:	:	:	:
12/31/2019 23:55	2521.544	80.88539	0

Anticipated Data

The data for future conditions follows the same format and resolution as that shown in **Error! Reference source not found.**. However, this data corresponds variables that are yet to occur. In our example this could be for the 1st of January 2020.

Typical Problems with Historical Data

When generating and recording data there might be some issues that need to be carefully corrected before using the DynADOR tool. These issues include incorrect timestamp, missing data, corrupted measurements, metering noise, among others. The user of the DynADOR tool should carefully revise the data before feeding it to the tool, since DynADOR could interpret these errors as genuine needs for reserve and produce erroneous results. For this work with LADWP, EPRI carefully revised the data sent by LADWP and interpolated for missing data and corrupt measurements.

Descriptive Statistics of Historical Data

The raw historical data from LADWP is shown in Fig. 4. This information presents useful information such as load and VRES trends for different seasons. However, it is hard to extract useful information about the actual instantaneous penetration of VRES as a function percentage of the load, or net demand profiles that would need to be accommodated by conventional dispatchable generation (e.g., nuclear, coal, gas, etc.).




Fig. 4 Raw historical data for 2019 in LADWP

From Fig. 4 it can also be seen that LADWP had a peak load of 5,663 MW and a valley of 1,763 MW in 2019. In that same year LADWP also had an instantaneous hourly penetration of VRES of 61.1% on 08/28/2019

(wind 15.12% and solar 45.97%) at 0900 h with a load of 2316.78 MW, Fig. 5. The average VRES hourly penetration was of 36.1% with a standard deviation of 8.9%. The majority of the periods of high percentual VRES production are in line with the solar irradiance, due to the deeper solar penetration compared against installed wind.




Fig. 5 Hourly VRES penetration in LADWP for 2019

Given the penetrations of wind and solar, the system net demand is different from the system-wide demand. For the sake of illustration, four clusters of net demand are produced using k-means, [7], and shown in Fig. 6.




Fig. 6 Generic profiles of demand and system net demand (i.e., demand after wind and solar generation) in LADWP for 2019 clustered using k-means

Fig. 6 shows that, as expected, the production from VRES offsets the energy needs for the system. In particular for LADWP in 2019 these were reduced by 14.94%, with 3.43% due to wind power production and 11.51% due to solar power production. Besides the reduced energy needs, the system needs to accommodate a significantly different profiles of residual demand. The average resulting net demand features steeper ramps over shorter periods of time. For instance, for cluster “1” which contains 38.63% of the total population of profiles in 2019, the average net demand ramps are of 128.9 MW/h sustained over 9 h, which are steeper than the 64.3 MW/h that would have been needed for a period of 15 h for the average load profile alone.

In terms of peak values, the net demand including VRES production reduces, but the peak demand is still relatively close to the system load. The valley on the other hand, reduces approximately by half of the load valley. It is important to note that the hour of the peak and valley changes with the integration of VRES. Both, the net demand valley and peak occur latter in the day due to the solar power production.

The descriptive statistics of by variable of the forecasted and actual quantities in this system for 2019 are shown in Fig. 7.




Fig. 7 Descriptive statistics of the historical forecasts and historical actuals for load (ℓ), wind (w) and solar (s)

In Fig. 7 the lighter areas cover materialization of the indicated variable in the 1% tails of the historical distribution at the indicated time. These figures show the trends and percentiles of each variable as a function of time. The slightly darker area covers the values that are within 1% and 99% of the historical values, and the even darker areas cover the values within 25% and 75%. The solid line denotes the median value of the variable. From these figures, it can be seen that the load increases from the early morning until a few hours after midday, and then tend decrease again. Wind is more uniformly distributed across the day, and solar follows a well-defined pattern aligned to solar irradiance. These figures are useful to make sure that there are no outliers in the forecasts of metered data. It also allows inspecting trends for the variables as well as timestamp alignment issues. If any of these were present in these time series, the resulting reserve requirements would be biased by the errors and produce unrealistic reserve requirements.

A similar inspection can be performed to the first order derivatives of the variables, which would denote the individual ramping needs by variable, Fig. 8.




Fig. 8 Ramping needs by forecasted and actual variables

Fig. 8 shows that the ramping needs of the load systematically increase during the early hours in the morning. These needs then widen-up across the day, and then tend to narrow at night. The ramping needs for solar are more homogenous across the day, in the sense that they do not follow a specific pattern. Finally, the ramping for solar is well defined, with positive ramps in the morning until midday, and then negative ramps until about 2000 h.

The difference between actuals and forecasts for the net demand (ε_n) are given as:

$$\varepsilon_t = n_{A,t} - n_{F,t} \quad (0.1)$$

In (0.1) the cases in which $\varepsilon_t > 0$ occur when the net demand has been under-forecasted, and the system needs to deploy up-reserves to maintain the balance. On the other hand, the cases in which $\varepsilon_t < 0$ occur when the net demand is over-forecasted, and the system needs to deploy down-reserves to maintain the balance. Fig. 9 shows the differences between forecasts and actuals (i.e., deviations) for LADWP in 2019. The figure in the upper left part shows these deviations as a function of time, and the figure in the upper right shows the densities of these deviations for each hour of the day in 5 min resolutions. The lower left figure presents the probability density function of the deviations for the whole year, with a calibrated Gaussian and calibrated Laplace distributions for reference. In this figure it can be seen that the aggregated deviations are closer to a Laplace distribution with a skewness of 0.717 and a kurtosis of 6.570. The lower right figure shows the distribution for at $t = 1200$ h. In this case it can be seen that the distributions at a specific time period are not that close to a parametric distribution.




Fig. 9 Deviations from forecasted quantities for LADWP in 2019

Fig. 9 shows that the deviations have relatively similar distributions for the night periods (between 2000 h and 0500 h). However, the distributions of the deviations widen substantially during the day (between 0500 h and 1800 h). The skewness towards the up-reserve requirements is notable during such periods. This indicates that the solar power production is the dominant source of deviations in the system, while the distributions of deviations due to wind are relatively constant across the day.

The forward rate of change of the net demand can be computed as follows:

$$\overline{\Delta n}_t = n_{t+1} - n_t \quad (0.2)$$

From equation (0.2) it can be seen that if $\Delta n_t > 0$, the net demand at forthcoming period is greater than the that at the current period and the system needs up ramps to make the transitions. On the other hand, if $\Delta n_t < 0$ it means that the net demand at the forthcoming period is lower than the actual, and the system needs down ramps to be able to make the transition. These needs are shown graphically in Fig. 10.




Fig. 10 Rate of change of the net demand for the LADWP system in 2019

In Fig. 10 the upper left part shows the ramping needs for the whole 2019 as a function of time, and the upper right shows the densities of these ramping needs for each hour of the day. The lower left shows the distribution the ramping needs for the whole year, while the lower right shows the ramping needs at hour 1200. From the upper right figure, it can be seen that the ramping needs start to be positive in the early hours in the morning, but around 0500 h these ramping needs start to reduce until they become negative. This is because of the solar production starting to pick up in the morning. During the middle of the day the ramping needs are mostly positive and relatively modest with a few percentiles of large magnitude. At 1600 the ramping needs are mostly positive, but the trend is to a rapid decline for the rest of the day, until these needs are fully negative at 2000 h. This is when the solar production starts reducing and fully disappears from the system.

Operating Reserves Calculation

LADWP Reserve Policy

LADWP developed its own operating reserve policy to accommodate variability from VRES. The method consists in computing the anticipated aggregated value of VRES by summing forecast for wind and solar and dividing it by half. A ceiling of 400 MW is enforced in order to avoid scheduling unnecessarily large reserves. Mathematically this method can be expressed as:

$$r_t^{\text{up}} = \min \left\{ \frac{1}{2} (w_t + s_t), 400 \right\} \quad (0.3)$$

While simple and practical, this method overlooks several important aspects that should be considering when determining the operating reserve requirements. For instance, this method allocates the amounts of reserve based directly on the anticipated amount of VRES in the system, but disregards entirely the quality of the forecasts. That is, in the cases in which VRES could be predicted with greater accuracy, the need for reserves would reduce; however, this method still would require large amounts of reserve specially as the penetrations of these resources deepen. Additionally, this method focuses only on the temporal aspect of

VRES production, but disregards other important variables like rate of change of production and deviation direction among others. Finally, there is no way of tailoring these needs to a desired confidence interval or derived reliability level that the system operator may want to maintain in the system.

The reserve requirements for LADWP using equation (0.3) for 2019 are shown in Fig. 11.




Fig. 11 Reserve requirements using LADWP policy for 2019

Exact Reserves for LADWP

The exact reserves to respond to variability and uncertainty in LADWP can be computed from the historical deviations. These reserves are the exact amounts of capacity that should have been held when drawing the operating plans of the system to respond to the deviations exactly, without having to resort to any other corrective action beyond reserve deployment. These requirements are shown in Fig. 12.




Fig. 12 Exact up and down-reserve requirements for LADWP in 2019

A visual inspection of Fig. 11 and Fig. 12 sheds light on important points on the actual needs for reserve in LADWP:

- It seems that the LADWP method under-estimates the needs for up-reserves during night and early mornings. This is particularly true during the night (i.e., 2000 – 2359 h), in which the exact reserve requirement seem to be substantially different from those obtained with (0.3), even for potentially low percentiles of deviations.
- The LADWP method also seems to under-estimate the actual needs for reserve during hours of high solar production (i.e., 0700 – 1900 h). This is particularly notable for cases in which the system operator may need to cover for a large percentile of the deviations.
- If on the other hand, the system operator was to choose to hedge the system against lower percentiles of the deviations, the LADWP method may be excessive.
- The LADWP method does not shed any light on the needs to carry reserve for downward deviations, which clearly cannot be disregarded by inspecting the exact down-reserve requirements shown in Fig. 12.

Comparison of Reserve Policies

In order to compare different reserve policies, it is useful to have a common metric to assess their expected performance. As such, the next section describes “risk” as a metric to assess the expected hedging that the reserves provision would bring forth in power system operation.

Risk as a Reserve Performance Metric

Given the historical distributions across time of the uncertainty (as shown in Fig. 9), it is then possible to estimate the risk of having insufficient reserves for any given reserve policy. Fig. 3 shows a hypothetical distribution of the historical deviations at an arbitrary time “ t ”, and arbitrary amounts of up (R_{up}) and down (R_{dn}) reserve requirements.




Fig. 13 Reserve requirements and risk of undergoing deviations

Given these amounts of reserve shown in Fig. 13, then three areas are identified in the deviations’ distributions:

- a. Green area: The area between the down and up-reserve requirements. Within this area, the system has sufficient resources to accommodate any deviation without having to resort to any other corrective action beyond the deployment of the reserves.

- b. Red area: This area denotes the cases in which the system has insufficient up-reserve (i.e., the system is “short”) and could undergo under-frequency, negative area control error (ACE), reliability violations, or even load shedding.
- c. Blue area: This area denotes the cases in which the system has insufficient down-reserve (i.e., the system is “long”) and could undergo over-frequency, positive ACE, or excessive VRE curtailment.

The risk of being short or long can be computed using the statistical information of the deviations and the reserve requirements. Risk (ρ) is given as the expectation of loss (insufficient reserve), [8], [9], and can be written as:

$$\rho = P \times \xi \quad (0.4)$$

where P is the probability of undergoing a deviation greater than the reserve requirement (i.e., expectation), and ξ is the extent of the deviation (i.e., loss). Therefore, the risk of being short and risk of being long are given as:

$$\rho_{\text{short}} = \left(\int_{R_{\text{up}}}^{\max(\varepsilon)} f(\varepsilon) d\varepsilon \right) \times (\max(\varepsilon) - R_{\text{up}}) \quad (0.5)$$

$$\rho_{\text{long}} = \left(\int_{\min(\varepsilon)}^{R_{\text{dn}}} f(\varepsilon) d\varepsilon \right) \times (-R_{\text{dn}} - \min(\varepsilon)) \quad (0.6)$$

where $f(\varepsilon)$ is the PDF of the historical deviations, R_{up} is the up-reserve requirement, R_{dn} is the down-reserve requirement, and $\max(\varepsilon)$ and $\min(\varepsilon)$ are the ends of the positive and negative tails of the probability distribution of the deviations. Then the first terms in equations (3) and (4) represent the probability of being short and long respectively, and the second term the extent of being short and long respectively.

To learn more about the relationship between reserve and risk, as well as risk-based methods to procure reserve, the interested reader is advised to see [5].

Reserve Procurement Comparison

In order to compare the reserve procurement policies a test day is chosen from the data set (July 1st, 2019) and excluded from the set fed for tool calibration (historical datasets). The forecasts for this day are shown in Fig. 14.




Fig. 14 Forecasts for the 1st of July 2019, used for reserve procurement comparison

Using the data shown in Fig. 14 with LADWP reserve policy and EPRI's dynamic reserve determination tool, the up and down-reserve requirements shown in Fig. 15 are obtained. For the sake of completeness, it was assumed that the amount of reserve obtained with la LADWP method is used in both directions (i.e., upward and downward).





Fig. 15 Upward and downward reserve procurement obtained with the LADWP policy and EPRI's dynamic reserve determination tool

From Fig. 15 it can be seen that the LADWP method estimates fewer reserve requirements for up-reserve during the night and during the early morning periods when compared against the requirements obtained using EPRI's dynamic reserve determination tool. It can also be seen, that except for the case in which a confidence interval of 99% is enforced, the LADWP method estimates that larger amounts of reserve are needed during the solar production hours. With respect to the down-reserve, the dynamic reserve requirements are lower compared against the up-reserve requirement for the EPRI's clustering method, and assumed exactly the same for the LADWP method. In this case, EPRI's method determine two “peak” values when the solar is either ramping up or down. These are the periods of greater uncertainty since the rate of change of VRES power is the largest (see Fig. 8).

While Fig. 15 allows performing a quantitative comparison on the amounts of reserve to procure at each time, it does not shed light on how well these reserves would perform in operations. In order to assess the expected performance of the reserves, it is useful to compute the risk associated with the reserve profiles, as shown in Fig. 16.






Fig. 16 Up and down-reserve profiles and associated risk for the LADWP system on the 1st of July 2019

Fig. 16 shows the up and down-reserves as a function of time, and the expected short and long risk in operations for: a) the LADWP reserve policy and b), c), d) for the dynamic reserves obtained with EPRI's DynADOR tool for different confidence intervals (CIs). In this figure it can be seen that the LADWP method results in a relatively low risk during the hours of high solar production. However, it does not guarantee such a positive outcome during the early morning and night periods. This is because the LADWP method only focuses on the VRES production itself, and not on its inherent uncertainty and variability. Therefore, this method disregards the frequent and potentially large deviations of wind power generation. EPRI's dynamic reserves result in a low risk at all time periods when a large CI is selected. That is because the tool is being requested to hedge against a large percentile of the possible deviations. As the CI is relaxed, the amount of reserves procured by this method reduce. However, such reduction in reserve procurement has an associated risk increment. An adequate selection of the CI interval is fundamental in order to guarantee a secure system operation. The selection of the CI could be chosen such that the resulting system operation would comply with any desired reliability standard.

Effect of Improved Forecasts

This subsection assesses the impact of improved solar forecasting on flexible reserve requirements. That is, given an improved system-wide PV forecast, by how much could the resulting reserve requirements be reduced, since the expected system uncertainty would effectively be reduced.

Method to Generate Improved Forecasts

The method used to generate improved forecasts simulates a desired percentage improvement in uncertainty applied uniformly to all forecasted values. First, the error between historical forecasts and realized measurements is computed. This error is then reduced by the desired percentage and added to the existing historical actual. A visual depiction of the process applied to a sample day is shown in Fig. 17[[[reference error]]]. In it, the resulting output (purple) lies a fixed fraction of the way between the original forecast (green) and the measured output (blue).




Fig. 17 Summary of the method to produce improved forecasts

For this analysis, an improvement of 50% was chosen. This relatively large forecast improvement was selected in order extoll the potential advantages of improved forecasts in general; it does not suggest that such improvements are easily attainable, nor does it provide suggestions as to how they would be attained. Likewise, the uniform improvement across all forecast values was done for simplicity; real, incremental forecast improvements would likely benefit select conditions more than others. Finally, please note that, at present, system-wide PV generation forecasts are generally not derived by aggregating production forecasts from individual PV plants like those analyzed elsewhere in this section. Unless this fact changes, improvements in individual plant forecasts may not influence bulk system forecasts.

Descriptive Statistics of Error

Fig. 18 shows the magnitude and direction of realized deviations from the DER forecast for both the original DER forecast and the DER forecast that contained the improved PV forecast. It can be seen that if the solar forecast is improved, the aggregated system error magnitude reduces significantly during hours of solar production.




Fig. 18 System deviations for a) the original forecast case and b) the improved solar forecast case; full year left, sample day right

To more clearly illustrate the reduction in error, the isolated deviations for just solar are shown in

Fig. 19.




Fig. 19 Isolated solar deviations for a) original forecast and b) improved forecasts; full year left, sample day right

In

Fig. 19 the effect of the improved forecasts on the magnitude of the deviations is quite evident. High-quality forecasts have the ability to reduce the magnitude of these deviations, and therefore to reduce the overall needs for reserve. The theoretical limit would be to have forecasts that are 100% accurate, and in such cases, there would be no need to allocate reserves for uncertainty.

Impact on the Reserve Requirements

The most valuable feature of the improved forecasts is that it allows drawing more precise operating plans for the system, and in particular, reduces the need for additional recourse to deviations (i.e., reserves). The figure below shows the impact on the LADWP policy for year 2019.




Fig. 20 LADWP reserve determination policy for a) the original solar forecast and b) the improved solar forecast

time, h

time, h

Fig. 20 shows that the amounts of reserve that the LADWP estimate with or without the improved solar forecasts are virtually the same. This is because the LADWP reserve policy focuses on the magnitude of VRES production, as opposed to the quality of the forecast of the VRES production. That means that even if it were possible to predict the VRES power with complete accuracy, the LADWP policy would still produce the same reserve requirements.

EPRI's dynamic reserve determination methods compute the reserve requirements based on the difference between historical forecast and historical actuals. That means that if the quality of the forecasts is improved, the magnitude of the deviations reduces, and this is translated into lower reserve requirements. The exact reserve requirements for the LADWP system using both, the original forecasts and the improved forecasts are shown in Fig. 21.




Fig. 21 Exact up and down-reserve requirements for the LADWP system for both, the original forecast and the improved forecast

Fig. 21 shows that as the quality of the forecasts improves, the exact needs for reserves in the system reduces. This is a desirable feature since it will translate into a more efficient utilization of the recourse resources while maintaining the same level of system reliability. The actual reserve requirements for the LADWP system for the test day (see figure Fig. 14) are shown in Fig. 22.

In Fig. 22, it is evident that the reserve requirements reduce as the quality of the forecast improves, even maintaining the same confidence interval (CI). For instance, for a CI = 99% (which means that 99% of the expected deviations would be captured by the reserve procured) the amounts of up-reserve for the original forecast are above those obtained with the LADWP policy. However, considering an improved forecast, the same CI of 99% results in a reserve profile that is below the LADWP policy during the high solar power production period. It is important to note here that the amounts of up-reserve for the night and early morning are the same using EPRI's dynamic reserve determination method. This is because these reserves are driven by wind power deviations, which remain the same for both cases. As mentioned earlier, the LADWP policy severely underestimates the needs for reserve during these periods. The same behavior can also be observed in the down-reserve requirements.







Fig. 22 Up and down-reserve requirements for the test day in LADWP using both, the original forecast and the improved forecast for solar generation

As in the previous section, in order to assess the expected performance of the reserve, it becomes necessary to quantify the expected number of potential violations (i.e., materializations above the reserve requirements) that this reserve profiles would not hedge against. Risk is a good metric since it takes into account the likelihood of the potential violation, as well as the potential extent/magnitude. Fig. 23 shows the reserve profiles and their associated risk for the LADWP policy (left pair of figures, a & c) and the dynamic reserves obtained with EPRI's clustering method for a CI of 90% (right pair of figures, b & d). The top row of Fig. 23 shows the reserves procured for the original forecast, and the bottom row shows the same but for the enhanced forecast.





c) Reserve and risk for the LADWP reserve policy, improved forecast



d) Dynamic reserves and risk for CI = 90%, improved forecast

Fig. 23 Up and down-reserve requirements and associated short and long risk for both, the original solar forecast and the improved solar forecast

In this prior figure, reserve profiles under the LADWP policy are virtually the same in both cases, with the only exception being a reduction in the already low, mid-day risk of being short. This risk is slightly lower in the case of the improved forecast, but the reduction is incidental since the amount of reserve is the same in both cases. The reason for this reduction is because the quality of the forecast is better, meaning that the deviations are lower in magnitude. The risk for this policy in the periods in which there is no solar production remains large and the same for both cases.

When using the dynamic reserves from EPRI's reserve determination tool, the profiles are similar in shape for both cases, but the overall amount is lower for the case of the improved forecast. Additionally, the risk profile is also lower for the improved forecast. This would translate in actual system operation as economic savings (lower reserves), while maintaining or potentially enhancing the system reliability.

The above reductions can be difficult to see in Fig. 23, but can be better appreciated in Fig. 24, which shows the average amounts of reserves over the whole period (top) as well as the average risks over the whole period (bottom).




Fig. 24 Average reserves (top) and average risk (bottom) assessment for the LADWP reserve policy and a collection of outputs from EPRI's reserve determination tool with different confidence intervals. Each reserve policy has one set of results for the original forecast, 'o', and one set for the improved/enhanced forecast, 'e'.

Reiterating, Fig. 24 shows the average amount of up and down-reserves as well as the average short and long risks for the whole period (the 24 h of the chosen day). Again, in this figure it can be seen that the reserves provided using the LADWP policy are the same in both cases, and the short risk is slightly lower. As for the dynamic reserve requirements obtained using EPRI's reserve determination tool, it can be seen that the average reserves for any CI are lower when using the improved forecasts. Additionally, the average risks are also lower considering the improved forecasts. This, again points to the conclusion that if used in actual system operation, the results would be economic savings while attaining the same or even greater reliability levels.

Conclusions

This study assessed the performance of the existing LADWP method to determine operating reserves to respond to VRES (wind and solar power production) vagaries and compares it against a method developed by EPRI to dimension such requirements dynamically taking into account the historic deviations from forecasts (variability and uncertainty), as well as the anticipated forecasts for future conditions. The LADWP method dimensions the needs for reserve with respect to the anticipated VRES production, however it does not take into account the quality of the forecasts associated with the random variables. The dynamic reserves method used by EPRI explicitly accounts for the magnitude of the deviations, direction, type (uncertainty or variability) as well as source (load, wind or solar generation), when computing the needs for reserve.

The LADWP system had a total penetration of 14.9% of energy from VRES (3.4% wind and 11.5% solar). Having solar dominating the VRES production is not surprising that the LADWP method results in larger amounts of reserve during hours of high solar irradiance. However, this disregards the high variability and uncertainty induced by wind power generation during the night and the early hours of the morning. As such, if the system only had the reserves dictated by the LADWP method, the night periods could be risky in terms of potential reserve violations. On the other hand, EPRI's dynamic reserve determination method strikes a better balance between reserve production and risk management over the whole study horizons. The risk level could be even adjusted indirectly by changing the desired confidence interval. Additionally, from 2021 the new version of the reserve determination tool (DynADOR) will feature a risk-based engine that can be effectively used to determine reserves to meet a specific risk ceiling chosen by the operator.

The study also presents an assessment of the impact on the reserve requirements if improved forecasts were to become available. It is shown that the improved forecasts reduce the magnitude of the expected deviations and therefore the needs for reserve. When assessing the LADWP method it was clear that this method remains agnostic to the quality of the forecasts, and only concentrates on the anticipated amount of VRES. The dynamic reserve determination method results in lower average reserve requirements and improvements in the risk profile of the system. Mirroring this to power system operation it would be expected to attain lower operating costs while improving system reliability.

References

- [1] "An Enhanced Dynamic Reserve Method for Balancing Areas - A Dynamic Reserve Method for Balancing Areas with High Levels of Variability and Uncertainty on Their System and Case Studies Showing the Performance of These Methods," Electric Power Research Institute (EPRI), Palo Alto, CA, 3002010941, Dec. 2017, Available: <https://www.epri.com/#/pages/product/000000003002010941/?lang=en-US>.
- [2] "Probabilistic Dynamic Reserve Determination and Risk Assessment," Electric Power Research Institute (EPRI), Palo Alto, CA. 3002016312, Dec. 2019, Available: <https://www.epri.com/#/pages/product/000000003002016312/?lang=en-US>.
- [3] "Dynamic Assessment and Determination of Operating Reserve (DynADOR) v3.0 " Electric Power Research Institute (EPRI), Palo Alto, CA Dec. 2019, Available: <https://www.epri.com/#/pages/product/000000003002016311/?lang=en-US>.

- [4] M. A. Ortega-Vazquez, "Dynamic Assessment and Determination of Operating Reserves (DynADOR) - Video," Electric Power Research Institute (EPRI), , Palo Alto, CA, 2020, Available: <https://epri.box.com/s/xk7niq4vgptpp3kmaq6dvdxcrg3h2owfn>.
- [5] M. A. Ortega-Vazquez, N. Costilla-Enriquez, E. Ela, and A. Tuohy, "Risk-Based Reserve Procurement," in *2020 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, 2020, pp. 1-7.
- [6] M. A. Ortega-Vazquez, "Risk-Based Reserve Procurement " Electric Power Research Institute (EPRI), , Palo Alto, CA, 2020, Available: <https://epri.box.com/s/9c2z6s2qjyf56pxefvpvw7fc06bimn9p>.
- [7] D. Arthur and S. Vassilvitskii, "k-means++: the advantages of careful seeding," presented at the 18th annual ACM-SIAM Symposium on Discrete Algorithms, New Orleans, LA, 2007.
- [8] A. Wald, "Statistical Decision Functions," (in en), *Ann. Math. Statist.*, vol. 20, no. 2, pp. 165-205, 1949/06 1949.
- [9] M. H. DeGroot, *Optimal Statistical Decisions*. John Wiley & Sons, 2004.

The Electric Power Research Institute, Inc. (EPRI, www.epri.com) conducts research and development relating to the generation, delivery and use of electricity for the benefit of the public. An independent, nonprofit organization, EPRI brings together its scientists and engineers as well as experts from academia and industry to help address challenges in electricity, including reliability, efficiency, affordability, health, safety and the environment. EPRI also provides technology, policy and economic analyses to drive long-range research and development planning, and supports research in emerging technologies. EPRI members represent 90% of the electricity generated and delivered in the United States with international participation extending to nearly 40 countries. EPRI's principal offices and laboratories are located in Palo Alto, Calif.; Charlotte, N.C.; Knoxville, Tenn.; Dallas, Texas; Lenox, Mass.;

Program:
DER Integration

© 2021 Electric Power Research Institute (EPRI), Inc. All rights reserved.
Electric Power Research Institute and EPRI are registered marks of the
Electric Power Research Institute, Inc. in the U.S. and worldwide.

3002023056

EPRI

3420 Hillview Avenue, Palo Alto, California 94304-1338 • PO Box 10412, Palo Alto, California 94303-0813 USA
800.313.3774 • 650.855.2121 • askepri@epri.com • www.epri.com