

Enhanced Ramp Event Forecasting Using Upstream Wind Measurements

Craig Collier - craig.collier@garradhassan.com
Jeremy Parkes - jeremy.parkes@garradhassan.com
Jonathan Collins - jonathan.collins@garradhassan.com
Lars Landberg - lars.landberg@garradhassan.com

GL Garrad Hassan
St Vincent's Works
Silverthorne Lane
Bristol, BS2 0QD, UK
Tel: +44 (0)117 972 9900

1 Introduction

Extreme volatility of the wind occasionally results in "ramps" in wind farm production. Wind power ramps are notoriously difficult to predict not only in terms of their occurrence but also in terms of their timing, amplitude, ramp rate, and duration, descriptive characteristics which guide the system operator in management of reserves.

An important category of ramps consists of those driven by rapidly changing meteorological conditions that result in correspondingly sharp increases or decreases in wind energy production (so-called upramps or downramps). In the geographical area considered in this investigation many meteorologically-driven ramps result from the passage of mid-latitude frontal systems driven by upper-level winds in a mid-latitude jet.

Using offsite and upstream measured data, an upstream learning model (ULM) is developed to augment an existing state-of-the-art wind power forecasting system, whose primary inputs are local to the wind farm site. The upstream learning model is then evaluated according to the following criteria:

- Its ability to improve the forecasting of whether or not a ramp event is going to occur (**ramp capture**)
- Refinements to forecasting the characteristics of the ramp event - **ramp rate**, **start time** and **duration**

It is found that conventional forecasting methods tend to under-predict ramp rate and over-predict ramp duration. Although the ULM in its current form only makes limited improvements to ramp capture, where ramps were forecast correctly it is able to improve the accuracy of the prediction of the profile of the ramps (ramp rate, duration and amplitude).

2 Methodology

2.1 Test Region

The geographical area chosen to test the upstream learning model (ULM) developed in this study is the region of the Pacific Northwest of the USA covered by the Bonneville Power Administration (BPA) system. Due to sharp seasonality in near-surface air temperatures at the latitudes of the Pacific Northwest, a maximum in upper-level wind (otherwise known as a jet) commonly traverses the BPA region in the mean climatological state. This results in lower level weather systems that propagate from west to east across the region and are responsible for large fluctuations in wind generation and so this region is a good test case for the value of using upstream measurements to predict ramp events at a downstream wind farm site.

For simplicity, only ramps caused by west to east propagation are considered in testing the model in this region. Up and down ramps are considered separately and ramps due to high wind speed shutdown are excluded with the result that only those ramps that are caused by the movement of large scale weather systems across the region are considered.

As shown in Figure 1, the wind farm sites whose data is being used in this study are oriented west-to-east along the Columbia River separating Oregon and Washington.

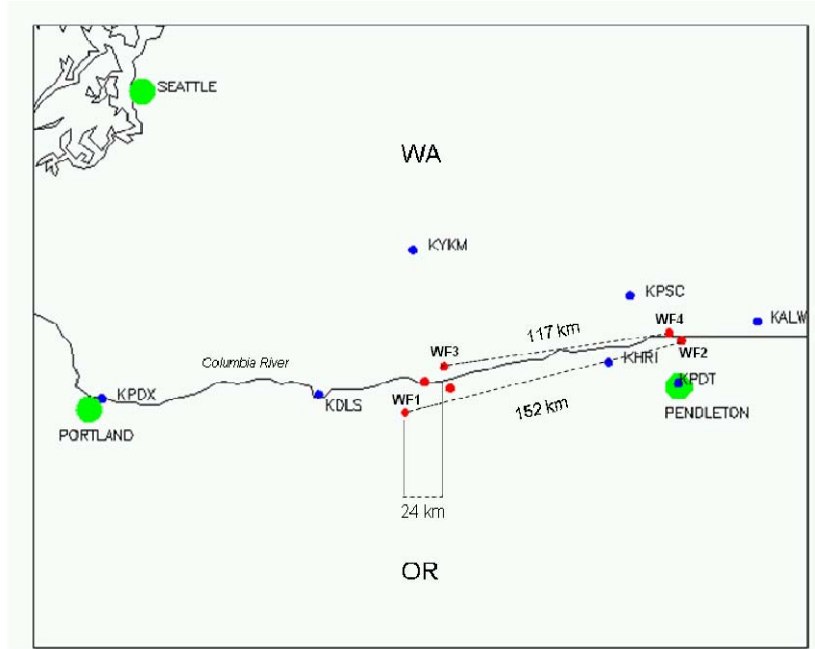


Figure 1: Map of BPA region showing locations of wind farms whose data is used in this study (red circles WF1 to WF4) and meteorological observing stations (blue circles)

Forecasts and measured data from upstream wind farms WF1 and WF3 are used in this study to enhance ramp event prediction at downstream wind farms WF2 and WF4. Figure 2 shows a ramp event occurring at upstream WF3 and then shortly after at WF4 caused by a weather system passing west to east across the test region. Data used in this study are from the period 2008-01-01 to 2008-12-31.




Figure 2: Measured power generation at two wind farms (WF3 and WF4 in Figure 1) separated by 117km showing a weather induced ramp at the upstream and then the downstream location

2.2 The Upstream Learning Model

According to the model developed herein, for a given pair of well-separated sites (S_u, S_d), the power forecast at downstream site S_d for a specific horizon h_d is given as a function of

- the error (ϵ) of the meteorological forecast at upwind site S_u for horizon h_u ,
- the numerical weather prediction (NWP)-derived power forecast specific to site S_d ,
- and the S_d power forecast as refined by onsite SCADA.

For simplicity, the meteorological forecast at the upstream site shall be denoted M_u and the total power forecast for the downstream site F_d (with component forecasts $f_{d,N}$ and $f_{d,S}$ for the respective NWP and SCADA - based models) such that

$$F_d(h_d) = G(\varepsilon(M_u(h_u)), f_{d,N}(h_d), f_{d,S}(h_d)) \quad (1)$$

Horizons h_u and h_d are determined by the speed of propagation of a site-impacting weather system (u_s). This speed may be a unique value for any particular system. Given the distance of separation between the upstream and downstream wind farms evaluated in this study (100 – 150 km), and a typical range of speeds assumed for eastward propagating disturbances, (e.g. 10 - 20 m/s), information from upstream farms (or stations) is likely to only be useful for predictions at horizons shorter than 6 hours for downstream farms (more properly, 3-4 hours at the 10 m/s rate and 1-2 hours at the 20 m/s rate). It is important to note that these horizon assumptions are based on the premise that disturbances propagate between sites at a constant rate. Determining with any precision the translation speeds of all systems tracked in the experiment period is beyond the scope of this study. The speeds cannot be known at all points; so for the purposes of this study, it was useful to measure value added for a range of forecast horizons.

For each set of simulations, h_u and h_d varied between 1 and 6 hours, such that 36 individual horizon combinations were tested to measure the value added by upstream information for a range of likely propagation speeds and rates of change. The upstream-learning model (ULM) described in Equation (1) operates as an add-on to the **GH Forecaster** system, and all experimental simulations are carried out at an hourly time-step using **GH Forecaster + ULM**. In order to appraise the value of the ULM for ramp prediction, control simulations on the downstream sites (using just the GH Forecaster system) were conducted, and the results were differenced. The processes and inputs of the basic GH Forecaster system are depicted in Figure 3.

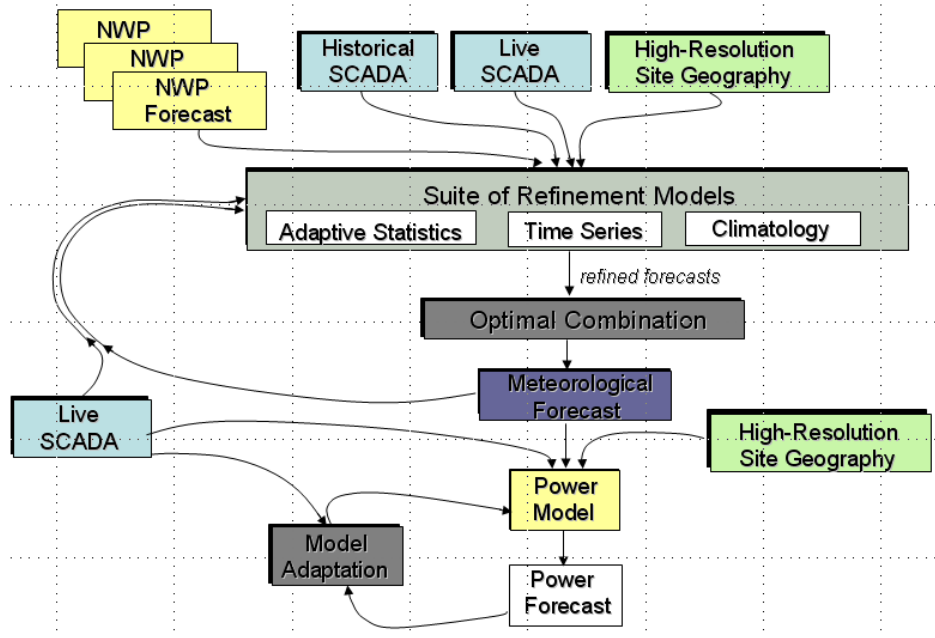


Figure 3: Illustration of the GH Forecaster setup, including flow of inputs and chronology of model processes

GH Forecaster draws upon raw NWP forecasts from several unique atmospheric models operating at a range of horizontal and temporal scales with updating frequencies varying from 1

to 6 hours. These meteorological forecasts are combined with historical and real-time site meteorological and generation data as well as high-resolution geographic datasets and collectively passed to a suite of models that refine the meteorological predictions for the site. Via various automatic learning techniques, adaptive statistics are used to remove systematic error, or bias, from an individual model; the most recent time series of production is used to eliminate nonsystematic error; and climatological data is used to adjust forecasts at very long horizons, for which there may be little confidence (or a large spread) in NWP solutions. The refined predictions are then optimally combined to reduce a cost function dependent on root mean square error, and the optimized meteorological forecast is then passed through a detailed wind farm power model that simulates the flow of wind through a wind farm site and generates a production forecast at each individual turbine. Live adjustments to the wind farm power model solution can be made with access to the most recent SCADA production.

The experimental setup is very similar to the control setup shown in Figure 3 except that the ULM serves as an additional input to the suite of refinement models, providing data from the upstream site. For each upstream-downstream site combination, four unique configurations of the ULM were tested, each assuming a different level of confidence in the upstream data.

Configuration 1 – The ULM-derived forecast is offered to a downstream forecast if and only if the downstream forecast production is demonstrating a monotonically-increasing divergence from the actual production over the most recent 3-hour period; the ULM-derived forecast is given 80% weight, with the remaining 20% of weight on the control forecast for the downstream site.

Configuration 2 – As Configuration 1 but the ULM-derived forecast is given 100% weight.

Configuration 3 – The ULM-derived forecast is offered to a downstream forecast regardless of downstream forecast performance and is given 80% weight.

Configuration 4 – The ULM-derived forecast is offered to a downstream forecast regardless of downstream forecast performance and is given 100% weight.

Note that configurations 1 and 2 may be considered selective-deployment configurations as the ULM model solution is only used given forecast error growth at the downstream site. Configurations 3 and 4 shall be considered nonselective-deployment configurations. For the 36 upstream/downstream horizon combinations used, 4 configurations of the ULM were tested to provide 144 simulations for each of the site pairings. For the results following, "Experiment 1" shall refer to the set of simulations using configuration #1, "Experiment 2" shall refer to the set of simulations using configuration #2, and so on.

2.3 Ramp Event Definition and ULM Evaluation Criteria

A systematic and automated method was employed for identifying meteorological ramp events over the experimental period. Herein a ramp event was defined as **any fluctuation in wind power generated over a period of no greater than 4 hours that is equal to or greater than 50% of the capacity of the wind farm.**




Figure 4: Example of a ramp event in real measured wind generation data

In order to assess the quality of a ramp forecast, it was necessary to define specific criteria with regards to its capture. A given ramp is captured if and only if a predicted fluctuation **equal to or greater than 50% of the actual ramp fluctuation** begins and ends **within 6 hours** of the starting and ending times of the actual ramp. This was measured by assuming a window around the actual ramp and identifying within it all non-overlapping ramps of the same direction as the actual ramp. The simulated ramp whose amplitude is closest to that of the actual ramp is selected as the model's representation. As several small ramps may be simulated within a larger ramp, these were carefully screened out so as not to penalize the model on its fidelity in predicting the large ramp.

Only ramps driven by meteorological conditions were considered. High-wind speed shut down events were not evaluated, though such events are also important in the management of electrical distribution systems. Of the meteorologically-driven ramps, only ramps which occur within the resolution of the forecast model are evaluated. Thus, if a ramp occurs over a period as short as 30 minutes, this event will be aliased by the model.

The goal of this study was improved prediction of ramp events by a state-of-the-art wind power forecasting system. Therefore, the addition of the ULM was evaluated on both improvement in ramp capture (predicting whether or not a ramp occurs) and on the prediction of key ramp characteristics – start time, rate of change, magnitude and duration.

3 Results

3.1 The Trial for Existing Wind Farm Sites

3.1.1 Control Simulations

Initial simulations without the ULM showed that prediction of ramp events is highly sensitive to forecast horizon. Consider Tables 2a and 2b below, which report the frequency of up and down ramp captures by GH Forecaster at WF1 to 4 as a function of horizon.

Farm Name		h = 1	h = 2	h = 3	h = 4	h = 5	h = 6	N
WF1	Cap	64.5	58	53	37	27	26	62
	RR	-25	-33	-41	-48	-55	-62	62
	Dur	+0.4	+0.5	+0.6	+0.8	+1.3	+1.9	62
WF2	Cap	73	70	67	49	37	36	90
	RR	-20	-30	-39	-44	-49	-50	90
	Dur	+0.3	+0.46	+0.5	+0.5	+0.8	+1	90
WF3	Cap	62.5	62.5	57.5	40	25	23	40
	RR	-19	-26	-36	-43	-49	-50	40
	Dur	+0.16	+0.08	+0.17	+0.38	+0.5	+1	40
WF4	Cap	80	76	69	56	39	33	75
	RR	-24	-30	-38	-41	-47	-48	75
	Dur	+0.3	+0.3	+0.3	+0.5	+0.66	+0.76	75

Table 1: Upramp capture rate (Cap), ramp rate bias (RR, % of capacity per hour), and duration bias (Dur, hours) for control simulations at forecast horizons between 1 and 6 hours. "N" represents the total number of upramps detected.

Farm Name		h = 1	h = 2	h = 3	h = 4	h = 5	h = 6	N
WF1	Cap	77	72	58	43	24	14	75
	RR	-30	-38	-49	-55	-65	-71	75
	Dur	+0.3	+0.26	+0.43	+0.5	+0.78	+2	75
WF2	Cap	77	73	70	53	39	23	109
	RR	-27	-36	-44	-54	-59	-60	109
	Dur	+1.2	+1.2	+1.2	+1.8	+1.7	+1.8	109
WF3	Cap	70	70	54	39	29	20	41
	RR	-16	-28	-36	-44	-50	-61	41
	Dur	+0.28	+0.27	+0.2	+0.25	+0.5	+0.88	41
WF4	Cap	73	71	62.5	50	35	21	48
	RR	-20	-27	-36	-43	-46	-43	48
	Dur	+0.23	+0.15	+0.37	+0.25	+0.47	+0.3	48

Table 2: Downramp capture rate (Cap), ramp rate bias (RR, % of capacity per hour), and duration bias (Dur, hours) for control simulations at forecast horizons between 1 and 6 hours. "N" represents the total number of downramps detected.

Note that differences in total numbers of ramps partly can be attributed to differences in record lengths; the period of record available at WF3 is 7 months while that at WF1, WF2 and WF4 is 1 year. However, it is interesting to note that WF2 production shows considerably more meteorological ramps than any of the other farms of similar record length.

In general, and as expected, forecast performance measures deteriorate for increasing forecast horizon. That is, ramp capture frequency decreases and ramp rate and duration biases increase for more distant prediction times. Ubiquitously, simulated ramp rates are biased low and durations are biased long, though generally fall within an hour of actual ramp duration. For upramps, ramp capture drops by a factor of 2 - 2.5, while ramp rate biases more than double from the 1-6 hour horizons. Deterioration is more pronounced for downramps, where capture rate drops by a factor of 3 - 3.5 between the next-hour and 6-hour horizons.

3.1.2 WF4 ULM Simulations Driven by WF1 and WF3

Ramp prediction performance was evaluated for use of the ULM at WF4 driven by data at upstream sites WF1 and WF3. WF4 straddles the Columbia River far east of the Columbia River Gorge, approximately 117 km downwind of WF3 and approximately 150 km downwind of WF1. For simulations driven by WF1, the forecast period spans all of 2008, while that for those driven by WF3 spans the first 7 months of the year. Ramp capture results for the WF4 simulations driven by WF1 indicate reduced ramp capture upon use of the ULM for downstream horizons shorter than 5 hours. There is neither significant improvement nor degradation in upramp capture at the 6-hour horizon and only modest degradation in downramp capture at this horizon. Ramp capture is reduced slightly in the nonselective deployment configurations of the ULM.

Ramp capture is also reduced for the simulations driven by WF3, with largest sensitivity at the 2-3 hour horizons. Downramp capture is particularly sensitive to the WF3-driven ULM, as downramp capture is reduced by nearly half for experiment configurations 3 and 4.

In contrast to the ramp capture results, ramp rate bias was reduced for a significant number of ramps predicted with the upstream learning model. In the WF1-driven simulations, the largest improvements were realized for upramp predictions at the 1, 2, and 3 hour horizons, where for the selective-deployment experiments, just under 50% of captured ramps were improved. For the nonselective-deployment experiments, over 50% were improved at the same horizons, with over 70% improved for the 1 and 2 hour horizons. Improvements of downramp simulations were not

quite as impressive for these simulations, but still noteworthy in the sense that use of the non-selective deployment configuration yielded improvements for an additional 20% of downramp predictions at horizons from 1 to 3 hours ahead (i.e., over 60% of the total events were improved). Qualitatively similar results were achieved with the WF3-driven ULM simulations.

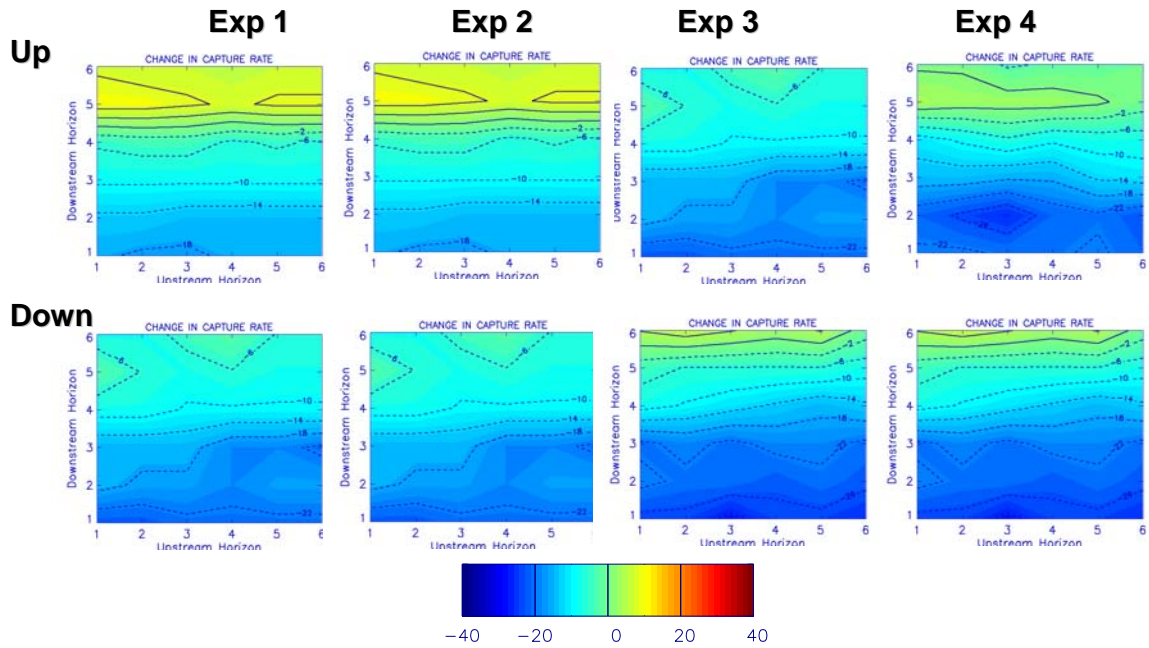


Figure 5: WF4 driven by WF1 – Change in ramp capture frequency (%) as a function of upstream forecast horizon and downstream forecast horizon upon application of ULM in the four different configurations (Exp. 1-4)

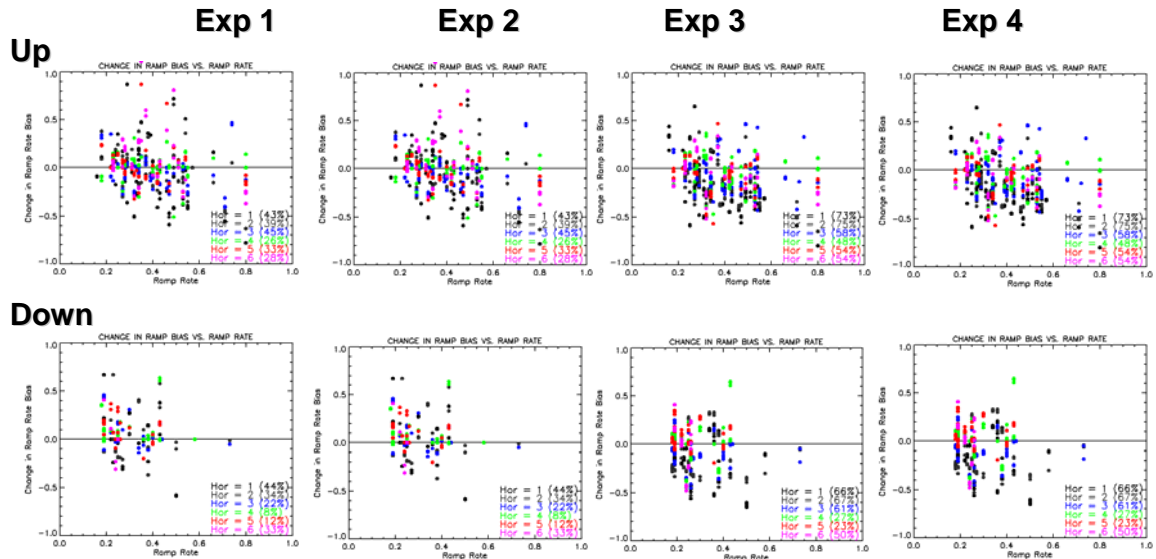


Figure 6: WF4 driven by WF1 – Change in ramp rate bias as a function of actual ramp rate. Downstream horizons shown in different colours. Points below $y=0$ line represent improved ramp rate forecasts. ULM applied in the four different configurations (Exp. 1-4)

3.1.3 WF2 ULM Simulations Driven by WF1 and WF3

The WF2 simulations reveal that ramp capture demonstrates similar sensitivity to the upstream data as for simulations at WF4. Employment of the WF1-driven ULM on WF2 forecasts results in moderately reduced capture rates (on the order of 2-5% for downramps and by about double for upramps) for the selective deployment configurations but significant reductions in capture for experimental configurations 3 and 4. A similar result was observed for simulations using a WF3-driven ULM. Once again, changes to ramp capture demonstrate sensitivity to downstream forecast horizon, with shortest horizon forecasts showing largest changes and relatively less significant changes at the longer horizons (5-6 hours).

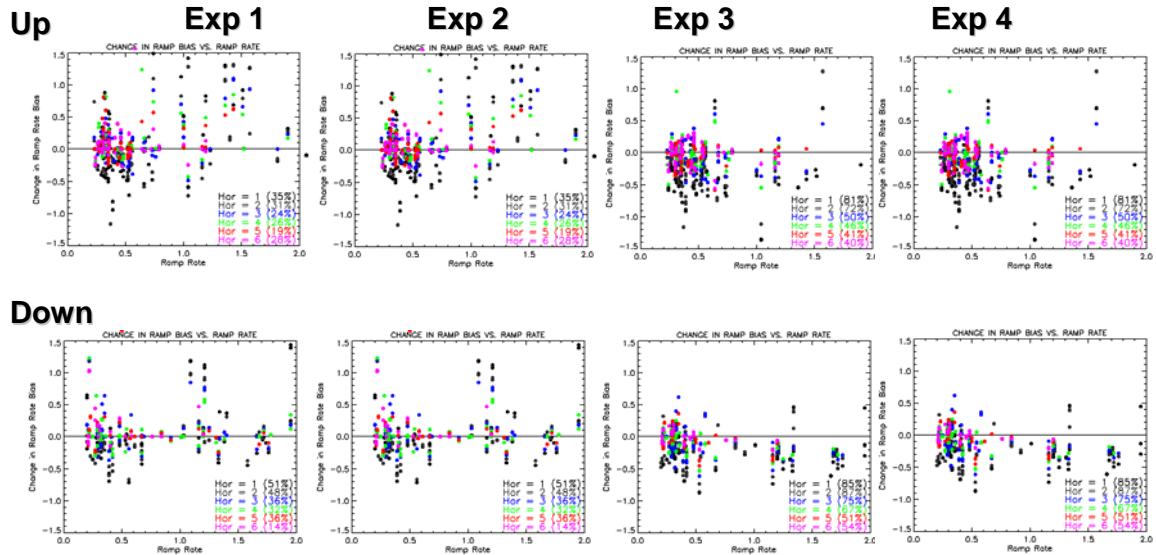


Figure 7: WF2 driven by WF1 – Change in ramp rate bias as a function of actual ramp rate. Downstream horizons shown in different colours. Points below y=0 line represent improved ramp rate forecasts. ULM applied in the four different configurations (Exp. 1-4)




Figure 8: Time history of 5-hour horizon forecasts for wind power generation at the WF1 for a week in February. Actual production is shown in black. The forecast from the control simulation is shown in red while the forecast from the WF1 driven ULM (configuration 4) simulation in is shown in blue

Of the ramps that are captured, use of the upstream model results in reduced forecast ramp rate bias for many upramps and downramps (Figure 7). In fact, simulation of over half of the ramps predicted is improved for horizons less than 4 hours for experimental configurations 3 and 4. While forecasts at the 5 and 6 hour horizons do not demonstrate the most impressive improvement with upstream data, a fraction of the ramp events show more realistic simulation by the ULM. Consider the significant upramps on 2/23, and the large downramp on 2/28 in Figure 8.

Important features such as ramp start time, duration, and ramp rate are better simulated with the ULM. It is evident that this refined solution, while certainly not consistently better, shows promise that upstream measurements can provide significant value for timing and amplitude of significant wind power fluctuations.

4 Conclusions

This study has attempted to evaluate the utility of upstream data for improvement of ramp detection and prediction quality at short forecast horizons. A conservative test was employed through an upstream learning model, by which wind forecasts at a given site were adjusted based on forecast errors at sites upwind over a range of horizon combinations. The adjustments from the upstream model were deployed under four sets of conditions to further test the sensitivity of downstream prediction enhancement to the degree to which the upstream-forecast-derived adjustment was applied. Ramp capture conditions were stringent, with a requirement that the forecasts simulate a fluctuation in power with start and end times within 6 hours of the actual start and end times of the ramp. Control simulations, which draw inputs from onsite data only, showed strong dependence of ramp capture and simulation quality on horizon. Forecasts at horizons of 4 hours or less tended to capture the majority of ramps observed in the data at the sites tested. However, beyond the 4-hour horizon, the ramp capture frequency decreased significantly while forecast quality statistics such as ramp rate bias increased. It is unclear whether the scoring statistics derived from the control simulations could be indicative of generalized ramp prediction performance for most state-of-the-art forecasting systems, but it is likely since all wind power forecasts ultimately rely on the same synoptic forecast model inputs, though differences between mesoscale modeling, statistical post-processing, and artificial intelligence techniques would result in some spread in performance amongst forecast systems. As particular models and techniques are deployed over unique forecast horizon bands, it is likely that ramp detection and simulation quality could improve at longer horizons, a subject of future study. Equally, it is likely that performance is geographically-dependent, as particular models are well known to perform more satisfactorily in particular wind regimes. The Pacific Northwest is a challenging regime for most forecast models due to the relative shortage of upstream high-fidelity observations over the Pacific Ocean.

Use of the upstream learning model for the site pairings tested yielded mixed results for forecasts for ramps at short horizons. Ramp capture was observed to either remain unchanged or to decrease for most horizons tested. This was likely due to the fact that forecast adjustments made to the wind levels at the downstream farm may have reduced the production forecast such that ramp "capture" failed to satisfy the amplitude criterion, thereby causing downstream forecasts to miss ramps that otherwise would have been captured. Additionally, missed ramps may have included a subset that did not propagate from the upstream site at all. Nonselective deployment of the upstream learning model tended toward increasingly reduced capture. Such a result suggests that standard forecasting techniques, relying on onsite data only, may be most reliable ramp predictors for shortest horizons but there may be increased value to this data at longer horizons, including for ranges not considered in this study.

The upstream learning models showed reduced predicted ramp rate bias for a large number of ramps detected. While such biases may have been reduced by as much as 8-12% of capacity per hour, some level of bias reduction was observed to occur for as many as 60-70% of captured ramps for horizons below 4 hours at the site pairings studied. When the upstream learning models were deployed without condition, the fractions of ramps for which ramp rate biases decreased grew larger.

Therefore, given the mixed results of this study, it should be concluded that increased upstream vigilance holds some benefit for wind power forecasting, though further work is needed to evaluate how best to use upstream forecast (or measured) data to produce robust improvement over forecast models that derive ground-truth inputs only from the site. This study has revealed that adjustments to wind forecasts based on forecast performance at upstream sites have the

possibility of helping or degrading a forecast, even when such adjustments are made very conservatively. Further investigation of the value added by upstream measurements, in particular exploration of the types of forecast scenarios in which upstream measurements will add the most value, is a necessity. This work highlights the complexity to be encountered with such an investigation, as forecasts do not respond well to broad-brush approaches and are likely to be sensitive to location, time of year, and upstream measurement site placement. Therefore, future work should be focused on completing the understanding of such sensitivities and construction of models that use high-fidelity upstream data in the most site-optimized ways.

5 References

1. Greaves B, Collins J, Parkes J, Tindal A. Temporal Forecast Uncertainty for Ramp Events, *EWEC 2009*
2. Colle BA, Westrick KJ, Mass CF. Evaluation of MM5 and Eta-10 Precipitation Forecasts over the Pacific Northwest during the Cool Season, *Weather and Forecasting 1999* (14-2)
3. McMurdie LA, Casola JH. Weather Regimes and Forecast Errors in the Pacific Northwest, *Weather and Forecasting 2009* (23-3)