Operationally Perfect Solar Power Forecasts: A Scalable Strategy to Lowest-Cost Firm Solar Power Generation

Richard Perez[1], Marc Perez[2], Marco Pierro[3], Jim Schlemmer[1], Sergey Kivalov[1], John Dise[2], Patrick Keelin[2], Mark Grammatico[2], Agata Swierc[2], Jorge Ferreira[2], Andrew Foster[2], Morgan Putnam[4] and Thomas Hoff[2]

[1] Atmospheric Sciences Research Center, SUNY, Albany, New York, USA

[2] Clean Power Research, Napa, California, USA

[3]University of Rome Tor Vergata, Rome, Italy, [4]RESurety, Boston, Massachusetts, USA

***Abstract* — The SUNY solar irradiance forecast model is implemented in the SolarAnywhere® platform. In this article, we evaluate its latest version and present a fully independent validation for climatically distinct individual US locations as well as one extended region.**

In addition to standard performance metrics such as mean absolute error or forecast skill, we apply a new operational metric that quantifies the lowest cost of operationally achieving *perfect forecasts*. This cost represents the amount of solar production curtailment and backup storage necessary to correct all over/under-prediction situations. This *perfect forecast metric* applies a recently developed algorithm to optimally transform intermittent renewable power generation into firm power generation with the optimal – least-cost – amount of curtailment and energy storage.

We discuss how *perfect forecast* logistics can gradually evolve and scale up into firm solar power generation logistics, with the objective of cost-optimally displacing conventional [dispatchable] power generation.

***Index Terms* —Irradiance, solar forecasting, PV fleet forecasting, firm power generation.**

I. Forecast Models

We analyze the latest version of the SUNY model as well as its four underlying Numerical Weather Prediction (NWP) models: HRRR [1], NDFD [2], GFS [3] and ECMWF [4]. In addition to these NWP models, the SUNY model blend includes satellite-derived cloud motion vector forecasts (CMM). The blend is a function of time horizon, solar geometry and predicted conditions. The current operational version of this model was independently evaluated by EPRI in 2017 and found to perform best among thirteen US forecast providers [5].

Here, we evaluate a beta version of the model that evolves the model's blend over time to locally capture the evolution of the relative performance of its underlying models.

II. Performance Evaluation

Performance is evaluated for point-specific locations, as well as for a fleet of sixteen locations in California. The evaluation period spans 16 months from January 2016 to April 2017. Out of this period, we analyzed ~ 11.5 months' worth of data when all models were present.

We consider three logistically important time horizons: 1, 3, and 24 hours ahead.

All validations are fully independent – i.e., for the SUNY model, the validation data are entirely distinct from the data used for blend optimization.

The seven SURFRAD network stations that span a wide range of climatic conditions [6] are used for the site-specific validations.

For the regional validation we consider the aggregated output of a 16 identical plants located in each of the state's climatic regions [7] (figure 1), using GHI as proxy for PV production. The considered points are located at the barycenter of each region. We use SolarAnywhere satellite-derived historical irradiances for performance benchmarking. We have previously shown that using satellite irradiances is acceptable to validate forecasts, yielding error metrics comparable to ground measurement validations [8]. In a recent article [9] we further discuss the appropriateness of satellite data for forecast validations: we show that while satellite data may be a suboptimal reference for single points (under-representing short-term variability) they are appropriate for intercomparing models, especially as the footprint evolves from single points (individual plants) to regions (PV fleets).

A. *Standard Metrics*

Standard metrics include MBE, MAE, RMSE, their relative (percent) counterparts normalized to either mean [daytime or 24-hour] value or nominal capacity, as well as Forecast Skill. The latter is determined with respect to *smart persistence* defined in accordance to IEA-SHCP Task 46 proposed practice [10] viz. (1) persistence of clear sky index, and (2) definition of current conditions as time-integrated over a period equal in length to the considered time horizon.

Here we focus our attention on the often-preferred absolute MAE metric and the Forecast Skill. The absolute MAE can be easily interpreted in terms of %MAE normalized to nominal capacity conditions by dividing by 1,000.

Absolute MAEs for SURFRAD sites and the California fleet are reported in Table I. Forecast skills are reported in Table II.

TABLE I
MEAN ABSOLUTE ERRORS

Location	SUNY	Smart Persitence	GFS	NDFD	ECMWF	HRRR
One Hour Ahead MAE (Wm-2)						
GoodwinCreek	43	51	73	78	66	83
Boulder	61	58	79	90	76	81
Sioux Falls	50	44	67	81	63	78
Penn State	53	53	73	78	70	96
Fort peck	47	44	67	79	60	77
Desert Rock	39	39	46	50	43	62
Bondville	47	49	69	78	64	84
SURFRAD MEAN	**49**	**49**	**68**	**76**	**63**	**80**
California Mean	34	39	51	53	45	58
California Fleet	**12**	**17**	**25**	**30**	**19**	**30**
3 Hours Ahead MAE (Wm-2)						
GoodwinCreek	57	95	74	80	67	83
Boulder	68	108	79	90	75	86
Sioux Falls	55	86	69	83	63	83
Penn State	62	99	75	81	70	98
Fort peck	55	79	69	80	60	80
Desert Rock	43	73	48	52	44	64
Bondville	58	95	71	82	66	84
SURFRAD MEAN	**57**	**91**	**69**	**78**	**64**	**83**
California Mean	42	65	52	54	46	66
California Fleet	**16**	**37**	**26**	**31**	**20**	**35**
24 Hours Ahead MAE (Wm-2)						
GoodwinCreek	64	144	79	87	68	na
Boulder	77	122	83	96	78	na
Sioux Falls	66	130	72	84	70	na
Penn State	67	129	77	85	71	na
Fort peck	60	94	72	79	65	na
Desert Rock	47	89	50	55	47	na
Bondville	71	134	81	90	75	na
SURFRAD MEAN	**64**	**120**	**73**	**82**	**68**	**na**
California Mean	46	87	53	57	46	na
California Fleet	**19**	**60**	**27**	**32**	**20**	**na**

TABLE II
FORECAST SKILL

FORECAST SKILL	SUNY	GFS	NDFD	ECMWF	HRRR
One Hour Ahead					
GoodwinCreek	23%	-43%	-41%	-17%	-61%
Boulder	4%	-37%	-36%	-19%	-34%
Sioux Falls	-1%	-55%	-69%	-32%	-79%
Penn State	8%	-28%	-40%	-19%	-68%
Fort peck	4%	-52%	-68%	-31%	-71%
Desert Rock	12%	-18%	-18%	-3%	-55%
Bondville	4%	-40%	-59%	-30%	-69%
SURFRAD MEAN	**8%**	**-39%**	**-46%**	**-21%**	**-61%**
California Mean	21%	-25%	-33%	-4%	-44%
California Fleet	**21%**	**-46%**	**-113%**	**-22%**	**-118%**
3 Hours Ahead					
GoodwinCreek	37%	13%	13%	29%	3%
Boulder	36%	17%	18%	29%	15%
Sioux Falls	35%	11%	4%	26%	-5%
Penn State	38%	23%	14%	30%	-1%
Fort peck	33%	8%	0%	22%	-5%
Desert Rock	35%	21%	21%	33%	-6%
Bondville	39%	19%	8%	31%	5%
SURFRAD MEAN	**36%**	**16%**	**11%**	**29%**	**2%**
California Mean	33%	11%	6%	27%	-16%
California Fleet	**54%**	**33%**	**4%**	**45%**	**-26%**
24 Hours Ahead					
GoodwinCreek	48%	33%	30%	46%	na
Boulder	32%	17%	15%	29%	na
Sioux Falls	45%	35%	30%	41%	na
Penn State	44%	32%	24%	40%	na
Fort peck	35%	15%	14%	27%	na
Desert Rock	41%	32%	30%	39%	na
Bondville	41%	29%	25%	37%	na
SURFRAD MEAN	**41%**	**28%**	**24%**	**38%**	**na**
California Mean	41%	26%	20%	38%	na
California Fleet	**66%**	**55%**	**34%**	**64%**	**na**

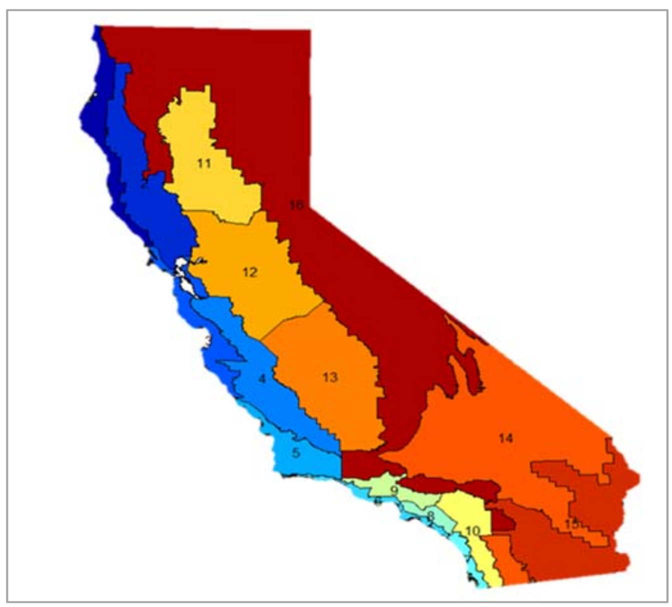

Fig. 1. Sixteen California Climatic Regions.

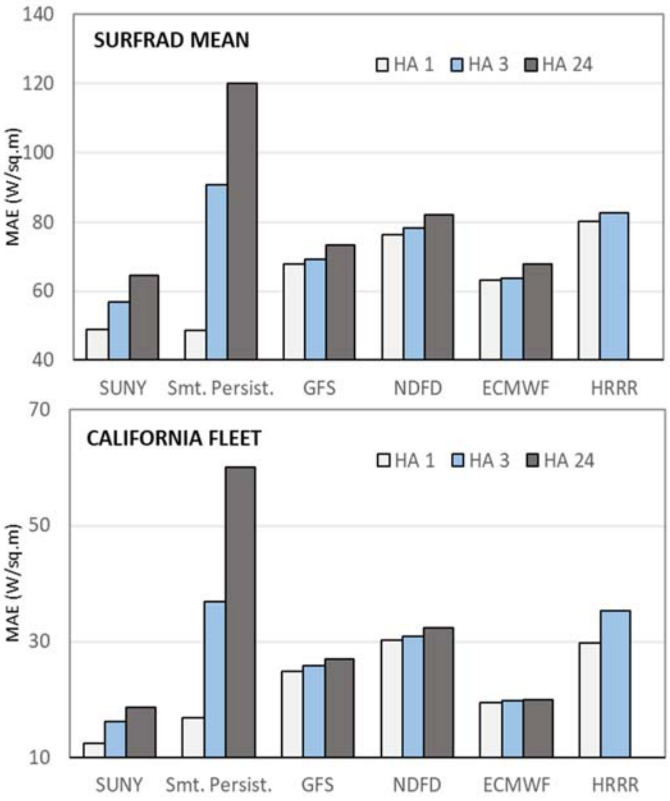

Fig. 2. Mean Absolute Errors for the individual SURFRAD stations (top) and for the California Fleet.

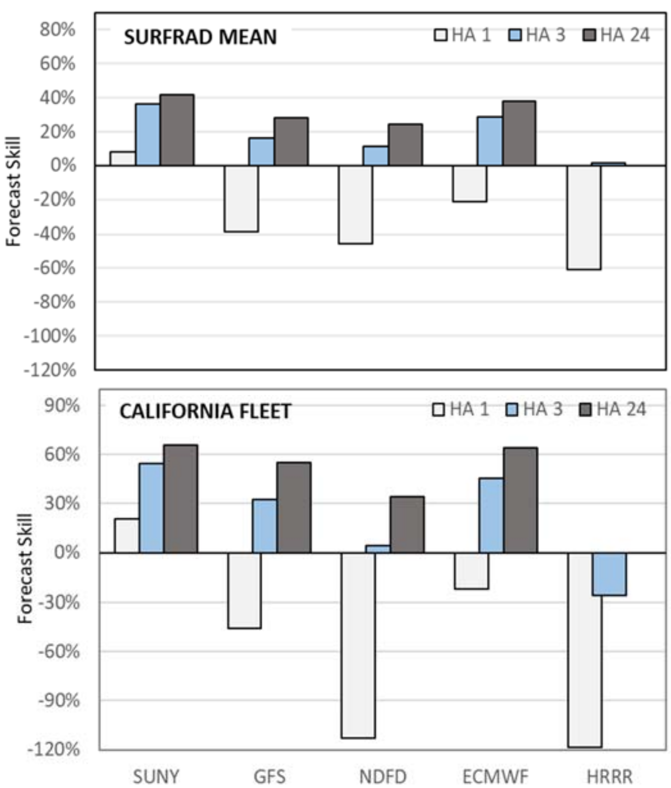

Fig. 3. Mean forecast skill for individual SURFRAD locations (top) and for the California Fleet (bottom)

A sample of the results in Table I and II are respectively illustrated in Figures 2 and 3.

These results are consistent with previous evaluations, with ECMWF exhibiting the best performance among the underlying NWP models, followed by GFS, NDFD and HRRR. The SUNY model is well ahead of the NWPs for short horizons (cloud motion advantage), and slightly better than ECMF (its major blend component) for longer time horizons.

The California fleet exhibits considerably reduced MAEs compared to individual sites. Whereas individual points in California are comparable to Desert Rock (see California mean of the 16 points in Table 1), the fleet MAE is reduced by a factor of 3 for the SUNY model, achieving 12 Wm^{-2} for 1 hour ahead and 19 Wm^{-2} for 24 hours ahead (i.e., respectively 1.2% and 1.9% of installed capacity)

Results for the forecast skill metric are also consistent with our previous findings: the underlying NWP models exhibit a negative skill for one hour ahead while this skill becomes positives beyond 3-hour time horizons. The SUNY model exhibits a positive skill for all horizons, reaching over 40% for 24 hours ahead forecasts. Interestingly, the skill differential between models is amplified for the regional fleet compared to individual locations: higher skill for the best models (SUNY is 66% for 24 hours ahead), lower skill for worst models.

B. Perfect Forecast Metric

In a previous article, we had introduced an initial version of this metric as the cost of storage necessary to offset any over-predictions [8]. This initial definition allowed nighttime storage recharge (i.e., implying low demand and low-cost electricity available at night).

The metric we apply in this article derives from an operationally more robust strategy built on a new algorithm to transform intermittent PV or wind generation into firm production at lowest cost [11]: this algorithm seeks the optimum (least-cost) combination of storage and PV oversizing/curtailment to meet a specified load profile with 100% certainty. This optimum combination depends on the relative costs of storage and PV. We consider two scenarios for PV and storage; (1) a current utility scale cost scenario with PV at \$1200 kWac turnkey and storage at \$200/kWh of storage capacity, and (2) a future (2045-50) utility-scale scenario with PV at \$400 kWac turnkey and storage at \$50/kWh of storage capacity.

The perfect forecast metric can either be expressed in terms of additional \$/kW above and beyond the cost of unconstrained PV, or in terms of levelized cost of energy (LCOE) premium above and beyond the LCOE of unconstrained PV. The LCOE metric requires one additional input: the weighted average cost of capital (WACC) – assumed here to be 3% (representative of utility industry [11, 12]).

In Table III, we present $/kW perfect forecast metric results for two sample SURFRAD sites and for the 16-points California Fleet.

TABLE III
PERFECT FORECAST $/KW PREMIUM

PERFECT FORECAST METRIC	SUNY	Smart Persitence	GFS	NDFD	ECMWF	HRRR
One Hour Ahead Perfect Forecast Metric ($/kW -- current)						
Goodwin Creek	$ 414	$ 182	$ 1,145	$ 1,441	$ 1,281	$ 2,330
Desert Rock	$ 398	$ 161	$ 968	$ 898	$ 692	$ 1,197
California Fleet	**$ 118**	**$ 89**	**$ 285**	**$ 169**	**$ 246**	**$ 595**
One Hour Ahead Perfect Forecast Metric ($/kW -- future)						
Goodwin Creek	$ 115	$ 47	$ 328	$ 365	$ 343	$ 627
Desert Rock	$ 110	$ 44	$ 261	$ 234	$ 192	$ 328
California Fleet	**$ 33**	**$ 23**	**$ 77**	**$ 50**	**$ 69**	**$ 163**
3 Hours Ahead Perfect Forecast Metric ($/kW -- current)						
Goodwin Creek	$ 589	$ 489	$ 1,180	$ 764	$ 892	$ 2,166
Desert Rock	$ 560	$ 434	$ 1,017	$ 912	$ 691	$ 1,076
California Fleet	**$ 172**	**$ 255**	**$ 316**	**$ 184**	**$ 262**	**$ 629**
3 Hours Ahead Perfect Forecast Metric ($/kW -- future)						
Goodwin Creek	$ 164	$ 130	$ 339	$ 220	$ 251	$ 567
Desert Rock	$ 149	$ 111	$ 275	$ 237	$ 191	$ 291
California Fleet	**$ 47**	**$ 71**	**$ 86**	**$ 54**	**$ 77**	**$ 173**
24 Hours Ahead Perfect Forecast Metric ($/kW -- current)						
Goodwin Creek	$ 835	$ 1,645	$ 1,234	$ 949	$ 1,016	na
Desert Rock	$ 711	$ 1,208	$ 1,166	$ 1,203	$ 772	na
California Fleet	**$ 199**	**$ 629**	**$ 441**	**$ 363**	**$ 224**	**na**
24 Hours Ahead Perfect Forecast Metric ($/kW -- future)						
Goodwin Creek	$ 227	$ 419	$ 356	$ 266	$ 277	na
Desert Rock	$ 189	$ 309	$ 309	$ 308	$ 205	na
California Fleet	**$ 52**	**$ 177**	**$ 115**	**$ 96**	**$ 62**	**na**

The perfect forecast metric results are interesting on two fronts.

First, operationally perfect 24-hour forecasts for the 16-plant fleet can be achieved at a cost of $200/kW with current hardware cost conditions and will be reduced to ~ $50/kW with anticipated future PV/Storage cost conditions. Perfect forecast thus amount to a small financial burden to guarantee operational certainty for the grid operator.

Second, the performance ranking of models is different from the standard metrics'. Particularly noteworthy is the better performance of persistence relative the underlying NWPs when benchmarked with the perfect forecast instead of the MAE metric. This ranking difference is illustrated in Figure 4. In this figure, the relative performance of each model is gauged against the average performance of all the model models across all considered locations and time horizons – a relative performance below 100% is better than the mean, and vice versa. Whereas persistence scores poorly when using the standard MAE as a metric, it bests the reference NWPs when using the perfect forecast metric. The blended SUNY blend scores very well with both metrics.

This observed ranking difference between the two metrics can be explained as follow: whereas the MAE is driven by the error of individual (hourly) forecast events, the perfect forecast metric is driven by the accumulation of under- or over-forecast conditions. The persistence is better balanced in this respect with shorter periods of enduring over/underpredicted conditions than the reference NWPS.

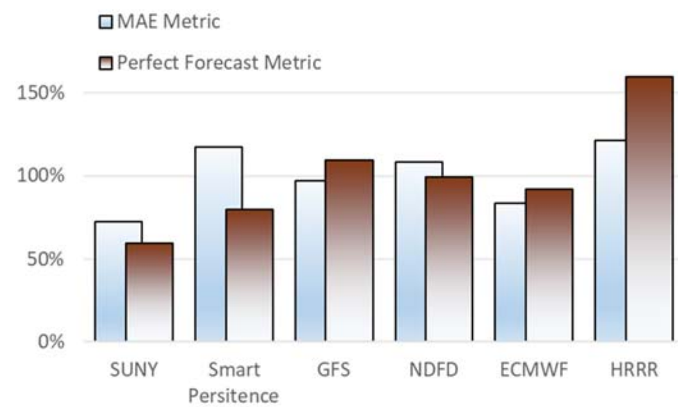

Fig. 4. Comparing model performance ranking across all locations and time horizons for standard and perfect forecast metrics. The value of 100% amounts to the mean error metric of all models/locations/time-horizons.

V. SCALABILITY TO FIRM POWER GENERATION

In addition to its use as a metric, perfect forecasting constitutes an economically attractive operational strategy for both solar producers and grid operators.

For grid operators, operationally delivering perfect forecasts removes all supply-side load imbalance uncertainty. It bypasses all missed forecast situations and associated costs, in effect eliminating the need for spinning reserves. For producers, this would amount to replacing imperfect administrative/regulatory penalties – that can evolve rapidly overtime -- by tangible hardware costs (PV overbuild and battery).

However, the real value of perfect forecast lies in its scalability to enable firm, effectively dispatchable PV power generation. Firm PV power generation (i.e., meeting the grid demand 24/7/365 regardless of time of day, time of year and weather conditions) is a prerequisite to very high PV penetration. The landmark Minnesota Solar Pathway project demonstrated that the least-cost means to deliver effectively dispatchable PV generation requires overbuilding and proactively curtailing PV generation to minimize multi-day storage requirements. Results show that below parity firm PV power generation is achievable in this not particularly sunny northern state if PV systems are overbuilt by 50-100% (hence curtailing 33-50% of their output) [11-15].

Figure 5 (from [13, 14]) illustrates the results of the algorithm applied to the objective of meeting Minnesota's MISO load with 100% certainty using an optimized mix of variable wind and solar resources and allowing for a maximum utilization of natural gas generation of 5%. The figure assumes utility-scale PV wind and battery future cost projections.

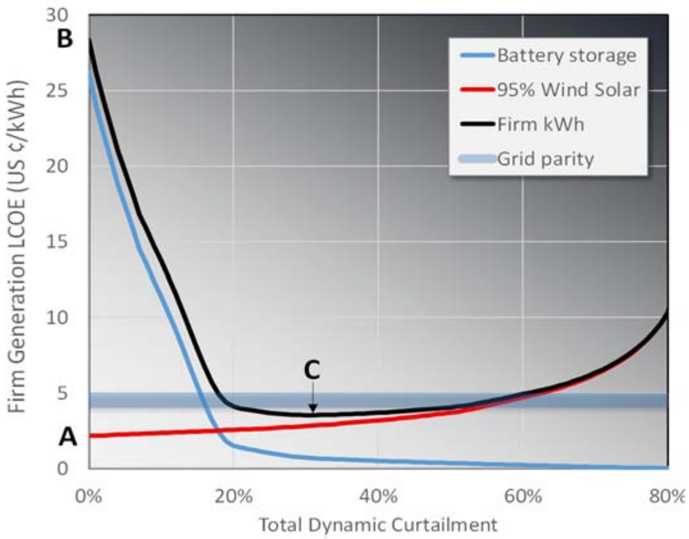

Fig. 5. Illustrating the catalyst role of overbuilding/curtailment in achieving firm least-cost power generation. Without overbuilding/curtailment, unconstrained variable wind/solar electricity will certainly achieve [apparent] grid parity (A). However transforming this variable renewable generation in effectively dispatchable firm power generation capable of meeting demand 100% of the time will remain well in excess of grid parity if overbuilding/curtailment is avoided because of the quantity of storage required to make up for multi-day and seasonal production gaps (B). With optimally overbuilt renewables, storage requirements can be reduced to the point where firm renewable power generation can achieve real grid parity (C), hence effectively displace conventional power generation

Importantly, the operational logistics of least-cost dispatchable PV generation – optimized PV overbuilding and storage associated with proactive operational PV curtailment -- are the same as the logistics of delivering operationally perfect forecasts, but on a larger scale (i.e., more storage and curtailment).

Thus, an operational perfect forecasts strategy constitutes a low-expense entry-level step and a learning curve toward enabling large-scale dispatchable PV power generation capable of meeting load demand 24/7/365.

The transition from perfect forecast to fully dispatchable PV can be gradual over time following grid operators' learning curve, PV penetration, and storage/PV costs decreases.

Figure 6 compares the duty cycles of meeting perfect forecast requirements and firm power generation requirements over a 10-day period. While this duty cycle is considerably heavier in the latter case, both involve a transformation of the solar resource into a predicted output in the case of perfect forecast and into the grid's load shape in the case of firm power generation. Both involve an optimization of storage and overbuild/curtailment requirements.

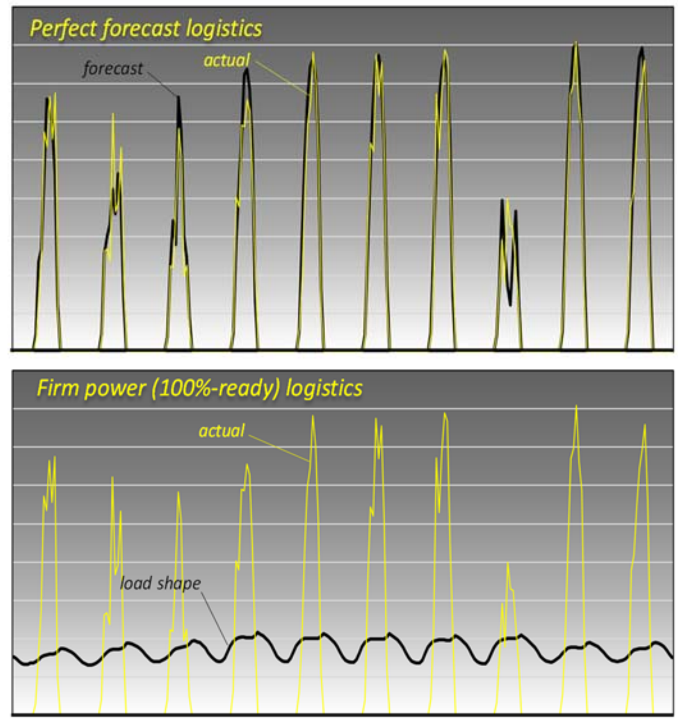

Fig. 6. Comparing a perfect forecast duty cycle (top) to a firm power generation duty cycle (bottom).

In Figure, we 7 illustrates a possible gradual transition of load requirements from perfect forecast requirements to firm power generation requirements.

The figure shows the same ten days' worth of forecasted PV production (black line) and the regional TSO load shape (red line). A gradual transition from perfect forecast logistics (guaranteed load = forecast) to firm power generation logistics (guaranteed load = regional load) could occur progressively as PV penetration increases while keeping the same operational storage/curtailment control logistics, in effect moving from the optimum storage/curtailment cost reported in Table III, to the optimum firm power generation cost illustrated in Figure 5. For the California fleet this optimum operating point would amount to an LCOE target of 3 cents per kWh for 100% penetration, fully dispatchable PV. This target cost could be further reduced by optimally blending PV with other renewables as was shown in Minnesota, and allowing for a small fraction of gas (<5%) to provide load target flexibility.

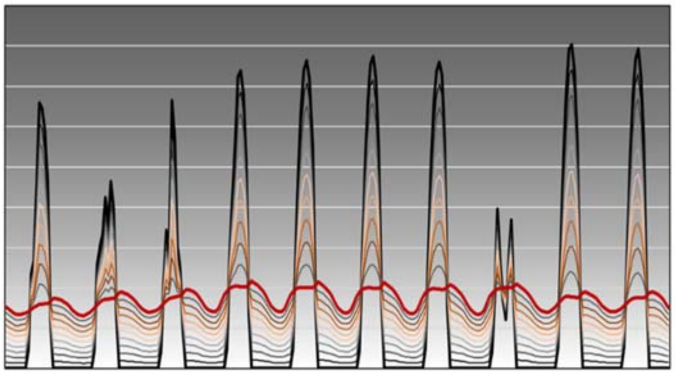

Fig. 7. Illustrating a gradual duty-cycle transition from a perfect forecast PV target (black line) to a regional load demand profile target (red line).

References

[1] High-Resolution Rapid Refresh (HRRR) - NOAA

[2] National Digital Forecast Database (NDFD) – NOAA

[3] Global Forecast System (GFS) – NOAA

[4] European Centre for Medium-Range Weather Forecasts (ECMWF)

[5] EPRI (2107): Solar Power Forecasting for Grid Operations: Evaluation of Commercial Providers. EPRI Technology Insights, Nov. 2017.

[6] Surface Radiation Budget network (SURFRAD) -- NOAA

[7] California Building Climate Zone Areas – California Energy Commission.

[8] Perez, R., J. Schlemmer, K. Hemker, Jr., S. Kivalov, A. Kankiewicz and J. Dise, (2016): "Solar Energy Forecast Validation for Extended Areas & Economic Impact of Forecast Accuracy". In 43th IEEE PV Specialists Conference, 2016

[9] Yang, D. and R. Perez (2019): Can we gauge forecasts using satellite-derived solar irradiance. Journal of Renewable and Sustainable Energy (under review)

[10] International Energy Agency, Solar Heating & Cooling Program, Task 46, Subtask C-1: Short-term solar forecasting .

[11] Perez, M., R. Perez, K. Rabago & M. Putnam, (2019): Overbuilding & curtailment: The cost-effective enablers of firm PV generation. Solar Energy 180, 412-422

[12] Perez, M., 2014. PHD Dissertation: A model for optimizing the combination of solar electricity generation, supply curtailment. Columbia University.

[13] Minnesota Department of Commerce: (2018): Solar Potential Analysis Report. < http://mnsolarpathways.org/wp-content/uploads/2018/11/solar-potential-analysis-final-report-nov15-2.pdf >

[14] Perez R. and M. Rabago, (2019): A radical idea to get a high-renewable electric grid: Build way more solar and wind than needed. TheConversation, May 29, 2019. https://theconversation.com/a-radical-idea-to-get-a-high-renewable-electric-grid-build-way-more-solar-and-wind-than-needed-113635

[15] Perez, M., R. Perez, K. Rabago & M. Putnam, (2019): Achieving 100% renewables: supply-shaping through curtailment. PVTECH Power Vol 19., May 2019. www.pv-tech.org