Conference Proceedings

ASES National Solar Conference 2016
San Francisco, CA, USA July 10-13, 2016

Forecasting Solar Power and Irradiance – Lessons from Real-World Experiences

Stephen D. Jascourt[1], Daniel Kirk-Davidoff[1] and Christopher Cassidy[1]

[1] MDA Information Systems LLC, Gaithersburg, MD (United States)

Abstract

MDA Information Systems, LLC developed a solar irradiance and power forecasting system based on a first principles science foundation employing high-quality scientific datasets such as AERONET and SURFRAD and utilizing the REST2 clear sky model as an underlying basis for the full-sky forecast. Real-time inputs include a diverse multi-model ensemble of numerical weather prediction (NWP) forecasts, ground-based solar monitoring observations and proprietary observations of solar power from client sites, and visible satellite imagery. Forecasts were made for challenging locations where daytime cumulus clouds and occasional storm systems passing through resulted in variability on time scales of minutes, hours, and days.

This paper focuses on lessons learned from our experience with real-world data and real-world power and irradiance forecasts. Topics include quality control of irradiance and power observations, sub-hourly variability and inverter-limited sites, tracking angles for single-axis trackers, and situational bias of NWP forecasts.

Keywords: *Data, ensemble, forecast, irradiance, model, numerical weather prediction, observations, power, quality control, real-time, tracking, variability*

1. Introduction

As the installed capacity of solar power has been rising exponentially, integration on the grid and the effect of solar power on power prices and markets has stimulated interest in forecasts of solar power and irradiance at solar farms and aggregated across collections of utility-scale and behind-the-meter distribution side installations. In some regions, this interest has reached the point where interested parties are procuring forecasts, while in other regions, the need for forecasts is anticipated to be coming soon, stimulating forecast trials and other assessments of how or when forecasts may provide value.

Anticipating this need for forecasts, MDA Information Systems, LLC began developing a forecast system for power and irradiance several years ago based on first principles and has continued to improve and expand its capabilities. Relationships with clients as well as participation in trials and collaborative projects has allowed us to obtain and analyze proprietary site observations of power and co-located or nearly co-located GHI and/or plane-of-array irradiance at time resolutions of five minutes or finer, at mid-latitude and tropical locations, in arid and humid climates, in continental and coastal locations, in all seasons. Likewise, we have obtained and analyzed ground-based irradiance measurements from publicly available sources, some of which are from high-quality well-maintained networks and others which are not. Quality control of both proprietary and public data is essential to using it in the forecast system as well as for validation of the forecast. After rigorous quality control including consistency checks among related parameters, the data can be used to investigate interesting questions such as deriving the actual operating tracking angles for sun-tracking arrays and examining sub-hourly variability.

This short paper highlights lessons learned from our analysis of real-world data and from our forecasts. We

doi:10.18086/solar.2016.01.15 Available at http://proceedings.ises.org

begin with an overview of the forecast system and then delve into some of the issues we encountered and how we addressed those issues.

2. Overview of MDA solar power and irradiance forecast system

The state-of-the-science MDA solar forecasting system is based on predicting irradiance, parsing the irradiance into direct and diffuse components, projecting it onto plane-of-array irradiance corresponding to a photovoltaic panel installation at a particular orientation which can be a function of time of day or sun angle, and running that through an empirical power curve based on that site or similar sites to obtain a power forecast. Global horizontal irradiance (GHI) is used from a multi-model ensemble of weather forecast models which includes the European Center for Medium Range Weather Forecasts (ECMWF), the NOAA High Resolution Rapid Refresh Model (HRRR), and others. For each individual model forecast, the GHI is nonlinearly bias-corrected through tuning against quality-controlled ground-based GHI measurements and other parameters over a recent history period, then prorated as a fraction of clear-sky conditions to match the diurnal curve down to 1-minute intervals. Because the model output typically represents hour or longer time averages but high-amplitude variability on a time scale of a few minutes can decrease the hour average power output by several percent for the same average irradiance, stochastic variability is added at 1-minute time scales. The stochastic variability is added using asymmetric distributions corresponding to the clear sky fraction and having appropriate temporal coherence. Then, the 1-minute GHI is parsed into direct and diffuse components and projected onto the panels to generate a plane-of-array (POA) irradiance, accounting for the extra circumsolar diffuse irradiance and ground-reflected light appearing on tilted panels. The POA irradiance is converted to power using multivariate empirical relationships derived from quality-controlled site data if available, otherwise using simple assumptions or applying the relationships found for other similar sites. This process is repeated for each individual model run and the results are blended using skill-based weights to produce the optimal forecast and the results are used collectively to generate forecast probability distributions. Additionally, satellite and real-time site data are employed to refine or correct the first few hours of the forecast.

The clear sky basis fundamental to this forecast approach employs the well-validated REST2 clear sky model (Gueymard, 2008), which calculates the clear sky transmissivity for GHI and for the direct beam. The accuracy of REST2 relies on good inputs of various scatterers and absorbers, including aerosol loading and Angstrom exponents and column water vapor, among others. The column water vapor comes from the model forecast. MDA analyzed years of sun photometer data from the NASA Aerosol Robotic Network (AERONET) together with weather model data to derive relationships between the weather parameters and the aerosol parameters. These relationships vary geographically and seasonally, allowing us to generate an aerosol parameter forecast tied to the weather forecast, resulting in better irradiance agreement with observations than by using persistence or static climatology. The parsing of all-sky (when not clear) GHI to direct and diffuse components combines the clear sky analysis with years of data from the Surface Radiation Network (SURFRAD), the gold standard in ground-based irradiance data, to yield relationships allowing us to derive the all-sky direct and diffuse components. Likewise, the 1-minute stochastic distributions of clear sky fraction were derived from a combination of SURFRAD data for all-sky irradiance and REST2 applied using our methodology and model data to yield the corresponding clear-sky irradiance.

More information about the forecast system is available in Jascourt et al. (2013, 2014, 2015, and 2016).

An example illustrating the quality of the forecast is shown in Figure 1. Fifteen-minute averaged forecast power at 1-hour lead time (blue) and metered power output (red) show remarkable agreement every day over a week, including clear and cloudy days with low and high variability despite no site data at all (neither real-time nor delayed) available from the preceding five weeks up to forecast time.

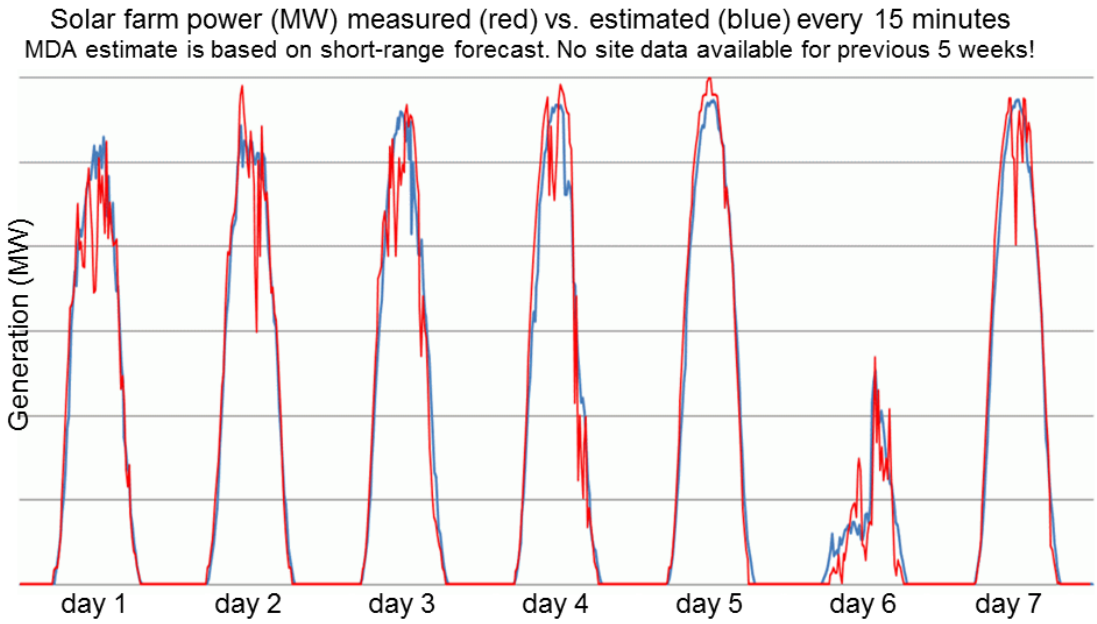

Fig. 1: One week of one-hour-ahead power forecasts for 15-minute blocks vs. actual power. No site observations were available for the preceding month up to forecast time.

3. Lessons learned from real-world experience

3.1. Quality control of PV site observations

PV farms always have power observations. Often the data recorders get stuck, for periods ranging from a few minutes to a few days. The latter are easy to detect but the former not, because values can also be stable for short periods and even at peak output for long periods at inverter-limited sites. Values can be cross-checked against calculated clear sky estimates to flag values which are unrealistically too high for the time of day. Also, in our experience so far, it is rare for a PV farm to produce exactly zero output even on a cloudy day when the sun is more than around 5 degrees above the horizon (accounting for terrain), so those zeroes are often spurious.

PV sites which report POA irradiance offer many more possibilities for quality control. If only GHI is reported, POA can be calculated. Then, the consistency between POA and power can be calculated. We have found many occasions at many sites when there were large discrepancies between POA and power. Sometimes this occurs at isolated times but more often in contiguous blocks of time, and it can help identify whether stuck power values are plausible. However, sometimes the problem is with the irradiance monitor. For example, sometimes shadowing occurs due to power poles or other objects. This can be detected by looking at irradiance as a fraction of clear sky irradiance versus azimuth and zenith angle to see if the fraction is consistently small at the same azimuth for a range of zenith angles. We have even detected brief shadows caused by wires using data at 1-minute intervals and highly accurate sun position calculations.

PV sites which report both GHI and POA irradiance offer even more cross-check possibilities. An example is shown in Figure 2 for a fixed tilt site at low latitude. On this day, the clear sky POA (green) was slightly less than the clear sky GHI (red) but the two were nearly identical. The morning was mostly clear; clouds developed by mid-day, intermittently blocking the sun, and a cloud deck moved in front of the sun towards the end of the day. Measured POA (white) was close to the clear sky curve in the morning while GHI (yellow) was lower by almost 200 W/m2. Similar discrepancies occurred over several weeks only during the mid to late morning. Thus, there must have been something partly shadowing the GHI sensor while the POA sensor had good exposure.

3.2. Determining actual angles for sun-tracking PV arrays

The forecast of POA irradiance and power is highly sensitive to the panel orientation during the morning up-

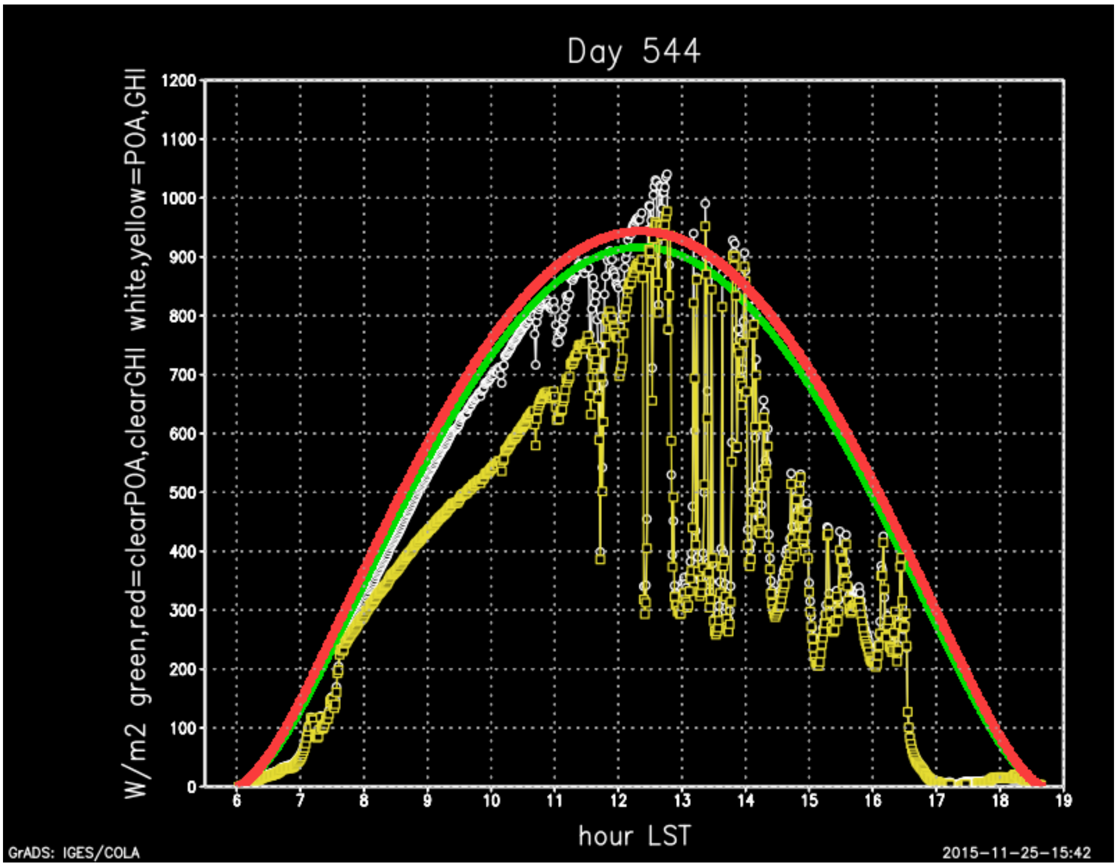

Fig. 2: GHI sensor shadowed from around 8:00 AM to noon, based on measured POA (white) indicating nearly clear sky conditions (green, calculated) while measured GHI (yellow) is far less than for clear sky (red, calculated)

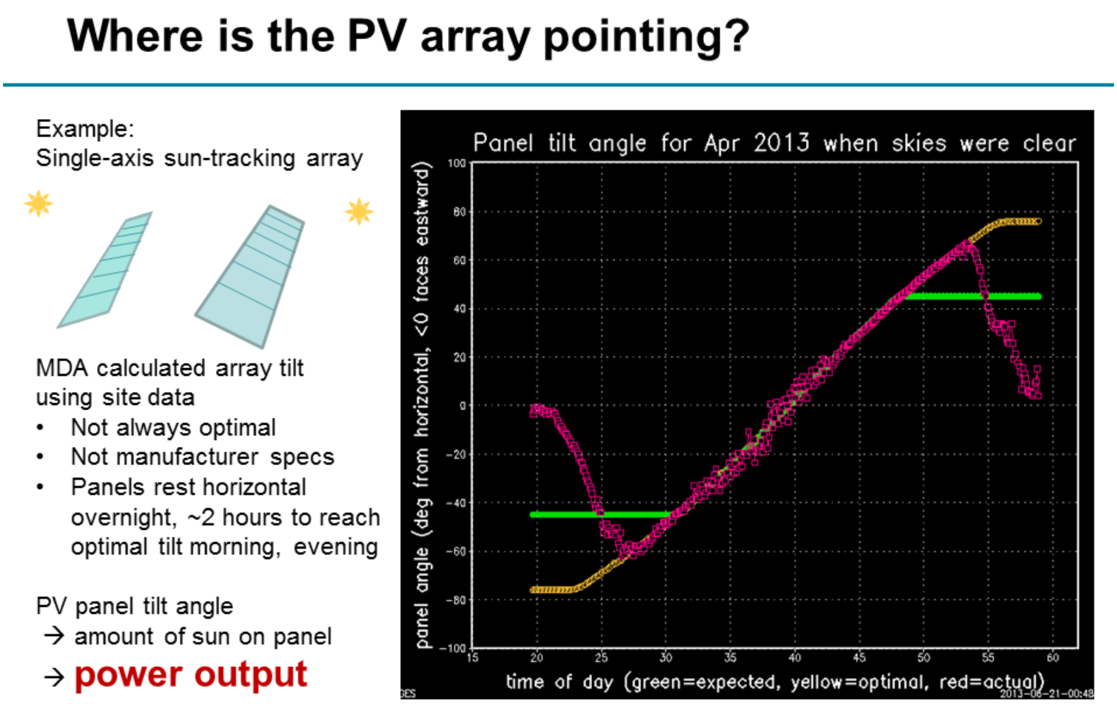

Fig. 3: Single-axis tracking angles for north-south axis tilting toward east (negative values on vertical axis) in morning (left) and toward west (positive) in afternoon (right). Yellow is optimal, green cuts off at manufacturer specifications (maximum tilt 45 degrees) and red is calculation from site data

ramp and evening down-ramp. This would be easy to deal with if we were to assume the tracking followed manufacturer specifications for tilt angles and followed the sun to the maximum extent the equipment can handle. However, in all sun-tracking systems we have encountered in all different geographic areas, the panels rest horizontal at night and can take up to a few hours to reach optimal orientation in the morning, then start heading down again at approximately the same rate to reach horizontal at sunset. MDA calculates the actual tracking positions based on site data. We have found that the rate of transition between the resting position and the optimal position varies from one site to another and the maximum tilt angle from horizontal often exceeds the manufacturer specifications, sometimes by a large amount. An example is shown in Figure 3 for a single-axis tracking array with a north-south axis. The panels take around 2 hours to reach optimal tilt toward the east in the morning, then start heading back to horizontal around 2 hours before sunset, reaching peak tilts of around 60 degrees although the manufacturer specifications indicate a maximum tilt of 45 degrees.

3.3. Quality control of public irradiance monitoring data

There are a variety of publicly available irradiance data sources, some of which report every few minutes and some only hourly.

Quality varies widely. For example, RAWS sites are abundant but are rarely serviced and, as they are intended to provide information in forests for the US Forest Service, they are located in forest clearings which still leave substantial shadows during morning and evening. Because they are abundant, some prominent research and private sector organizations building gridded GHI products tune their output to match the RAWS observations, claiming excellent fit to observations despite actually having poor fit to reality.

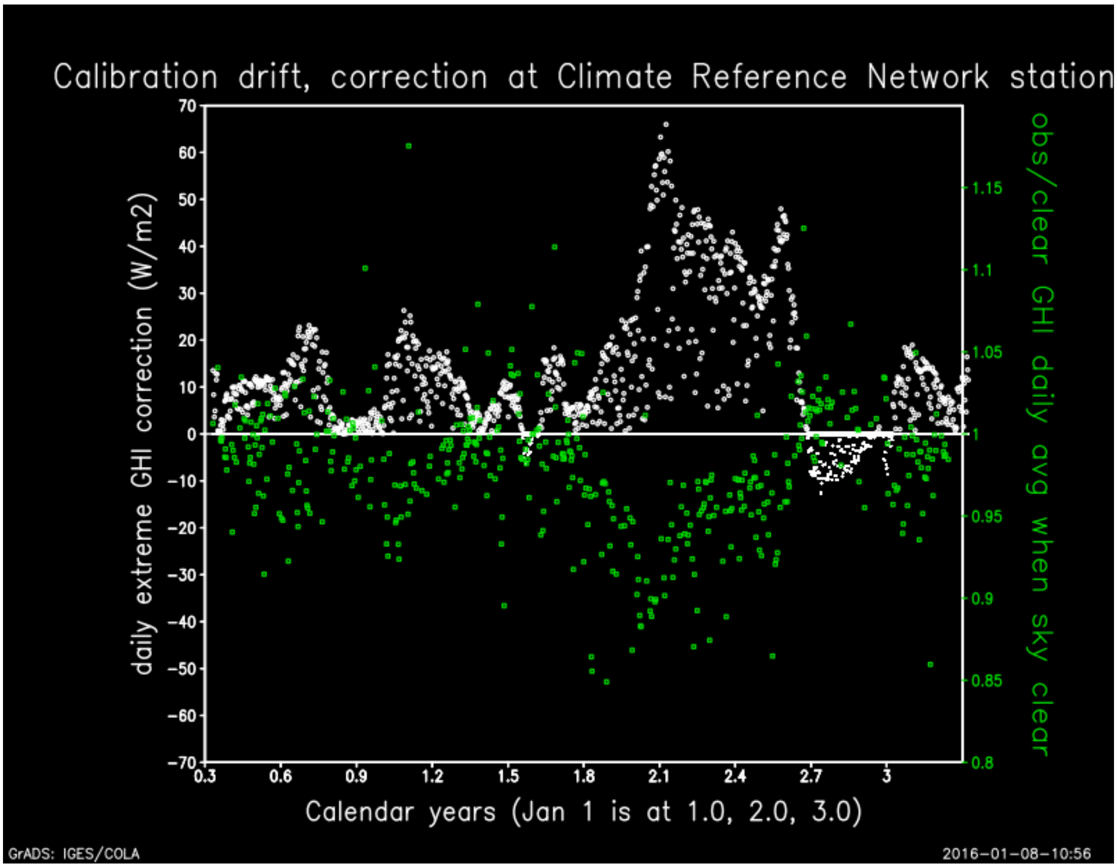

Fig. 4: Calibration drift and correction at a US Climate Reference Network station. Green shows the daily average of the ratio of GHI to clear sky GHI for times of day when skies were clear, over a 3-year period with date progressing from left to right (scale on right) . Estimated GHI corrections were applied and the largest amplitude corrections for any time of day are in white (axis on left)

However, even good quality observations at annually maintained sites can have issues. Figure 4 shows calibration drift or sensor soiling and corrections we applied for a Climate Reference Network site. This shows the value of having a good clear sky model to check against observations during identifiable clear sky times. The site tends to drift toward low values until it is serviced, then it is better for a while. Seeing this, we make corrections to level out the clear sky fraction at 1.0 and prorate the corrections also to times when the sky is not clear. The corrections are usually small but on some days the peak corrections can be rather large. These data after correction are then used for tuning forecast model GHI values.

3.4. Sub-hourly variability

We are finding that while we cannot predict the minute when an individual cloud will pass in front of the sun tomorrow, we can predict which hours will have rather steady cloud conditions and which hours will have rapid fluctuations. Our method involves careful statistical analysis of years of research-quality data. Figure 5 shows an example. The white dots are measured 1-minute GHI averaged over 5 minutes. The green dots come from averaging the observations over an hour and then applying the statistical method to synthesize one-minute values and then taking 5-minute averages. This simulates a perfect one hour forecast where we have no information about details during each hour. The statistical method recovers the wild fluctuations at the correct time even though the values for each minute are not correct during the period of high variability. We did this because the fluctuations affect the hourly average power generation, so it improves our forecast of total power generation for the hour. When the peak values during periods of high-amplitude fluctuations exceed the maximum which the PV-inverter system can output for, the power output is capped. However, the downward spikes are matched in amplitude in the power output. Thus, the average power is lower than the power based on the average irradiance. We calculate that this difference can reach a few percent of capacity at times although it is usually smaller. This stochastic method also provides a side benefit because the forecast amount of sub-hourly variability may also be of interest to the electric industry.

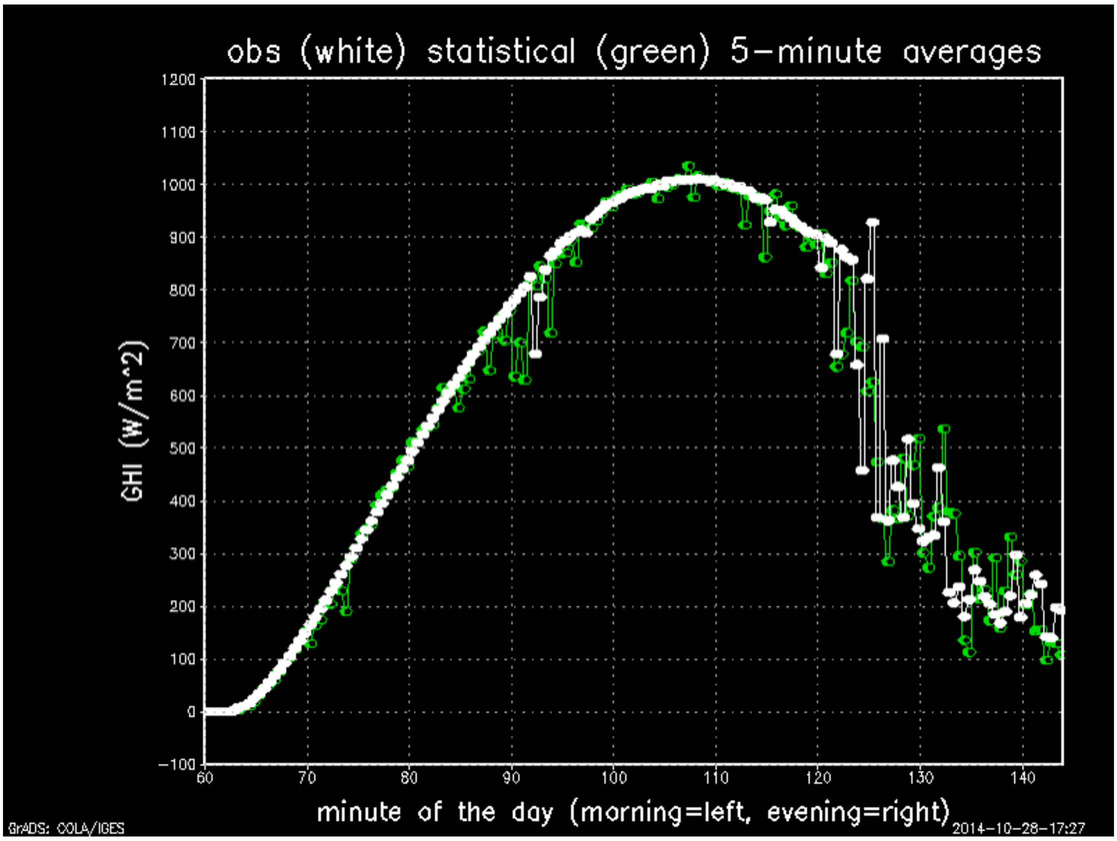

Fig. 5: Stochastic sub-hourly variability (green) versus actual (white). Observed and statistical values are 5-minute averages of 1-minute values. Statistical values receive only information about hour averages and attempt to recover the actual variability. Therefore, hourly averages of the two should match but 5-minute values would only match by chance. The goal is to match the observed variability.

3.5 Forecast bias

The Numerical Weather Prediction (NWP) model forecasts tend to be too sunny, particularly in winter and spring, in all different regions we have examined, and at all lead times including the first hour after the model is available (a few hours after model initialization due to latency for data ingest and computation and dissemination). We found this for ECMWF, GFS, NAM, RAP, and HRRR and will examine others. Power forecasts derived from passing the model forecasts through the MDA solar forecast system verify with little error on clear days, but on cloudy days, many of the model forecasts show nearly clear conditions.

The cause of the too-sunny forecasts are varied. Some cases involved poor forecasts of the movement of cut-off lows, others involved low-level moisture trapped under inversions that did not mix out as much or as soon as predicted, and there were cases of mesoscale cloud features associated with convection, sea breeze and other convergence zones, and other situations.

While model blends reduce error, bias remains. Even skill-weighting the contributions from each model does not improve this situation much. However, giving additional weight to cloudier forecasts does help.

Figure 6 shows 3-month bias in forecast power derived from different models (colors) at different lead times (different lines of the same color) as a percentage of AC capacity (vertical axis) throughout the day (horizontal axis). Figure 7 shows the forecast from various models and lead times for a clear day at one site, illustrating that correctly predicted clear days are not contributing most of the bias. Rather, the bias is due to predicting too many sunny days, with problems even in short range forecasts for later the same day.

The MDA forecast accuracy was improved and bias reduced by applying heavier weighting to models predicting lower irradiance in the MDA forecast blend.

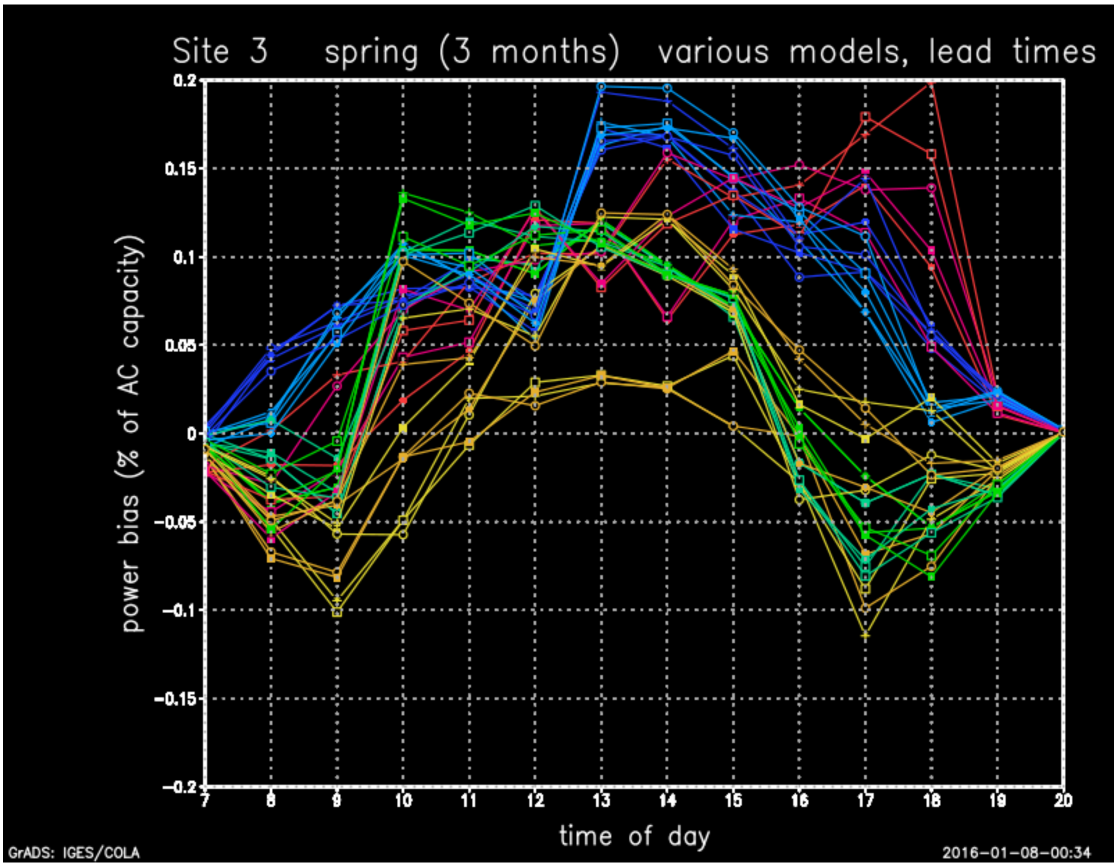

Fig. 6: Bias of power derived from NWP model irradiance forecasts over a three month period as a function of time of day (hour in local standard time). Each color is a different model. Each line is for forecasts of different lead times. Most of the forecasts are showing large mid-day to afternoon bias of 5 to 15 percent.

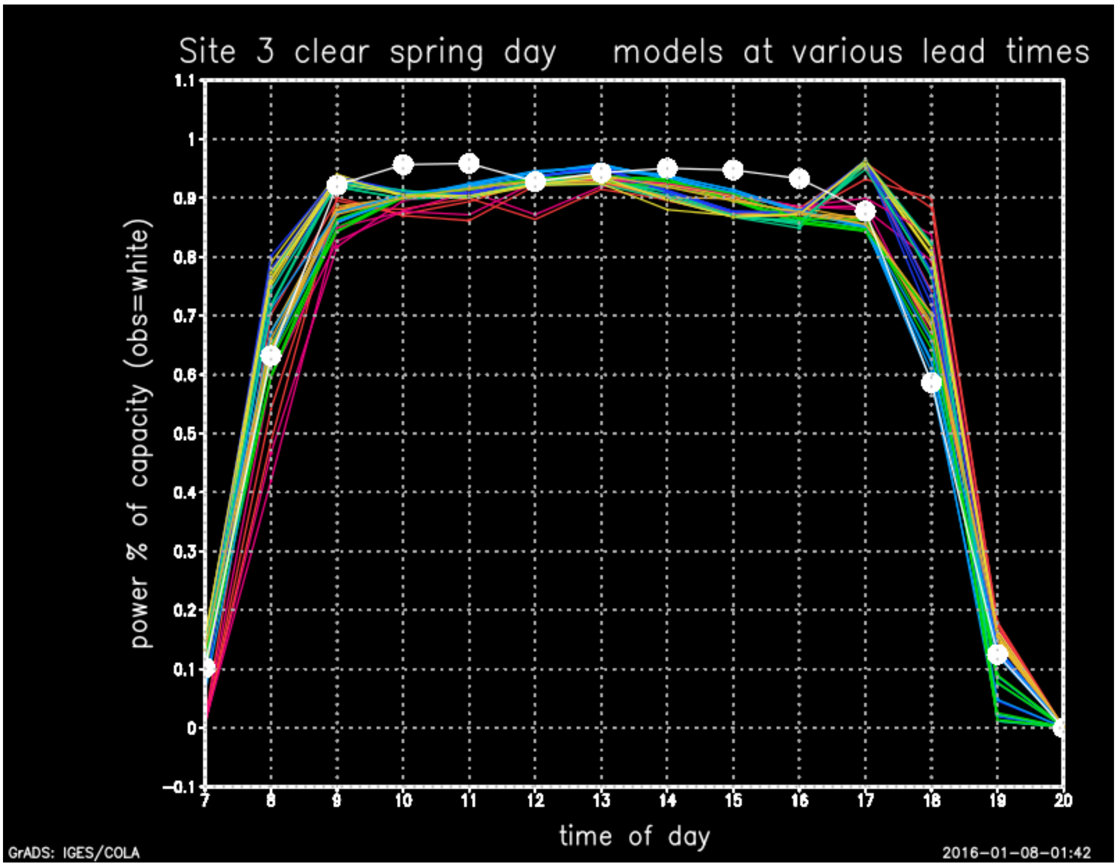

Fig. 7: Forecasts for one day for the same solar farm as in Figure 6 using the same colors for the same underlying NWP models and different lines for the same sets of lead times. This was a clear day, with observed values plotted in white. It does not show the high bias in the three-month average values shown in Figure 6, indicating that those high values are not due to overpredicting power on sunny days.

4. Conclusion

MDA has developed a sophisticated state-of-the-science solar power and irradiance forecasting system. The forecast system even simulates sub-hourly variability. Experience analyzing both proprietary site power and irradiance measurements and public irradiance monitoring data have led to emphasis on data quality control to filter an extensive variety of erroneous and suspect measurement reports and correct those which are correctable and to ascertain actual operating conditions such as orientations of sun-tracking arrays when those have differed from manufacturer specifications. Better results could be obtained if actual tracking were directly and accurately reported and if observing and monitoring systems were better maintained. Additionally, most numerical weather prediction models predict higher irradiance than observed on cloudy days, even at rather short lead times. Improvements in the underlying model forecasts might result from better parameterization of boundary layer mixing and other boundary layer physics as well as improvements in microphysics affecting cloud optical thickness. Meanwhile, MDA mitigates against model bias through the manner in which model forecasts are weighted in the multi-model ensemble.

5. References

Gueymard, C. A., 2008: REST2: High-Performance Solar Radiation Model for Cloudless-Sky Irradiance, Illuminance, and Photosynthetically Active Radiation – Validation with a Benchmark Dataset. Solar Energy, 82, 272-285

Jascourt, S.D., D. B. Kirk-Davidoff, C. Cassidy, and T. Hartman, 2016: The forecast is too suny and the observations are biased. 7th Conference on Weather, Climate, and the New Energy Economy, New Orleans, LA, Amer. Meteor. Soc. [Available online at https://ams.confex.com/ams/96Annual/webprogram/Paper288708.html]

Jascourt, S.D., D. Kirk-Davidoff, C. Cassidy, and T. Hartman, 2015: The MDA Solar Power Forecast System: Sub-hourly variability and behind-the-meter generation. 6th Conference on Weather, Climate, and the New Energy Economy, Phoenix, AZ, Amer. Meteor. Soc. [Available online at https://ams.confex.com/ams/95Annual/webprogram/Paper269331.html]

Jascourt, S.D., D. Kirk-Davidoff, C. Cassidy, and T. Hartman, 5 February 2014: Forecasting Electric Generation from Solar Energy. 5th Conference on Weather, Climate, and the New Energy Economy, Atlanta, GA, Amer. Meteor. Soc. [Available online at https://ams.confex.com/ams/94Annual/webprogram/Paper236483.html]

Jascourt, S.D., D. Kirk-Davidoff, C. Cassidy, and L. Heitkemper, 9 January 2013: Solar Energy Irradiance and Electric Generation Forecasting. 4th Conference on Weather, Climate, and the New Energy Economy, Austin, TX, Amer. Meteor. Soc. [Available online at https://ams.confex.com/ams/93Annual/webprogram/Paper222941.html]