Probabilistic ramp detection and forecasting for wind power prediction

C. Ferreira
LIAAD/INESC TEC and
ISEP/IPP, Polytechnic
Institute of Porto
Porto, Portugal
cfg@isep.ipp.pt

J. Gama
LIAAD/INESC TEC
and
FEP, University of Porto
Porto, Portugal
jgama@liaad.up.pt

V. Miranda
INESC TEC (formerly
INESC Porto), and
FEUP, University of Porto
Porto, Portugal
vmiranda@inescporto.pt

A. Botterud
Argonne National Laboratory
Decision and Information
Sciences Division
Argonne, IL, USA
abotterud@anl.gov

Abstract – **This paper presents a new approach to the critical problem of detecting or forecasting ramping events in the context of wind power prediction. The novelty of the model relies on departing from the probability density function estimated for the wind power and building a probabilistic representation of encountering, at each time step, a ramp event according to some definition. The model allows the assignment of a probability value to each possible magnitude of a predicted ramp and its worth is assessed by several metrics including ROC curves.**

Keywords – wind power, uncertainty, ramping, forecasting.

I. INTRODUCTION

This paper proposes a new way to detect and represent the probability of ramping events in short term wind power forecasting. Ramping is one notable characteristic in a time series associated with a drastic change in value in a set of consecutive time steps. Two properties of a ramping event, i.e. slope and phase error, are important from the point of view of the System Operator (SO): they have important implications in the decisions associated with unit commitment or generation scheduling, especially if there is thermal generation dominance in the power system. Unit commitment decisions, generally taken some 12-48 hours in advance, must prepare the generation schedule in order to smoothly accommodate forecasted drastic changes in wind power availability. A comprehensive analysis of ramp modeling and prediction may be found in [1]. Some important works in this area are mentioned in the following paragraphs.

The authors of [2][3] define *direction, duration* and *magnitude* for ramps in two basic types: upward (or *ramp-ups*), and downward (or *ramp-downs*), related with meteorological phenomena [4]. To consider a ramp event, the minimum duration is assumed in [3] to be of 1 hour; however in [5] one finds events in intervals of 5 to 60 minutes. The magnitude of a ramp is given as a percentage of the wind farm nominal power.

In [2], the authors define a ramp event to be a power output change higher than 50% of the wind farm nominal power, occurring over a period of 4 hours or less. They define two metrics: forecast accuracy and ramp capture - these metrics correspond to Precision and Recall defined in Section V below.

In [6], the authors combine feature selection and five data-mining algorithms to predict power ramp rates 10 to 60 minutes ahead. They use the *mean standard deviation, maximum* and *minimum* wind speed over all turbines and also the measured wind farm *power* and *power ramp rate*. Considering that the huge number of predictors can degrade the performance, *a boosting tree* algorithm was adopted to select the most interesting features, which are used to train five data mining algorithms: Multilayer Perceptron, Support Vector Machines, Random Forest, Classification and Regression Trees and Pace Regression.

In [7], a research project in Alberta (Canada) named WEPROG developed a special purpose tool providing real-time uncertainty weather forecasts for wind ramp prediction. The tool uses data from their Multi-Scheme Ensemble Prediction System (MSEPS), taking 75 forecasts to represent uncertainty for several weather parameters. An extreme ramp event would be a change in power of more than 80 MW over an hour.

In [8] the authors present the development by AWS Truepower to predict wind ramps between 0 and 6 hours ahead. The system capabilities include: a probabilistic ramp forecast module that can predict ramp rate probabilities for different time resolutions; a hybrid deterministic-probabilistic ramp event forecast that outputs deterministic values, a confidence interval for the events satisfying the ramp event definition and also the average power production for 15-min intervals. The ramp detection methodology associates algorithms with ramp types. The system learns models for significantly different weather conditions. To access the performance of probabilistic ramp rate forecasts and to compare two forecast methodologies, the Critical Success Index (CSI) [9] and the Ranked Probability Skill Score (RPSS) [10] were computed.

In [11] the authors identify ramps by mapping the initial wind power series into a signal that results from computing the average of time power differences. The authors propose two

probabilistic forecasting methods aiming to predict wind power output using ramp information, as well as to predict ramp timing. One method uses information extracted from the translated space, ramp intensity and ramp forecast time information. A method predicting ramping translates the signal of an ensemble of wind power curves and then uses the ensemble votes to define a confidence interval for the ramp timing. The quality of results depends on the number of ensemble members predicting the interval time of ramp occurrence and the size of the time intervals.

In [1] hard constraints are added to decision-making applications: for instance, in stochastic optimization or risk assessment. In [3] a probabilistic ramp event forecasting system is presented. It has been shown that by using probabilistic forecasting systems, higher economic benefits can be obtained.

The approach proposed in this paper requires and departs from a probabilistic wind power forecasting model. This model must allow a form of representation (implicit or explicit) of a joint probability density function (pdf) for the wind power, taking as variables the wind power prediction at different time stamps within the forecasting horizon. Such a representation already includes or takes in account the possible cross time step dependencies (a designation more general than just cross-correlations, which are called autocorrelation in the context of time series). In theory, given such a pdf, a Monte Carlo sampling process may allow the generation of scenarios, each consisting of a sequence of predictions for a number of successive time steps, in a time series fashion.

Our new model departs from a scenario generation procedure and builds a new ramp detection process and a new probabilistic assessment phase. The scenario generation model acts as a sampling mechanism in the Monte Carlo sense:

1. Generate N wind power scenarios (over a set of time steps), sampled from the wind power model.

2. For each scenario, detect in each hour the possibility of having a ramp event using a filter model.

3. Count the no. of detections in each hour for the whole set of scenarios, associated with each ramp type.

4. Based on the sample ratio of number of detected events n_i of type i over the size N of the sampled set, define probabilities for each ramp event in each hour of the forecasting horizon.

A refinement of this procedure allows one to build a histogram, for each hour, of the probability of having a ramping event above a certain magnitude: ramping becomes represented as a random variable associated to a probability distribution. This leads to the use of the results for decision making because the risk (or probability) of having a ramp exceeding a given threshold is quantified.

The robustness of the method, its capacity to avoid false alarms or missed alarms and its tuning are addressed in this paper. We present a comparative evaluation of performance in the ROC space (i.e. receiver operating characteristic, plotting true positive vs. false positive rates) [15] – from a case study using data from a wind farm in the USA.

II. GENERATING WIND POWER SCENARIOS

The key piece of the method developed is the availability of an estimate for the probability density function (pdf) of the wind power prediction. This paper will not discuss the methods to obtain such pdf estimate nor their validity. This pdf is taken as a multivariate function in a high dimension space (the number of dimensions equal to the number of time steps represented in the forecasting horizon). The probability of a ramping event of a given nature becomes then assimilated to the calculation of a special marginal distribution associated with some matching filter describing, in the multivariate space, the ramping event defined. In order to do this, a discrete representation of the pdf by wind power scenarios is necessary.

By a scenario, one understands a sequence of predicted wind power values spanning the entire prediction horizon, and by a discrete representation of the pdf one means a set of scenarios with an empirical density similar to the original pdf.

This is equivalent to some method of generating scenarios from a known pdf by using a sampling technique and turns the ramp event model into a Monte Carlo generating process.

Fig. 1 shows a usual representation of the uncertainty associated to wind power in short term 24 hour ahead prediction, based on quantiles [12], and a discrete representation. In the work reported in this paper, a state of the art method to generate scenarios according to a Monte Carlo sampling was followed [1]. We should note that we use both historical data and weather forecasts as input to our scenario generator. Moreover, our ramp detection model is independent of the scenario generation method used.
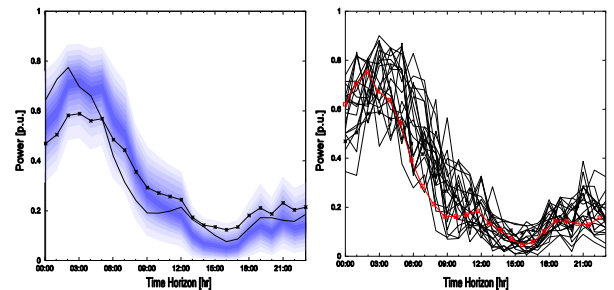

Fig. 1. Representation of wind power forecast as intervals (quantiles gathered by pairs and centered in the median), on the left, and scenarios of wind power generation, on the right.

III. DETECTING RAMP EVENTS

A. Defining ramps

A *ramp* is a change in power output (from a wind farm) with large enough amplitude and over a relatively short period of time. Fig. 2 illustrates this concept. Ramps may be up or down; in both cases, when not predicted, they may cause serious problems in system operation and dispatch, with high costs and additional risks incurred in.

There is no consensually accepted formal definition of a ramp. The ramp concept is related with the power signal P(t), a defining threshold ΔP_{ramp} and a time interval Δt. Some definitions adopted by different authors follow.
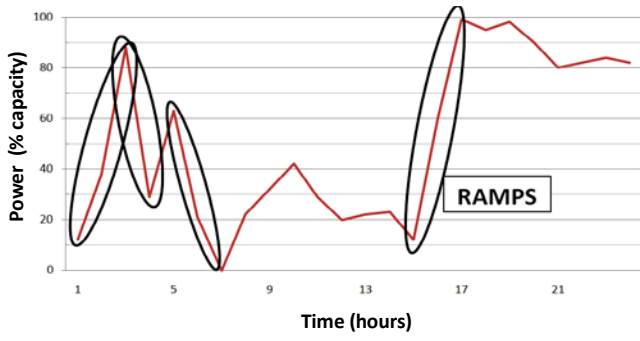
Fig. 2. Illustration of ramp events, defined as a change of at least 50% in power in an interval of 4 hours.

Definition 1 [5]: A ramp event is considered to occur at the beginning of an interval, if the magnitude of the increase or decrease in the power signal, at time Δt ahead of the interval, is greater than the ramping threshold value, ΔP_{ramp}:

$$|P(t+\Delta t)-P(t)| > \Delta P_{ramp} \qquad (1)$$

Definition 2 [5]: A ramp is considered to occur in a time interval Δt if the difference between the maximum and the minimum power output measured in that interval is greater than the threshold value, ΔP_{ramp}:

$$\max(P[t, t+\Delta t])-\min(P[t, t+\Delta t]) > \Delta P_{ramp} \qquad (2)$$

Definition 3 [6]: A ramp occurs if the difference between the power measured at the initial and final points of a time interval Δt is greater than a predefined reference value to the Power Ramp Rate, ΔP_{ramp}:

$$|P(t+\Delta t)-P(t)|/\Delta t > \Delta P_{ramp} \qquad (3)$$

The definitions above work directly with the wind power signal. Other approaches transform the signal into a more appropriate representation, e.g., considering k-order differences in the power amplitude (see [11]). Let P_t be the wind power time series and P_t^f the associated transformed signal that was obtained according to

$$P_t^f = \text{mean}\{P_{t+h} - P_{t+h-n_{am}}; h = 1,...,n_{am}\} \qquad (4)$$

where n_{am} stands for the number of averaged power differences to consider. Then,

Definition 4 [11]: A ramp event is said to occur in an interval, if the absolute value of the filtered signal P_t^f exceeds a given threshold value, ΔP_{ramp}:

$$| P_t^f | > \Delta P_{ramp} \qquad (5)$$

Definition 5: This is a new definition developed under the project. It uses a high-pass filter, i.e. a filter that passes high-frequency signals and attenuates (reduces the amplitude of) signals with frequencies lower than the cutoff frequency. The simpler high-pass filter can be formulated as:

$$y[i] = \alpha(y[i-1] + x[i] - x[i-1]) \qquad (6)$$

It can only pass relatively high frequencies because it requires large (i.e., fast) changes and tends to quickly forget its prior output values (see Fig. 3). The parameter α takes values in the interval $]0;1]$. Values near 1 imply that the output will decay very slowly but will also be strongly influenced by small changes in the input signal.

A constant input (i.e., an input with $(x[i] – x[i - 1] = 0)$) will always lead to an output decay to zero. A small α implies that the output will decay quickly, requiring large changes in the input (i.e., $(x[i] – x[i - 1])$ is large) for the output to vary considerably.

Fig. 3 illustrates the concept. The y signal may be further treated by a band filter, removing small peaks and only keeping the values above a certain threshold, compatible with the ramp definition accepted.
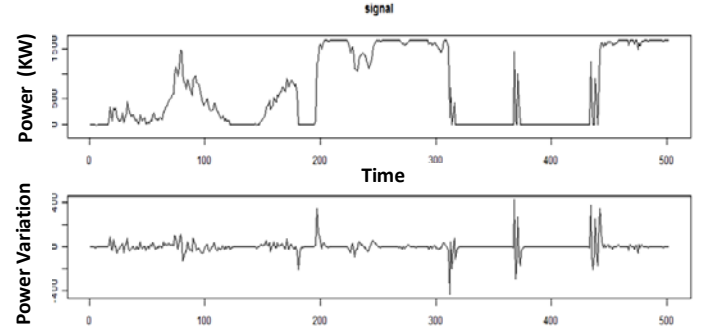

Fig. 3. Top: wind power. Bottom: high-pass filtered signal ($\alpha = 0.25$).

B. Building a Probabilistic Ramp Representation

The detection mechanisms outlined above serve as indicator function for each time step in one scenario. When applied in a set of scenarios as a sample of the wind power pdf, we have in place a Monte Carlo process that can give as a result the probability of a specific ramp event and also an estimate of the ramping probability distribution at each time step as a function of the ramp amplitude.

The general algorithm is:

- Generate a large set of N wind power scenarios, sampled with the wind power forecasting model.

- For each scenario, detect in each time step if there is a ramp event of each type defined.

- Count the total of ramp event detections in each time step for the whole sampled set of scenarios, associated with each ramp type.

- Based on the sample ratio of number of detected events n_i of type i over the size N of the sampled set, define probabilities for each ramp event in each time step of the forecasting horizon.

C. Building cumulative ramp probability diagrams

A histogram may be built by defining a set of bins $\Delta^b P$, ranging from ΔP_{ramp} to a user specified maximum power change, and define the Vote Counting for each histogram bin b

$$V_k^b = \sum_{j=1}^{N} [\Delta^{bl}P < F(\Delta P_j^k) < \Delta^{bu}P] \qquad (7)$$

where the lower and upper bound of the histogram bins are bl and bu, and F is the ramp definition in terms of power change.
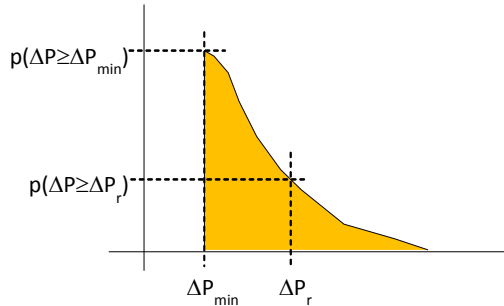

Fig. 4. A cumulative ramp probability diagram allowing risk evaluation for a ramping value ΔP_r. ΔP_{min} is the minimum acceptable power variation that does not trigger a ramp event alarm.

This allows the definition of a cumulative ramp probability diagram such as in Fig. 4. From diagrams such as these, a measure of risk can be associated with the probability $p(\Delta P \geq \Delta P_r)$ of having a ramp event with a change equal or greater than P_r.

D. Deciding that a ramp event should be declared

From the set of scenarios one may thus define an empirical probability of exceeding a given threshold ΔP_{ramp}:

$$P(E_k) = \frac{1}{N} \sum_{j=1}^{N} [F(\Delta P_j^k) > \Delta P_{ramp}] \qquad (8)$$

By setting a cut-off threshold *thr* on the probability $P(E)$, we may declare the occurrence of an event (ramp) at time step *k* if

$$P(E_k) > thr \qquad (9)$$

This declaration converts the probabilistic ramp prediction into a binary choice that may lead to a decision to act and is needed in the operation context. Section V presents a way to define the threshold *thr* in an optimal way.

IV. APPLICATION OF THE NEW MODEL

To verify the quality of the model, we organized an experiment built from real data from a large wind farm in the U.S., in a time period of 12 weeks (21/10/2009 to 18/02/2010). We generated 5000 scenarios (possible predictions of power) using [1]. The algorithm described above was run for each 24 hour ahead, and cumulative ramp probability diagrams were built for windows of 3 hours, counting and classifying possible ramps through the use of the high-pass filter to detect possible ramps and a band filter to eliminate small changes.

Fig. 5 shows the case for one day. One can observe a point forecast, produced by the model in [14], and the actual values measured. Below, one has ramp cumulative probability diagrams, for 3 hour steps, which provide information about the probability of having a ramp (given a definition) and also about the probability of the magnitude of such a ramp.

These diagrams are rotated, relative to Fig. 4, and the vertical axis is associated with a ΔP value, while the horizontal axis corresponds to the probability of having a ramp event of a magnitude equal or greater than a given power threshold. It is possible to have at the same time some probability of having either ramp up or ramp down, representing a variety of behavior observed in the generated wind power scenarios for the same type of occurrence..
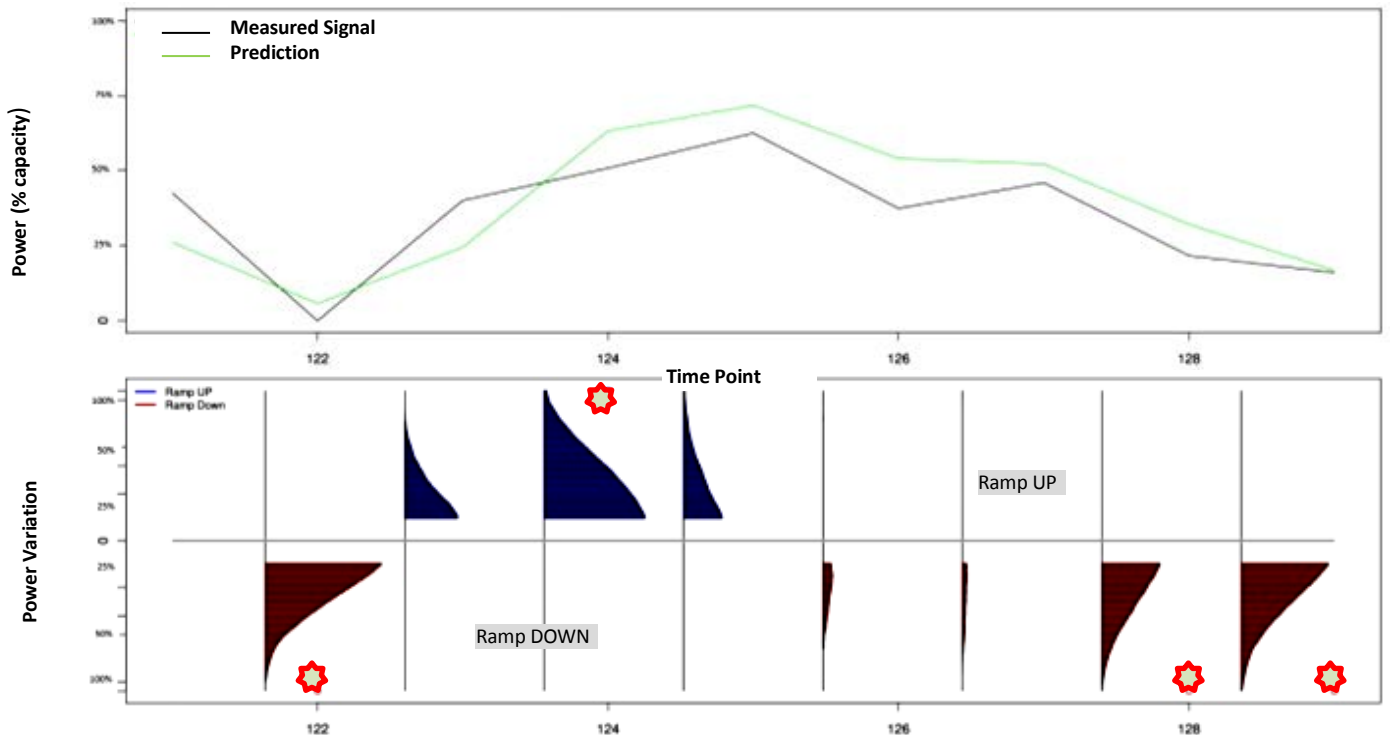

Fig. 5. Top: actual measured wind power and the point forecast for 24 hours (slightly over-forecasting in most time steps). Bottom: cumulative probability diagrams for ramp-up (above) and ramp-down (below) events. A star marks intervals with a prediction/declaration for a ramp to occur.

V. QUALITY ANALYSIS

Ramp event detection is a process where one may define hits (TP – true positives or TN – true negatives) and misses (FP – false positives, or false alarms, when a ramp is predicted but does not occur, and FN – false negatives or missed detections, when no ramp is predicted but occurs). The results below were obtained using a three-hour aggregation defining a ramp UP or DOWN of magnitude change higher than 25% of the wind farm nominal power. This is the ΔP_{ramp} threshold value set for the time period Δt (= 3 hours, in this case). With this Δt value, definitions one, two and three lead to equal results. As for definition 4 (a moving average), the results come from setting n_{am} = 2. With n_{am}=1 one gets almost the same results as when running definitions one, two and three.

Some widely used statistics to assess the quality of deterministic event detections defined in [0, 1] are: precision (or sensitivity, or true positive rate TPR), recall and specificity. *Precision* is defined as the ratio between the number of true positives and of positive forecasts. *Recall* is defined as the ratio between the number of true positives and of observed positives. *Specificity* is the fraction of true negatives and the quantity (1 – Specificity) may be called the *false positive rate FPR*.

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \; ; \; \text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \; ; \; \text{Specificity} = \frac{\text{TN}}{\text{TN} + \text{FP}} \quad (10)$$

These concepts can be used to assess the effect of the actions resulting from probabilistic information. It is evident that in probabilistic forecasts a new degree of freedom is introduced: a threshold to define the occurrence of the event.

A technique that may help in choosing a threshold level that optimizes event detection is the *receiver operating characteristic* or *ROC* curve, which is a graphical plot in the plane FPR × TPR and domain [0,1]×[0,1]. This plot is achieved by changing progressively a cut-off value that defines the detection of an event – this value is associated with the probability of observing a ramp, from examining all scenarios. In this ROC domain, the main diagonal (0,0) to (1,1) defines a prediction like a random guess. The optimum corresponding to perfect discrimination is the point (0,1) where all positives are detected and no negatives are taken as (false) positives.

Fig. 6 displays the ROC curves that we obtain in the classification of ramp-up events using definition 1, 4 and 5 and setting ΔP_{ramp} (amplitude of the band filter that eliminates small events) equal to 25% of the nominal power. The model produces a better result than random guessing. The cut-off value that should be adopted depends [15] on the relative costs of missing a positive and of assuming a positive when there is none. If these costs are equal, and assuming a uniform event distribution, then the slope of TPR/FPR equals 1 and the cut-off value to be adopted, to accept or reject an alarm, should be the one that leads to the tangent to the ROC curve with slope 1 that is closest to the optimum point (0,1). This is relevant in the application of the method to wind power ramp prediction: the cost of missing a positive may be related with emergency ramping of generators or power purchases at high spot prices or load disconnection; the cost of accepting a false alarm is related with higher unit commitment costs or higher cost of operational reserve allocated.
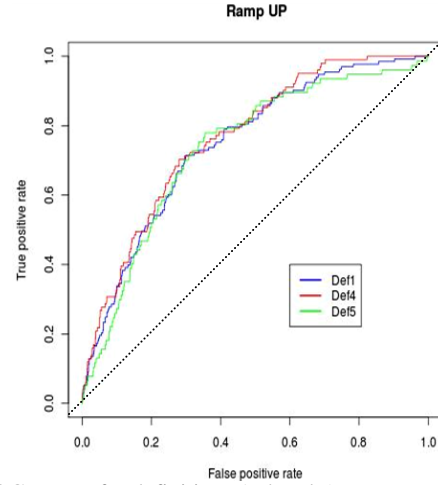

Fig. 6. ROC curves for definitions 1, 4 and 5.

Assume binary forecasts, where each example can be labeled using one of two classes in the set {y,n}, and a forecast can output the corresponding {Y,N}. Consider that we know the distribution of yes and no events, i.e., the probabilities P(y) and P(n), and that we define the costs *cost*(Y;n) and *cost*(N;y) to be, respectively, the costs of predicting a event when no event occurs (False Positive) and the cost of predicting no event when an event actually occurs (False Negative). The slope of the line (a tangent line) that touches the ROC curve at the optimum operating point – a point with coordinates (FPR$_0$, TPR$_0$) that is associated with a probability threshold *thr* – is given by P(n)cost(Y;n)/P(y)cost(N;y).

If this distribution is unknown, we can estimate the distribution from the observations. The point (FPR_0, TPR_0) where the tangent line and the curve touch is the optimum operating point, in the sense that this point minimizes the Expected Cost given by the following expression:

$$\text{EC=P(Y;n)} \times \text{cost(Y;n)+P(N;y)} \times \text{cost(N;y)} \quad (11)$$

where P(Y;n) is the probability of predicting an event when it does not occur (probability of a false positive) and P(N;y) is the probability of predicting that an event does not occur when it really occurs (probability of a false negative).

Fig. 7 presents the tangent lines, and identify the associated optimum operating point, including the optimal *thr$_0$* associated to Eq. (9), that we get by running definition 1 and setting two error cost configurations to predict ramp-up events. In Conf. 1 we define the error costs to be *cost* (N;y) = 200, i.e. the cost of a FN (cFN), and *cost* (Y;n) = 10, i.e. the cost of a FP (cFP). In Conf. 2 we consider *cost* (N;y) = 10 and *cost* (Y;n) = 200.

The ROC curve is annotated with the corresponding empirical probability values, observed on the set of scenarios. The point obtained from optimizing Eq. (11) is associated to the optimal threshold *thr$_0$* value. If the probability calculated is above *thr$_0$*, one should declare the prediction of occurrence of a ramp; if not, a prediction of no ramp. Referring to Fig. 5, this declaration was produced for the intervals marked with a star, where the probability of a ramp of magnitude above 25% of the wind farm nominal power exceeds the optimal *thr$_0$* $\approx$ 0.2 obtained from optimizing Eq. (11) over the ROC curve.
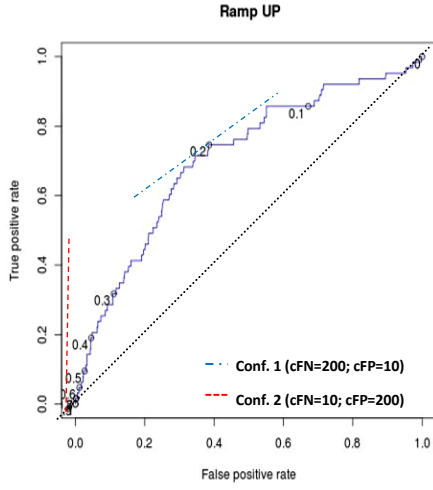
Fig. 7. ROC curve and tangent lines for definition 1 and two cost configurations.
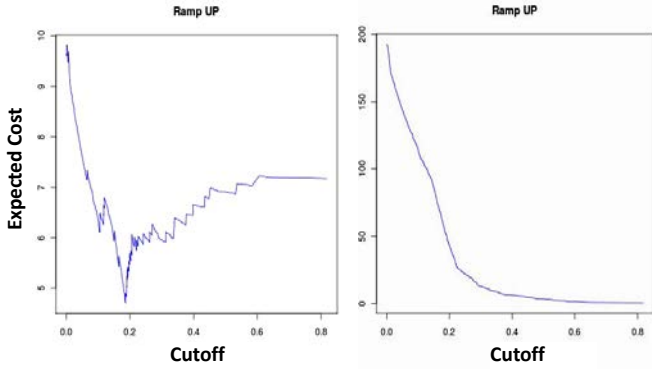

Fig. 8. Expected Cost using definition 1: cFN=200; cFP=10, on the left, and cFN=10 and cFP=200, on the right.

Fig. 8 plots the expected cost for a set of probability thresholds. These plots were generated by running definition 1 to identify ramp-ups. In these figures we can easily identify the minimum expected costs that define the cut-off probability threshold corresponding to the optimum operating point.

Forecasters that assign a probability to each event often use the Brier Score (BS) [10], which is a score function that measures the average squared deviation between predicted probabilities and the actual outcomes. It is computed as

$$BS = \frac{1}{N}\sum_{t=1}^{N}(X_t - Y_t)^2 \qquad (11)$$

where X_t is the event forecast probability; Y_t is the actual outcome (0 if not happened, 1 if happened) and N – number of forecasting instances. It is obvious that the optimum Brier Score would be of 0, for perfect sharp predictions. The Brier Score obtained in the experiment being described gave the values described in Table I. By inspecting these results we can see that by using definition 1 we get a lower BS when detecting ramp-ups. In contrast, definitions 4 and 5 get lower BS for detection of ramp down events.

	RAMP UP			RAMP DOWN		
	D1	*D4*	*D5*	*D1*	*D4*	*D5*
BS	0.078	0.091	0.082	0.080	0.086	0.076

Another useful metric is the critical success index (CSI), defined as:

$$CSI = \frac{TP}{TP+FN+FP} \qquad (12)$$

The CSI metric takes values in the interval [0;1], where 1 means correct prediction. CSI measures the fraction of observed and/or forecast events that were correctly predicted.

In Table II we present CSI values obtained in experiments to predict upward ramp events, for definitions 1, 4 and 5 and considering phase errors for time lags of 2 and 4 periods. The probability cut-off value is 0.3. The results for ramp-down display a similar performance.

Regarding other definitions and parameters, we can say that the performance of our model improves when we consider large sizes of the time-step (Δt), large aggregation windows and, obviously, consider phase error. Overall, we can say that we obtain the best results of our experiments by running the detection process for ramp definitions 4 and 5.

TABLE II.   CSI COMPARISON AGAINST A POINT FORECAST SYSTEM: CSI FOR THREE RAMP DEFINITIONS AND CONSIDERING PHASE ERRORS WITH A LAG OF 2 AND 4 HOURS.

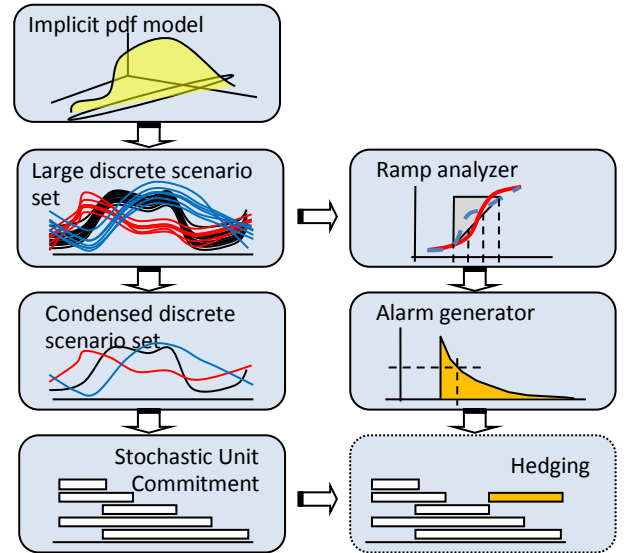|	RAMP UP – CSI					
	Probabilistic Forecast			*Point Forecast*		
Phase Error	*D1*	*D4*	*D5*	*D1*	*D4*	*D5*
---	0.12	0.20	0.15	0.09	0.17	0.08
2	0.30	0.36	0.32	0.18	0.31	0.24
4	0.38	0.45	0.38	0.26	0.37	0.32

Fig. 9. Conceptual modules relating scenario generation, unit commitment and ramp event analysis.

VI. Ramp forecasting and unit commitment

Fig. 9 makes explicit the role of probabilistic ramp forecasting. By setting alarms at specific hours and by defining probabilities associated with ramp presence and amplitude, the model establishes risk levels for ramp events – in the form of a probability of having a ramp of a given amplitude or greater. A system operator may then decide, based on the risks he is willing to take, whether to accept a specific unit commitment or to hedge against the adverse event and plan for some extra reserve, at some cost. This also indicates how system operators may take full advantage of probabilistic wind power forecasting models – to adopt a stochastic model for the unit commitment exercise. However, even if a classical model is used, the reasoning about ramp risks and hedging applies.

VII. Conclusions

Ramp forecasting has been recognized as a difficult exercise, related with predictions about the derivative of a time series. But it is of the utmost importance to take in account the possibility of ramp events, especially in systems with high penetration of wind and where the remainder generation resource is dominantly of the slow thermal type, such as coal and nuclear power. The problem is somewhat less serious if the power system has significant hydro generation or gas turbines – but the issue is about system security and costs incurred in mitigating risks. i.e. the cost of adopting hedging policies.

To adequately assess costs and risks, a probabilistic approach is mandatory. This paper brings to light a new approach to define a probabilistic model for ramp forecasting, in a form that it may be useful for system operators that have to define generation unit commitment. The output is a histogram, at every hour, indicating the probability of a ramp exceeding a certain threshold, for all magnitudes above a minimum level defined. This does not provide dispatch decision suggestions – however, the declaration of a ramp event based on a probability threshold contributes to the decision process: it may serve as input to a hedging process, where the operator may decide if he runs the risk of being subject to a ramp event (e.g., in the case of a low probability for an event of a damaging magnitude) or if he hedges by changing the unit commitment in an appropriate (more costly) way to avoid problems in case the event materializes.

Tests with real data from a US wind farm have proved the validity and usefulness of the approach. The experimental results, evaluated under the ROC curve concept, show clear advantages of the probabilistic forecaster over point forecasts and random guesses. It must be said that the quality of ramp forecasting depends a great deal on the quality of meteorological forecasts, translated into Numerical Weather Models, which supply data to the general short term wind power forecasting problem.

The model does not take in account very short term meteorological/wind predictions, as they are not usually available. Nevertheless, the work presented here, by assigning a probability to each possible ramp magnitude, is a clear step forward, providing a methodology useful to the industry.

References

[1] C. Ferreira, J. Gama, V. Miranda, A. Botterud, "A survey on wind power ramp forecasting", Report ANL/DIS 10-13, Argonne National Laboratory, Dec. 2010.

[2] B. Greaves et al., "Temporal Forecast Uncertainty for Ramp Events", Proc. EWEC'09, Marseille, France, 2009.

[3] C.W. Potter, E. Grimit and B. Nijssen, "Potential Benefits of a Dedicated Probabilistic Rapid Ramp Event Forecast Tool," Proc. PSCE'09, Seattle, USA, 2009.

[4] J. Freedman, M. Markus and R. Penc, "Analysis of West Texas Wind Plant Ramp-up and Ramp-down Events," AWS Truewind, LLC, Albany, NY, 2008.

[5] C. Kamath, "Understanding Wind Ramp Events Through Analysis of Historical Data," Proc. IEEE PES Trans. and Dist. Conf. and Expo, New Orleans, LA, United States, 2010.

[6] H. Zheng and A. Kusiak, "Prediction of Wind Farm Power Ramp Rates: A Data-Mining Approach", J. Solar Energy Engineering, vol. 131, 2009.

[7] J.U. Jørgensen and C. Mörlen, "AESO Wind Power Forecasting Pilot Project", Ebberup, Denmark, Technical Report 2008.

[8] AWS Truewind-LLC, "AWS Truewind's Final Report for the Alberta Forecasting Pilot Project", Alberta, Canada, 2008.

[9] J. T. Schaefer, "The Critical Success Index as an Indicator of Warning Skill," Weather Forec., nº 5, pp. 570–575, 1990.

[10] A. P. Weigel, M.A. Liniger and C. Appenzeller, "Generalization of the Discrete Brier and Ranked Probability Skill Scores for Weighted Multimodel Ensemble Forecasts", Monthly Weather Review, vol. 135, no. 1, pp. 118-124, 2007.

[11] A. Bossavy, R. Girard and G. Kariniotakis, "Forecasting Uncertainty Related to Ramps of Wind Power Production", Proc.EWEC'10, Warsaw, Poland, 2010.

[12] H. A. Nielsen, H. Madsen, and T. S. Nielsen, "Using quantile regression to extend an existing Wind Power Forecasting system with probabilistic forecasts," Wind Energy, vol. 9(1-2), pp. 95–108, 2006.

[13] J. Juban, N. Siebert and G. N. Kariniotakis, "Probabilistic Short-term Wind Power Forecasting for the Optimal Management of Wind Generation", Proceedings IEEE PowerTech, Lausanne, France, pp. 683-688, 2007.

[14] R.J. Bessa, V. Miranda and J. Gama, "Entropy and Correntropy Against Minimum Square Error in Offline and Online Three-day Ahead Wind Power Forecasting", IEEE Trans. on Power Sys., vol. 24, no. 4, pp. pp. 1657-1666, Nov. 2009.

[15] F. Provost and T. Fawcett, "Analysis and visualization of classifier performance: Comparison under imprecise class and cost distributions", KDD'97, USA, 1997, pp. 43-48.