

SANDIA REPORT

SAND2015-0156

Unlimited Release

Printed January 2015

Comparison of Errors in Solar Power Plant Variability Simulation Methods

Matthew Lave

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

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Sandia National Laboratories

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

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

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

Telephone: (865) 576-2087
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

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

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: <http://www.ntis.gov/help/ordermethods.aspx#online>



Comparison of Errors in Solar Power Plant Variability Simulation Methods

Matthew Lave
Photovoltaic and Distributed Systems Integration
Sandia National Laboratories
P.O. Box 969, MS-9052
Livermore, CA 94551-0969

Abstract

Four PV power plant variability simulation methods – no-smoothing, time average, Marcos, and the wavelet variability model (WVM) – were compared to measured data from a 19MW PV power plant to test the relative accuracy of each method. Errors (simulated vs. measured) were quantified using five application-specific metrics: the largest down ramps, the largest up ramps, the mean absolute error in matching the cumulative distribution of large ramps, the total energy contained in down ramps over the entire period considered, total energy in down ramps on the worst day. These errors were evaluated over timescales ranging from 1-second to 1-hour and over plant sizes of 1 to 14MW and the total plant size of 19MWs to determine trends in model errors as a function of timescale and plant size.

Overall, the WVM was found to most often have the smallest errors. The Marcos method also often had small errors, including having the smallest errors of all methods at small PV plant sizes (1 to 7MWs). The no-smoothing method had large errors and should not be used. The time average method was an improvement over the no-smoothing method, but generally has larger errors than the WVM and Marcos methods.

CONTENTS

1.	Introduction	6
2.	PV Power plant Simulation Methods	7
3.	PV Power plant data	9
4.	Quantification of Errors	11
4.1.	Largest up and down ramps	11
4.2.	Mean absolute error in matching the cumulative distribution of ramps with magnitude larger than 10% of capacity	12
4.3.	Energy in down ramp rates greater than 10% of capacity per minute	12
5.	Results	15
5.1.	Model Performance by Timescale	15
5.1.1.	Qualitative Analysis as a Function of Timescale	15
5.1.2.	Quantitative Analysis as a Function of Timescale	17
5.2.	Model Performance by Plant Size	21
5.2.1.	Qualitative Analysis as a Function of Plant Size	21
5.2.2.	Quantitative Analysis as a Function of Plant Size	23
6.	Conclusions	29
	References	31

FIGURES

Figure 1: Data availability. Black vertical lines indicate that <i>both</i> irradiance and plant power output measurements were available.	9
Figure 2: Measured irradiance (left axis) and measured power (right axis) on January 16, 2012.	10
Figure 3: Power output from individual 500kW inverters (different colored lines) over the whole day (left) and zoomed in on the afternoon clear period (right).	10
Figure 4: Cumulative distributions of ramp rates at various timescales for measured and simulated 19MW PV power plant output. The x-axis is the absolute value of the ramp rate, such that both positive and negative ramps are included. The y-axis is zoomed in to show ramps that occur 10% of the time or less to help illustrate the differences between the simulation models.	16
Figure 5: Errors (modeled minus measured) in the cumulative distributions of ramp rates for measured and simulated 19MW PV power plant output.	17
Figure 6: Measured and simulated largest down ramps (top left) and up ramps (top right) at various timescales. The bottom plots show the errors: modeled minus measured.	18
Figure 7: Mean absolute error in matching the measured cumulative distribution for ramp rates larger than 10% of capacity.	19
Figure 8: Measured and simulated energy in down ramp rates greater than 10% per minute over the whole time period of 349 days (top left) and on the single day with the most energy in the down ramps (top right) at timescales of 1 to 60 seconds. The bottom plots show the errors: modeled minus measured.	20
Figure 9: Plot showing the method with the smallest error for each metric and timescale.	21

Figure 10: Cumulative distributions of ramp rates over various plant sizes (MWs) for measured and simulated 19MW PV power plant output. The x-axis is the absolute value of the ramp rate, such that both positive and negative ramps are included. The y-axis is zoomed in to show ramps that occur 10% of the time or less to help illustrate the differences between the simulation models.	22
Figure 11: Errors (modeled minus measured) in the cumulative distributions of ramp rates for measured and simulated PV power plant output at 60s over various PV plant sizes (MWs).	23
Figure 12: Measured and simulated largest down ramps (top left) and up ramps (top right) at various plant sizes (MWs). The bottom plots show the errors: modeled minus measured.	24
Figure 13: Mean absolute error in matching the measured cumulative distribution for ramp rates larger than 10% of capacity for various plant sizes (MWs).	25
Figure 14: Measured and simulated energy in down ramp rates greater than 10% per minute over the whole time period of 349 days (top left) and on the single day with the most energy in the down ramps (top right) over various plant sizes (MWs). The bottom plots show the errors: modeled minus measured.	26
Figure 15: Plot showing the method with the smallest error for each metric and plant size (MWs). All error metrics were computed based on 1-minute ramps.	27

1. INTRODUCTION

As opposed to conventional generation sources such as coal or nuclear power plants, the output from solar PV power plants is variable due the movement of the sun through the sky and clouds obscuring the sun. This variability introduces new challenges to operation of the electric grid.

At short timescales, (seconds to minutes), quick changes in solar power output can cause frequency or voltage to exceed allowable limits, which may cause additional tap change or switching operations on transformers or capacitors that mean additional required maintenance or early replacement of these devices. At longer timescales (minutes to hours), solar variability can increase the amount of regulating and ramping reserves required to balance the transmission system. Uncertainty about the magnitude of long-term (hours to days) solar fluctuations can increase electricity production cost due to sub-optimal generation unit commitment and dispatch.

The variability of large solar PV power plants is particularly of concern since such fluctuations at such plants can have a significant impact to the electric grid. This is especially important in small balancing areas, such as islands. For example, concern over large PV plant variability in Puerto Rico led the Puerto Rico Electric Power Authority (PREPA) [1] to institute strict regulations requiring PV plant operators to limit changes in output to a rate of 10% per minute and to provide frequency support. Large amounts of storage are required to comply with these requirements, increasing project costs.

In order to understand the impact of solar variability to the electric grid, and to correctly size storage for ramp rate control, it is necessary to have an accurate understanding of PV plant variability. However, in most cases grid integration studies are run to determine the effects of many different hypothetical PV scenarios, or PV plant developers need to size their plant storage before the plant is operational. Thus, PV power plant simulation methods must be used to model the variability of yet-to-be-built PV power plants.

Since often irradiance measurements from a point sensor are available at the location of the PV power plant to be simulated, in this work we consider PV power plant simulation methods that start with a single irradiance point sensor and scale the output to simulate a large PV power plant. These models have not previously been compared to one another using the same validation data set and error metrics, making it difficult for grid integration studies and PV developers to understand the relative errors in each method.

In a previous work [2], we compared the methods, data requirements, and ease of implementation of a variety of PV plant simulation models. Here, we quantitatively compare the simulated PV power variability to the measured PV power variability at a large PV power plant. We compare model performance both as a function of timescale (1-second to 1-hour) and as a function of PV plant size (1MW to 14MW) to determine the best models for each application.

2. PV POWER PLANT SIMULATION METHODS

Four different variability simulation methods were considered. Each method starts with the measurements from an irradiance point sensor timeseries and then applies smoothing to the timeseries to simulate the spatially-averaged irradiance over the footprint of a PV plant. All four methods are described in detail in [2], but, briefly, are:

- 1) No smoothing: the irradiance point sensor timeseries was used directly.
- 2) Time averaging: the point sensor timeseries was smoothed by taking a moving average with time window $\Delta t = \frac{\sqrt{A}}{CS}$, where A is the PV plant area and CS is the cloud speed.
- 3) Marcos: The method described in Marcos, et al. [3] of a low-pass filter where the cutoff frequency, $f_c = \frac{0.02}{\sqrt{A}}$.
- 4) WVM: The wavelet variability model (WVM) [4], which applies different smoothing at different timescales based on the distance across the plant and the cloud speed.

Additional PV plant simulation methods based on satellite-derived irradiance were described in [2], but we were unable to obtain the required data to include these models in our analysis.

Since each method's output is the plant-averaged irradiance, an irradiance to power model must be used to allow for direct comparison to plant AC power measurements. For all simulation methods, we used the Sandia Array Performance Model (SAPM) [5] to convert the plant-average irradiance into plant-total DC power. Since module temperature measurements were also available at the PV plant, both plant-average irradiance and module temperature were input to the SAPM (module temperature was assumed constant across the whole plant). The plant DC power was converted to plant AC power using a linear DC to AC derate (e.g., as used in [6]). While a more advanced DC to AC model such as the Sandia Inverter Model [7] could have been used, the inverter manufacturer was not known, and acceptable performance was achieved with the linear model. If the simulated AC power exceeded the rated capacity, then it was set equal to the rated capacity to simulate inverter clipping.

3. PV POWER PLANT DATA

For this analysis, we used data measured at a 19MW PV power plant in southern Colorado. The plant uses monocrystalline silicon PV modules mounted on single-axis trackers. For this study, total AC power output of the plant and measured irradiance from a point sensor mounted on one of the trackers (i.e., POA irradiance) at 1-second intervals were used. The PV plant area A was computed from the plant footprint using satellite imagery to be approximately 0.79 km^2 .

Data was available from September 2011 until October 2012, though intermittent days were missing data, mainly due to communication issues with the irradiance sensor, which reported through a wireless mesh network that was at times unreliable. Figure 1 shows the data availability. Even with the intermittent outages, data was available for 349 days with no strong seasonal bias, so should accurately represent annual trends.

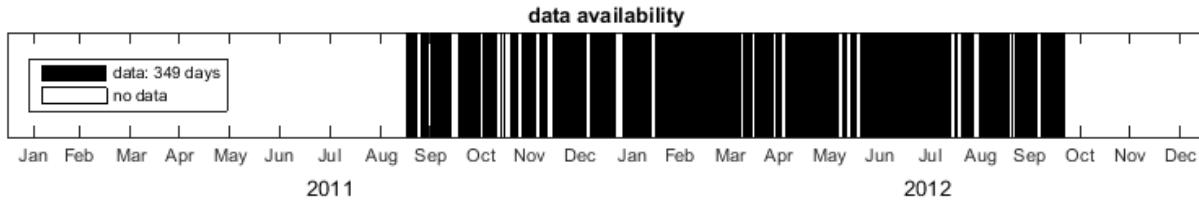


Figure 1: Data availability. Black vertical lines indicate that *both* irradiance and plant power output measurements were available.

All of the PV plant simulation methods considered here use the measured irradiance to simulate the plant power output. This assumes that the measured irradiance is representative of the general trends in irradiance received by the PV modules. Figure 2 plots the POA irradiance and power output on one day to test this assumption. During the middle of the day (e.g., 0900 to 1500), clear periods and cloudy periods are indeed generally well-correlated between the irradiance and the power output. However, during clear periods in mornings and evenings, the measured power was at times significantly less than the measured irradiance in a way that leads to large positive or negative ramps in the power that did not exist in the irradiance.

While we originally suspected that this morning and evening “step” behavior might have been caused by different turn on/off times of the inverters, Figure 3 plots the 500kW inverters separately and shows that this is not the case. Rather, the separate inverters have similar timeseries during the clear periods, including a notable ramp down around 15:45. We are unsure of the cause of this behavior.

Similar morning and evening ramps in power were observed on most other days with clear morning or evening periods. Since these ramps are not related to cloud-caused variability but rather seem to be a feature of this specific PV plant, we do not feel that they represent predictable variability that the models could reasonably simulate, so we chose to eliminate times before 0900 and after 1500 on all days to eliminate this irradiance to power mismatch.

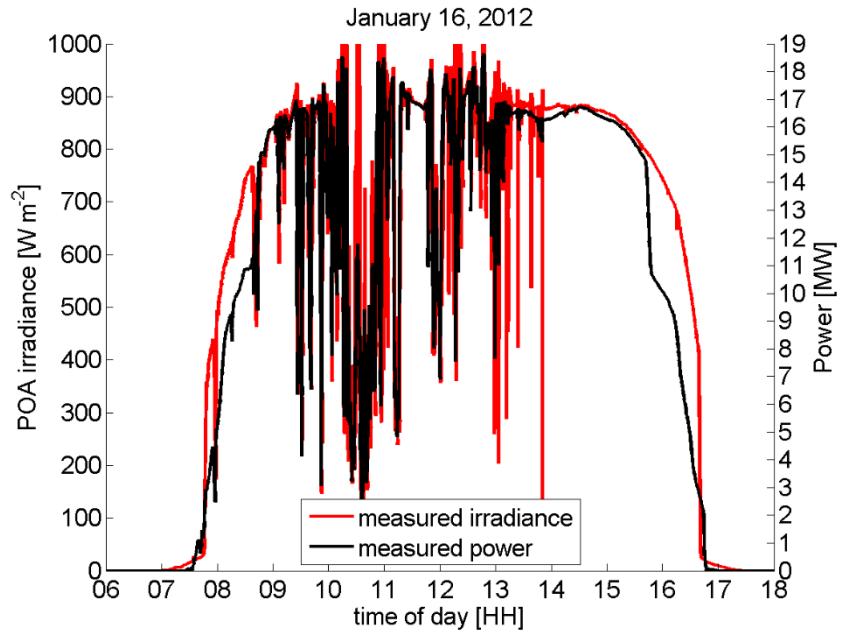


Figure 2: Measured irradiance (left axis) and measured power (right axis) on January 16, 2012.

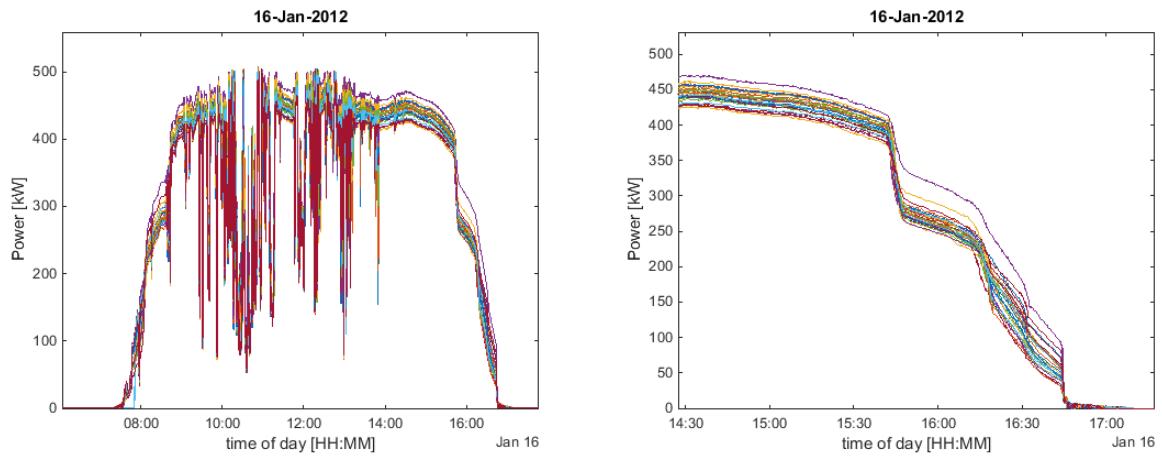


Figure 3: Power output from individual 500kW inverters (different colored lines) over the whole day (left) and zoomed in on the afternoon clear period (right).

4. QUANTIFICATION OF ERRORS

Since all of the methods start with only a single irradiance point sensor, the simulation methods will not be able to exactly match the power output timeseries as the onset timing of clouds on the point sensor may not exactly match the onset timing of cloud fronts covering a significant amount of the PV modules leading to reduced plant power output. However, in this study we evaluate the PV plant simulation models' accuracy at matching the true PV plant power variability. Thus, an accurate description of the changes in power output, rather than an exact timeseries match, is desired.

The simulation models are likely to be used to understand the typical range of fluctuations to advise grid operations (e.g., reserve requirements), and to determine storage capacity and control needs to comply with ramp rate restrictions (e.g., as imposed in Puerto Rico). Thus, it is important for the simulations to accurately match the distribution of power ramps and ramp rates, and so our evaluation focuses on comparing the simulated distributions of PV plant power ramps to the measured distributions of PV plant power ramps at various timescales. We use the metrics described in the sections below to quantify the model performance.

4.1. Largest up and down ramps

The largest up and down ramps relate to the amount of reserves that must be carried on the grid to ensure that electric service can be maintained when the PV ramps. If the magnitudes of the largest ramps are over-predicted by the simulation models, excessive reserves may be procured by the utility, increasing system operating cost. Conversely, if the magnitudes are under-predicted, the lack of sufficient reserves may lead to a system imbalance.

The largest down ramp is also important to storage applications, as the ramp rate relates to the rate at which the storage must inject power. Overestimating the largest ramp would lead to inefficient control strategies and underestimating could lead to unintended ramp rate violations.

Ramp rates for each timescale were computed as:

$$RR(t, \bar{t}) = \frac{1}{\bar{t}} \sum_{t'=t}^{t+\bar{t}-1} P(t') - P(t' - \bar{t}),$$

where t increases in increments of \bar{t} : $t = 1 + n\bar{t}$, where n is a positive integer.

The largest down ramp is computed as:

$$LDR(\bar{t}) = \min(RR(\bar{t})),$$

where \bar{t} is the timescale and $RR(\bar{t})$ is the ramp rate timeseries, and negative values in $RR(\bar{t})$ indicate down ramps. Similarly, the largest up ramp is:

$$LUR(\bar{t}) = \max(RR(\bar{t})).$$

Both LDR and LUR retain the RR units of power per timescale (e.g., watts per second). Note that in Figures 6 and 12, the y-axis units are the generic units "% per timescale" to allow for easy

comparison between timescales. If normalized to the same units (e.g., % of capacity per second), the short-timescale ramp rates would be much larger than the long-timescale ramp rates.

4.2. Mean absolute error in matching the cumulative distribution of ramps with magnitude larger than 10% of capacity

In addition to matching the maximum up and down ramp rates, it is also important to match the cumulative distribution of all large ramps to give an accurate description of not just the worst case but also the more common, large ramp rates. While this comparison can be done separately at distinct intervals by e.g., testing the simulated number of down ramps larger than 10% of capacity, then the simulated number of down ramps larger than 20% of capacity, etc, we instead compute the difference between the cumulative distribution of measured versus simulated ramp rates for ramps larger than 10% of capacity.

We intend this metric to apply to balancing reserve decisions which are concerned with large changes in power output. Thus, we focus on large magnitude ramps at all timescales and evaluate the mean absolute error in matching the cumulative distribution of ramps larger than 10% of capacity. This is distinct from the metric described in the Section 4.3, which focused on the ramp *rates*. We define ramps as changes in power output:

$$\text{ramp}(t, \bar{t}) = \bar{t} \text{RR}(t, \bar{t})$$

such that, for example, a 10% per min ramp rate and a 10% per second ramp rate are both equivalently a 10% ramp.

The error, e , in matching the distribution of ramps at each ramp magnitude, M , is given by:

$$e(M, \bar{t}) = P(|\text{ramp}_{\text{sim}}(t, \bar{t})| > M) - P(|\text{ramp}_{\text{meas}}(t, \bar{t})| > M)$$

For example, if the simulated probability of a 1-minute ramp larger than 15% of capacity was 1.5% and the measured probability of a 1-minute ramp larger 15% of capacity was 2%, then:

$$e(15\% \text{ of capacity}, 1\text{min}) = 2\% - 1.5\% = 0.5\%.$$

The mean absolute error in matching the cumulative distribution of ramps larger than 10% of capacity is then:

$$\text{MAE}_{\text{cdf}}^{\text{ramp} > 10\%}(\bar{t}) = \sum_{M = 10\% \text{ of capacity}}^{\text{M = maximum ramp}} e(M, \bar{t})$$

$\text{MAE}_{\text{cdf}}^{\text{ramp} > 10\%}$ is unitless since it is the difference in the frequency of occurrence (i.e., the unitless probability) of ramps. It is expressed as a percentage of occurrence during daytime in Figures 7 and 13.

4.3. Power in down ramp rates greater than 10% of capacity per minute

The power in down ramps at a rate faster than 10% of rated capacity per minute is a proxy for the amount of storage required to comply with the Puerto Rico regulation that ramps must not exceed a rate of 10% per minute. We assume that up ramps can be controlled through other means (e.g., inverter control), but that down ramps will have to be mitigated with energy storage.

Although it is expected that the Puerto Rico ramp restriction will be enforced at 2-second intervals (e.g., 0.33% per 2-seconds), we evaluate the energy in down ramps (with rate faster than 10% per minute) at a variety of timescales.

The power in down ramps PDR is computed as:

$$PDR(t, \bar{t}) = \begin{cases} 0, & RR(t, \bar{t}) > -\frac{10\%}{\text{minute}} \\ \left(RR(t, \bar{t}) + \frac{10\%}{\text{minute}}\right) \times \bar{t}, & \text{else} \end{cases},$$

When the ramp rate is not less than 10% per minute down, the PDR is zero. For a down ramp faster than 10% per minute, the PDR is the amount of power by which the ramp exceeds the 10% per minute rate. Thus, the magnitude of PDR is amount by which the down ramp exceeds a rate

of 10% per minute. For example, a 1.9MW per 30-second down ramp, $RR = \frac{1.9\text{MW}}{30\text{sec}}$, at the 19MW PV plant (i.e., a 20% per minute ramp rate) would be calculated to have a down ramp power

$$PDR = \left(\frac{-1.9\text{MW}}{30\text{sec}} + \frac{1.9\text{MW}}{\text{minute}}\right) * 30\text{sec} = \frac{-0.95\text{MW}}{30\text{sec}} * 30\text{sec} = -0.95\text{MW}.$$

By definition, PDR is always negative (indicating a loss of power due to a down ramp). In Figures 8 and 14, $|PDR|$ is plotted to give a positive sign, and thus corresponds to the power that must be injected by storage to reduce the ramp to a rate of 10% of capacity per minute.

We apply two metrics based on the power in the down ramps: (1) the average daily power in down ramps over the whole time period, relating to the total amount of power required to be injected by the energy storage, and (2) the most energy in down ramps on a single day, which, assuming that the storage can be recharged overnight, will relate to the capacity of storage required.

The power in down ramp rates greater than 10% per minute over all days is

$$PDR_{\text{avg. per day}}(\bar{t}) = \frac{1}{\text{number of days}} \sum_{\text{all days}} PDR(t, \bar{t}),$$

On the worst day:

$$PDR_{\text{worst day}}(\bar{t}) = \sum_{\text{worst day}} PDR(t, \bar{t})$$

Since the PV output cannot be perfectly predicted, and due to battery state of charge limitations, the actual required storage will be larger than, but related to, these calculated values for power in fast down ramp rates.

5. RESULTS

Two separate analyses were undertaken to (a) quantify the models' performance as a function of timescale, and (b) quantify the models' performance as a function of PV plant size.

5.1. Model Performance by Timescale

In this section, we test model performance by timescale. In all cases, the total PV plant output (19MW rated capacity) was compared to the simulated output using the four plant simulation models. Each model was run to simulate the 19MW PV power plant, creating a simulated power output timeseries over the whole period of record.

5.1.1. Qualitative Analysis as a Function of Timescale

As a precursor to computing the quantitative metrics described in Section 4, we first calculated the cumulative distributions of ramp rates at timescales ranging from 1-second to 1-hour for each simulated timeseries, as well as for the measured power output timeseries. The comparison of the cumulative distributions at 1-second to 15-minute timescales is shown in Figure 4, and can be used to qualitatively compare the modeled variability to the measured variability.

In Figure 4, the x-axis is the absolute value of the ramp rate, meaning that both positive and negative ramps were combined. In general, both positive and negative ramps will occur with equal probability, so combining the ramps using the absolute value is common. For specific applications such as determining energy storage requirements, it may be of interest to separate the up and down ramps, as we present later in some of the quantitative metrics. The y-axis in Figure 4 is the percentage of ramps during the 0900 to 1500 interval on all 361 days that exceeded the ramp magnitude on the x-axis. For example, the measured cumulative distribution shows that approximately 4% of all ramp rates were larger than 10% of capacity per minute, such that there were approximately 5200 such ramps over the 349 day period.

It can be seen in Figure 4 that the no smoothing case overestimated the probabilities of large ramps. This overestimation occurs because the no smoothing case does not account for the spatial smoothing that occurs across the 19MW PV plant. At longer timescales, the effect of spatial smoothing across the plant is small, and all four methods converge to have close agreement with the measured cumulative distribution. Beyond these general observations, though, it is difficult to differentiate between the models based on Figure 4.

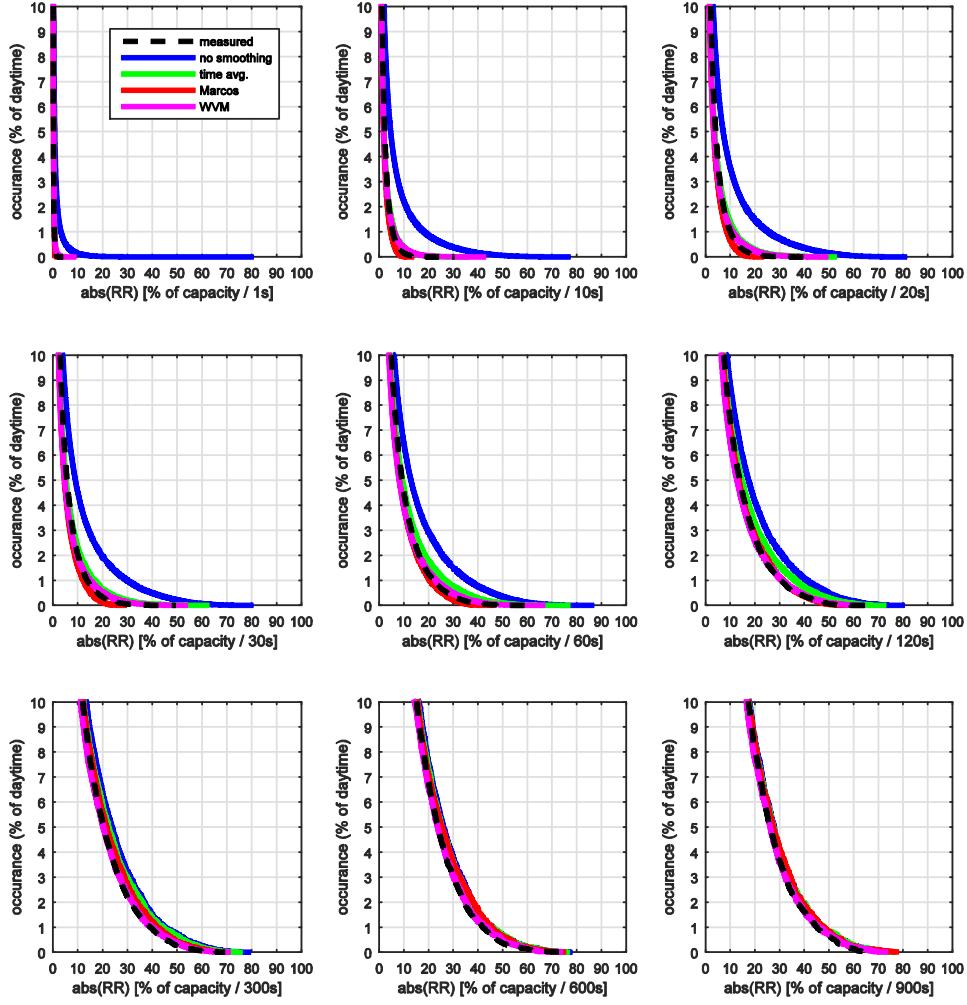


Figure 4: Cumulative distributions of ramp rates at various timescales for measured and simulated 19MW PV power plant output. The x-axis is the absolute value of the ramp rate, such that both positive and negative ramps are included. The y-axis is zoomed in to show ramps that occur 10% of the time or less to help illustrate the differences between the simulation models.

To more directly compare these differences in modeled versus measured cumulative distributions, Figure 5 shows the errors of each model in matching the cumulative distributions of ramp rates. The error was computed as the simulated probabilities of ramp rates minus the measured probabilities of ramp rates. In this way, positive y-axis values in Figure 5 indicate the model overestimated the ramp probability for the ramp magnitude listed on the y-axis, and conversely negative values mean the ramp probability was underestimated. For example, the no-smoothing case had a 2% error at predicting 10% per minute ramps, meaning that it estimated 2% more ramps than actually occurred in the power output.

As seen in the cumulative distributions directly (Figure 4), the no smoothing case always has positive errors, meaning that it always overestimates the ramp rate probabilities. The time

average method was similarly found to overestimate the probability of large ramps (e.g., greater than 10% of capacity) at all timescales, likely due to its simplistic spatial smoothing model.

The Marcos and WVM models had less consistent behavior across the various timescales. Both methods typically had the smallest errors at each timescales. The WVM method had negative errors at small ramp magnitudes at all timescales. The Marcos method similarly had negative errors at shorter timescales, though these errors turned positive at the longest timescales (>10 -minutes).

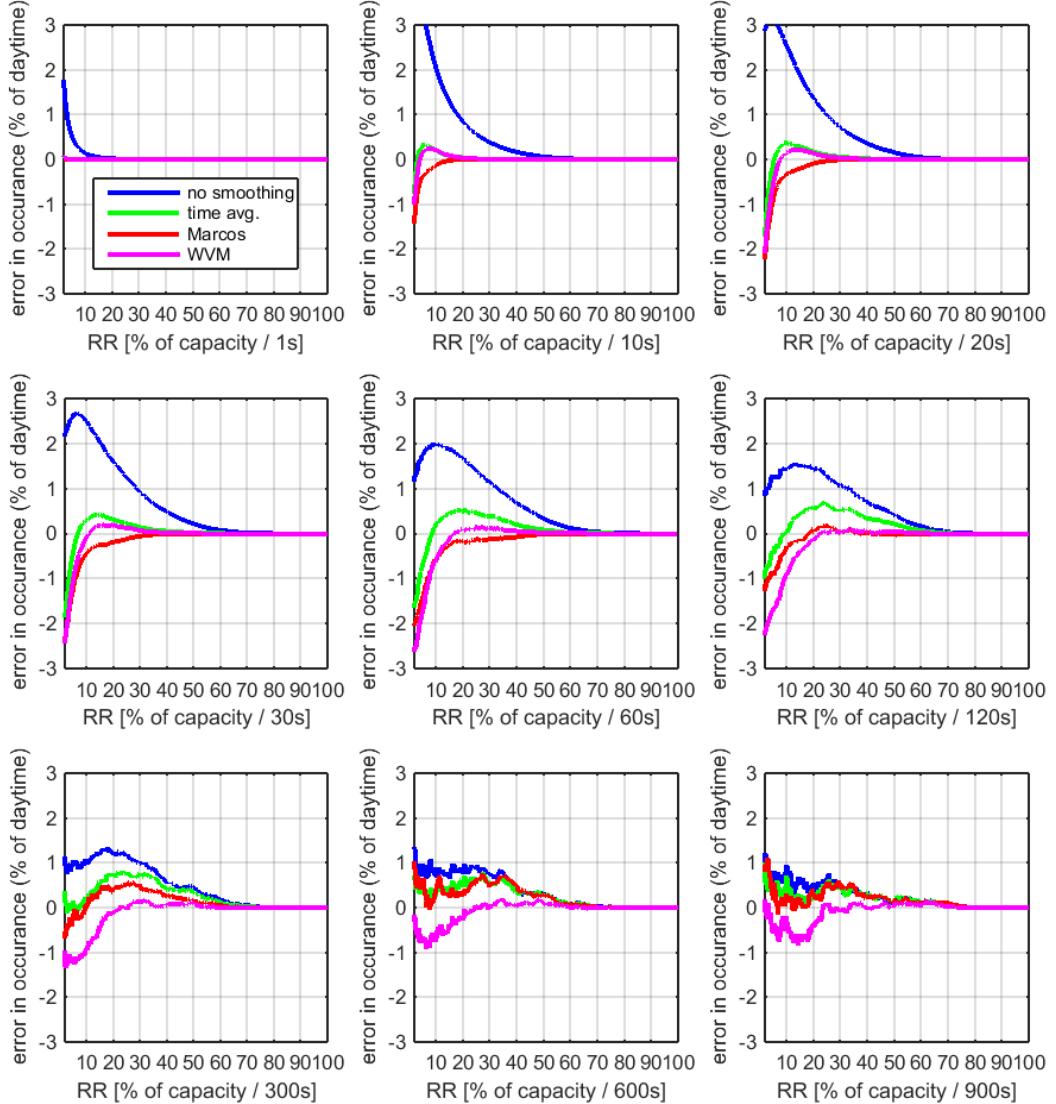


Figure 5: Errors (modeled minus measured) in the cumulative distributions of ramp rates for measured and simulated 19MW PV power plant output.

5.1.2. Quantitative Analysis as a Function of Timescale

To quantify our observations from the cumulative distribution comparisons in section 5.1.1, we applied the quantitative metrics defined in section 4. The largest down ramps (LDRs) and largest

down ramps (LURs) for the measured AC power and the simulated AC powers are shown in Figure 6. Note that down ramps are plotted as negative values, such that the ramps largest magnitude ramps are plotted at the bottom of the largest down ramps plot. In general, the largest ramps (both up and down) increase in magnitude as the timescale increases, since the longer time allows for further deviation from the previous value (note the ramp *rates* generally decrease with increasing timescale). The time average, Marcos, and WVM simulation methods all capture the increasing ramp magnitude with increasing timescale behavior. The no smoothing simulation method, however, predicts roughly constant largest ramp magnitudes at all timescales. At the longest timescales (~ 300 -seconds and longer), all simulation methods converge to have similar accuracy.

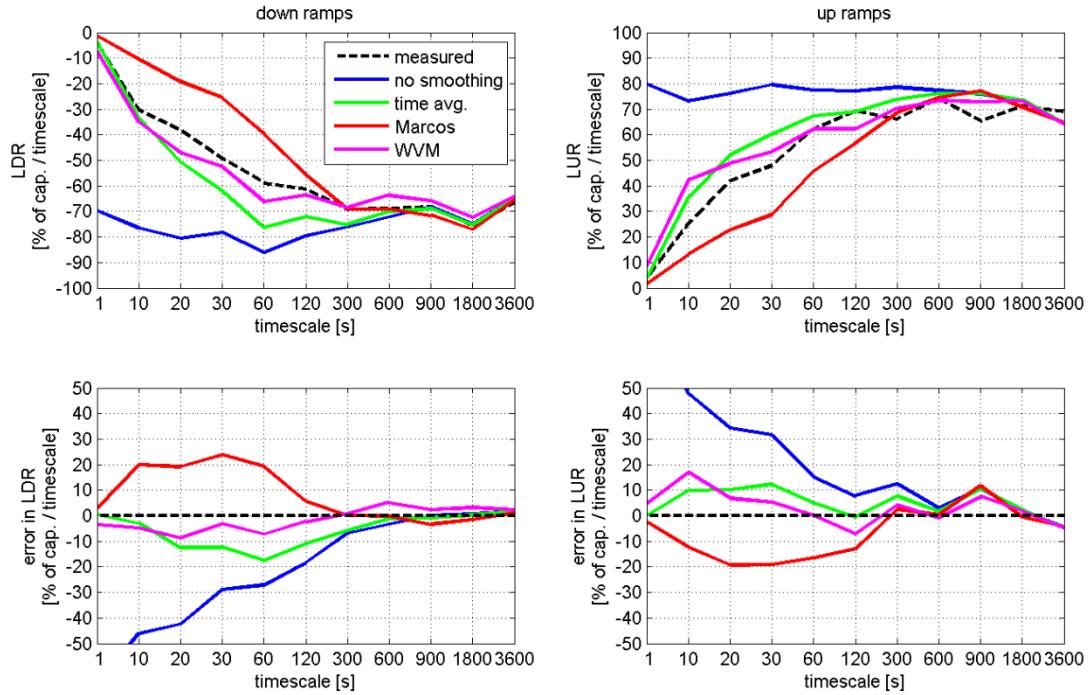


Figure 6: Measured and simulated largest down ramps (top left) and up ramps (top right) at various timescales. The bottom plots show the errors: modeled minus measured.

The simulation methods are most different at short timescales. Spatial smoothing has the biggest effect at the 1-second timescale; such that the 19MW plant measured largest ramp rates are much smaller than the no-smoothing simulated ramp rates. However, the other 3 simulation methods do a good job of modeling the spatial smoothing, and have small errors at the 1-second timescale. As the timescale increases, the no-smoothing simulated largest ramps gradually approach the measured largest ramps. The Marcos method under predicts the largest ramp magnitudes at timescales of 10-seconds to 2-minutes. The time-average method over predicts the largest ramp magnitudes at these same timescales, but with smaller magnitude errors than the Marcos method. The WVM has the smallest magnitude errors at these mid-range timescales, and most often over predicts the largest ramp magnitudes. At timescales of 5-minutes and longer, all methods have similar performance.

While the largest ramps have direct applications to reserves and storage, they describe only a single point in the distribution of ramp rates. To account for all large ramps, Figure 7 shows the mean absolute error in matching the measured cumulative distribution of ramps larger than 10% of capacity ($MAE_{cdf}^{ramp > 10\%}$).

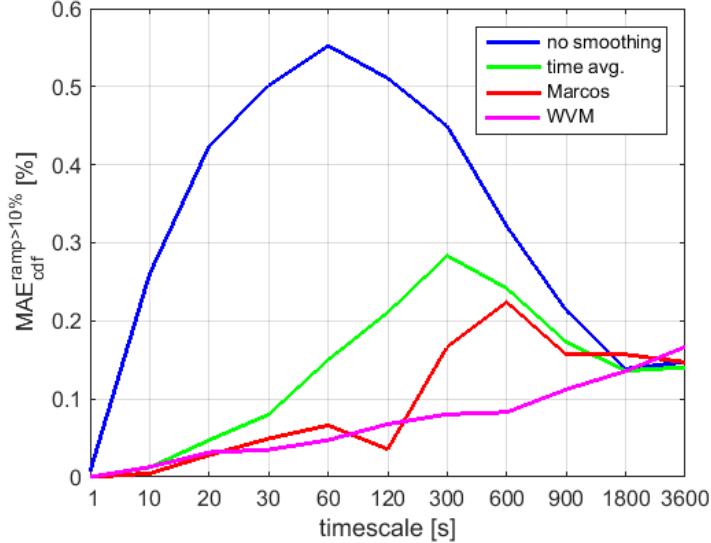


Figure 7: Mean absolute error in matching the measured cumulative distribution for ramp rates larger than 10% of capacity.

The no smoothing method has the largest errors at matching the cumulative distribution of large ramps at timescales of 30-minutes and shorter; at 1-hour all methods have similar performance.

The $MAE_{cdf}^{ramp > 10\%}$ for the no-smoothing method follows a parabolic shape: errors are highest at mid-timescales (maximum error at the 60-second timescale) and lowest at the shortest and longest timescales. While the no smoothing method most overestimates the magnitudes of ramps at the shortest timescales (as seen in Figure 6), the short-timescale ramp magnitudes are small, meaning there are few ramps greater than 10% of capacity and $MAE_{cdf}^{ramp > 10\%}$ is small. As the timescale increases, the ramp magnitudes increase while the error in the no smoothing method is still substantial, leading to large errors at mid-timescales. At the longer timescales, the errors in the no-smoothing method decrease, and $MAE_{cdf}^{ramp > 10\%}$ becomes small again.

The $MAE_{cdf}^{ramp > 10\%}$ values for the other three simulation methods show less pattern. The time averaging and Marcos methods have a weak parabolic shape (with maximums at 5-minutes and 10-minutes, respectively), while the WVM method increases in $MAE_{cdf}^{ramp > 10\%}$ with increasing timescale. The WVM and Marcos methods have comparable performance at short timescales (1 to 60-seconds), while the WVM has better performance at longer timescales (5-minutes to 15-minutes). At the longest timescales (30-minutes and 1-hour), all four methods have similar $MAE_{cdf}^{ramp > 10\%}$ values.

To test the simulation methods' accuracies at matching the down ramps frequency and magnitudes, Figure 8 shows the power in down ramp rates larger than 10% of capacity per minute. Both the power in down ramps over whole time period (349 days) and on the single worst day are shown.

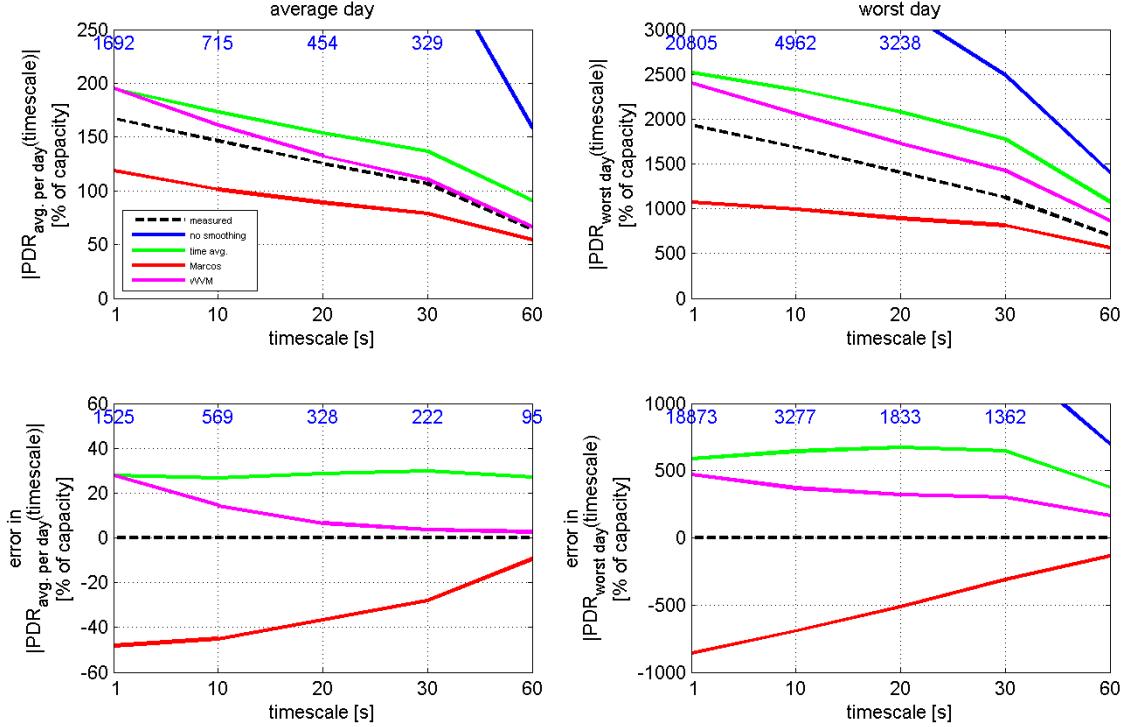


Figure 8: Measured and simulated power in down ramp rates (PDR) at a rate faster than 10% per minute averaged over the whole time period of 349 days (top left) and on the single day with the most energy in the down ramps (top right) at timescales of 1 to 60 seconds. The bottom plots show the errors: modeled minus measured.

The magnitude of power in measured down ramps, both over the whole time period and over the single day, decreases slightly with increasing timescale. This occurs because ramp rates faster than 10% at short timescales may be canceled out by ramps the other direction such that at longer timescales they do not exceed 10% of capacity. For example, a 1-second down ramp of 20% of capacity followed by 9 1-second up ramps of 2% would have some PDR at the 1-second timescale but zero PDR at the 10-second timescale (since the 10-second ramp would be only -2%). All four methods correctly predict this decrease, but with varying levels of accuracy. The no smoothing method overestimates the energy in down ramps at all timescales, much more so than the other methods. This would lead to an expensive oversizing of storage!

The time average, Marcos, and WVM methods all predicted down ramp power values much better than the no-smoothing method. The time averaging method overestimated the magnitude of power in the down ramps, while the Marcos method underestimated this power. The WVM slightly overestimated the magnitude of power in down ramps, but generally had the best performance of the four models.

To summarize the results, Figure 9 shows the method with the smallest error for each timescale and error metric pair. In general, the WVM is found to have the smallest errors at mid-timescales (e.g., 20 to 60-seconds). Although other trends may exist in Figure 9, such as the time averaging method best matching the LDR and LUR at the shortest timescales (1-second and 10-seconds), and the Marcos method best matching the LDR and LUR at long timescales (e.g., >5-minutes), the error magnitudes at these shortest and longest timescales were generally of similar magnitude among the time average, Marcos, and WVM methods.

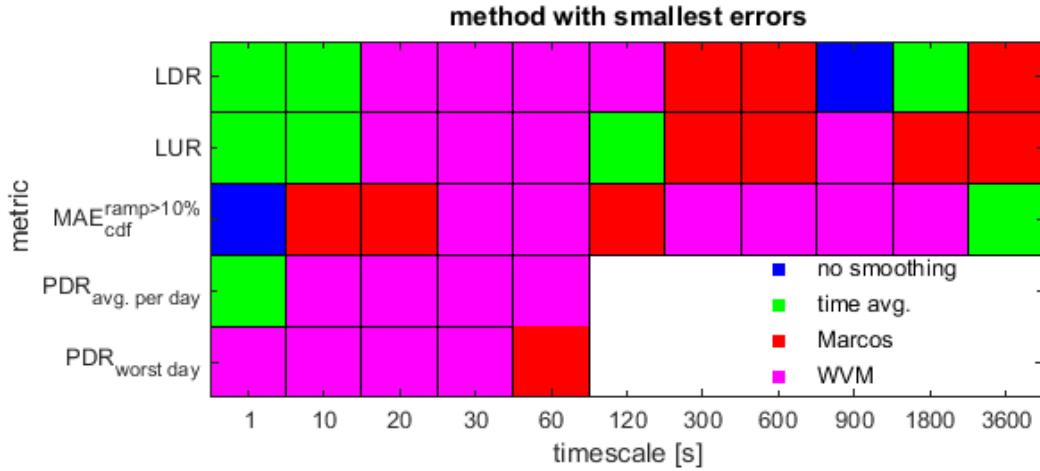


Figure 9: Plot showing the method with the smallest error for each metric and timescale.

5.2. Model Performance by Plant Size

The analysis in section 5.1 showed the models' performance as a function of timescale at matching the total 19MW PV plant output. Since inverter-level PV plant output data was also available at the same plant, in this section we test the model performance as a function of plant size. This granular inverter data was only available from 28 500kW inverters at the plant (rather than the full 38), limiting the plant size analysis to a maximum of 14MW. Thus, we analyze each model's performance at matching the 60-second ramp rates for plant sizes ranging from 1MW to 14MWs.

5.2.1. Qualitative Analysis as a Function of Plant Size

Figure 10 shows the cumulative distribution of ramp rates at a selection of the plant sizes considered. The ramp rates on the x-axis in Figure 10 are shown as a percent of plant capacity to better allow for comparison between plant sizes. We term these ramps as a percent of capacity "relative" ramps. For the measured distributions, the probabilities of large relative ramps always decrease with increasing plant size. It is important to remember, though, that the magnitude of ramps (in MW) will increase as the plant size increases.

At small plant sizes, all four methods closely match the cumulative distribution of ramp rates. Small plant sizes have small amounts of smoothing, analogous to the long timescales section 5.1,

which leads to this similar performance among all models. As the plant size increases, the no-smoothing method becomes worse at matching the measured distribution, always overestimating the probability of large ramps. The no-smoothing method produces a nearly identical distribution of relative ramp rates at every plant size since, since by definition it doesn't account for smoothing across the plant. The time averaging method closer matches the distributions of larger plant sizes than the no-smoothing method, but still overestimates the ramp probabilities. The Marcos and WVM methods both underestimate the large ramp probabilities for large plant sizes.

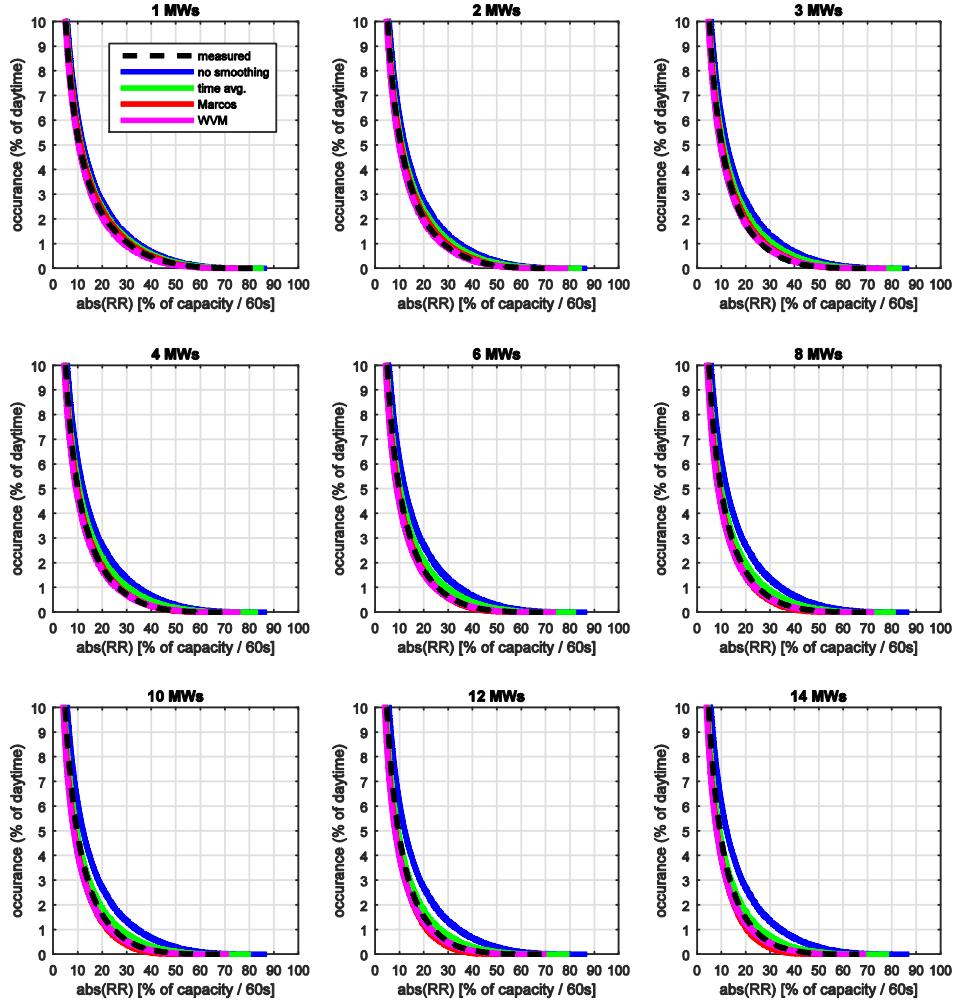


Figure 10: Cumulative distributions of ramp rates over various plant sizes (MWs) for measured and simulated 19MW PV power plant output. The x-axis is the absolute value of the ramp rate, such that both positive and negative ramps are included. The y-axis is zoomed in to show ramps that occur 10% of the time or less to help illustrate the differences between the simulation models.

Figure 11 shows the errors at matching the cumulative distributions of ramp rates. As in the timescale analysis, the no smoothing method has the largest errors, and its errors tend to peak at ramp rates of around 10% of capacity per minute. The other three methods typically under predict the occurrence of small and medium ramp rates; the upper threshold for this over

prediction changes by model and by plant size. At the largest plant sizes, the Marcos and WVM methods both under predict the occurrence of nearly all relative ramp rates. This is consistent with the 19MW total plant analysis in section 5.1, except that for the 19MW plant the WVM was found to have slightly positive errors at relative ramp rates above approximately 20% of capacity per minute.

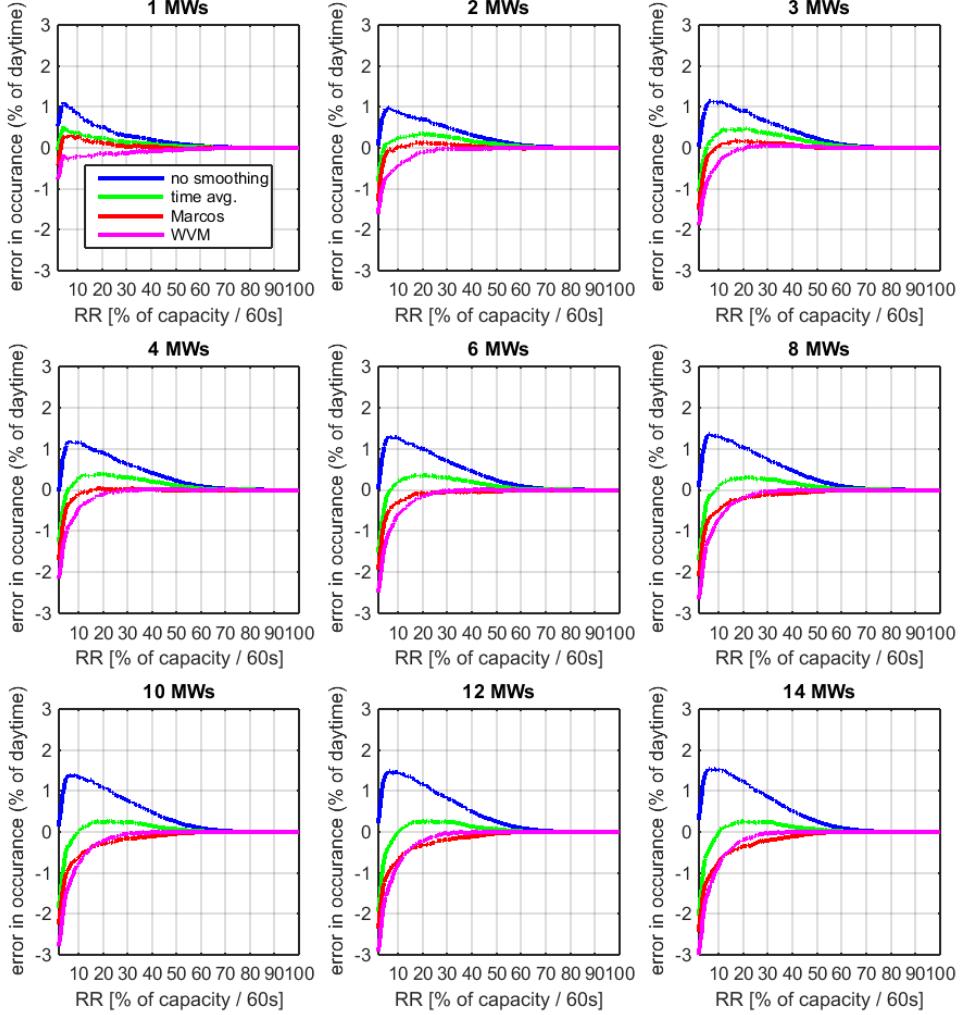


Figure 11: Errors (modeled minus measured) in the cumulative distributions of ramp rates for measured and simulated PV power plant output at 60s over various PV plant sizes (MWs).

5.2.2. Quantitative Analysis as a Function of Plant Size

The largest down ramps (LDR) and largest up ramps (LUR) are plotted in Figure 12 as relative ramp rates. The measured relative ramp rates roughly follow a trend of decreasing as capacity increases, as expected due to the increase spatial smoothing across larger plants. This trend is not linear, though, as the largest relative ramps decrease significantly from 1MW to 3MW, then level

off and have little decrease from 3MW to 14MWs. This behavior may due to the order in which inverters were added to create plants larger than 3MWs. The 3MW plant consisted of six inverters which combined spanned the total east-west extent of the 19MW PV plant. Additional inverters south of these six were added to create larger plant sizes, increasing the north-south diversity but leaving the east-west diversity unchanged.

Consistent with section 5.1, the Marcos method always underestimates the largest relative ramps, while the other three methods always overestimate the magnitudes of these largest relative ramps. At the 1MW plant size, the Marcos and WVM methods overestimate the largest ramp magnitudes; the time average and no-smoothing methods closer match the measured largest ramps at this smallest plant size. As the plant size increases, the WVM generally matches the largest ramps best.

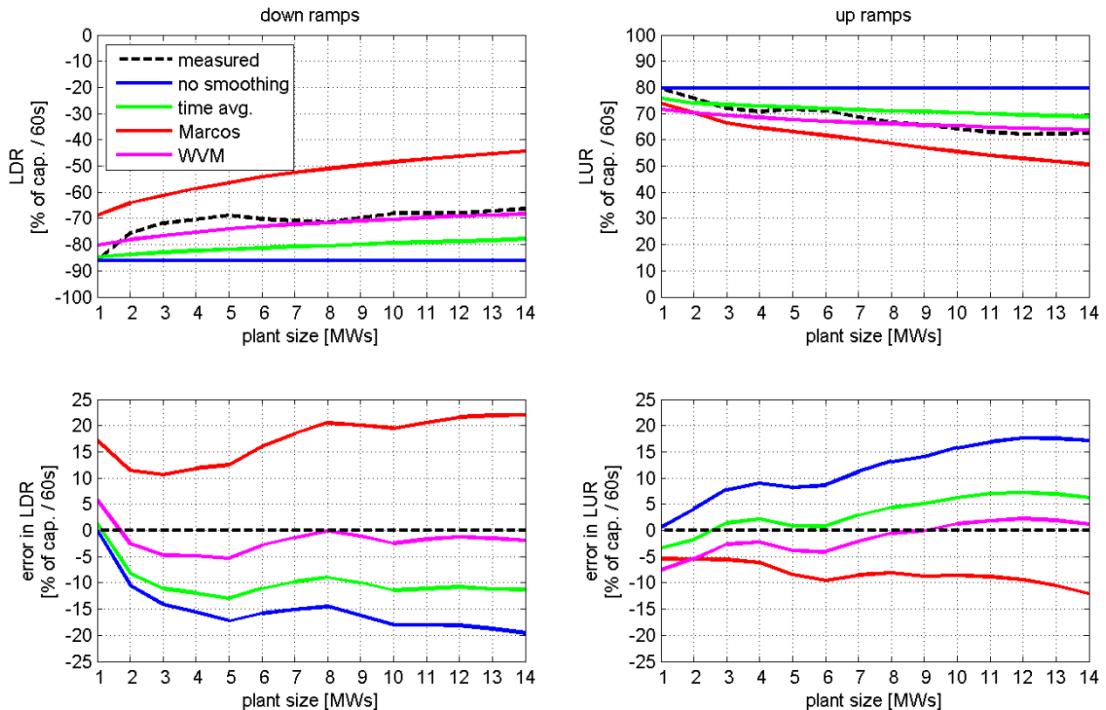


Figure 12: Measured and simulated largest down ramps (top left) and up ramps (top right) at various plant sizes (MWs). The bottom plots show the errors: modeled minus measured.

The mean absolute errors in matching the ramps larger than 10% of capacity are shown in Figure 13. The no-smoothing method's errors become larger as the plant size (and spatial smoothing) increases. The error increases approximately linearly with increasing plant size for plant sizes between 3MW and 14MW. The time average method, conversely, has decreasing $MAE_{cdf}^{ramp > 10\%}$ in this 3MW to 14MW range. The Marcos and WVM methods have similar small errors at small plant sizes (1 to 7MWs), though the Marcos method error increases to larger than the time averaging method at larger plant sizes. This (Marcos $MAE_{cdf}^{ramp > 10\%}$ larger than time average $MAE_{cdf}^{ramp > 10\%}$) is not consistent with the 19MW analysis in section 5.1, where the Marcos method had smaller error than the time average method, suggesting that the Marcos error may reduce and the time average error may increase as the plant size is further increased. The WVM has

consistently small $MAE_{cdf}^{ramp > 10\%}$ (less than 0.07%) at all considered plant sizes, and these values are consistent with the $MAE_{cdf}^{ramp > 10\%}$ value found for the 19MW plant in section 5.1.

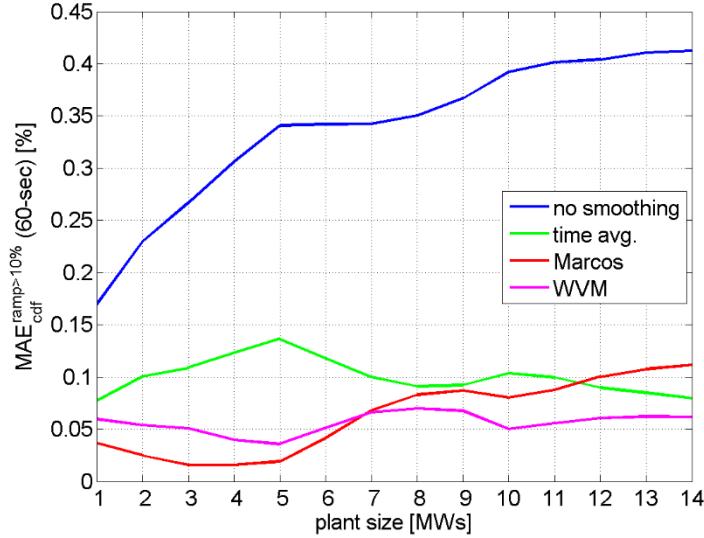


Figure 13: Mean absolute error in matching the measured cumulative distribution for ramp rates larger than 10% of capacity for various plant sizes (MWs).

Figure 14 shows the measured and simulated power in the down ramp rates faster than 10% per minute averaged per day over the 349 days ($PDR_{avg. per day}^{(t)}$), and for the worst day ($PDR_{worst day}^{(t)}$). The no-smoothing method predicts identical power in down ramps for all plant sizes since it does not account for spatial smoothing. The very slight deviation in predicted energy across the different plant sizes is due to missing data: when data was not available for a certain inverter, all plant sizes that required that inverter were removed from the analysis for that day. This is the cause of the slight step between 13MW and 14MW in the no smoothing average daily PDR shown in the top left of Figure 14.

The measured PDR decreases as plant size increase. All methods besides the no smoothing method correctly capture this behavior. For the average over the whole time period, the time average, Marcos, and WVM methods all have similar magnitudes of errors, though the Marcos method is slightly better at matching small plant (1-5MW) PDRs than the other methods. For the worst day, the WVM most closely captures the shape of the measured PDR. The time average method always over predicts the PDR by a approximately the same amount at each plant size, while the Marcos method always under predicts the PDR, especially at large plant sizes.

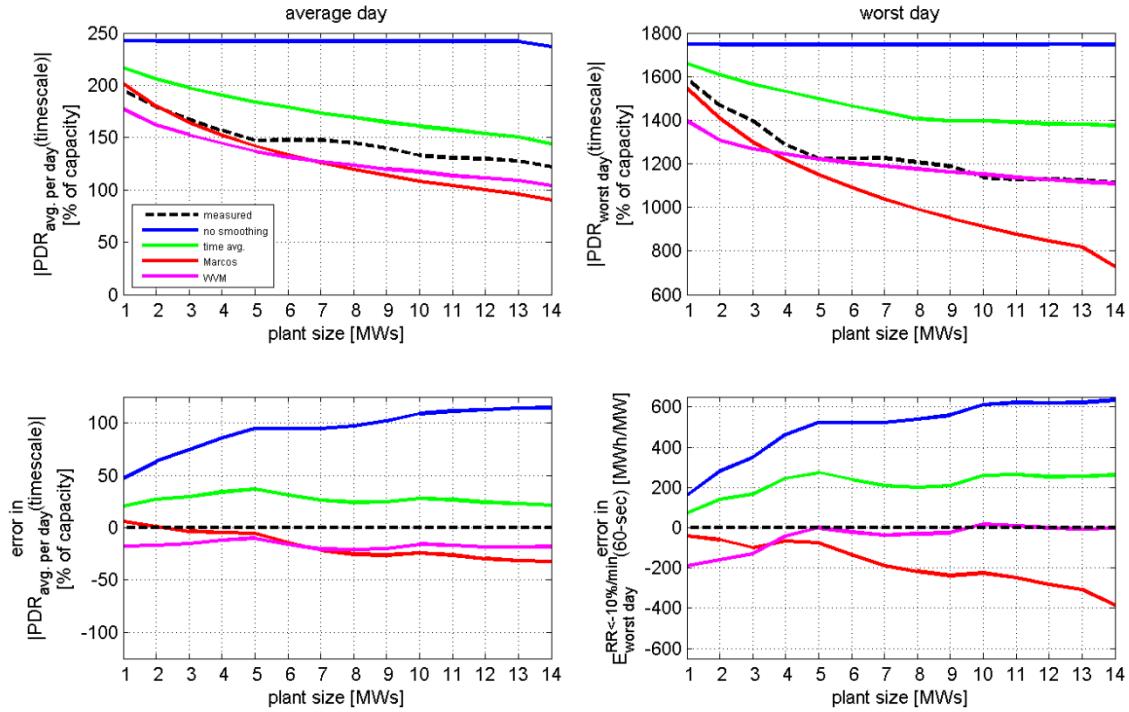


Figure 14: Measured and simulated power in down ramp rates greater than 10% per minute average per day over the whole time period of 349 days (top left) and on the single day with the most energy in the down ramps (top right) over various plant sizes (MWs). The bottom plots show the errors: modeled minus measured.

Finally, Figure 15 displays the best performing method for each plant size and error metric combination. It is generally observed that the Marcos method most often has the smallest error for small plant sizes (1-6MW), while the WVM most often has the smallest errors for larger plant sizes (7-14MW). This is consistent with the 19MW total plant comparison in section 5.1, where the WVM was mostly found to have the smallest errors at the 1-minute timescale. The Marcos method was developed using data from PV plants sized between 0.1MW and 9.5MW [3]; this may be why it performs best at the smaller plant sizes. Similarly, the WVM was primarily validated at a 48MW PV plant [4], which may explain its best performance at the larger plants.

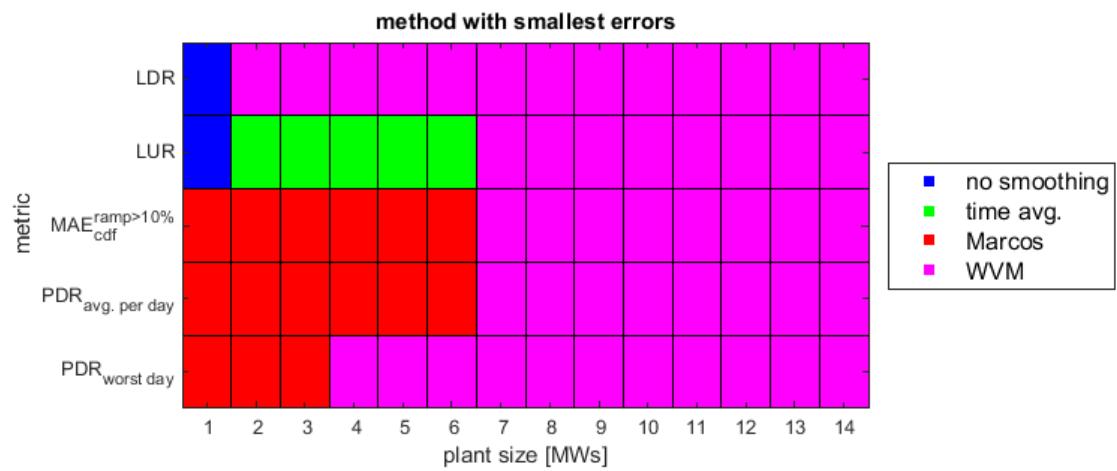


Figure 15: Plot showing the method with the smallest error for each metric and plant size (MWs). All error metrics were computed based on 1-minute ramps.

6. CONCLUSIONS

Errors in four PV power plant variability simulation methods – no-smoothing, time average, Marcos, and WVM – were analyzed as both a function of timescale and as a function of PV plant size. Five application-specific error metrics were created to test the performance of each method at matching the measured power variability.

The no-smoothing method was found to have the largest errors at nearly all timescales and plant sizes since it does not account for the spatial smoothing of irradiance across the plant and so always overestimated the magnitude and frequency of large ramps. The no-smoothing errors were the largest at short timescales and large PV plant sizes, when spatial smoothing has the strongest effect on plant power variability.

The time average method had reduced errors compared to the no-smoothing method, since the time averaging method accounts for spatial smoothing. However, the time average method still typically overestimated the probability and frequency of large ramps. Especially, at mid-timescales (30-seconds to 10-minutes) and small plant sizes (1-9MWs), the time average method had larger errors than both the Marcos and WVM methods.

The time average method assumes that clouds travel linearly across the length of the PV plant at a fixed speed without distortion; in reality clouds do not travel in such fixed patterns. Thus, the time average method likely overestimates the correlations between locations within the PV plant and, hence, overestimates the ramp rates.

The Marcos method underestimated the magnitude of the largest down and up ramps at short and mid-timescales (1-second to 2-minutes), and at all plant sizes (1 to 14MWs, 19MWs). However, for small plant sizes (1 to 6MWs), the Marcos method had the smallest mean absolute error in matching the cumulative distribution of large ramps. The Marcos method also had the smallest error in estimating the average daily power in large down ramps for small plant sizes (1 to 6MWs), and had the smallest error in estimating the worst daily power in large ramps for very small plant sizes (1 to 3MW).

The Marcos method uses a low pass filter defined from a cutoff frequency seen in a Fourier Transform to smooth fluctuations. While using a cutoff frequency in this manner appears to appropriately smooth the total variability, it also appears to overly smooth the largest ramps, leading to the Marcos model's poor performance at simulating the largest ramps. We also note that the Marcos method also applies the same amount of smoothing each day, since the smoothing is based on the cutoff frequency which depends on the plant size but not the daily cloud speed. This will cause the Marcos method to have different performance in locations with different cloud speeds. For example, slow clouds speeds in locations such as Hawaii lead to more spatial smoothing than locations with fast cloud speeds. The Marcos method may underestimate the smoothing (and hence, overestimate the ramp magnitudes and frequencies) in such slow cloud speed locations. The Marcos method was developed using data from small PV plants (0.1MW to 9.5MW [3]), which explain its better performance for small PV plants.

The WVM method most often had the smallest errors. Across all timescales for the 19MW plant, the WVM had errors that were the smallest, or comparable to the smallest errors, for all error metrics. Similarly, for plant sizes of 7MW to 14MWs, the WVM had the smallest errors for all metrics

The WVM simulates correlations between PV modules within the PV plant as a decaying function of distance and timescales: as distance between modules increases or timescale decreases, pairs of modules within the plant become less correlated. The average correlation among all module pairs in the PV plant is inversely related to the amount of smoothing. For consistency among all simulation methods, we only input the plant size, not the plant shape, to the simulation methods. The exponential shape of the decay of total energy in down ramps simulated by the WVM occurs since the smoothing increases exponentially as the plant size increases. However, the measured total energy in down ramps takes a less continuous form which may be due to order in which inverters were added to create the larger plant sizes: the 3MW plant size consisted of six inverters which covered the total east-west extent of the 19MW plant, larger plant sizes were created by adding inverters which were south of these original six.

The WVM is the only method considered that can account for the plant shape, but it was not tested in this way to allow for equal comparison between the simulation models. It is possible that better WVM model performance could be obtained by accounting for the plant shape. Further improvement may also be achieved through an anisotropic correlation model (such as [8]) which accounts for the cloud motion direction relative to the PV plant.

Overall, among the models tested, the WVM most often had the smallest errors both with varying timescale and with varying PV plant size. The Marcos method also typically had small errors, and especially had small errors at small PV plant sizes (1 to 6MWs). The no-smoothing method should only be used as a rough approximation for very small plants or at very long timescales. The time average method is an improvement over the no-smoothing method, and is easy to implement, but generally has larger errors than the WVM and Marcos methods.

REFERENCES

- [1] PREPA. (2012, April 30, 2013). *Puerto Rico Electric Power Authority Minimum Technical Requirements for Photovoltaic Generation (PV) Projects* http://www.fpsadvisorygroup.com/rso_request_for_quals/PREPA_Appendix_E_PV_Minimum_Technical_Requirements.pdf
- [2] M. Lave, A. Ellis, and J. S. Stein, "Simulating Solar Power Plant Variability: A Review of Current Methods," Sandia National Laboratories, Albuquerque, NM; Sandia National Laboratories (SNL-CA), Livermore, CA (United States)2013.
- [3] J. Marcos, L. Marroyo, E. Lorenzo, D. Alvira, and E. Izco, "From irradiance to output power fluctuations: the PV plant as a low pass filter," *Progress in Photovoltaics: Research and Applications*, vol. 19, pp. 505-510, 2011.
- [4] M. Lave, J. Kleissl, and J. S. Stein, "A Wavelet-Based Variability Model (WVM) for Solar PV Power Plants," *Sustainable Energy, IEEE Transactions on*, vol. PP, pp. 1-9, 2012.
- [5] D. King, W. Boyson, and J. Kratochvil, "Photovoltaic Array Performance Model," SAND2004-3535, 2004.
- [6] PVWatts. National Renewable Energy Laboratory, <http://www.nrel.gov/rredc/pvwatts/about.html>.
- [7] D. King, S. Gonzalez, G. Galbraith, and W. Boyson, "Performance Model for Grid-Connected Photovoltaic Inverters," Sandia National Laboratories SAND2007-5036, 2007.
- [8] E. Arias-Castro, J. Kleissl, and M. Lave, "A Poisson model for anisotropic solar ramp rate correlations," *Solar Energy*, vol. 101, pp. 192-202, 2014.

DISTRIBUTION

1	MS1033	Abraham Ellis	6112
1	MS1033	Joshua Stein	6112
1	MS1033	Robert Broderick	6112
1	MS1033	Clifford Hansen	6112
1	MS9052	Matthew Lave	6112
1	MS0899	Technical Library	9536 (electronic copy)

