EPIC Solar Forecasting

TASK 3 Final Report

Grid-Connected and Embedded PV Fleet Forecasting Accuracy

March 2018

Prepared for

Itron, Inc.

Prepared by

Clean Power Research, LLC

Legal Notice from Clean Power Research

Contents

Executive Summary

Under this project, solar power forecasts were provided to the California ISO (CAISO) via two mechanisms. First, behind-the-meter "embedded" systems offset loads at the customer level. Aggregated forecasts of these "fleets" are used in combination with load forecasts to develop the forecast at the utility level to be served by the wholesale market. Second, utility scale "grid-connected" systems feed the wholesale markets directly within CAISO, so solar forecasts inform what units need to be scheduled to supplement the non-dispatchable renewable sources.

This report describes new methods for improving both of these solar production forecasts under Task 3 of the Itron EPIC project EPC-14-001 of the California Energy Commission (CEC), for which Clean Power Research (CPR) was a major subcontractor. Task 3 covers a broad range of activities that led directly to improvements in CPR's ability to efficiently and accurately produce solar production forecasts for the California electric grid. Task 3 also included work related to a dashboard information system (to be evaluated by Itron) and that work is not covered here.

Enhancements Using Embedded System Production Data

Several forecast improvements relied upon a data set of individually metered production systems from a third party source. Prior to this project, SolarAnywhere FleetView treated systems as if newly installed: they were always available for service, they operated as new, there was no age-induced module degradation, and they were free from dust and dirt. In reality, none of these assumptions are true. This project developed and evaluated methods for incorporating these real-world effects into the forecasts.

Methods were advanced for determining system specifications and shading based on measured production inputs. FleetView has historically relied upon data provided by the installers, including PV module make/model, array orientation (tilt and azimuth angles), inverter make/model, and shading profiles. However, this data is not always accurate. Installers may input incorrect information (e.g., using a different azimuth angle convention than the convention assumed by FleetView) or they may not include shading even if obstructions exist that would significantly impact production.

The project included the development of methods to use measured production data in order to confirm or correct system specifications. System specifications are adjusted to minimize error in production profiles relative to actual output.

Other Forecast Improvements

Other forecast improvements include the following. The first relates to inverter efficiency, which is a function of power level. Prior forecasts used a fixed, default power curve applied to all inverters for simplicity. Under this project, the use of model-specific inverter power curves was evaluated and built into software.

Additional work focused on improving the operational SolarAnywhere forecast models at both the short-term (hour ahead) and longer-term (day ahead) time horizons by using advanced ensemble methods leveraging forecasts from multiple sources.

Methods were evaluated for increasing the performance of the forecasts, such as through the use of representative, rather than actual, fleets. These methods reduce the number of systems for simulation, thereby reducing computation time. This may be important as the number of systems increases with higher penetration levels.

Historical simulations were developed and provided to Itron for training by their neural network. This activity improved the net load forecasts that were delivered to CAISO.

Data on system sizing and installation dates were collected from the three California IOUs. A process was developed to automatically update FleetView to take advantage of this and use it to scale up the fleet forecasts.

Project Partnerships

Utility partner meetings were held to gather input on applicability beyond the ISO. Project meetings were held with partners, Sacramento Municipal Utility District (SMUD), Southern California Edison (SCE) and Pacific Gas & Electric (PG&E). Together, we identified the use of the PV simulation tools for quantifying the impact of distributed PV on the distribution grid, and more regionalized BTM PV forecast for utility load modeling as key areas of interest.

In addition, several of the forecast improvements were demonstrated in SCE's Preferred Resources Pilot (PRP) program. This demonstration covered two feeders in their service territory: Johanna and Santiago. These feeders have roughly 4,500 PV installations tied to both residential and commercial buildings. To support the PRP, CPR and Itron are validating their set of modeling tools against measured PV production data from these two regions. The regions sit within a challenging climate that blends coastal fog formation and pushes inland towards hot and arid conditions. This report describes the validation of the CPR PV model output against measured PV meter data for a number of unique PV systems in the Johanna and Santiago regions.

Introduction and Background

The key challenge facing the California ISO (CAISO) and the electric utilities as they integrate higher and higher concentrations of PV into the grid is the uncertainty associated with PV generation profiles. PV is inherently a variable resource and utilities are charged with maintaining high system reliability at low costs. The uncertainty in PV is reflected in conservative scheduling of regulation and spinning reserves.

The work described in this report was done under California Energy Commission (CEC) contract EPC-14-001. The work covers Task 3 of the project related to improvents in forecasting accuracy for behind-the-meter (embedded) PV and utility-scale (grid-connected) systems. This work covers a broad range of activities that led directly to improvements in CPR's ability to efficiently and accurately produce solar production forecasts for the California electric grid.

Forecasts are used in two ways: (1) embedded system fleet forecasts are delivered to Itron as inputs to net-load forecasts; and (2) grid-connected system forecasts can be used by CAISO to schedule units for delivering the net forecasted load. Both of these forecasts use CPR's SolarAnywhere FleetView software product, into which the improved methods are incorporated.

Project Partnerships

As part of the project, utility partner meetings were held to gather input on applicability beyond the ISO. Project meetings were held with partners SMUD, SCE and PG&E. Key areas of interest were identified to be the use of the PV simulation tools for quantifying the impact of distributed PV on the distribution grid and more regionalized BTM PV forecast for utility load modeling.

The specific use case for PG&E was the PV modeling for distribution grid planning. As the number and capacity of distributed PV continues to grow in PG&E territory, the cost of and uncertainty around operating the distribution grid is growing. Distributed PV can create a number of problems at the distribution level. The problems largely arise when the PV capacity becomes a significant portion of the regional load. PG&E was seeking to quantify the regional, feeder-level capacity and energy contribution of distributed PV. This project demonstrated that the PV modeling tools were useful in demonstrating capacity and energy contribution Challenges were encountered, however, when system shading information was not recorded. Activities included:

- Held meetings with project partners SMUD, SCE, and PG&E
- Refined utility partner BTM/utility-scale PV fleet grouping capabilities
- Performed in-depth PG&E and SCE sub-fleet modeling analysis.

Enancements Using Embedded System Production Data

Several forecast improvements relied upon a data set of individually metered production systems from a third party source. This data was used to gauge the effectiveness of the new methods in improving simulation accuracy and, by extension, forecast accuracy, which relies on the simulations. Improvements include the incorporation of module degradation effects, module soiling, system availability, and the accuracy of system design specifications.

Prior to this project, SolarAnywhere FleetView treated systems as if they were newly installed: they were always available for service (i.e., they were on-line), they operated as newly installed, there was no age-induced module degradation, and they were free from dust and dirt. In reality, none of these assumptions are true. Methods were developed and evaluated under this project methods to incorporat these real-world effects into the forecasts.

Module Degradation

Effect of Module Degradation on Modeling Accuracy

Although CPR's PV modeling tools have long applied module degradation, the effect has generally been applied starting at the beginning of the simulation period. In other words, if you specified a degradation rate of 0.5% per year, and simulated the period from January 1, 2015 to January 1, 2016, the module's rating as of January 1, 2016 would be reduced by exactly 0.5% from its value on January 1, 2015. When working with newly installed or hypothetical systems, this is exactly the behavior desired. However, when modeling output for a 10-year-old system, the module rating should already be reduced by 5% on January 1, 2015 and should be reduced by an additional 0.5% by January 1, 2016.

CPR added a commissioning date to all system specifications, not previously available in FleetView. The module degradation can now be calculated at the specified rate beginning on that date, regardless of the time period being simulated.

To estimate the effect that this change might have in a real-world application, CPR simulated the SDG&E BTM fleet of approximately 14,000 systems for a single day (July 4, 2013) using a 0.5% annual degradation rate and a per-system commissioning date based on the system's California Solar Initiative (CSI) incentive payment date. Systems were installed as early as 2008, but typical age was about 2-3 years. Note that 2013 was a transition year where new installations were no longer being funded under CSI, so this analysis was performed for mid-2013 where reliable data was available.

Total daily fleet energy production without degradation was 883 MWh and peak power was 113 MW AC. With degradation, total daily energy production dropped to 874 MWh and peak power was 112 MW

AC. The relative Mean Absolute Error in power production over the course of the day was 1.03%. Figure 1 shows SDG&E fleet production for July 4, 2013 with and without the effects of module degradation.

Figure 1. Effects of Module Degradation

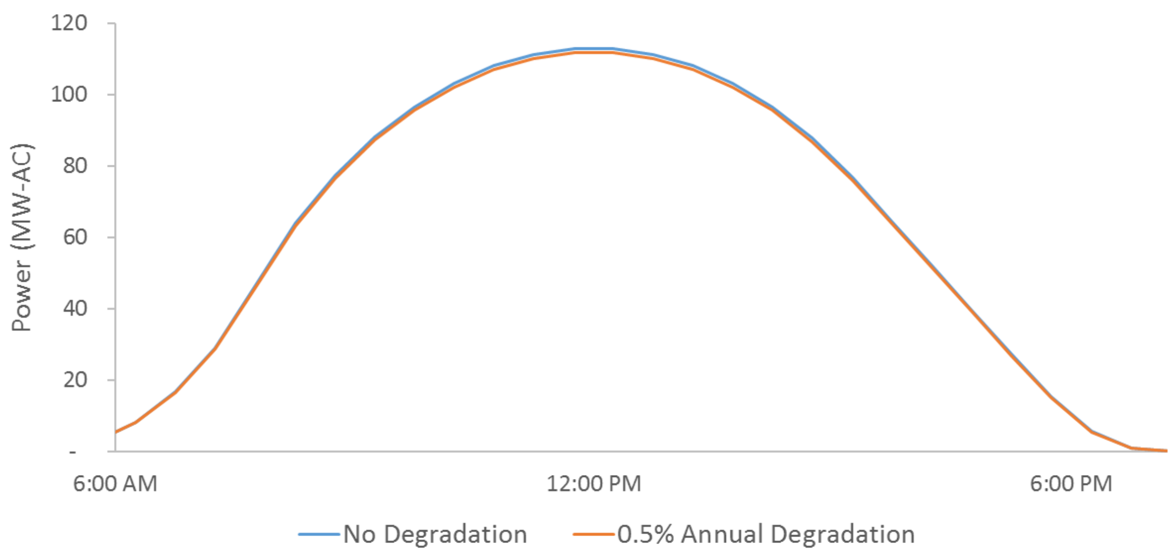

Methodology

The following process was applied the data Itron collected on behalf of the California Solar Initiative to identify a degradation signal.

1. Obtain SolarAnywhere 1 km V2.5 data.
2. Obtain PV specifications from PowerClerk.
3. Use PVSimulator to model PV power production using data from steps 1 and 2.
4. Align hourly ground power data and hourly SolarAnywhere historic power for each site. Remove negative ground values and non-zero nighttime ground values from the data.
5. Calculate annual MAE for each year between SolarAnywhere power and ground power, from 2011-2015, for each site. Periods of time when ground data was not present or zero were not included in the calculations.
6. Discard any MAE value greater that 20% from consideration
7. Discard any site that did not have an MAE value for all 5 years (2011-2015).
8. Average the remaining 207 sites' MAE for each year, to produce one average MAE for each year (2011-2015)

Results and Discussion

From 2011 to 2015 an increase in average MAE for all systems was observed. This increase, when only taking the 2011 and 2015 into account, results in a 0.42% per year increase in average MAE for all systems. This rate is lower than the rough estimate of 1% per year, going by the typical 80% of capacity after 20 years of use. Comparing the results to the measured degradation of solar panels of about 0.5% year, puts the result of 0.42% annually - well within a reasonable range.

Figure 2: Relationship between average MAE for 207 sites selected from the Itron data that both have 5 years of data and annual MAE less than 20%

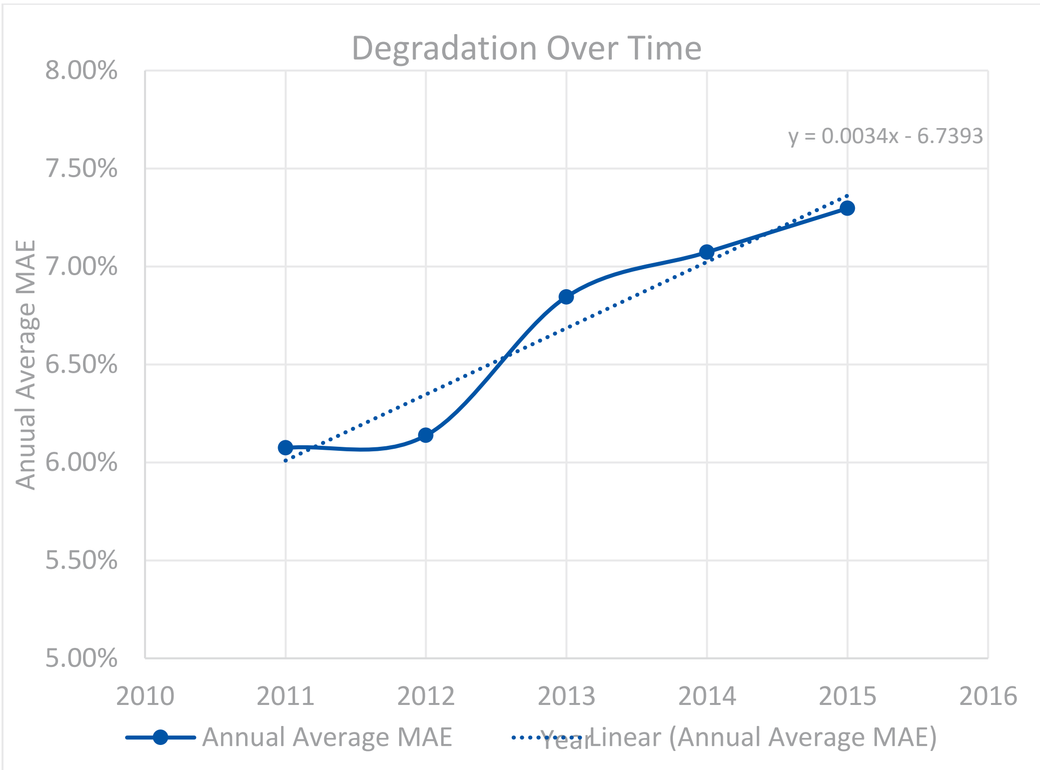

The annual increase in average MAE for all systems is interpreted as degradation. The increase is on the expected order of magnitude and in expected direction. The increase in MAE would be expected because the SolarAnywhere power simulations do not currently take degradation into account. This would lead to a small increase in simulation error over time as PV panels degrade. A linear line of best fit has a slope of +0.0034 which, when divided by the average power of the 207 systems, results in an annual degradation rate of 0.32%.

CPR employed a second approach to identify degradation. CPR averaged the monthly maximum energy generation over 5 years for the same 207 filtered systems, resulting in 60 average monthly maximum values. A linear line of best fit trend line was then applied. The resulting slope was -0.0057 kWh. This is then divided to by the average system power output, of 1.053 kWh, for an annual degradation rate of 0.54%.

Figure 3. Degradation via Monthly Max Values

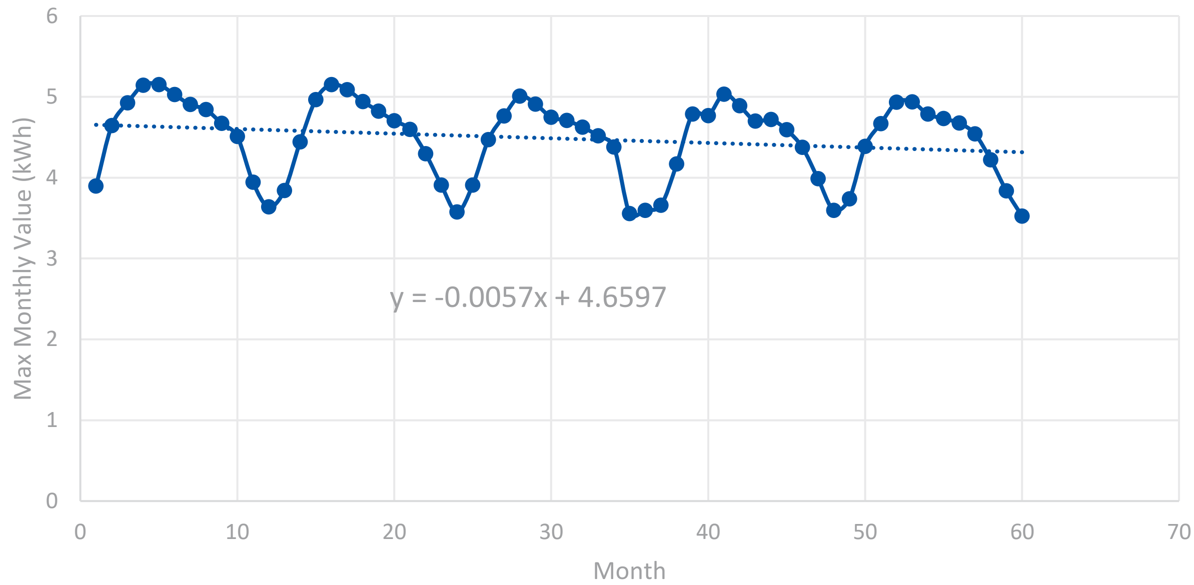

Both approaches result in similar results, and are consistent with industry studies reporting 0.5% degradation. Comparable results from both the satellite simulations relative to ground and using the ground information indicates consistency in the satellite data prior to applying degradation, and builds confidence in applying a modeled degradation approach to better predict the real world PV fleet output.

Soiling

Algorithm Summary

The soiling algorithm allows SolarAnywhere power simulation to take soiling of PV panels into account. Not considering module soiling losses during PV simulations can lead to systemically high biases in PV power. The soiling algorithm is a function of time and precipitation. It assumes that soiling increases at a constant temporal rate and is reduced by precipitation events. There are 2 categories of precipitation events; major and minor events. Major precipitation events remove more soil from PV modules than minor events do. This soiling algorithm is custom-designed to work with daily precipitation data from the SNODAS (Snow Data Assimilation System) dataset, which is produced using a reanalysis with measured input to the base numerical model. Results and improvement are shown with and without the soiling model applied.

Results

Results shown are based on five years of hourly data from the 500-system Itron metered PV fleet. One year of BTM data from a major solar installers was also used to firm up the soiling rate calculations reported above.

Table 1. Soiling Analysis Results

Soiling Results		
Average MAE Unsoiled	Average MAE Soiled	Absolute % Difference
6.62%	5.95%	0.67%

Relative Percent Improvement	**10.12%**

Yearly stats	2011	2012	2013	2014	2015
MBE	1.03%	1.06%	3.40%	4.48%	4.43%
MBE Soiled	-1.19%	-1.07%	0.50%	1.26%	2.47%
MAE	6.00%	6.03%	6.76%	7.15%	7.15%
MAE Soiled	5.91%	5.78%	5.79%	5.96%	6.32%

Figure 4. Soiling by Year

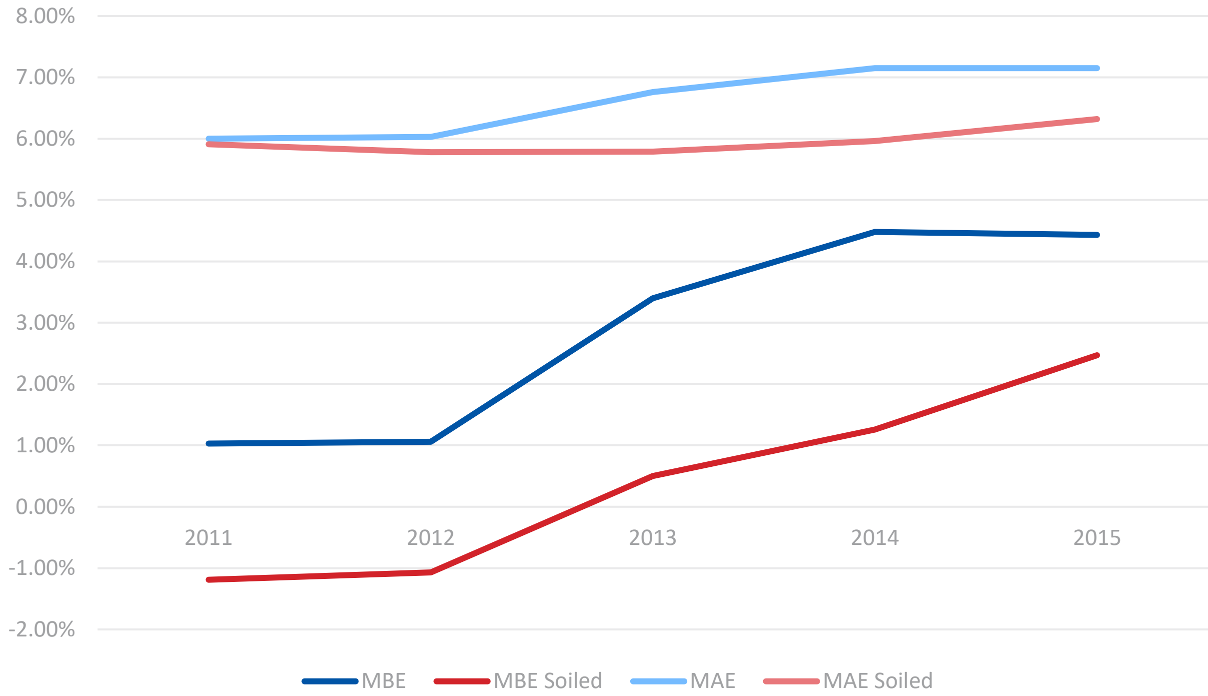

System Availability

Forecasts should account for the fact that not all systems are on-line at any given time. Some may be unavailable due to any number of factors, such as fuse/breaker trips, maintenance, and line power disturbances (which cause the units to trip offline). It is not possible to monitor every system in the fleet for availablility, so an overall factor could be used to represent average outage rates. The factor would then be applied to the fleet as a whole.

Methodology

To determine the availability of the fleet of CSI systems that Itron metered, the following steps were performed:

1. Produce historical power simulations using SolarAnywhere 1 km V2.5 data in combination with PVsimulator and site specifications from PowerClerk.
2. Obtain ground data from 500 well-maintained and monitored PV rooftop systems dispersed throughout the state of California. Data from each system covered approximately 5 years, ranging from 2010 to 2016.
3. Align hourly ground power data with hourly SolarAnywhere historical power data for each site. Remove negative and non-zero nighttime ground values from the data.
4. Discard time periods with missing ground data to eliminate them from availability calculations.

5. Treat periods of time when the systems reported zero and SolarAnywhere simulations were non-zero during daylight hours as times when the system was not available.
6. Determine % availability for each system by dividing the number of periods when the system was available by the total periods not discarded , then take the average to get fleet availability.

Results and Discussion

Overall, systems had high availability. The average of all 476 systems was 98.27% availability. 319 systems had 100% availability.

Figure 5. Histogram of System Availability

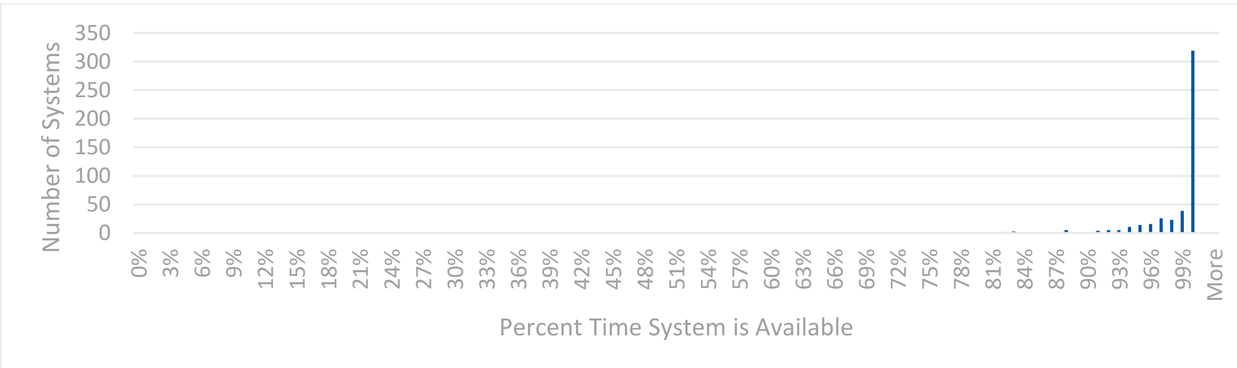

The figure shows the distribution of system availability in the fleet. The trend is highly biased towards near 100% availability. Below are additional statistics on the fleets availability.

Table 2. Fleet Availability Statistics

Statistics	
Average	98.27%
Max	100.00%
Min	64.01%
Mode	100.00%
Median	99.84%
Standard Deviation	3.56%

One aspect of availability that was not considered is partial system availability. This can occur when a single module or an inverter in the system is not functioning properly. This would result in decreased power output from the system but would not result in the system be reported as unavailable, rather it might be interpreted as degradation or soiling. Since the system is not reporting 0s, this aspect of availability remains unmeasured. This partial availability would be seen in increased error between the power simulations, and actual ground data. This error would be difficult to systematically identify.

Additionally, the relationship between system size and availability was investigated. The hypothesis was that it would be more likely that larger systems are carefuly monitored and maintained and would therefore have a higher availability. Figure 6 shows that relationship, however the sample size is small enough that this relationship may be obscured. Finally, the available data skewed heavily towards single family residential systems.

Figure 6. Availability as a function of system size

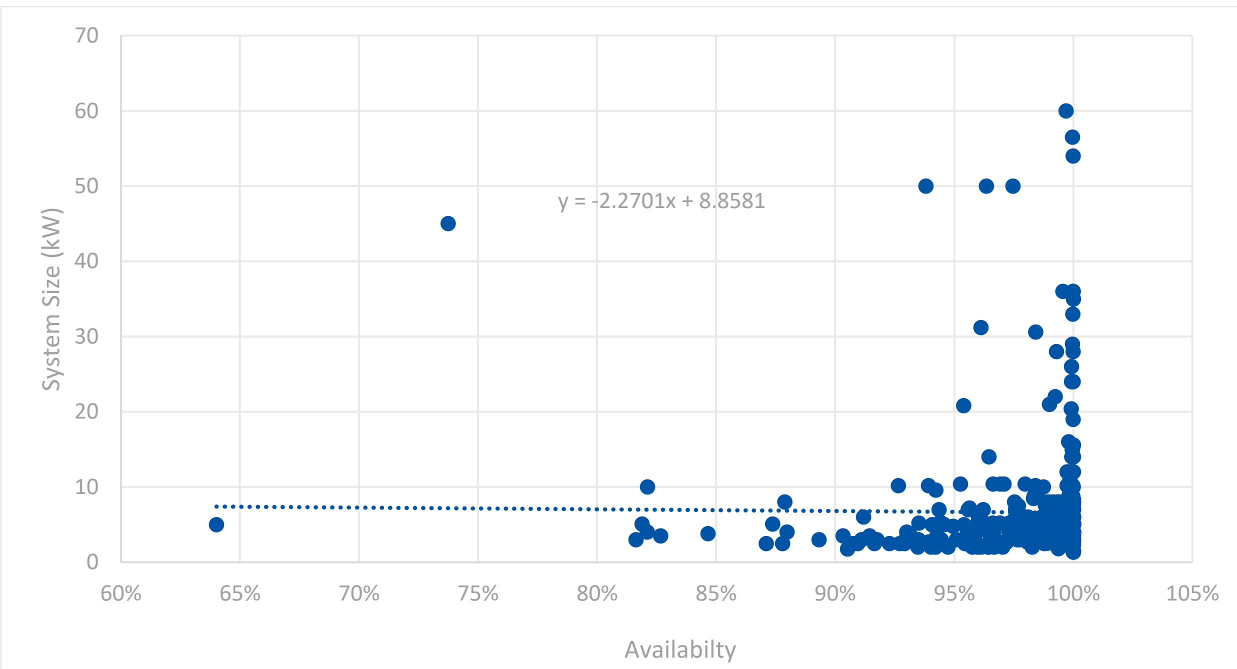

Improving System Specifications by Inference

With more than 5 GW of utility-scale PV capacity, forecasting output from large PV plants is becoming increasingly important to the California ISO.[1] Unfortunately, detailed system specifications, which would

[1] http://www.eia.gov/todayinenergy/detail.php?id=24852

improve the accuracy of modeled PV output, are difficult to obtain because most of these plants are privately owned. It should be possible, CPR hypothesized, to use historical weather data and measured system output to infer some or all of a PV system's specifications automatically. The same approach might also be used to determine BTM system specifications.

The goal for this part of the project was to develop a command-line tool that would compare measured PV production data with the simulated output from candidate systems with various tilt and azimuth combinations. It would then identify the candidate system whose output resulted in the best fit to the measured data.

To simplify the problem somewhat, exact system location (latitude and longitude) is a required input. Also, the first version of the tool would only attempt to infer tilt, azimuth, and AC and DC system ratings for fixed (i.e. non-tracking) systems. Furthermore, to reduce the overall error in candidate system output, a baseline specification would be used to provide any known system details such as commissioning date, row count, row spacing, or solar obstructions.

Methodology

The overall approach used to infer PV system specifications from measured production data was as follows:

1. Transform measured production time series data, as necessary, through compression and time shifting, to simplify comparison with modeled data. For example, 10-minute interval data is compressed to 30-minute interval data with time stamps at 15 minutes and 45 minutes after the hour, to match SolarAnywhere native weather data.
2. Identify clear sky days for the system's location. We selected days where the sum of the simulated system power from a horizontal system was >= 98% of the sum of the simulated clear sky power for that system. By examining only clear sky days, it becomes easier to identify inverter clipping and differences caused by a mismatch in system azimuth.
3. Identify and track reporting errors in the measured data. Periods with reporting errors can be filtered out so as not to affect data comparison.
    a. Gaps in time stamps indicate missing data.
    b. Days where the maximum power is < 80% of previous day's maximum power can indicate a problem with the system (system outage) or a reporting error.
    c. Days with periods where clear sky power is zero and measured power > 3% of the max measured power (nighttime production) can indicate either a calibration error or a reporting error.
    d. Days with periods where measured power is zero and clear sky power > 5% of max clear sky power can indicate either a system outage or reporting error.
4. Create a time-correlated list of AC ratings that defines the time periods to be modeled
5. Simulate candidate systems, filter and compare modeled with measured production

a. Create a set of candidate systems with the same basic attributes as the baseline system, but with varying combinations of DC rating, tilt, and azimuth.
b. For each candidate system, set the AC rating to the average of maximum daily measured power, then simulate production for the most recent rating time period
c. Filter the measured and simulated production data and compare only clear sky days with no reporting errors.
d. Select the candidate system for each rating period with the lowest relative mean absolute error (rMAE). If selected candidate systems differ in tilt and/or azimuth, then manually review candidate systems. Differences in DC and/or AC rating are expected between rating periods.

Results

Measured PV production data for the period from March 1, 2013 through March 31, 2014 was obtained for a number of utility-scale PV plants. Here, we'll focus on three of the smaller plants, which we've designated Plant A, B, and C. CPR used publicly available information in an attempt to determine actual plant specifications manually. This information was supplemented by satellite imagery to determine the approximate number of rows of modules, number of inverters, and the array orientation. That information was used as the basis for a baseline system to be used as a template for each of the candidate systems whose simulated output would be compared to measured.

The figures below show measured output for two of the PV plants studied. In Figure 7, based on the increases in maximum power output, it appears that the plant was undergoing construction from March through July 2013. Starting in August 2013, maximum power output remains flat, in spite of seasonal changes that would normally cause a drop in output. From this, we can deduce that the plant has a DC to AC ratio that's high enough to allow its maximum power output to remain relatively constant throughout the year.

Figure 7. Measured PV Output, Example One

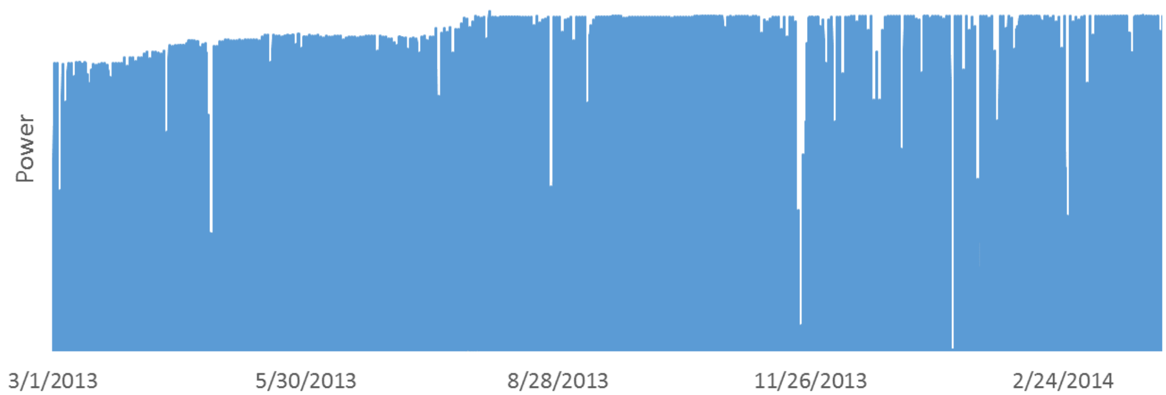

The plant output shown in Figure 8 is also clipped due to a high DC to AC ratio. However, the ratio is not high enough to permit its maximum power output to remain constant throughout the year and we see a drop in maximum power output from October through January. Also worth noting is the lack of data in mid-January 2014. This could be either a reporting error or a plant outage. It's impossible to know from the measured data alone.

Figure 8. Measured PV Output, Example Two

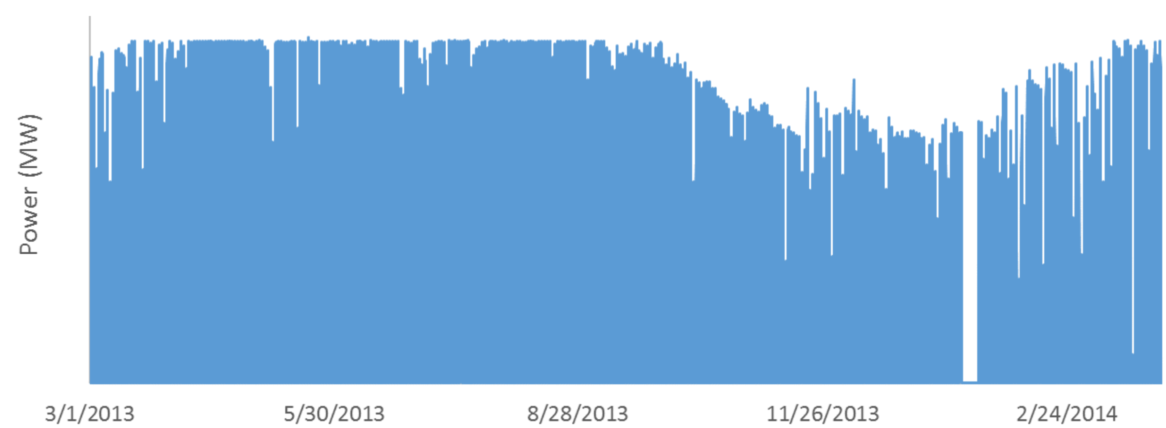

These two measured data examples illustrate some of the challenges in deriving system specifications from measured data: Changes in plant capacity, inverter clipping and lack of seasonal output changes due to high DC to AC ratios, missing data, and unknown PV plant operational status.

When automatically inferring specifications, the tool correctly identified the gross capacity changes over time, but had difficulty during transitional periods where capacity changed on an almost daily basis. Once capacity had stabilized, the simulated output from the selected candidate system matched the measured output reasonably well (see Figure 9).

Figure 9. Measured and Simulated PV Output for Selected Days, Example One

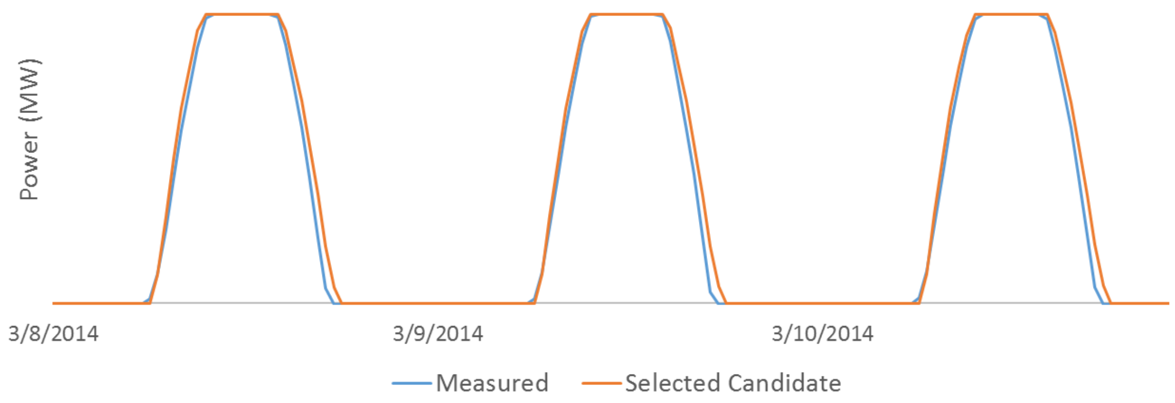

For the output shown in Figure 8, the AC capacity was slightly underestimated by the spec inference tool (see Figure 10) and the DC to AC ratio was overestimated (see Figure 11).

Figure 10. AC Capacity Underestimated, Example Two

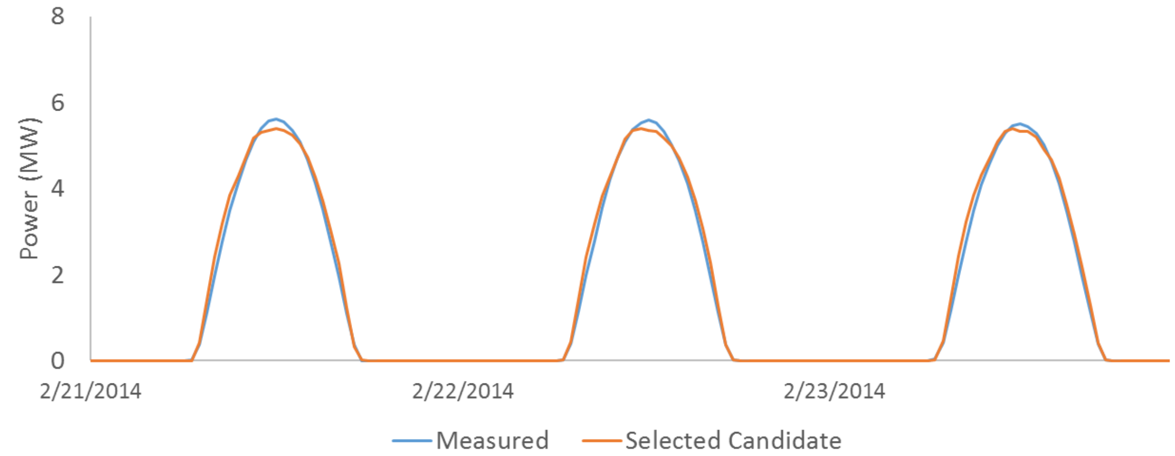

Figure 11. DC to AC ratio Overestimated, Example Two

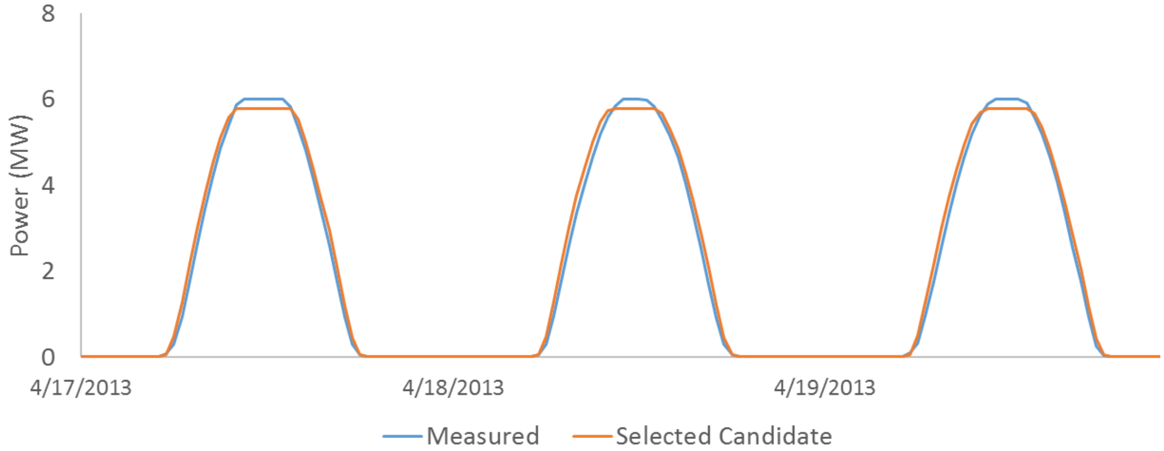

Table 3 summarizes the system specifications inferred for three sets of measured data. Information about these three systems was readily available online for these systems and in two cases, the baseline system yielded the lowest relative Mean Absolute Error. However, for Plant A, a much better candidate was identified and the error was reduced significantly.

Table 3. Summary of Results for Three Systems

		Plant A	Plant B	Plant C
Baseline				
	Rating (MW DC-PTC)	7.733	19.668	20.703
	Rating (MW AC with Losses)	6.821	17.347	18.260
	Tracking	none (fixed)	none (fixed)	none (fixed)
	Azimuth	180	180	180
	Tilt	20	20	20
	DC to AC Ratio	1.22	0.98	0.98
	rMAE	23.0%	12.6%	12.8%
Selected Candidate				
	Rating (MW DC-PTC)	7.772	19.668	20.703
	Rating (MW AC with Losses)	7.402	17.347	18.260
	Tracking	none (fixed)	none (fixed)	none (fixed)
	Azimuth	180	180	180
	Tilt	25	20	20
	DC to AC Ratio	1.05	0.98	0.98
	rMAE	6.2%	12.6%	12.8%

Clean Power Research made significant progress in creating an automated tool for inferring PV system specifications using measured PV output data. Due to the complexities inherent in interpreting such data, we believe that additional accuracy is possible. For example, we did not account for solar obstructions, soiling, module degradation and other factors that decrease DC output. Consideration of such details were outside of the scope of this project, but we hope to continue development of this tool and the algorithms it implements. Preliminary versions of the tool,when implemented in software, could be applied to a system that improves the quality of reported PV system specifications in PowerClerk, by incorporating measured production data from PV systems within a utility or ISO territory.

Other Forecast Improvements

Inverter Power Curve

Historically, when modeling PV system output, CPR has relied on the CEC weighted average efficiency rating to determine the output of an inverter relative to its DC input. This single number has been used in conjunction with an inverter power curve that is the same regardless of inverter make or model. In an effort to improve model accuracy, CPR has implemented two new ways to specify the inverter power curve.

The first method allows the system specification to contain a list of power level/efficiency pairs. Using this method, you could, for example, specify the five power levels for which the CEC publishes inverter test results and the inverter efficiency at each level. The following XML code snippet shows an example of this method for specifying the inverter curve.

Figure 12. Actual Inverter Power Curve versus Existing Default

```xml
<EfficiencyRatings>
	<PowerLevelEfficiency PercentMaxAcPower="10" PercentEfficiency="93.4" />
	<PowerLevelEfficiency PercentMaxAcPower="20" PercentEfficiency="96.9" />
	<PowerLevelEfficiency PercentMaxAcPower="30" PercentEfficiency="97.2" />
	<PowerLevelEfficiency PercentMaxAcPower="50" PercentEfficiency="97.5" />
	<PowerLevelEfficiency PercentMaxAcPower="75" PercentEfficiency="97.4" />
	<PowerLevelEfficiency PercentMaxAcPower="100" PercentEfficiency="97.2" />
</EfficiencyRatings>
```

The second method we've implemented for specifying the inverter curve is to list a set of coefficients and exponents used in a formula to calculate AC power output for a given DC input. This method facilitates precise mathematical control over the shape of the output curve.

To estimate the effect of a more accurate inverter curve on modeled output, we simulated output for a one-year period from two 5 kW systems that were identical in every way except for the inverter curve. Using the CPR default inverter curve yielded a maximum power output of 4.51 kW and a total of 9,042 kWh for the year, while using an inverter specified by a list of power level/efficiency pairs yielded a maximum power output of 4.388 kW and a total of 8,786 kWh for the year. The relative Mean Absolute Error was 2.9%.

Error! Reference source not found. shows the results of using this approach versus the default inverter power curve used in CPR's simulation model.

Figure 13. Actual Inverter Power Curve versus Existing Default

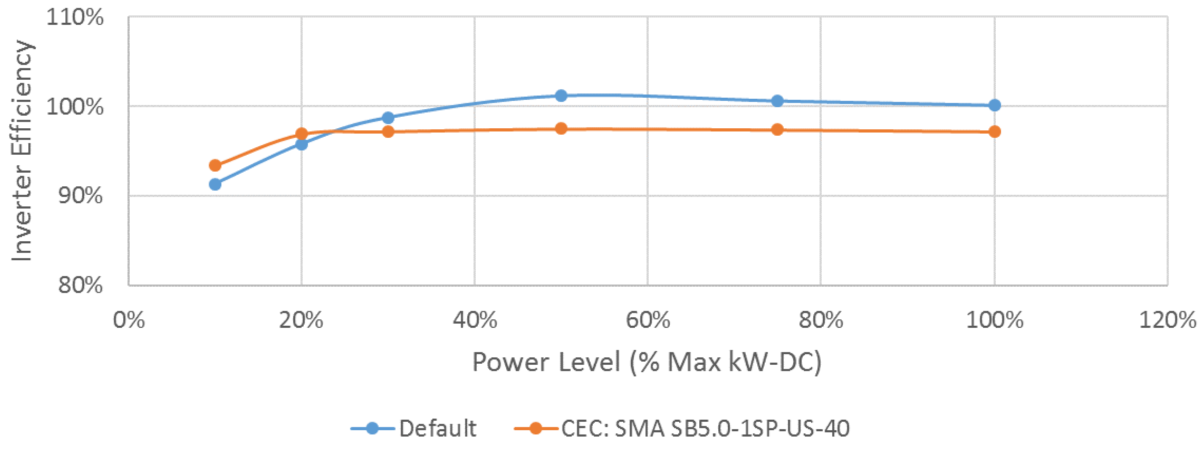

Ensemble Methods

Additional work focused on improving the operational SolarAnywhere (SA) forecast models at both the short-term (hour ahead) and longer-term (day ahead) time horizons by using advanced ensemble methods leveraging forecasts from multiple sources.

Representative PV System Fleets

As the number of behind-the-meter (BTM) PV systems in California continues to grow, tracking the capacity and forecasting the output from those systems becomes more important to grid operators and balancing authorities. At the same time, while simulating and aggregating power output from individual systems provides greater accuracy, it also requires ever-increasing computing resources. However, system capacity from multiple systems in nearby locations with similar orientations can be combined to create representative systems, thereby reducing the number of distinct systems in a fleet to be simulated while retaining the diversity of locations and orientations that characterize the fleet's power production. While these "Representative PV Fleets" introduce some level of error into the modeling process, the decrease in simulation time may prove to be a worthwhile trade-off. In addition, representative fleet concepts can be applied in a top-down manner to extrapolate PV fleet production in areas where the detailed specification of individual resources is unknown.

Within this project, Clean Power Research produced behind-the-meter (BTM) PV fleet power forecasts every 30 minutes for five load regions[2] in the territories of California's three investor-owned utilities (IOUs). These five load regions were identified by the California ISO. These forecasts are produced using satellite-derived irradiance values from SolarAnywhere® at 1 km x 1 km spatial resolution. System specifications such as latitude, longitude, tilt and azimuth, PV module and inverter efficiency ratings, obtained from IOUs, the California Energy Commission (CEC) and the California Solar Initiative (CSI), are used to model power output from approximately 186,000 systems. The power output from these individual systems is then aggregated to provide fleet power output. These systems, however, only represent about 43% of the total systems online.

According to the California Public Utility Commission's (CPUC) California Solar Statistics web site,[3] more than 440,000 behind the meter PV systems are currently installed in California IOU territories and that number appears to be growing steadily, with more than 30% of the systems installed in 2015 (see Figure 14). In addition to the increased computing horsepower required to model such large numbers of systems, specifications for the systems in the publicly available data are inexact or missing altogether. For example, locations are anonymized by providing only the systems' zip code. Furthermore, system orientation (tilt and azimuth) is only available for 30% of the systems. Rather than creating generic systems and guessing at their orientation and exact location, we currently scale the modelled PV fleet power output on the assumption that the locations and system orientations of new systems will have a distribution similar to that of the current fleet captured in PowerClerk. In a way, this is one of the simplest methods for creating a bottom-up representative PV fleet.

[2] These regions are San Diego Gas and Electric (SDG&E), Southern California Edision (SCE) Inland, Southern California Edision (SCE) Coastal, Pacific Gas & Electric (PG&E) Bay Area, and Pacific Gas & Electric (PG&E) Non-Bay Area

[3] https://www.californiasolarstatistics.ca.gov/

Figure 14. Cumulative Installed BTM Systems in California

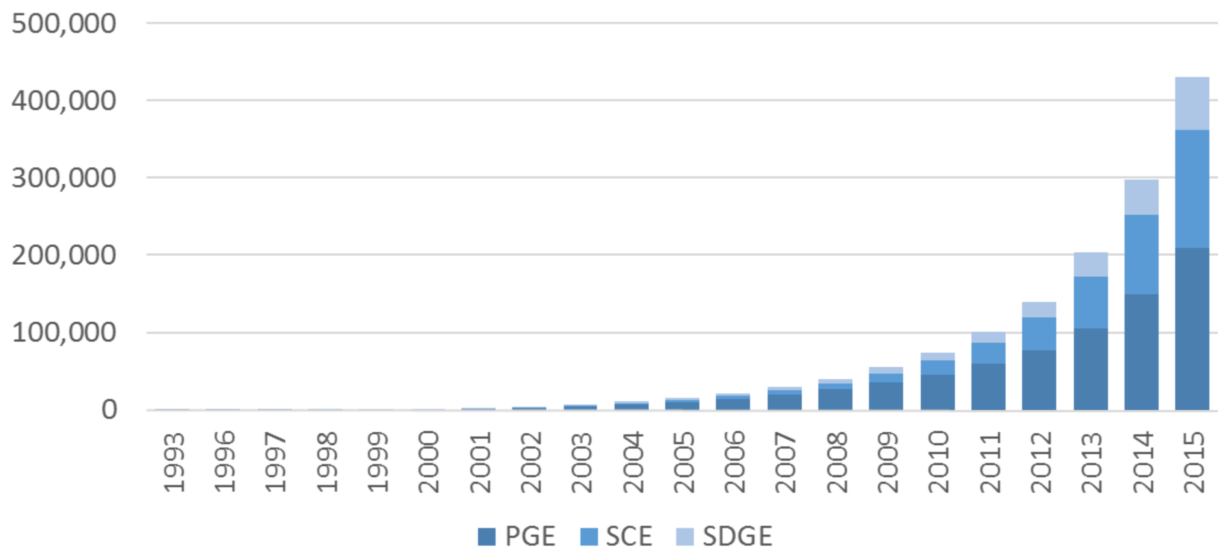

Methodology

To investigate methods for reducing the computational resources required for modelling large PV fleets and to better quantify the margin of error that might be introduced by generalizing the locations and orientation of systems in such fleets, we created various representative fleets using the CSI systems in the PG&E Non-Bay Area load region behind-the-meter fleet as a baseline for comparison. This baseline fleet, consisting of the 34,562 PV systems mapped in Figure 15, is spread across a large portion of California and has a wide variety of system orientations. The 30-minute power output from each of these systems was simulated for a one-year period from January 1, 2014 through December 31, 2014 and the results were aggregated to produce the baseline fleet output every 30 minutes during the period. "Bottom-up" representative fleets were created by binning the capacity in a baseline fleet with known system specifications. The relative Mean Absolute Error (rMAE)[4] for each of the representative (test) fleets was calculated compared to the output of this baseline (reference) fleet.

[4] Thomas E. Hoff, J. K. (2012). REPORTING OF IRRADIANCE MODEL RELATIVE ERRORS. Proc. ASES Annual Conference. Raleigh, NC: American Solar Energy Society.

Figure 15. PG&E Non-Bay Area CSI Systems

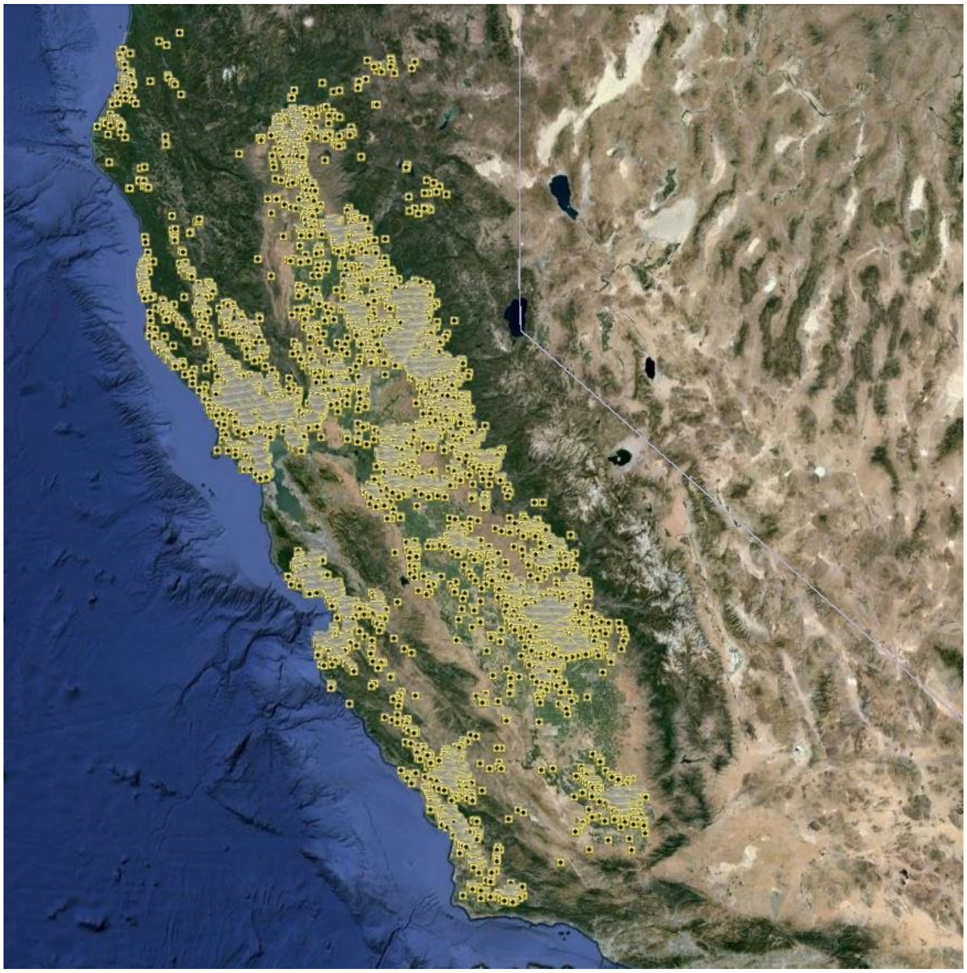

Geographic Bin Selection

Site locations for the Representative fleets were selected using one of two methods. In the first method, a grid was created, by evenly dividing the rectangle bounding the systems in the baseline fleet. Six different spatial resolutions were used with this method. The first five spatial resolutions tested were 1.6° latitude and longitude (approximately 160 x 160 kilometers), 0.8°, 0.4°, 0.2°, and 0.1° latitude and longitude (approximately 10 x 10 kilometers).

Capacity for each system was mapped to the nearest location on the grid, then further binned by orientation (tilt, azimuth, and tracking). Table 1 shows the selected geographic bins and how they were combined with the orientation bins (described in the next section) to create the systems in each representative fleet. The map in Figure 16 shows the locations for one representative fleet. Multiple systems were created at each location sized to represent the capacity of the actual systems in each orientation bin.

For example, in the fleet shown in Figure 16, among the 32 systems created at 37.942° latitude, -120.593° longitude, there would be a south-facing system, with a 22.5° tilt, rated at 798.8 kW AC.

Figure 16. Representative Fleet with system locations at 0.4° latitude/longitude spacing.

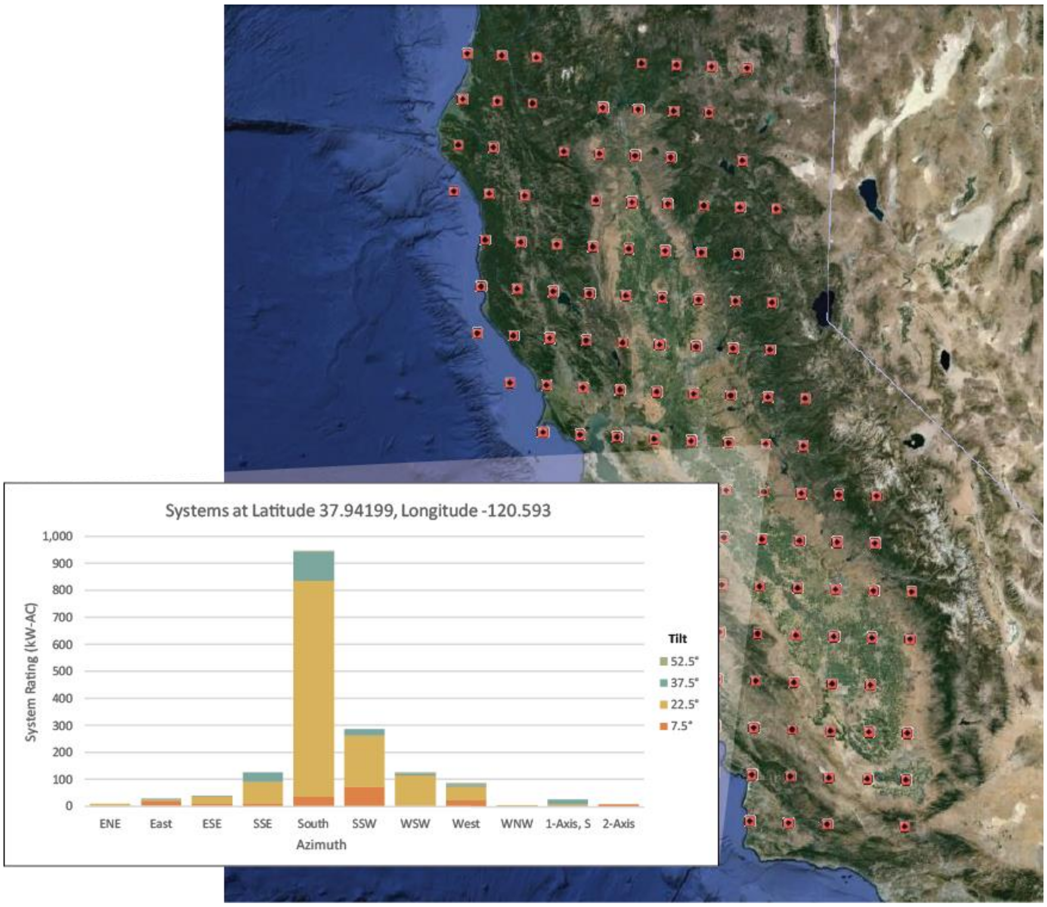

For the sixth spatial resolution, the fleet we created had all capacity mapped to a single location, then binned by orientation. The location selected was the capacity-weighted geographic center of the baseline fleet.

Note that SolarAnywhere Enhanced Resolution data has a spatial resolution of 1 km x 1 km, so the systems in the baseline fleet are already implicitly binned by location to the nearest 1 km, with no binning by orientation. This implicit binning has the effect of reducing the number of actual locations from 35,562 to 10,866.

With the second method for representative fleet creation, we mapped each system's capacity based on the zip code of the PV site and used the geographic center of the zip code as the location, then further binned the capacity based on the system's orientation. System locations for these fleets are shown in

Figure 17. Finally, in a variation of the zip code based method, we mapped each system's capacity based on the zip code, but created a single system with all of the zip code's capacity, located it at the geographic center of the zip code, and used the baseline fleet's capacity-weighted azimuth and tilt (17° and 175°, respectively) as that system's orientation.

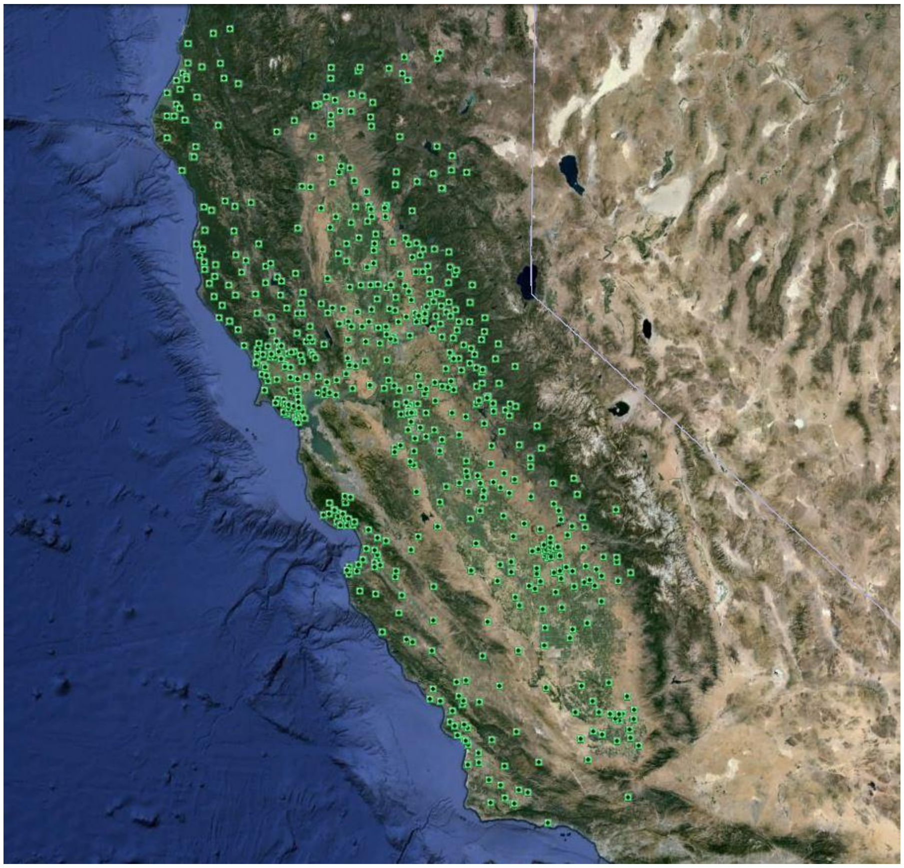

Figure 17. Zip code based locations

Orientation Bin Selection

In addition to binning system capacity by location, when creating the representative fleets, the capacity of the actual systems was binned by tilt, azimuth and tracking to capture the diversity in system orientations at each location typical in large PV fleets. System orientation bins were based on 10° azimuth and 5° tilt increments (648 bins), 20°, azimuth and 10° tilt increments (162 bins), or 30° azimuth and 15° tilt increments (72 bins). Dual-axis tracking systems, where azimuth and tilt vary continuously throughout the day, constituted an additional bin.

For each location, the capacity for each array[5] was assigned to the bin that most closely matched the azimuth and tilt of that array. For example, in the case where 30° azimuth and 15° tilt increment bins were used, capacity for arrays with azimuths that were +/- 15° from south (165° to 195°) with tilts between 22.5° and 37.5° would have been added to the 180° azimuth/30° tilt capacity bin.

System Creation and Simulation

Once the total capacity was determined for each location/orientation bin, systems were created with the appropriate capacity. Table 4 shows the number of systems in each of the 22 representative fleets created by combining spatial and orientation bins.

Table 4. Number of systems in Representative Fleets by spatial resolution and orientation bin

	Number of systems			
Spatial resolution	Azimuth/Tilt Increments			
	10°/5°	20°/10°	30°/15°	Single Orientation
Single Location	362	130	73	-
Zip Codes	15,305	8,824	6,302	601
160 x 160 km	1,926	730	418	-
80 x 80 km	4,020	1,707	1,025	-
40 x 40 km	6,818	3,329	2,077	-
20 x 20 km	11,119	6,091	4,022	-
10 x 10 km	16,276	9,986	6,841	-

[5] Capacity is analyzed at the array level rather than the system level in order to properly account for systems with multiple arrays.

The maximum power rating for the inverter used for each system was calculated based on the capacity-weighted DC to AC ratio for the baseline fleet of 1.027 as recorded for acutal systems. Inverter efficiency for each system was set to 96.2% - also based on the capacity-weighted inverter efficiency rating of the baseline fleet – and other DC losses were set to 11% - once again using the capacity-weighted DC losses of the baseline fleet. After creating the systems, power output for each system was simulated for every 30-minute period from January 1, 2014 through December 31, 2014 and the results were aggregated to produce 30-minute interval fleet power. Those results were then compared to the output from the baseline fleet.

Effect of Spatial Resolution and Orientation Bin Count on Relative Mean Absolute Error

The amount of error introduced by using bottom-up representative PV fleets with regular geographic dispersion, rather than fleets consisting of individual systems with exact system specifications varied from 4.2% for the coarsest spatial resolution and smallest number of orientation bins, to 1.1% for the fleet with 10 km x 10 km spatial resolution and the largest number of orientation bins. As shown in Figure 18, the greatest impact on error was due to spatial resolution, rather than the number of orientations considered. However, further increases in spatial resolution would likely have proportionally less impact as we approach the native resolution of the satellite data.

Figure 18. Relative Mean Absolute Error for Representative Fleets

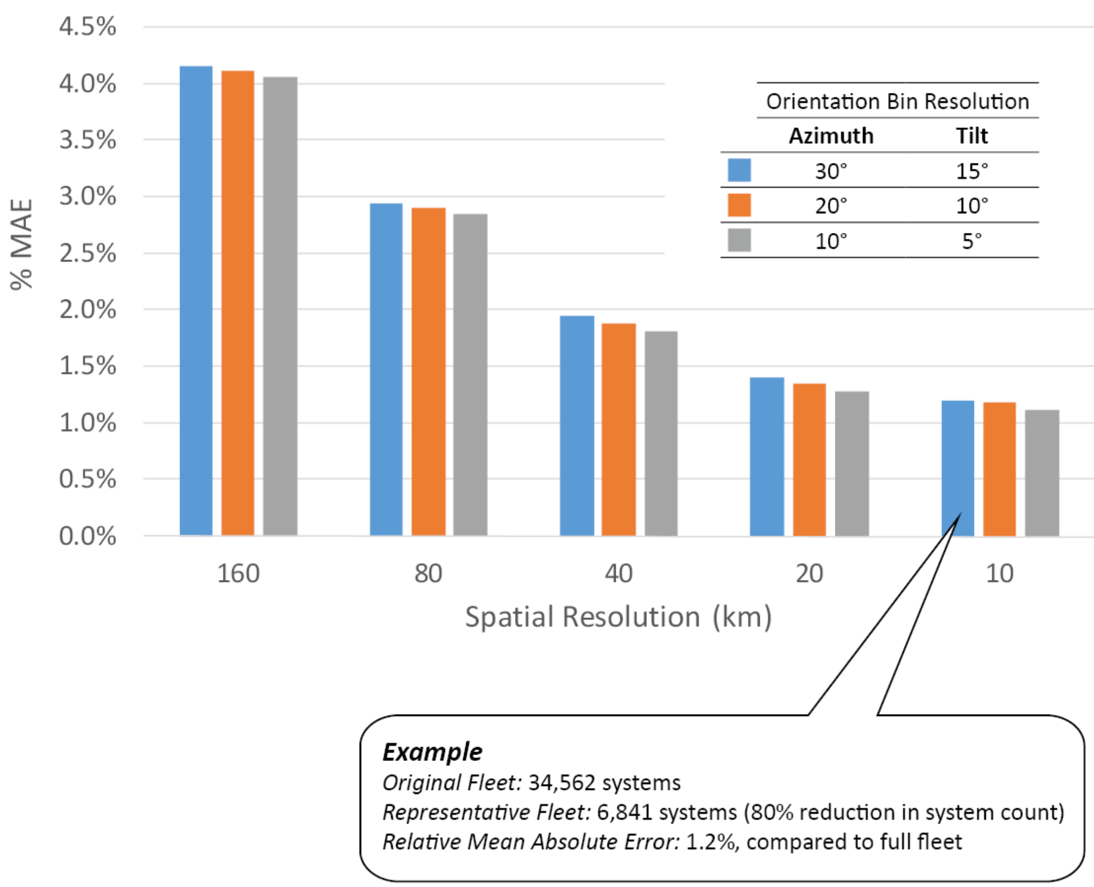

The relative Mean Absolute Error for the representative fleets with a single location had significantly higher error than the other fleets, ranging from 10.1% to 10.7%.

The representative fleet based on multiple orientations at each zip code fared reasonably well with rMAE ranging from 2.2% to 2.4%. However, the zip code based fleet that used a single orientation at each zip code had a much higher rMAE at 6.6%. While a multi-orientation zip code based fleet may be appropriate when exact system locations are unknown, performance is only slightly better than the representative fleets with the highest number of orientation bins and spatial resolution, and error is approximately double.

The graphs in Figure 19 show the correlation between the 30-minute power values for selected representative fleets versus the baseline fleet. At the same spatial resolution, there was little difference between fleets with different numbers of orientation bins, so these were omitted.

Figure 19. Thirty-minute Power Value Correlation

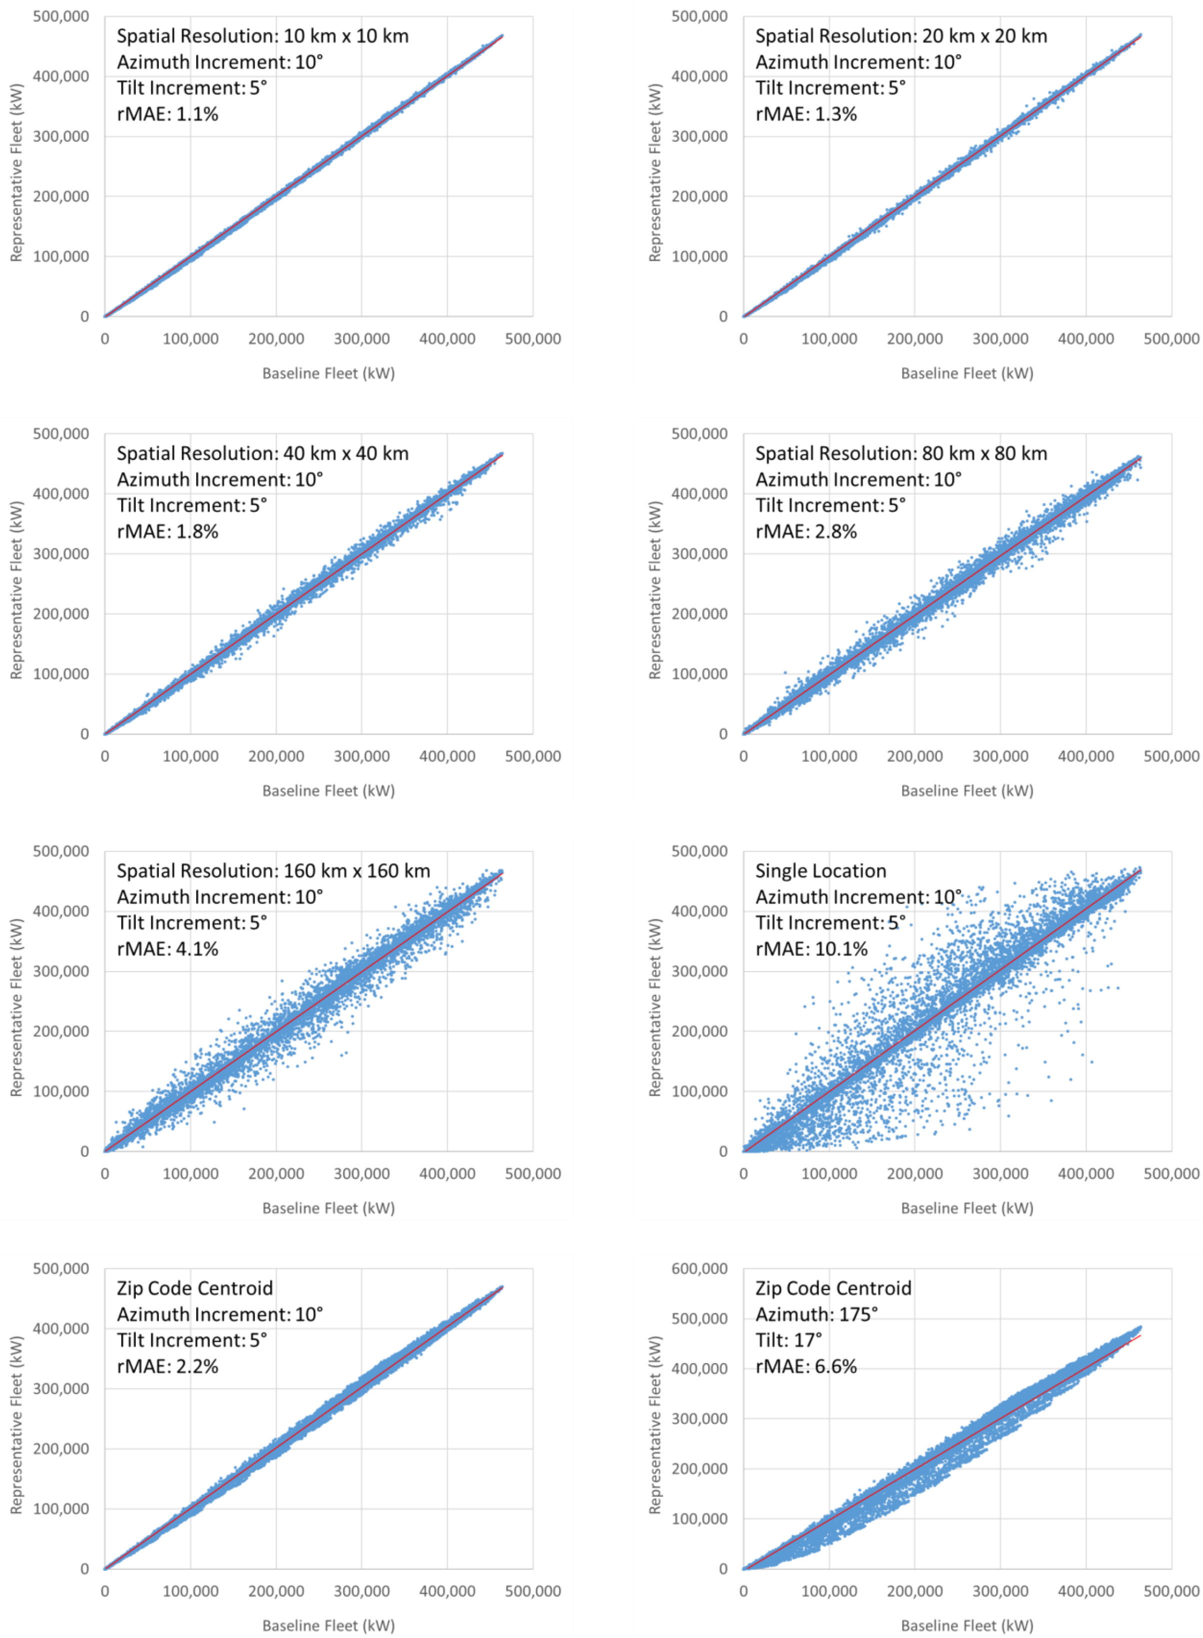

Performance Benefits

In general, the amount of time it takes to simulate a PV fleet scales linearly with the number of systems to be simulated with approximately an 8% additional reduction in time when there are many systems to be simulated at a single location. Although simulation times for the representative and baseline fleets were tracked, we feel that the actual times, which vary greatly depending on computer system load and data transmission speeds over the Internet, should not be considered when evaluating performance. Instead, based on the number of locations and systems, Table 5 shows the hypothetical reduction in time required to simulate these representative fleets relative to the baseline fleet. Even at the highest spatial resolution and the largest number of orientation bins evaluated, simulation times could be reduced by 63.5%. However, fewer orientation bins at the same high spatial resolution adds only 0.1% error, while yielding an 83.4% reduction in simulations time.

Table 5. Estimated Reduction in Simulation Times for Representative Fleets Relative to Baseline Fleet

	Azimuth/Tilt Increments			
	10°/5°	**20°/10°**	**30°/15°**	**Single Orientation**
Single Location	99.2%	99.7%	99.8%	
Zip Codes	66.3%	80.0%	85.3%	98.6%
160 x 160 km	95.9%	98.4%	99.1%	
80 x 80 km	91.4%	96.3%	97.7%	
40 x 40 km	85.5%	92.9%	95.5%	
20 x 20 km	75.7%	86.3%	90.7%	
10 x 10 km	63.5%	76.8%	83.4%	

Examining Individual Days

This study examines rMAE over an entire year. However, when forecasting PV fleet output, much shorter time scales are involved and it's useful to look at individual days to get an idea of some of the shortcomings of certain types of representative fleets.

In Figure 20 , for example, the June 21 output from the single location fleet drops dramatically in the afternoon due to clouds at that location, while the single-orientation zip code fleet shows the higher, narrower curve that's typical when fleets do not include non-optimal orientations. Apart from their higher rMAE, these types of representative fleets should be avoided if daily production shape is important.

Figure 21, which shows fleet output for December 21, 2014 exhibits similar problems for these two fleets. Finally,

Figure 22 illustrates the importance of spatial resolution on partially cloudy days.

Figure 20. Baseline and Representative Fleet Output for June 21, 2014

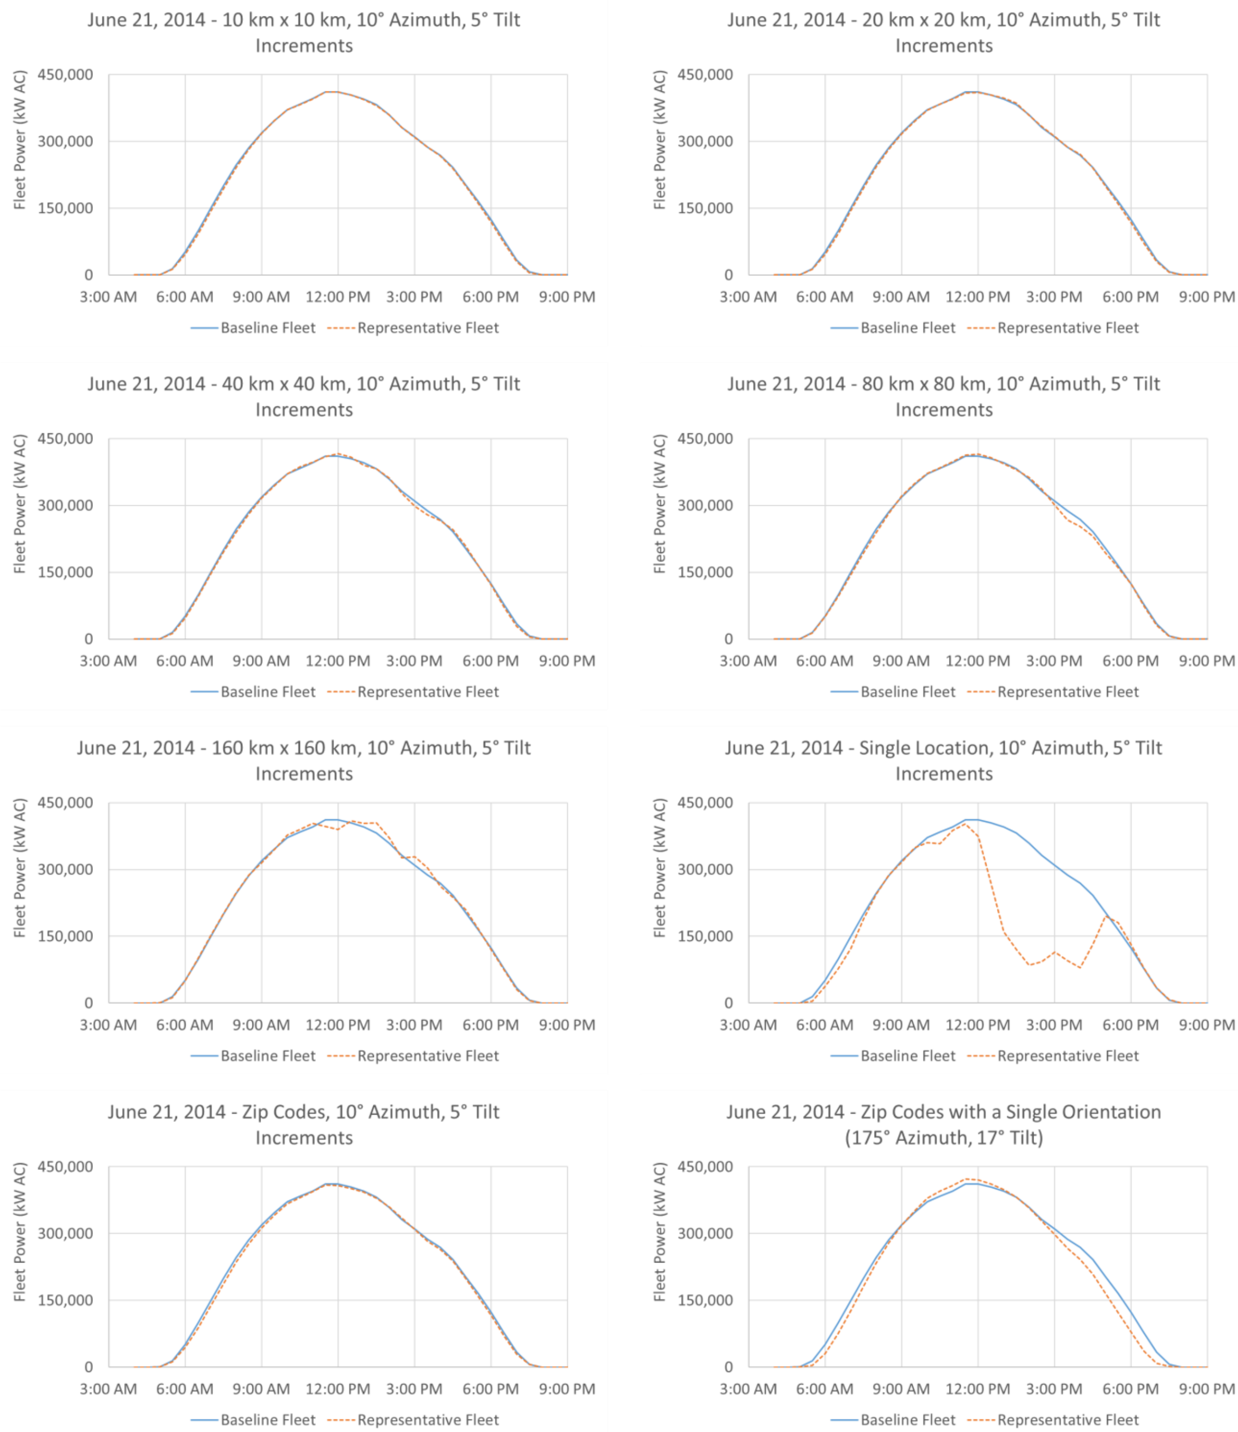

Figure 21. Baseline and Representative Fleet Output for December 21, 2014

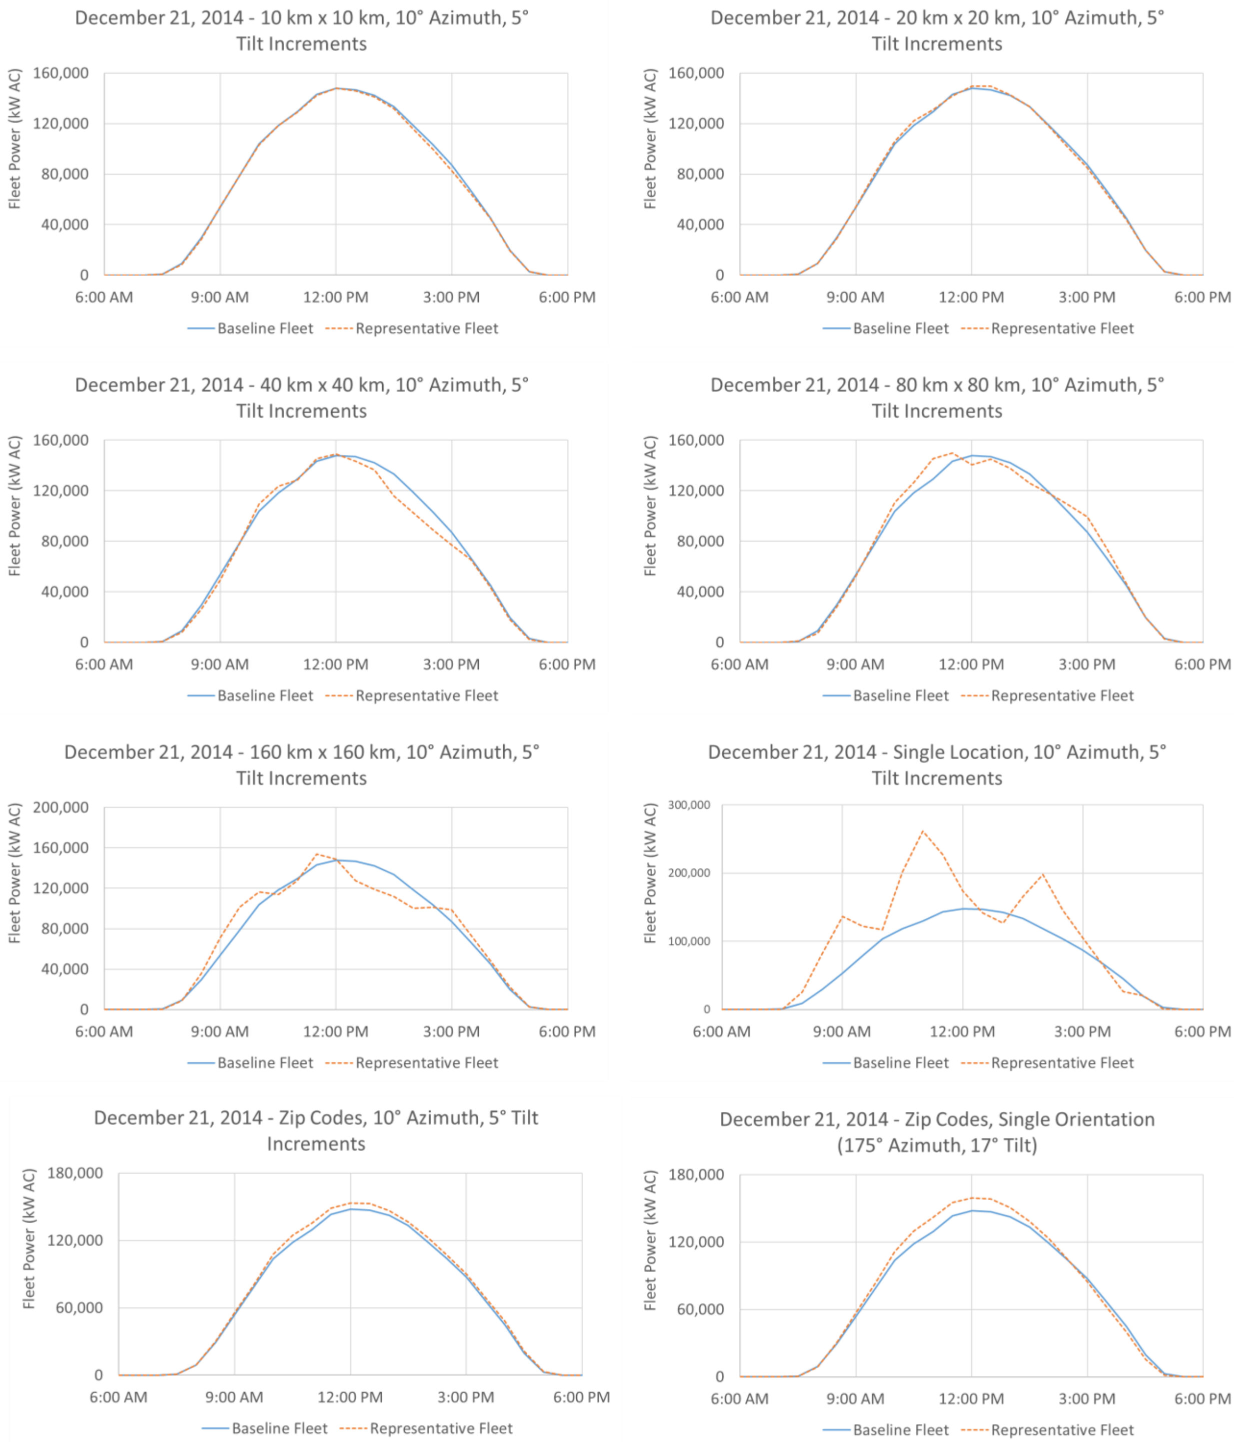

Figure 22. Baseline and Representative Fleet Output for February 26, 2014

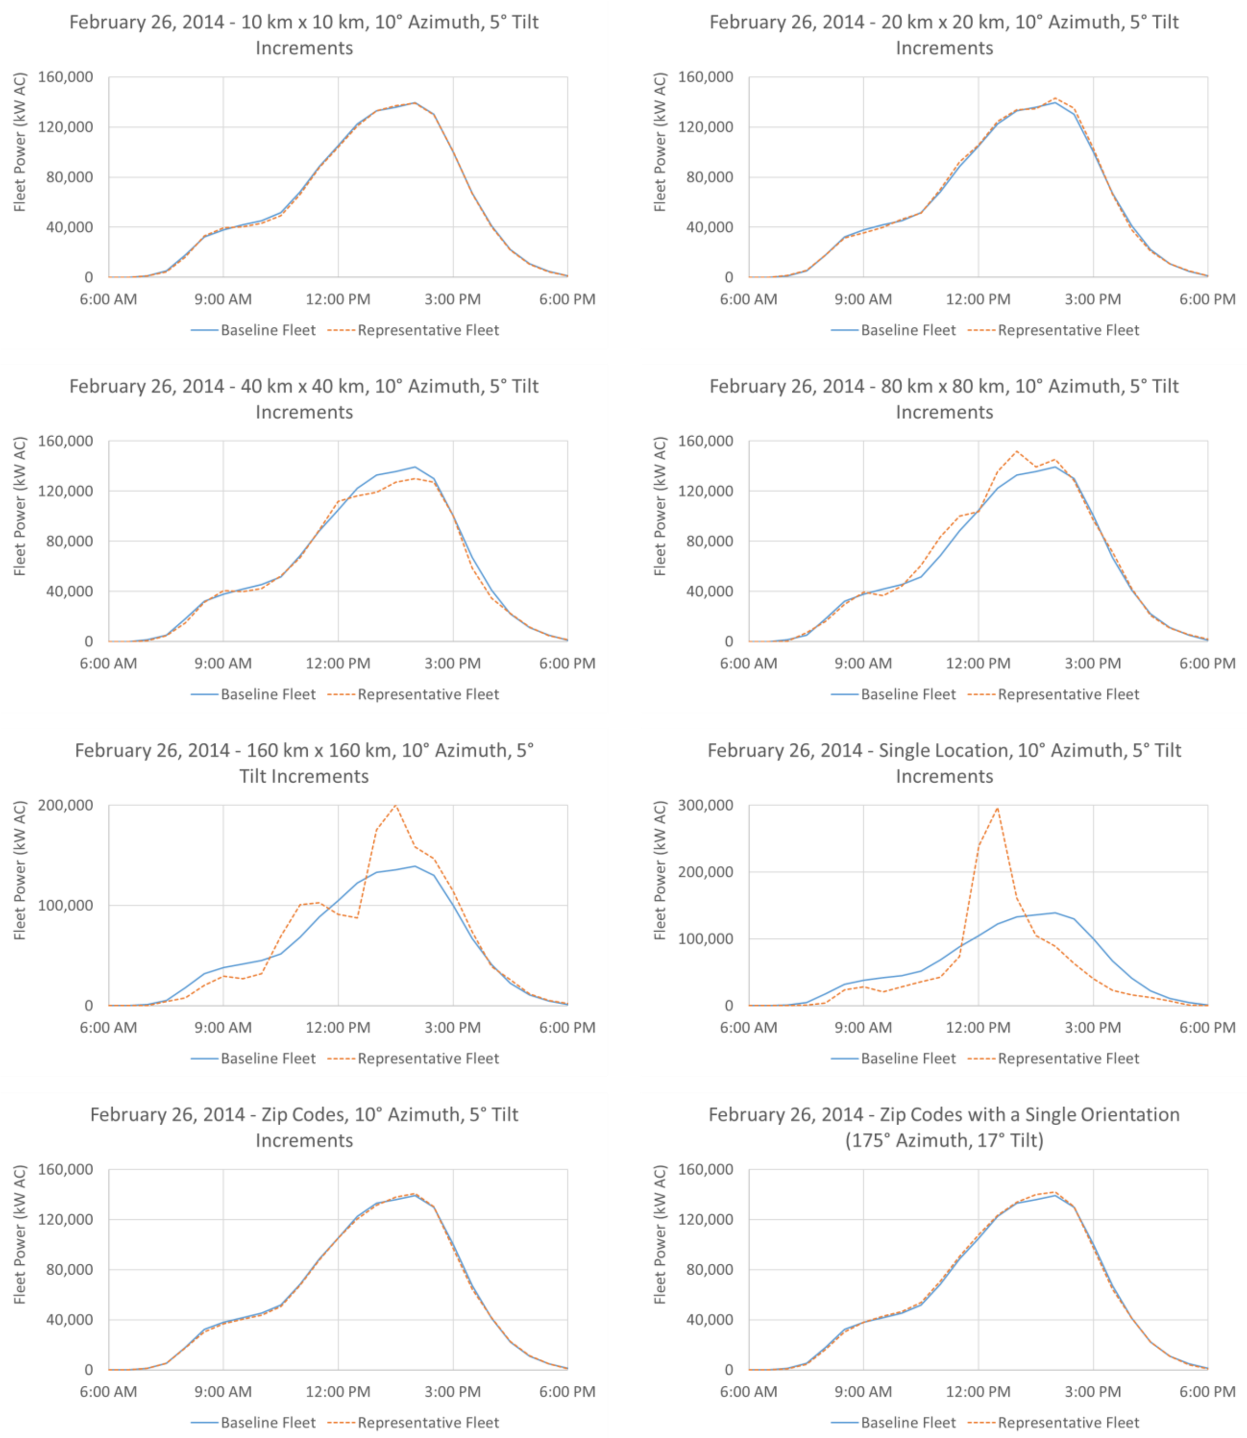

Bottom-up representative PV fleets, created by generalizing the location and orientation of a set of individual systems with known specifications, can be used in modelling fleet output to reduce computing resource requirements by more than 80%, while introducing as little as 1.2% rMAE on an annual basis. By applying scaling factors to the known historical California BTM PV fleet, this was the the approach used for the trial of this project with CAISO.

Zip code based representative fleets, which make use of known individual system orientation data can reduce computing resource requirements by more than 65%, while introducing as little as 2.2% rMAE on an annual basis.

Although they can be simulated very quickly, representative fleets that make use of a single location exhibit more than 10% rMAE and have a fairly inaccurate power production curve on a daily basis. At 6.6% rMAE, zip code fleets that use a single system orientation have less error than single location fleets, but typically exhibit a narrower daily production curve with a higher peak. Applying this approach correctly as the size of the PV fleet continues to grow will need to account for loss of accuracy and ensure that any avoidable error isn't introduced.

Dynamic Regional Fleet Capacity Updates

The equipment comprising BTM PV systems do not always remain in service on a continuous basis. Owners sometimes replace system components such as the inverter. They also may add or remove modules. Utility-scale systems are sometime built in phases, with capacity growing over time. In addition, outages – both planned and unplanned – can cause capacity to drop. As part of this project, CPR has implemented the ability to track changes in system capacity over time and use that information when simulating system output.

California leads the nation in BTM PV installations – systems whose production, for the most part, is not tracked by utilities. In order to provide an accurate estimate of the power produced by these systems, it's important to have detailed information about each system's configuration - where it's located, the date it was installed, the orientation of each array in the system, the models and quantities of modules and inverters that are installed, and the elevation of any solar obstructions, such as buildings and trees surrounding the system that are above the bottom edge of the panels. When combined with accurate weather data, whether historical or forecast, this system configuration information can be used with modeling software to produce a reasonably accurate estimate of the system's production.[6]

Beginning in 2007, detailed specifications for systems incentivized under the California Solar Initiative (CSI) were collected using PowerClerk, an online software service from Clean Power Research. Clean

[6] https://www.nrel.gov/analysis/sam/pdfs/2008_sandia_ieee_pvsc.pdf

Power Research also collected specifications for non-CSI BTM systems incentivized under Self Generation Incentive Program (SGIP), the Emerging Renewables Program (ERP). By the end of 2014, CPR team had collected detailed specifications for more than 140,000 CSI systems and 43,000 non-CSI systems as shown in Figure 18. These system specifications, with a combined capacity over 2.1 GW, were used in the creation of five of the BTM fleets used in the CAISO forecast.

Figure 23. Capacity of California's BTM PV Fleet as tracked by Clean Power Research

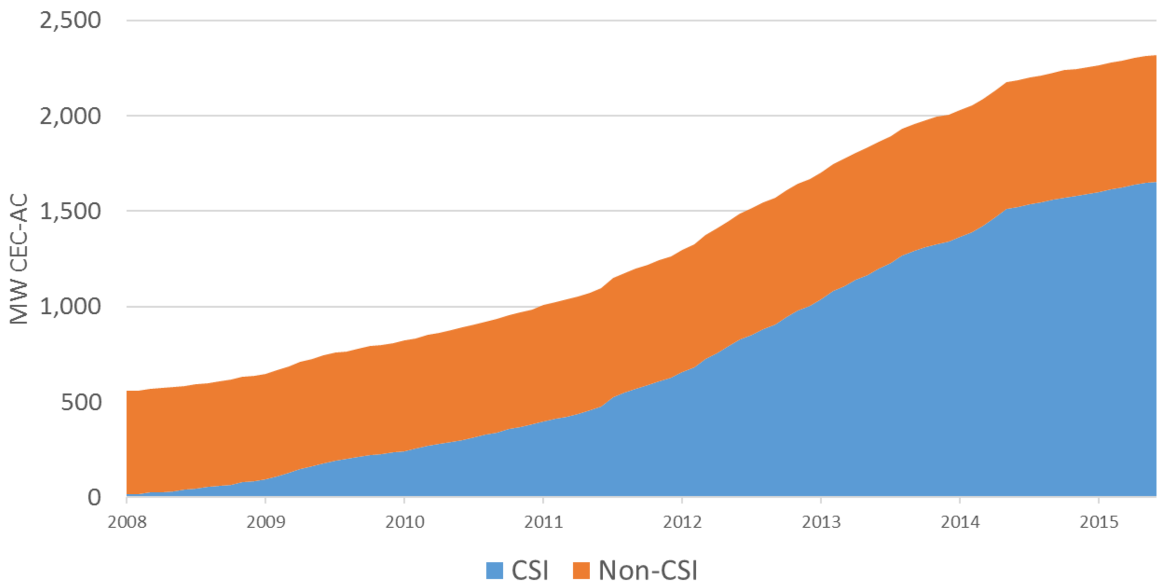

As incentives available through the California Solar Initiative began to run out, an increasing number of PV systems were being installed that were not tracked in PowerClerk. This was especially true in PG&E and SDG&E territories, and CPR began to look for ways to keep the capacity of its PV fleets up to date.

Initially, CPR experimented with capacity data provided to CAISO by the IOUs. However, the poor quality (e.g. redundant and missing data) of data prevented its use. Next, we obtained market research data from GTM Research. Although the data quality was better, reporting was only by quarter for the entire state, rather than by IOU or region.

In July 2015, the CPUC began posting monthly editions of the NEM Currently Interconnected Data Set (CIDS) on the California Solar Statistics web site.[7] Although this data set is useful for estimating total installed capacity, its usefulness as a source for detailed system specifications is limited due to anonymized locations (only zip code is provided) and missing data (tilt and azimuth) is available for only 30% of the systems.

[7] http://www.californiadgstats.ca.gov/downloads/

Figure 24. California BTM Fleet Capacity 2008 to 2016

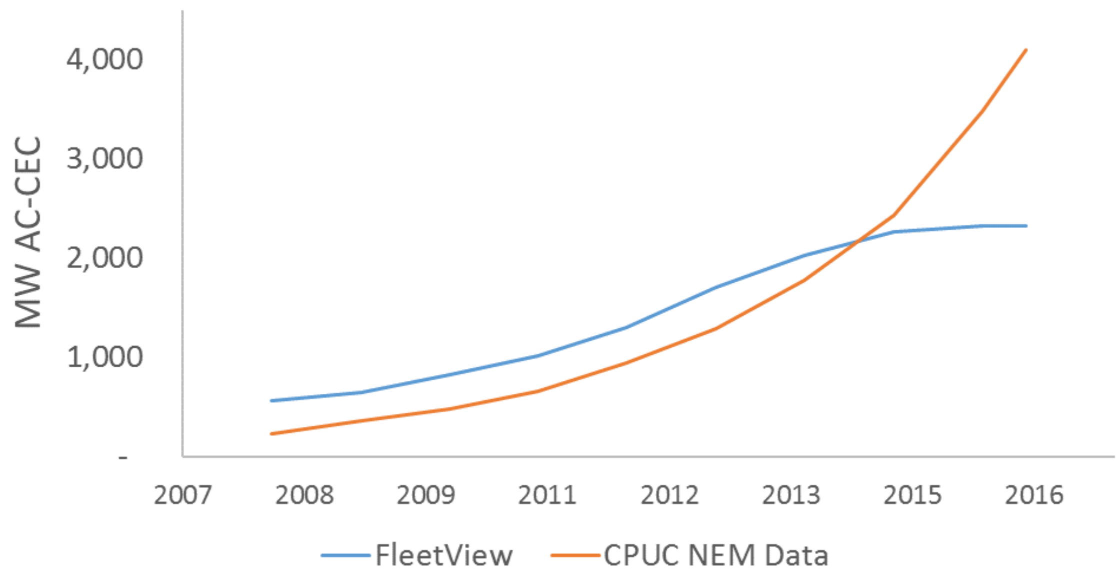

By using capacity data from the NEM Interconnection Applications Data Set in conjunction with the detailed system specifications in FleetView, CPR was able to develop the time-dependent scaling factors that were applied to the historical simulations of the five CAISO behind-the-meter fleets whose output is used to train Itron's load forecasting software. Furthermore, the scaling factors are now automatically recalculated when CIDS updates are published, then projected to future dates and applied to the CAISO BTM fleet forecasts.

Determination of Fleet Historical Scaling Factors

Since, the NEM Currently Interconnected Data Set contains only those systems that are thought to be currently online, and not those that have been decommissioned, we must look to the NEM Interconnection Applications Data Set to get a complete picture of historical capacity over time. Figure 19 depicts the complete picture of historical capacity.

Determination of Fleet Forecast Scaling Factors

For forecasting, a linear formula for the scaling factor growth trend is derived for each CAISO zonal fleet, defined by this project, using CIDS and FleetView capacities for the most recent two months. Monthly scaling factors are then calculated dynamically using those formulas, based on the time elapsed since the beginning of the growth trend period and applied to the PV production forecast. Figure 25 shows a plot of the trend line for September through December 2015. Technically, coefficients only need to be updated if the growth rate of the scaling factors change. However, Clean Power Research monitors the

Currently Interconnected Data Set and automatically updates the coefficients whenever the CIDS is updated.

Figure 25. Projected PG&E Non-Bay Area Scaling Factors

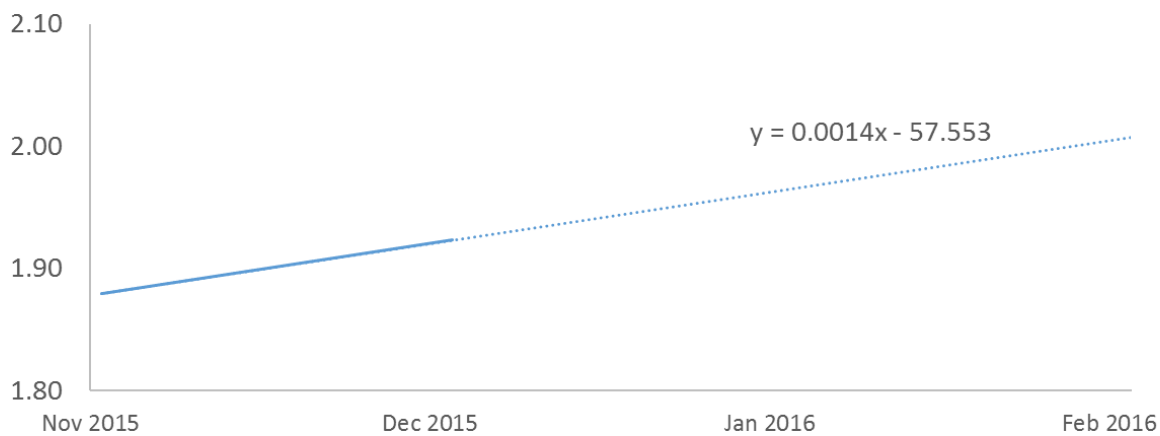

Table 4 shows the capacity (MW_{CEC-AC}) of each of the five CAISO BTM fleets from both the Currently Interconnected Data Set (blue) and FleetView (green) for the 6-month period from April to September 2015. The yellow columns show the ratio of CIDS to FleetView capacity.

Table 6. Installed PV Capacity (MWCEC-AC) Currently Installed Data Set vs FleetView

	Apr 2015			May 2015			Jun 2015			Jul 2015			Aug 2015			Sep 2015		
	CIDS	FleetView	Ratio	CIDS	FleetView	Ratio	CIDS	FleetView	Ratio	CIDS	FleetView	Ratio	CIDS	FleetView	Ratio	CIDS	FleetView	Ratio
PG&E Bay Area	455.8	396.8	1.14868	465.2	396.9	1.17212	474.4	398.9	1.18924	482.4	400.0	1.20576	493.2	400.0	1.23282	502.8	400.0	1.25681
PG&E Non-Bay Area	961.7	760.0	1.26540	992.2	761.1	1.30363	1,027.5	762.4	1.34770	1,056.7	765.0	1.38126	1,089.1	765.0	1.42372	1,123.9	765.0	1.46914
SCE Coastal	477.6	460.7	1.03671	490.1	462.4	1.05993	501.8	462.7	1.08463	514.8	465.0	1.10707	524.8	465.0	1.12868	538.0	465.0	1.15714
SCE Inland	459.3	437.8	1.04908	477.6	441.3	1.08224	494.2	442.1	1.11770	509.8	442.6	1.15185	524.4	442.6	1.18493	543.3	442.6	1.22761
SDG&E	364.4	245.8	1.48251	374.3	250.6	1.49344	390.0	252.8	1.54251	402.0	254.0	1.58285	402.3	254.0	1.58395	413.0	254.0	1.62583

The scaling factors and consequent fleet ratings calculated for November 20, 2015 6:00 PM, for example, would be as follows:

Fleet	Scaling Factor	Scaled CEC-AC Capacity (MW)
Non-Metered: PG&E Bay Area	1.3002497	520.2
Non-Metered: PG&E Non-Bay Area	1.5414869	1,179.2
Non-Metered: SCE Coastal	1.1985151	557.3
Non-Metered: SCE Inland	1.2929136	572.2
Non-Metered: SDG&E	1.66500428	422.0
Total		3,251.8

Historical Behind-the-Meter PV Fleet Production Modeling

One of the available inputs to the load forecasting model produced by Itron for the California ISO, is estimated PV production. For this project, the California ISO has identified five zonal fleets for which they require separate forecasts: PG&E Bay Area, PG&E Non-Bay Area, SCE Coastal, SCE Inland, and SDG&E.

Using CSI data, CPR had previously created fleets of individual systems for each of the five zones. As more non-CSI systems began to come online, those fleets were no longer representative of the actual California PV fleet capacity. However, the large number of systems in these fleets did provide a representative sample of geographic distribution, diversity of orientations, and other system characteristics such as DC to AC sizing and inverter efficiency.[8] Therefore, CPR simulated historical PV fleet production for each of those five fleets for the period from Jan. 1, 2010 through December 31, 2015 using SolarAnywhere Enhanced Resolution data, which has a temporal resolution of 30 minutes. Interpolation was used to calculate 15-minute interval values. PV production was then scaled to match monthly capacity derived by combining the non-CSI portion of the NEM Currently Interconnected Data Set (CIDS) capacity as of June 30, 2015 with CSI capacity data obtained from PowerClerk. Scaling factors

[8] This method assumes that geographic diversity of capacity and other system characteristics remained unchanged beyond the time at which systems incented under CSI began to comprise a smaller share of the total fleet.

were linearly interpolated for the periods between each month. The CSV-format PV production data files for each fleet were made available to Itron via a File Transfer Protocol (FTP) server. The charts below show the scaled versus unscaled PV production for each of the five CAISO fleets. As indicated in the charts, CPR's estimate of non-CSI capacity before 2014 was higher than that reported in the CIDS. These results make the assumption that the CIDS is the best source for capacity available.

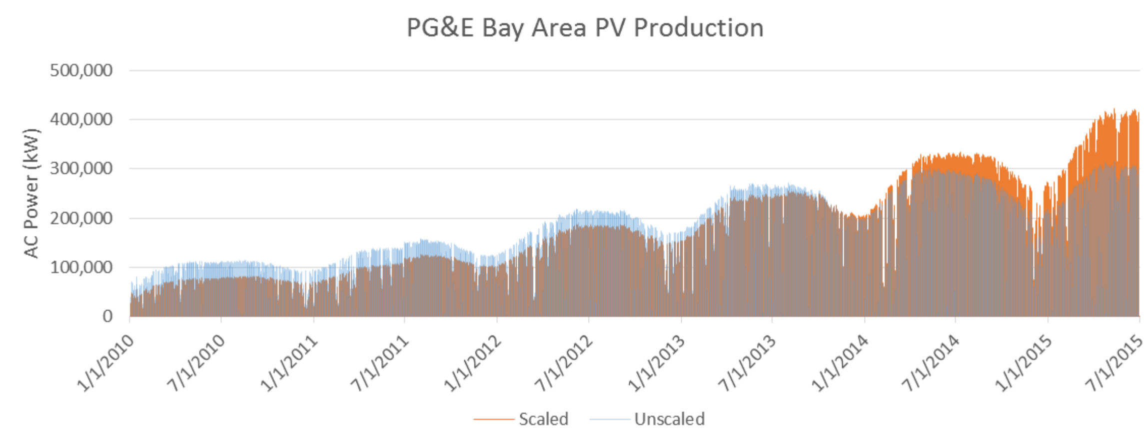

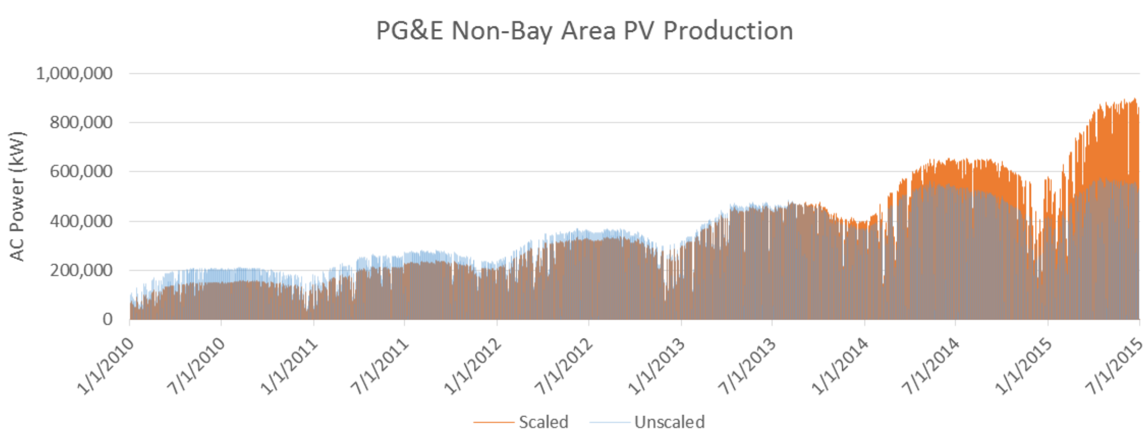

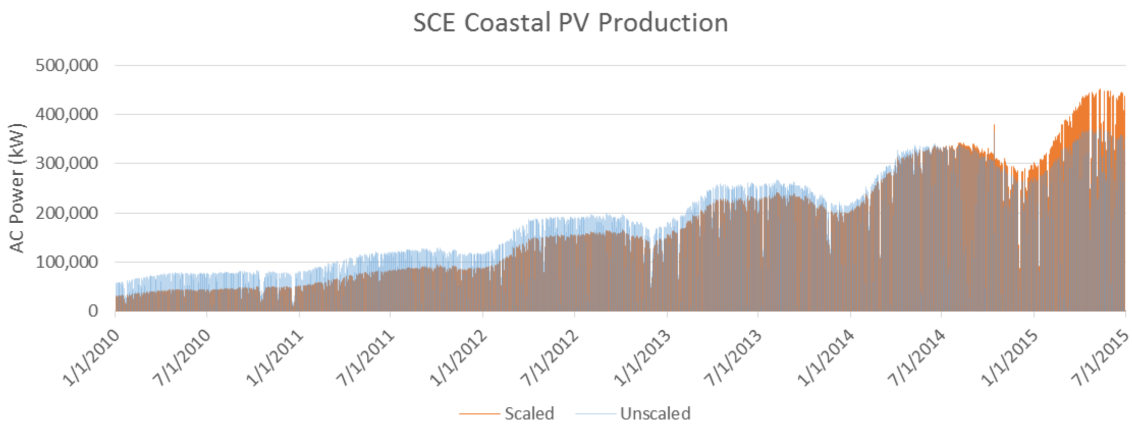
SCE Coastal PV Production
AC Power (kW)
500,000
400,000
300,000
200,000
100,000
0
1/1/2010
7/1/2010
1/1/2011
7/1/2011
1/1/2012
7/1/2012
1/1/2013
7/1/2013
1/1/2014
7/1/2014
1/1/2015
7/1/2015
Scaled
Unscaled

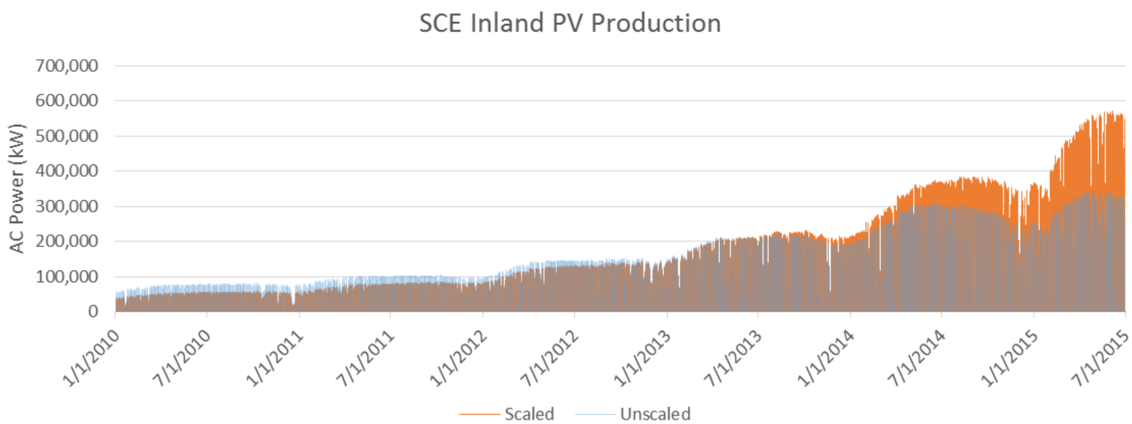
SCE Inland PV Production
AC Power (kW)
700,000
600,000
500,000
400,000
300,000
200,000
100,000
0
1/1/2010
7/1/2010
1/1/2011
7/1/2011
1/1/2012
7/1/2012
1/1/2013
7/1/2013
1/1/2014
7/1/2014
1/1/2015
7/1/2015
Scaled
Unscaled

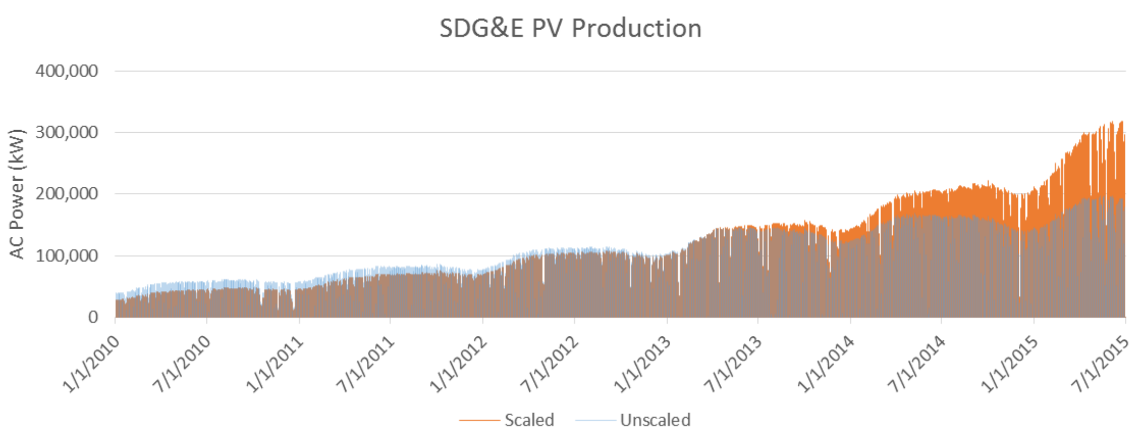
SDG&E PV Production
AC Power (kW)
400,000
300,000
200,000
100,000
0
1/1/2010
7/1/2010
1/1/2011
7/1/2011
1/1/2012
7/1/2012
1/1/2013
7/1/2013
1/1/2014
7/1/2014
1/1/2015
7/1/2015
Scaled
Unscaled

Robustness of the CAISO Forecast Delivery

To improve availability and accuracy of the forecasts provided to CAISO during this project, Clean Power Research made the following operational changes to the manner in which forecasts are produced and delivered:

1. Add distributed processing support so that larger fleets could be processed more quickly.
2. Run forecasts on a scalable cloud-based platform that permits better monitoring and increases reliability.
3. Automatically detect changes to the NEM Currently Interconnected Data Set and, trigger an automatic update of the scaling factors applied to the baseline fleet output.

Real-time Data Feedback

The use of feedback from real time production data have the potential to also improve forecasts. By using the current conditons and knowledge of the clear sky profile, it would be possible to advance the current observed clearness index along the clear sky profile to produce a "persistence forecast." This is based on the assumption that the cloud conditions will not vary from the current conditions, or in other words that the current conditions will persist. However, obtaining real time data, fast enough for a forecast to be produced and dissemintated for decision making may prove difficult.

The approach taken in this project was to focus on the use of production data from distributed rooftop systems. A large number of systems would needed, particularly if systems did not report data reliably, if they were out of service, or if they were reporting bad data. All of these possibilities were belived to be a factor for distributed systems.

Itron provided near real-time access to data from approximately 30 systems in the Bay Area. Some time delay in the readings was inevitable: the process required data transmission from the PV system itself to Itron's database, followed by ingesting into the forecast system. A proof-of-concept system for retreiveing and ingesting the data was developed and demonstrated, and the process typically took 20 minutes when working correctly.

The process was demonstrated for a period of about one month. Real time data was being collected, evaluated, and used to modify the CPR solar forecast. During this period the modified forecasts were evaluated and small forecast improvements were observed in the modified forecast of PV system power for each of the 30 PV systems. This was expected since persistence, at least for short time scales, has been proven to add skill to the forecast. However, for forecasts beyond three hours no improvement in forecast skill was observed.

After a month, the cellular carrier phased out support of the modems that were used to collect the data. This prevented a comprehensive evaluation. The use of data from these distributed systems is also costly, so a more complete evaluation would not only have to determine whether a forecast improvement were possible on a consistent basis, but also whether any such improvement would justify ongoing maintenance costs at scale. It is not clear how many systems would be required to have a meaningful impact state-wide and if the meaningful impact is driven by a relative number of sites, the solution could become cost-prohibitive.