

Chance Constrained Optimal Power Flow: Risk-Aware Network Control under Uncertainty

Daniel Bienstock^{*}, Michael Chertkov[†], Sean Harnett^{* †}

^{*}Department of Industrial Engineering and Operations Research and Department of Applied Physics and Applied Mathematics, Columbia University, 500 West 120th St. New York, NY 10027 USA, and [†]Theoretical Division and Center for Nonlinear Studies, Los Alamos National Laboratory, NM 87545 USA

Submitted to Proceedings of the National Academy of Sciences of the United States of America

When uncontrollable resources fluctuate, Optimum Power Flow (OPF), routinely used by the electric power industry to re-dispatch hourly controllable generation (coal, gas and hydro plants) over control areas of transmission networks, can result in grid instability, and, potentially, cascading outages. This risk arises because OPF dispatch is computed without awareness of major uncertainty, in particular fluctuations in renewable output. As a result, grid operation under OPF with renewable variability can lead to frequent conditions where power line flow ratings are significantly exceeded. Such a condition, which is borne by simulations of real grids, would likely result in automatic line tripping to protect lines from thermal stress, a risky and undesirable outcome which compromises stability. Smart grid goals include a commitment to large penetration of highly fluctuating renewables, thus calling to reconsider current practices, in particular the use of standard OPF. Our Chance Constrained (CC) OPF corrects the problem and mitigates dangerous renewable fluctuations with minimal changes in the current operational procedure. Assuming availability of a reliable wind forecast parameterizing the distribution function of the uncertain generation, our CC-OPF satisfies all the constraints with high probability while simultaneously minimizing the cost of economic re-dispatch. CC-OPF allows efficient implementation, e.g. solving a typical instance over the 2746-bus Polish network in 20s on a standard laptop.

Significance statement. A critical component of modern power grids is the computation that determines constant generator output levels to use during the next short term (15-minute) time window, relying on static demand forecasts. In normal operation these fixed output levels are automatically adjusted in real-time in proportion to demand changes. The aggressive incorporation of wind generation renders this setup problematic due to large random wind fluctuations: it causes equipment overloads thus impacting grid stability. We present an effective and scientifically rigorous methodology for setting generator outputs in a risk aware manner; our experiments indicate that our methodology is both fast and effective.

Power Flows | Uncertainty | Wind Farms | Networks

Abbreviations: OPF, Optimal Power Flow; CC-OPF, Chance Constrained Optimum Power Flow; UC, Unit Commitment; AC, Alternating Current; DC, Directed Current; PF, Power Flow; GB, Generator Bounds; TL, Thermal Limits; SOCP, Second Order Conic Programming

The power grid can be considered one of the greatest engineering achievements of the 20th century, responsible for the economic well-being, social development, and resulting political stability of billions of people around the globe. The grid is able to deliver on these goals with only occasional disruptions through significant control sophistication and careful long-term planning.

Nevertheless, the grid is under growing stress and the premise of secure electrical power delivered anywhere and at any time may become less certain. Even though utilities have massively invested in infrastructure, grid failures, in the form of large-scale power outages, occur unpredictably and with increasing frequency. In general, automatic grid control and regulatory statutes achieve robustness of operation as conditions display normal fluctuations, in particular approxi-

mately predicted inter-day trends in demand, or even unexpected single points of failure, such as the failure of a generator or tripping of a single line. However, larger, unexpected disturbances can prove quite difficult to overcome. This difficulty can be explained by the fact that automatic controls found in the grid are largely of an engineering nature (i.e. the flywheel-directed generator response used to handle short-term demand changes locally) and are largely *not* of a data-driven, algorithmic and distributed nature. Instead, should an unusual condition arise, current grid operation relies on human input. Additionally, only some real-time data is actually used by the grid to respond to evolving conditions.

All engineering fields can be expected to change as computing becomes ever more enmeshed into operations and massive amounts of real-time data become available. In the case of the grid this




Fig. 1. Bonneville Power Administration [1] shown in outline under 9% wind penetration, where green dots mark actual wind farms. We set standard deviation to be 0.3 of the mean for each wind source. Our CC-OPF (with 1% of overload set as allowable) resolved the case successfully (no overloads), while the standard OPF showed 8 overloaded lines, all marked in color. Lines shown orange are at 4% chance of overload. There are two dark red lines which are at 50% of the overload while other (dark orange) lines show values of overload around 10%.

Reserved for Publication Footnotes

change amounts to a challenge; namely how to migrate to a more algorithmic-driven grid in a cost-effective manner that is also seamless and gradual so as not to prove excessively disruptive – because it would be impossible to rebuild the grid from scratch. One of the benefits of the migration, in particular, concerns the effective integration of renewables into grids. This issue is critical because large-scale introduction of renewables as a generation source brings with it the risk of large, random variability – a condition that the current grid was not developed to accommodate.

This issue becomes clear when we consider how the grid sets generator output in “real time”. This is typically performed at the start of every fifteen-minute (to an hour) period, or time window, using fixed estimates for conditions during the period. More precisely, generators are dispatched so as to balance load (demand) and generator output at minimum cost, while adhering to operating limitations of the generators and transmission lines; estimates of the typical loads for the upcoming time window are employed in this computation. This computational scheme, called Optimum Power Flow (OPF) or economic dispatch, can fail, dramatically, when renewables are part of the generation mix and (exogenous) fluctuations in renewable output become large. By “failure” we mean, in particular, instances where a combination of generator and renewable outputs conspire to produce power flows that significantly exceed power line ratings. When a line’s rating is exceeded, the likelihood grows that the line will become tripped (be taken uncontrollably out of service) thus compromising integrity and stability of the grid. If several key lines become tripped a grid would very likely become unstable and possibly experience a cascading failure, with large losses in serviced demand. This is not an idle assumption, since firm commitments to major renewable penetration are in place throughout the world. For example, 20% renewable penetration by 2030 is a decree in the U.S. [2], and similar plans are to be implemented in Europe, see e.g. discussions in [3, 4, 5]. At the same time, operational margins (between typical power flows and line ratings) are decreasing and expected to decrease.

A possible failure scenario is illustrated in Fig. 1 using as example the U.S. Pacific Northwest regional grid data (2866 lines, 2209 buses, 176 generators and 18 wind sources), where lines highlighted in red are jeopardized (flow becomes too high) with unacceptably high probability by fluctuating wind resources positioned along the Columbia river basin (green dots marking existing wind farms). We propose a solution that requires, as the only additional investment, accurate wind forecasts; but no change in machinery or significant operational procedures. Instead, we propose an intelligent way to modify the optimization approach so as to mitigate risk; the approach is implementable as an efficient algorithm that solves large-scale realistic examples in a matter of seconds, and thus is only slightly slower than standard economic dispatch methods.

Maintaining line flows within their prescribed limits arises as a paramount operational criterion toward grid stability. In the context of incorporating renewables into generation, a challenge emerges because a nominally safe way of operating a grid may become unsafe – should an unpredictable (but persistent) change in renewable output occur, the resulting power flows may cause a line to persistently exceed its rating. It is natural to assess the risk of such an event in terms of probabilities, because of the non-deterministic behavior of e.g. wind; thus in our proposed operational solution we will rely on techniques involving both mathematical optimization and risk analysis.

When considering a system under stochastic risk, an extremely large variety of events that could pose danger might emerge. Recent works [6, 7, 8] suggest that focusing on instantons, or most-likely (dangerous) events, provides a practicable route to risk control and assessment. However, there may be far too many comparably probable instantons, and furthermore, identifying such events does not

answer the question of what to do about them. In other words, we need a computationally efficient methodology that not only identifies dangerous, relatively probable events, but also mitigates them.

This paper suggests a new approach for handling the two challenges, that is to say, searching for the most probable realizations of line overloads under renewable generation, and correcting such situations through control actions, simultaneously and efficiently in one step. Our approach relies on methodologies recently developed in the optimization literature and known under the name of “Chance-Constrained” (CC) optimization [9]. Broadly speaking, CC optimization problems are optimization problems involving stochastic quantities, where constraints state that the probability of a certain random event is kept smaller than a target value.

To address these goals, we propose an enhancement of the standard OPF to be used in the economic dispatch of the controllable generators. We model each bus that houses a power source subject to randomness to include a random power injection, and reformulate the standard OPF in order to account for this uncertainty. The formulation minimizes the average cost of generation over the random power injections, while specifying a mechanism by which (standard, i.e. controllable) generators compensate in real-time for renewable power fluctuations; at the same time guaranteeing low probability that any line will exceed its rating. This last constraint is naturally formulated as a chance constraint – we term our approach Chance-Constrained OPF, or CC-OPF.

This paper is organized as follows. In “Models” we first describe the various mathematical models used to describe how the grid operates, as well as our proposed methodology. We then present, in “Experiments” a number of examples to demonstrate the speed and usefulness of our approach. Finally, in “Methods” we provide in-depth technical details on our approach. Further details, and proofs, are provided in the “Supporting Information” (SI).

Models

Transmission Grids: Controls and Limits. The power systems we consider in this paper are *transmission grids* which operate at high voltages so as to convey power economically, with minimal losses, over large distances. This is to be contrasted with *distribution systems*; typically residential, lower voltage grids used to provide power to individual consumers. From the point of view of wind-power generation, smooth operation of transmission systems is key since reliable wind sources are frequently located far away from consumption.

Transmission systems balance consumption/load and generation using a complex strategy that spans three different time scales (see e.g. [10]). At any point in time, generators produce power at a previously computed base level. Power is generated (and transmitted) in the Alternating Current (AC) form. An essential ingredient toward stability of the overall grid is that all generators operate at a common frequency. In real time, changes in loads are registered at generators through (opposite) changes in frequency. A good example is that where there is an overall load increase. In that case generators will marginally slow down – frequency will start to drop. Then the so-called primary frequency control, normally implemented on gas and hydro plants with so-called “governor” capability will react so as to stop frequency drift (large coal and nuclear units are normally kept on a time constant output). This is achieved by having each responding generator convey more power to the system, proportionally to the frequency change sensed. (In North America the proportionality coefficient is normally set to 5% of the generator capacity for 0.5Hz deviation from the nominal frequency of 60Hz.) This reaction is swift and local, leading to stabilization of frequency across the system, however not necessarily at the nominal 60Hz value. The task of the secondary, or Automatic Gain Control (AGC), is to adjust the common frequency mismatch and thus to restore the overall balance between generation and consumption, typically in a matter of minutes. Only some of the generators in a local area may be involved in this step.

The final component in the strategy is the tertiary level of control, executed via the OPF algorithm, typically run as frequently as every fifteen minutes (to one hour), and using estimates for loads during the next time window, where base (controllable) generator outputs are reset. This is not an automatic step in the sense that a computation is performed to set these generator levels; the computation takes into account not only load levels but also other parameters of importance, such as line transmission levels. Tertiary control computation, which is in the center of this paper, thus represents the shortest time scale where actual off-line and network wide (in contrast to automatic primary and secondary controls of frequency) optimal computations are employed. The three levels are not the only control actions used to operate a transmission system. Advancing further in the time scale, OPF is followed by the so-called Unit Commitment (UC) computation, which schedules the switching on and off of large generation units on the scale of hours or even days.

A critical design consideration at each of the three control levels is that of maintaining “stability” of the grid. The most important ingredient toward stable operation is *synchrony* – ultimately, all the generators of the network should stabilize thus locking, after a perturbation followed by a seconds-short transient, at the same frequency. Failure to do so not only proves inefficient but, worse, it threatens the integrity of the grid, ultimately forcing generators to shut down for protective reasons – thus, potentially, causing a large, sudden change in power flow patterns (which may exceed equipment limits, see below) and possibly also an unrecoverable generation shortage. A second stability goal is that of maintaining large voltages. This is conducive to efficiency; lower voltage levels cause as a byproduct more generation (to meet the loads) and larger current values. Not only is this combination inefficient, but in an extreme case it may make impossible to meet existing loads (so-called “voltage collapse” is a manifestation of this problem). The third stability goal, from an operational perspective, is that of maintaining (line) power flows within established bounds. In long transmission lines, a large flow value will cause excessive voltage drop (an undesirable outcome as discussed). On a comparatively shorter line, an excessively large power flow across the line will increase the line temperature to the point that the line sags, and potentially arcs or trips due to a physical contact. For each line there is a given parameter, the *line rating* (or *limit*) which upper bounds flow level during satisfactory operation.

Of the three “stability” criteria described above, the first two (maintaining synchrony and voltage) are a concern only in a truly nonlinear regime which under normal circumstances occur rarely. Thus, we focus on the third – observing line limits.

OPF – Standard Generation Dispatch (tertiary control). OPF is a key underlying algorithm of power engineering, see the review in [11] and e.g. [12, 10]. The task of OPF, usually executed off-line at periodic intervals, is to re-dispatch generation over a control area of the transmission grid, for example over the Bonneville Power Administration (BPA) grid shown in Fig. 1. In outline the standard OPF can be stated as the following constrained optimization problem:

$$\begin{aligned} \text{OPF:} \quad & \min_p c(p), \quad \text{s.t.} \quad [1] \\ \text{PF:} \quad & \text{Power Flow Equations} \\ \text{TL:} \quad & \text{Line flow (capacity) Limits} \\ \text{GB:} \quad & \text{Generation Bounds} \end{aligned}$$

Here, $p = (p_i | i \in \mathcal{G})$ is the vector of controllable/re-dispatchable generation available at the subset \mathcal{G} of nodes of the full set of grid nodes, \mathcal{V} ; $\mathcal{G} \subset \mathcal{V}$. The above problem is endowed by three types of constraints: Power flow (PF), Line Limit (TL) and Generation Bound (GB) constraints. (PF) consists of AC Power Flow equations (AC-PF) which are simply Kirchoff’s circuit laws stated in terms of power flows and potentials (voltages). Here, for each node $i \in \mathcal{V}$ its voltage U_i is defined as $v_i e^{j\theta_i}$, where v_i and θ_i are the voltage magnitude

and phase angle at node i . See e.g. [12, 10]. The power flow equations are quadratic and thus can constitute an obstacle to solvability of OPF (from a technical standpoint, they give rise to nonconvexities). In transmission system analysis a linearized version of the AC equations is commonly used, the so-called “DC-approximation”. In this approximation (a) all voltages are assumed fixed and re-scaled to unity; (b) phase differences between neighboring nodes are assumed small, $\forall \{i, j\} \in \mathcal{E} : |\theta_i - \theta_j| \ll 1$, where \mathcal{E} stands for the set of the grid edges, or lines; (c) thermal losses are ignored (reactance dominates resistance for all lines). Then, the power flow over line $\{i, j\}$, with line susceptance β_{ij} ($= \beta_{ji}$) is related linearly to the respective phase difference, $f_{ij} = \beta_{ij}(\theta_i - \theta_j)$, while the balance of power can be stated in the following matrix form

$$\begin{aligned} \text{PF:} \quad & B\theta = p - d, \quad \text{where } B = (B_{ij} | i, j \in \mathcal{V}), \quad [2] \\ \forall i, j: \quad & B_{ij} = \begin{cases} -\beta_{ij}, & \{i, j\} \in \mathcal{E} \\ \sum_{k; \{k, j\} \in \mathcal{E}} \beta_{kj}, & i = j \\ 0, & \text{otherwise} \end{cases} \quad [3] \end{aligned}$$

where d is the vector of the exogenous (uncontrollable) demands at each node (possibly equal to zero at some nodes). We also model an uncontrollable demand as a negative load. The linearized approximation is often considered around a stationary solution of the full AC-PF system (stationary operational point); thus elements of the susceptance matrix B account for renormalization due to the base case solution. The GB and TL constraints in problem [1] are

$$\text{GB:} \quad \forall i \in \mathcal{G} : p_i^{min} \leq p_i \leq p_i^{max}, \quad [4]$$

$$\text{TL:} \quad \forall \{i, j\} \in \mathcal{E} : |f_{ij}| \leq f_{ij}^{max}, \quad [5]$$

where p_i^{min}, p_i^{max} are lower and upper generation bounds. f_{ij}^{max} represents the line limit (typically a thermal limit), which is assumed to be strictly enforced in constraint (TL). This conservative condition will be relaxed in the following. Finally, the objective $c(p)$ to minimize in problem [1] is a sum of convex quadratic functions of the components of p (fixed price curve per generator). In summary, problem [1] is a convex optimization problem solved for a fixed vector of demands, d . In practice, however, demand will fluctuate around d ; generators then respond by adjusting their output proportionally to the overall fluctuation as discussed above in relation to frequency control. When some of the generation is due to renewables, standard OPF would model their output as constant (in the time window of interest), and would manage their fluctuation by having controllable generators adjust their output in the same way used to handle demand variations.

Using modern optimization tools [1] is an easily solved problem. This scheme works well in current practice, as demands do not substantially fluctuate on the time scale for which OPF applies. Thus the standard practice of solving [1] in the feasibility domain defined by Eq. [2],[3],[4],[5], using demand forecasts based on historical data (and ignoring fluctuations) has produced a very reliable result - generation re-dispatch covering a span of fifteen minutes to an hour, depending on the system.

Chance constrained OPF: motivation. The separation of generation control into the hierarchy of primary, secondary and OPF has worked in the past because of the slow time scales of change in uncontrolled resources (mainly loads). That is to say, frequency control and load changes were well-separated. Clearly, an error in the forecast or an under-estimation of possible d for the next –e.g., fifteen minute– period may lead to an operational problem in standard OPF; see e.g. the discussions in [13, 14]. This was not considered a significant handicap until recently, however, simply because line trips due to overloading as a result of OPF-directed generator dispatch were (and still are) rare. The projected increase of renewable penetration in the future, accompanied by the decreasing gap between normal operation and limits set by line capacities, will make these overload events more frequent and generally increase risk (see [3]). One would also suspect

that the rare overload event in the grid of today is due to setting up the TL limits too conservatively. In general, lowering the TL limits succeeds in preventing overloads, but it also forces excessively conservative choices of the generation re-dispatch, potentially causing extreme volatility of the electricity markets. (See e.g. the discussion in [15] on abnormal price fluctuations in ERCOT and New Zealand markets, which are both heavily reliant on renewables.) CC-OPF, introduced next, is less conservative (it is probabilistic) and also offers an exact and efficient algorithm for balancing the cost of operation with the risk of overload. We will introduce uncertain power sources into the OPF. This will change the optimization problem from deterministic to probabilistic; we will seek to minimize average cost, with the previously hard (deterministic) constraints becoming *chance constraints* as explained next.

Redefining the line flow constraints. Power lines do not fail (i.e., trip) instantly when their flow limits are exceeded. A line carrying flow that exceeds the line’s thermal limit will gradually heat up and possibly sag, increasing the probability of an arc (short circuit) or even a contact with neighboring lines, with ground, with vegetation or some other object. Each of these events will result in a trip. The precise process is extremely difficult to calibrate (this would require, among other factors, an accurate representation of wind strength and direction in the proximity of the line)¹. Additionally, the rate at which a line overheats depends on its overload which may dynamically change (or even temporarily disappear) as flows adjust due to external factors; in our case fluctuations in renewable outputs. What *can* be stated with certainty is that the longer a line stays *overheated*, the higher the probability that it will trip – to put it differently, if a line remains overheated long enough, then, after a span possibly measured in minutes, it will trip. In summary, (thermal) tripping of a line is primarily governed by the historical pattern of the overloads experienced by the line, and thus it may be influenced by the status of other lines (implicitly, via changes in power flows); further, exogenous factors can augment the impact of overloads.

Even though an exact representation of line tripping seems difficult, we can however state a practicable alternative. Ideally, we would use a constraint of the form “for each line, the fraction of the time that it exceeds its limit within a certain time window is small”. Direct implementation of this constraint would require resolving dynamics of the grid over the generator dispatch time window of interest. To avoid this complication, we propose instead the following static proxy of this ideal model, a *chance constraint*: we will require that the probability that a given line will exceed its limit is small. Let f_{ij} be the flow on line $\{i, j\}$, where the bold face indicates that it is a random quantity. Denote the “small probability” above by ϵ_{ij} , and the flow limit on line $\{i, j\}$ by f_{ij}^{max} . Then the chance constraint for each line, $\{i, j\}$, is:

$$\text{CCTL} : \quad \forall \quad P(|f_{ij}| > f_{ij}^{max}) < \epsilon_{ij}. \quad [6]$$

Chance constraints [17], [18], [19] are but one possible methodology for handling uncertain data in optimization. Broadly speaking, this methodology fits within the general field of stochastic optimization. Constraint [6] can be viewed as a “value-at-risk” statement; the closely-related “conditional value at risk” concept provides a (convex) alternative, which roughly stated constrains the expected overload of a line to remain small, *conditional* on there being an overload (see [9] for definitions and details).

Yet another alternative would be to rely on *robust optimization*. In this setting we would view the output of a renewable source j as (for example) an unknown quantity constrained to lie in an interval $[l_j, u_j]$, and formulate an optimization problem that requires (for example) that the flow of each line stay within its limit no matter what the value of each renewable output j is, so long as it remains in its corresponding interval $[l_j, u_j]$. Other variations are possible, but always with the same deterministic flavor. See e.g. [20]. This general approach is attractive because it makes few assumptions on the un-

derlying uncertain process. However, we would argue that our chance constrained approach is reasonable (in fact: compelling) in view of the nature of the line tripping process we discussed above.

Uncertain power sources. We assume that there is a collection of wind sources (farms), with one wind source located at each node in a given subset \mathcal{W} of \mathcal{V} . For each $i \in \mathcal{W}$, the amount of power generated by source i is assumed to be of the form $\mu_i + w_i$, where μ_i is constant, assumed known from the forecast, and w_i is a zero mean independent random variable with standard deviation σ_i .

The physical assumptions behind this model of uncertainty are as follows. Independence of fluctuations at different sites is due to the fact that the wind farms are sufficiently far away from each other. For the typical OPF time span of 15 min and typical wind speed of 10m/s, fluctuations of wind at the farms more than 10km apart are not correlated. We also rely on the assumption that transformations from wind to power at different wind farms is not correlated.

To formulate and calibrate our models, we make simplifying assumptions that are approximately consistent with our general physics understanding of fluctuations in atmospheric turbulence; in particular we assume Gaussianity of w_i ². We will also assume that only a standard weather forecast (coarse-grained on minutes and kilometers) is available, and no systematic spillage of wind in its transformation to power is applied³.

We also have a strong additional –and purely computational– reason for the Gaussian assumption. As we will show below, under this assumption chance constraint [6] can be captured using a tractable deterministic optimization problem; our substitute for standard OPF (see below for details). As we will indicate below, in the case of non-Gaussian distributions our approach is easily modified so as to retain tractability, but at the cost of relying on a conservative approximation of the chance constraint.

Other fitting distributions considered in the wind-modeling literature, e.g. Weibull distributions and logistic distributions [21, 22], will be discussed later in the text as well⁴. In particular, we will demonstrate on out-of-sample tests that the computationally advantageous Gaussian modeling of uncertainty allows as well to model effects of other distributions. Our approach relies on a “robust” wind forecast; typically this would involve obtaining a reasonably accurate estimate on mean wind strength and a conservative variance estimation at each farm. This robust forecast will be used to compute our control as indicated in the next section. The overall edifice should be validated with actual data or out-of-sample simulations. (See Discussion and Methods for details.)

Affine Control. Since the power injections at each bus are fluctuating, we need a control to ensure that generation is equal to demand at all times within the time interval between two consecutive OPFs. We term the joint result of the primary control and secondary control the *affine* control. The term will intrinsically assume that all governors involved in the controls respond to fluctuations in the generalized load (actual demand which is assumed frozen minus stochastic wind resources) in a proportional way, however with possibly differ-

¹We refer the reader to [16] for discussions of line tripping during the 2003 Northeast U.S.-Canada cascading failure.

²Correlations of velocity within the correlation time of 15 min, roughly equivalent to the time span between the two consecutive OPF, are approximately Gaussian. The assumption is not perfect, in particular because it ignores significant up and down ramps possibly extending tails of the distribution in the regime of really large deviations.

³See [3], for some empirical validation.

⁴Note that the fitting approach of [21, 22] does not differentiate between typical and atypical events and assumes that the main body and the tail should be modeled using a simple distribution with only one or two fitting parameters. Generally this assumption is not justified as the physical origin of the typical and anomalous contributions of the wind, contributing to the main body and the tail of the distribution respectively, are rather different. Gaussian fit (of the tail) – or more accurately, faster than exponential decay of probability in the tail for relatively short-time (under one hour) forecast – would be reasonably consistent with phenomenological modeling of turbulence generating these fluctuations.

ent proportionality coefficients:

$$\forall \text{node } i \in \mathcal{G} : p_i = \bar{p}_i - \alpha_i \sum_{j \in \mathcal{W}} \omega_j. \quad \alpha_i \geq 0, \quad [7]$$

Here the quantities $\bar{p}_i \geq 0$ and $\alpha_i \geq 0$ are design variables satisfying (among other constraints) $\sum_{i \in \mathcal{G}} \alpha_i = 1$. Notice that we do not set any α_i to a standard (fixed) value, but instead leave the optimization to decide the optimal value. (In some cases it may even be advantageous to allow negative α_i but we decided not to consider such a drastic change of current policy in this study.) The generator output p_i combines a fixed term \bar{p}_i and a term which varies with wind, $-\alpha_i \sum_{j \in \mathcal{W}} \omega_j$. Observe that $\sum_i p_i = \sum_i \bar{p}_i - \sum_i \omega_j$, that is, the total power generated equals the average production of the generators minus any additional wind power above the average case.

This affine control scheme creates the possibility of requiring a generator to produce power beyond its limits. With unbounded wind, this possibility is inevitable, though we can restrict it to occur with arbitrarily small probability, which we will do with additional chance constraints for all controllable generators, $\forall g \in \mathcal{G}$,

$$\text{CCgen} : P(p_g^{min} \leq \bar{p}_g - \alpha_g \sum_{j \in \mathcal{W}} \omega_j \leq p_g^{max}) > 1 - \epsilon_g. \quad [8]$$

CC-OPF: Formal Expression. We can now formally state the main optimization problem we introduce in this manuscript. CC-OPF is given by the following modification of the standard OPF problem [1]:

$$\begin{aligned} \text{CC-OPF:} \quad & \min_{\bar{p}, \alpha} \mathbb{E}_w [c(\bar{p})] \quad \text{s.t.} \quad [9] \\ \text{PFAV:} \quad & \text{Power Flow Eqs. (2) under average wind} \\ \text{CCTL:} \quad & \text{Chance constrained line limits, Eq. (6)} \\ \text{CCgen:} \quad & \text{Chance constrained generator bounds, Eq. (8)} \end{aligned}$$

where $\mathbb{E}_w [c(\bar{p}, \alpha)]$ is the expected cost of generation over the varying wind power w . If the cost function c is convex quadratic (standard practice), then so is this expectation. The expected cost objective and the chance constrained condition in **CC-OPF** both account for fluctuations in wind and also for the standard generation adjusting to these fluctuations via the aforementioned proportional control.

[23] considers the standard OPF problem under stochastic demands, and describes a method that computes *fixed* generator output levels to be used throughout the period of interest, independent of demand levels. In order to handle variations in demand, [23] instead relies on the concept of a *slack bus*. A slack bus is a fixed node that is assumed to compensate for all generation/demand mismatches – when demand exceeds generation the slack bus injects the shortfall, and when demand is smaller than generation the slack bus absorbs the generation excess. A vector of generations is acceptable if the probability that each system component operates within acceptable bounds is high – this is a chance constraint. To tackle this problem [23] proposes a simulation-based local optimization system consisting of an outer loop used to assess the validity of a control (and estimate its gradient) together with an inner loop that seeks to improve the control. Experiments are presented using a 5-bus and a 30-bus example.

Chance constrained optimization has also been discussed recently in the Unit Commitment setting – discrete-time planning for operation of large generation units on the scale of hours-to-months accounting for the long-term wind-farm generation uncertainty [24, 25, 26].

At first glance CC-OPF constitutes a significantly more difficult problem than the original OPF. Not only does this formulation contain additional (affine control) variables not present in the standard OPF, but even more significantly the chance constraints and the objective render it into a stochastic optimization problem, requiring the evaluation of expectation and the probability of overload over the exogenous statistics of the wind. A direct computational approach to solving problem [9], as in [23], though universal (applicable to any

type of exogenous distribution), would require a number of technical assumptions and elaborations to guarantee convergence and feasibility and would be prohibitively expensive, e.g. as discussed in [23]. Our technical approach is given next.

CC-OPF: From Stochastic Formulation to Conic Programming.

Our methodology applies and develops general ideas of [9] to the power engineering setting of the generation re-dispatch under uncertainty. We show that under the assumptions of the basic power flow linearity, proportionality of the standard generation response to fluctuations, along with Gaussianity and statistical independence of the wind fluctuations at different sites, the CC-OPF [9] is reduced exactly to a deterministic optimization problem, see Eq. [10] of the Methods. Moreover this deterministic optimization over \bar{p} and α is a convex optimization problem, more precisely, a Second-Order Cone Program (SOCP) [27, 28], allowing an efficient computational implementation discussed in detail in the ‘‘Cutting Plane’’ Subsection of the Appendix.

Let us emphasize that many of our assumptions leading to the computational efficient formulation are not restrictive and allow natural generalizations. In particular, using techniques from [9], it is possible to relax the phenomenologically reasonable but approximately validated assumption of wind source Gaussianity (validated according to actual measurements of wind, see [21, 22] and references therein). For example, using only the mean and variance of output at each wind farm, one can use Chebyshev’s inequality to obtain a similar though more conservative formulation. And following [9] we can also obtain convex approximations to [6] which are tighter than Chebyshev’s inequality, for a large number of empirical distributions discussed in the literature. In any case, we will perform (below) out-of-sample experiments involving our controls; first to investigate the effect of parameter estimation errors in the Gaussian case, and, second, to gauge the impact of non-Gaussian wind distributions.

Experiments/Results

Here we will describe qualitative aspects of our affine control on small systems; in particular we focus on the contrast between standard OPF and CC-OPF, on problematic features that can arise because of fluctuating wind sources and on out-of-sample testing of the CC-OPF solution, including the analysis of non-Gaussian distributions. Some of our tests involve the BPA grid, which is large; in Supplementary Information we also present a second set of tests that address the scalability of our solution methodology to large cases.

Failure of standard OPF. Above (see eq. [1]) we introduced the so-called standard OPF method for setting traditional generator output levels. When renewables are present, the natural extension of this approach would make use of some fixed estimate of output (e.g., mean output) and to handle fluctuations in renewable output through the same method used to handle changes in load: ramping output of traditional generators up or down in proportion to the net increase or decrease in renewable output. This feature could seamlessly be handled using today’s control structure, with each generator’s output adjusted at a fixed (preset) rate. For the sake of simplicity, in the experiments below we assume that all ramping rates are equal.

Different assumptions on these fixed rates will likely produce different numerical results; however, this general approach entails an inherent weakness. The key point here is that mean generator output levels *as well as* in particular the ramping rates would be chosen *without* considering the stochastic nature of the renewable output levels. Our experiments are designed to highlight the limitations of this ‘‘risk-unaware’’ approach. In contrast, our CC-OPF produces control parameters (the \bar{p} and the α) that are risk-aware and, implicitly, also topology-aware – in the sense of network proximity to wind farms.

We first consider the IEEE 118-bus model with a quadratic cost function, and four sources of wind power added at arbitrary buses




Fig. 2. 118-bus case with four wind farms (green dots; brown are generators, black are loads). Shown is the standard OPF solution against the average wind case with penetration of 5%. Standard deviations of the wind are set to 30% of the respective average cases. Lines in red exceed their limit 8% or more of the time.



Fig. 3. Failure of the standard OPF shown for a snapshot of the standard OPF solution on Polish grid from MATPOWER [29]; full snapshot is shown in Section III of SI. Color coding and conditions of the experiment are equivalent to these of Fig. 2.

to meet 5% of demand in the case of average wind. The standard OPF solution is safely within the thermal capacity limits for all lines in the system. Then we account for fluctuations in wind assuming Gaussian and site-independent fluctuations with standard deviations set to 30% of the respective means. The results, which are shown in Fig. 2, illustrate that under standard OPF five lines (marked in red) frequently become overloaded, exceeding their limits 8% or more of the time. This situation translates into an unacceptably high risk of failure for any of the five red lines. This problem occurs for grids of all sizes; in Fig. 3 we show similar results on a 2746-bus Polish grid. In this case, after scaling up all loads by 10% to simulate a more highly stressed system, we added wind power to ten buses for a total of 2% penetration. The standard solution results in six lines exceed-

ing their limits over 45% of the time, and in one line over 10% of the time. For an additional and similar experiment using the Polish grid see the Supplementary Information.

Cost of reliability under high wind penetration. The New York Times article “Wind Energy Bumps Into Power Grid’s Limits” [30] discusses how transmission line congestion has forced temporary shutdown of wind farms even during times of high wind. Our methodology suggests, as an alternative solution to curtailment of wind power, an appropriate reconfiguration of standard generators. If successful, this solution can use the available wind power without curtailment, and thus result in cheaper operating costs.

As a (crude) proxy for curtailment, we perform the following experiment, which considers different levels of renewable penetration. Here, the mean power outputs of the wind sources are kept in a fixed proportion to one another and proportionally scaled so as to vary total amount of penetration, and likewise with the standard deviations.




Fig. 4. 39-bus case under standard solution. Even with a 10% buffer on the line flow limits, five lines exceed their limit over 5% of the time with 30% penetration (left). The penetration must be decreased to 5% before the lines are relieved, but at great cost (right). The CC-OPF model is feasible for 30% penetration at a cost of 264,000. The standard solution at 5% penetration costs 1,275,020 – almost 5 times as much.




Fig. 5. 39-bus case. Red lines indicate high probability of flow exceeding the limit under the standard OPF solution. Generators are shades of blue, with darker shades indicating greater absolute difference between the chance-constrained solution and the standard solution.

First, we run our CC-OPF under a high penetration level. Second, we add a 10% buffer to the line limits and reduce wind penetration (i.e., curtail) so that under the *standard* OPF solution line overloads are reduced to an acceptable level. Assuming zero cost for wind power, the difference in cost for the high-penetration CC-OPF solution and the low-penetration standard solution are the savings produced by our model (generously, given the buffers).

For the 39-bus case, our CC-OPF solution is feasible under 30% of wind penetration, but the standard solution has 5 lines with excessive overloads, even when solved with the 10% buffer. Reducing the penetration to 5% relieves the lines, but more than quadruples(!) the cost over the CC-OPF solution. See Figure 4. Note that this approach does not only show the advantage of the CC-OPF over standard OPF but also provides a quantitative measure of the advantage, thus placing a well-defined price tag on reliability.

Non-locality. We have established that under fluctuating power generation, some lines may exceed their flow limits an unacceptable fraction of the time. Is there a simple solution to this problem, for instance, by carefully adjusting (a posteriori of the standard OPF) the outputs of the generators near the violated lines? The answer is no. Power systems exhibit significant non-local behavior. Consider




Fig. 6. 39-bus case with four wind farms (green dots; brown are generators, black are loads). Lines in red are at the maximum of $\epsilon = .02$ chance of exceeding their limit. The three cases shown are left to right .1%, 8%, and 30% average wind penetration. With penetration beyond 30% the problem becomes infeasible.




Fig. 7. 30-bus case with three wind farms. The case on the left supports only up to 10% penetration before becoming infeasible, while the one on the right is feasible for up to 55% penetration.




Fig. 8. 9-bus case, 25% average penetration from two wind sources. With shifting winds, the flow on the orange line changes direction with a large absolute difference.

Fig. 5. In this example, the major differences in generator outputs between the standard OPF solution and our CC-OPF model's solution are not obviously associated with the different line violations. In general, it seems that it would be difficult to by-pass CC-OPF and make small local adjustments to relieve the stressed lines. On the positive side, even though CC-OPF is not local and requires a centralized computation, it is also only slightly more difficult than the standard OPF in terms of implementation.

Increasing penetration. Current planning for the power system in the United States calls for 30% of wind energy penetration by 2030 [2]. Investments necessary to achieve this ambitious target may target both software (improving operations) and hardware (building new lines, sub-stations, etc), with the former obviously representing a much less expensive and thus economically attractive option. Our CC-OPF solution contributes toward this option. A natural question that arises concerns the maximum level of penetration one can safely achieve by upgrading from the standard OPF to our CC-OPF.

To answer the question we consider the 39-bus New England system (from [31]) case with four wind generators added, and line flow limits scaled by .7 to simulate a heavily loaded system. The quadratic cost terms are set to $\text{rand}(0,1) + .5$. We fix the four wind generator average outputs in a ratio of 5/6/7/8 and standard deviations at 30% of the mean. We first solve our model using $\epsilon = .02$ for each line and assuming zero wind power, and then increase total wind output until the optimization problem becomes infeasible. See Figure 6. This experiment illustrates that at least for the model considered, the 30% of wind penetration with rather strict probabilistic guarantees enforced by our CC-OPF may be feasible, but in fact lies rather close to the dangerous threshold. To push penetration beyond the threshold is impossible without upgrading lines and investing in other (not related to wind farms themselves) hardware.

Changing locations for wind farms. In this example we consider the effect of changing locations of wind farms. We take the MATPOWER 30-bus case with line capacities scaled by .75 and add three

Distribution	Max. prob. violation
Normal	0.0227
Laplace	0.0297
logistic	0.0132
Weibull, $k = 1.2$	0.0457
Weibull, $k = 2$	0.0355
Weibull, $k = 4$	0.0216
t location-scale, $\nu = 2.5$	0.0165
Cauchy	0.0276

Fig. 9. Maximum probability of overload for out-of-sample tests. These are a result of Monte Carlo testing with 10,000 samples on the BPA case, solved under the Gaussian assumption and desired maximum chance of overload at 2.27%.




Fig. 10. BPA case solved with average penetration at 8% and standard deviations set to 30% of mean. The maximum probability of line overload desired is 2.27%, which is achieved with 0 forecast error on the graph. Actual wind power means are then scaled according to the x-axis and maximum probability of line overload is recalculated (blue). The same is then done for standard deviations (green).

wind farms with average power output in a ratio of 2/3/4 and standard deviations at 30% of the average. Two choices of locations are shown in figure 7. The first remains feasible for penetration up to 10% while the second can withstand up to 55% penetration. This experiment shows that choosing location of the wind farms is critical for achieving the ambitious goal of high renewable penetration.

Reversal of line flows. Here we consider the 9-bus case with two wind sources and 25% average penetration and standard deviations set to 30% of the average case and analyze the following two somewhat rare but still admissible wind configurations: (1) wind source (a) produces its average amount of power and source (b) three standard deviations *below* average; (2) the reverse of the case (1). This results in a substantial reversal of flow on a particular line shown in Figure 8. This example suggests that when evaluating and planning for grids with high-penetration of renewables one needs to be aware of potentially fast and significant structural rearrangements of power flows. Flow reversals and other qualitative changes of power flows, which are extremely rare in the grid of today, will become significantly much more frequent (typical) in the grid of tomorrow.

Out-of-sample tests. We now study the performance of our method when there are errors in the underlying distribution of wind power. We consider two types of errors: (1) the true distribution is non-Gaussian but our Gaussian fit is “close” in an appropriate sense, and (2) the true distribution is Gaussian but with different mean or standard deviation. The experiments in this section use as data set the BPA grid, which as noted before has 2209 buses and 2866 lines, and collected wind data; altogether constituting a realistic test-case.

We first consider the non-Gaussian case, using the following probability distributions, all with fatter tails than Gaussian: (1) Laplace, (2) logistic, (3) Weibull (three different shapes), (4) t location-scale with 2.5 degrees of freedom, (5) Cauchy. For the Laplace and logistic distributions, we simply match the mean and standard deviation. For the Weibull distribution, we consider shape parameters $k = 1.2, 2, 4$ and choose the scale parameter to match the standard deviations. We then translate to match means. For the t distribution, we fix 2.5 degrees of freedom and then choose the location and scale to match mean and standard deviation. For the Cauchy distribution, we set the location parameter to the mean and choose the scale parameter so as to match the 95th percentiles.

We use our model and solve under the Gaussian assumption, seeking a solution which results in no line violations for cases within two standard deviations of the mean, i.e. a maximum of about 2.27% chance of exceeding the limit. We then perform Monte Carlo tests drawing from the above distributions to determine the actual rates of violation. See Figure 9. The worst-performer is the highly-asymmetric (and perhaps unreasonable) Weibull with shape parameter 1.2, which approximately doubles the desired maximum chance of overload. Somewhat surprisingly, the fat-tailed logistic and Student’s t distributions result in a maximum chance of overload significantly less than desired, suggesting that our model is too conservative in these cases.

Next we consider the Gaussian case with errors. We solve with nominal values for the mean and standard deviation of wind power. We then consider the rate of violation after scaling all means and standard deviations (separately). While the solution is sensitive to errors in the mean forecast, the sensitivity is well-behaved. With a desired safety level of $\epsilon = 2.27\%$ for each line, an error in the mean of 25% results in a maximum 15% chance of exceeding the limit. The solution is quite robust to errors in the standard deviation forecast, with a 25% error resulting in less than 6% chance of overload. See Figure 10.

Discussions

This manuscript suggests a new approach to incorporating uncertainty in the standard OPF setting routinely used in the power industry to set generation during a time window, or period (typically 15 min to one hour duration). When uncertainty associated with renewable generation is quantified in terms of the probability distribution of output during the next period, we incorporate it through chance constraints - probabilistic conditions which require that any line of the system will not be overloaded for all but a small fraction of time (at most one minute per hour, for example). Additionally, the modeling accounts for local frequency response of controllable generators to renewable changes. The key technical result of this manuscript is that the resulting optimization problem, CC-OPF, can be stated as a convex, deterministic optimization problem. This result also relies on plausible assumptions regarding the exogenous uncertainty and linearity of the underlying power flow approximations/equations. In fact, our CC-OPF is a convex (conic) optimization problem, which we solve very efficiently, even on realistic large-scale instances, using a sequential linear cutting plane algorithm.

This efficient CC-OPF algorithm becomes an instrument of our (numerical) experiments which were performed on a number of standard (and nonstandard) power grid data sets. Our experimental results are summarized as follows:

- We observe that CC-OPF delivers feasible results where standard OPF, run for the average forecast, would fail in the sense that many lines would be overloaded an unacceptably large portion of time.
- Not only is CC-OPF safer than standard OPF, but it also results in *cheaper* operation. This is demonstrated by considering the optimal cost of CC-OPF under sufficiently high wind penetration solution (where standard OPF would fail) and the low penetration solution (corresponding to the highest possible penetration where standard OPF would not fail).
- We discover that solutions produced by CC-OPF deviate significantly from what amounts to a the naive adjustment of the standard OPF obtained by correcting dispatch just at those generators which are close to overloaded lines.
- We test the level of wind penetration which can be tolerated without upgrading lines. This experiment illustrates that, at least for the model tested, the 30% of wind penetration with rather strict probabilistic guarantees enforced by our CC-OPF may be feasible; but much lower wind penetration remains feasible under the standard approach.
- We experiment with location of wind farms and discover strong sensitivity of the maximum level of penetration on the choice of location - optimal choice of wind farm location is critical for achieving the ambitious goal of high renewable penetration.
- Analyzing fluctuations of line flows within CC-OPF solution admissible under high wind penetration, we discover that these fluctuations may be significant, in particular resulting in reversal of power flows over some of the lines. This observation suggests that flow reversals and other qualitative changes of power flows, which are extremely rare in the grid of today, will become significantly much more frequent (typical) in the grid of tomorrow.
- We also studied an out-of-sample test consisted in applying CC-OPF (modeling exogenous fluctuations as Gaussian) to other distributions. Overall these tests suggest that with a proper calibration of the effective Gaussian distribution our CC-OPF delivers a rather good performance. One finds that the worst CC-OPF performance is observed for the most asymmetric distributions.

The nature of the problem discussed in the manuscript – principal design of a new paradigm for computationally efficient generation re-dispatch that accounts for wind fluctuations – inevitably required incorporation of a number of assumptions and approximations. In particular, we made simplifying assumptions about static forecasts and general validity of power flow linearization. We have

also focused solely on failures associated with line congestion ignoring other possible difficulties, for example those associated with loss of synchronicity and voltage variations. However, we would like to emphasize that all of these assumptions made (admittedly natural for a first attack on the problem) also allow generalization within the approach just sketched:

- Accounting for time evolving forecast/loads/etc. Wind forecast, expressed in terms of the mean and standard deviation at the wind farm sites, changes on the scales comparable to duration of the generation re-dispatch interval. Loads may also change at these time scales. When the slowly evolving, but still not constant, wind and load forecasts are available we may keep the quasi-static power flow description and incorporate this slow evolution in time into the chance constrained scheme. These changes will simply result in generalizing the conic formulation of Eq. [10] by splitting what used to be a single time interval into sub-intervals and allowing the regular generation to be re-dispatched and parallel coefficients to be adjusted more often. Ramping rate constraints on the controllable generation may naturally be accounted in the temporal optimization scheme as well.
- Accounting for nonlinearity in power flows. Evolution of the base case invalidates the linearization (DC-style) hypothesis. However, if variations around one base case becomes significant one may simply adjust the linearization procedure doing it not once (as in the case considered in the manuscript) but as often as needed. Slow adjustment of the base case may also be included into the dynamical procedure mentioned one item above. Additionally, some interesting new methodologies for handling nonlinearities have recently emerged, see [32].
- Accounting for synchronization bounds. Loss of synchronicity and resulting disintegration of the grid is probably the most acute contingency which can possibly take place in a power system. The prediction of those conditions under which the power grid will lose synchronicity is a difficult nonlinear and dynamic problem. However, as shown recently in [33], one can formulate an accurate linear and static necessary condition for the loss of synchronicity. A chance-constrained version of the linear synchronization conditions can be incorporated seamlessly in our CC-OPF framework.

Finally, we see many opportunities in utilizing the CC-OPF (possibly modified) as an elementary unit or an integral part of even more complex problems, such as combined unit commitment (scheduling large power plants normally days, weeks or even months ahead) [25] with CC-OPF, planning grid expansion [34] while accounting for cost operation under uncertainty, or incorporating CC-OPF in mitigating emergency of possible cascades of outages [35, 36, 37, 38, 39, 40, 41, 42]. In this context, it would be advantageous to speed up our already very efficient CC-OPF even further. See, for example, [43], [44]. A different methodology, relying on distributed algorithms, can be found in [45].

Methods

CC-OPF as Deterministic (Conic) Programming. Eq. (9) states the problem of generation dispatch under uncertainty due to wind as a stochastic optimization problem with chance (probabilistic) constraints. The critical part of our approach consists in providing a convex expressions for the expectation of the objective and for the probabilities in Eq. (9), as a function of generation dispatch optimization variables. This gives rise to the following deterministic optimization

problem (see Section I of SI for derivations and proofs)

$$\min_{\bar{p}, \alpha} \sum_{i \in \mathcal{G}} \left(c_{i1} \left(\bar{p}_i^2 + \alpha_i^2 \sum_{j \in \mathcal{W}} \sigma_j^2 \right) + c_{i2} \bar{p}_i + c_{i3} \right) \quad [10]$$

$$\text{s.t.} \quad \sum_{i \in \mathcal{G}} \alpha_i = 1, \quad \alpha \geq 0, \quad \bar{p} \geq 0 \quad [11]$$

$$B\bar{\theta} = \bar{p} + \mu - d, \quad \sum_{i \in \mathcal{V}} (\bar{p}_i + \mu_i - d_i) = 0 \quad [12]$$

$$\text{for } 1 \leq i \leq n-1: \quad \sum_{j=1}^{n-1} B_{ij} \delta_j = \alpha_i, \quad \delta_n = 0, \quad [13]$$

$$\forall \{i, j\} \in \mathcal{E}:$$

$$s_{ij}^2 \geq \beta_{ij}^2 \sum_{k \in \mathcal{W}} \sigma_k^2 (B_{ik}^+ - B_{jk}^+ - \delta_i + \delta_j)^2, \quad [14]$$

$$|\beta_{ij}(\bar{\theta}_i - \bar{\theta}_j)| \leq f_{ij}^{max} - \phi^{-1}(1 - \epsilon_{ij}/2) s_{ij} \quad [15]$$

The objective here is simply $\mathbb{E}_w [c(\bar{p}, \alpha)]$ written explicitly; the vectors \bar{p} and α model our control methodology as described above; $\bar{\theta}$ is the average phase angle vector; the first equation in [12] amounts to the standard DC model equation relating phase angle to power injections. The second constraint in [12] is a flow balance statement: the sum of all power injections is zero. The nature of our control will guarantee that power is balanced for any configuration of wind power.

Constraints [14] and [15] constitute a (deterministic) representation of the chance constraint [6]. Here, β_{ij} is the susceptance on line (i, j) ; σ_k is the standard deviation of wind source k ; ϕ is the standard normal cumulative distribution function; B^+ is an appropriate, sparse generalized inverse of the bus susceptance matrix; the δ are auxiliary variables. It can be shown (see SI) that, under the independence assumption for the wind fluctuations w_i , but without requiring Gaussianity, the right-hand side of constraint [14] is the variance of the flow on line (i, j) . Thus, at optimality [14] and [15] will amount to

$$|\beta_{ij}(\bar{\theta}_i - \bar{\theta}_j)| \leq f_{ij}^{max} - \phi^{-1}(1 - \epsilon_{ij}/2) \sigma_{ij} \quad [16]$$

where σ_{ij} is the standard deviation of the flow f_{ij} on line (i, j) .

Further, it can be shown that the bus angles θ_i and line flows f_{ij} both are affine combinations of the w_i . Thus, assuming that the w_i are Gaussian, then so will be the θ_i and f_{ij} . Since the left-hand side of equation [16] is the absolute value of the expected flow on line (i, j) , it follows under the Gaussianity assumption that [16] is as claimed a valid representation of the chance constraint [6]: it states that the expectation of flow on line (i, j) is the right multiple of a standard deviation away from the maximum f_{ij}^{max} , as per the risk tolerance ϵ_{ij} .

In deriving [15] we are thus explicitly making use of the Gaussianity assumption. However, using the results in [9] one can show that in the case of arbitrary distributions of the w_i (but assuming finite variances) one can obtain a convex (conic) conservative approximation to the chance constraint by simply replacing the quantity $\phi^{-1}(1 - \epsilon_{ij}/2)$ in [15] with an appropriate coefficient Ω (effectively, this approach relies on estimates from the theory of large deviations). Thus, even in the non-Gaussian case, our general approach remains essentially the same.

Cutting-Plane Algorithm. The number of conic constraints [14] is equal to the number of lines, and in the case of a large grid this can prove challenging. For example, in the Polish 2003-2004 winter peak case, Cplex [46] reports over 6000 conic constraints after pre-solving. All of the commercial solvers [46, 47, 48] we experimented with reported numerical difficulties and were unable to solve a Second Order Conic Programming (SOCP) of this size.

In order to solve these large cases, we employed a *cutting-plane algorithm*. Refer to the Section 2 of SI for a detailed description. The algorithm repeatedly solves linear programs that include all the linear constraints in our formulation together with a finite set of first-order approximations to the conic constraints. Having solved one such problem the algorithm checks for a violated conic constraint and if found, it approximates that conic constraint with its first-order approximation at the point of violation. We repeat this process until the largest violation of a conic constraint is sufficiently small, in our case less than 10^{-6} . This method is able to solve cases with thousands of buses usually in just CPU seconds on a laptop computer. The following table shows typical behavior:

Iteration	Max rel. error	Objective
1	1.2e-1	7.0933e6
4	1.3e-3	7.0934e6
7	1.9e-3	7.0934e6
10	1.0e-4	7.0964e6
12	8.9e-7	7.0965e6

Each row of this table shows that maximum relative error and objective value at the end of several iterations. The total run-time was 25 seconds. Note the “flatness” of the objective. This makes the problem nontrivial – the challenge is to find a *feasible* solution (with respect to the chance constraints); at the onset of the algorithm the computed solution is quite infeasible and it is this attribute that is quickly improved by the cutting-plane algorithm.

Power Grid Models. The BPA grid data was extracted from [49] (available to the authors). Corresponding wind data was extracted from [1]. All other power grid cases, in particular Polish Grid, 118 bus, 39 bus and 30 bus systems are publicly available from [31]. The wind-related modifications for each example are explained in the text above.

CC-OPF as Conic Programming

Notation. The following notation and conventions will be used throughout. (i) n = number of buses, m = number of lines. (ii) Let \hat{B} be the $(n-1) \times (n-1)$ -matrix obtained by removing from B row and column n . (Assuming the grid is connected, \hat{B} is invertible.) (iii) In what follows we use bold font to indicate uncertain quantities. (iv) We will write $\mathbf{S} \doteq \sum_{i \in \mathcal{W}} \mathbf{w}_i$, $\sigma^2 \doteq \sum_{i \in \mathcal{W}} \sigma_i^2$. (v) For convenience of notation, here we assume that bus n does not hold a generator and does not hold a wind farm, in other words, $n \notin \mathcal{G} \cup \mathcal{W}$. This assumption is easily attained by adding a “dummy” bus (with zero generation, demand and wind output) and attaching it to the grid with a dummy line. (vi) For a bus i , let d_i be its *demand*. When $i \notin \mathcal{D}$ we write $d_i = 0$. Write $\bar{p}_i = \alpha_i = 0$ for each bus i that is not under the affine control discipline. For each i with $i \notin \mathcal{W}$ let $\mu_i = 0$ and let \mathbf{w}_i be the random variable with mean and variance equal to 0. Then, $\mathbf{b}_i \doteq \bar{b}_i + \bar{p}_i - \alpha_i \mathbf{S} + \mathbf{w}_i$, where $\bar{b}_i \doteq \mu_i - d_i$, is the net power injected into the network at bus i . Note also that we must always have

$$\begin{aligned} 0 &= \sum_i \mathbf{b}_i = \sum_{i \in \mathcal{G}} (\bar{p}_i - \alpha_i \mathbf{S}) + \sum_{i \in \mathcal{W}} (\mu_i + \mathbf{w}_i) - \sum_i d_i \\ &= \sum_{i \in \mathcal{G}} \bar{p}_i + \sum_{i \in \mathcal{W}} \mu_i - \sum_i d_i. \end{aligned} \quad [17]$$

Analysis. Let $\bar{\mathbf{b}} = (\bar{b}_1, \dots, \bar{b}_{n-1})^T$, $\alpha = (\alpha_1, \dots, \alpha_{n-1})^T$, $\bar{\mathbf{p}} = (\bar{p}_1, \dots, \bar{p}_{n-1})^T$, and $\mathbf{w} = (\mu_1 + \mathbf{w}_1, \dots, \mu_{n-1} + \mathbf{w}_{n-1})^T$. Thus, given a choice of the control variables \bar{p}_i and α_i , we have that the (random) phase angle vector $\boldsymbol{\theta} = (\theta_1, \dots, \theta_{n-1})^T$ satisfies:

$$\hat{B} \boldsymbol{\theta} = \bar{\mathbf{b}} + \bar{\mathbf{p}} - \mathbf{S} \alpha + \mathbf{w}. \quad [18]$$

The vector $\boldsymbol{\theta}$ has one degree of freedom, and so we set $\theta_n = 0$ and thus $E(\boldsymbol{\theta}_n) = 0$. So for all $1 \leq i \leq n$,

$$\boldsymbol{\theta}_i = \bar{\theta}_i - \delta_i \mathbf{S} + \sum_{j=1}^{n-1} \pi_{ij} \mathbf{w}_j, \quad \bar{\theta}_i \doteq \begin{cases} [\hat{B}^{-1}(\bar{\mathbf{b}} + \bar{\mathbf{p}})]_i, & i < n, \\ 0, & \text{otherwise.} \end{cases} \quad [19]$$

$$\pi_{ij} \doteq \begin{cases} [\hat{B}^{-1}]_{ij}, & i < n, \\ 0, & \text{otherwise.} \end{cases}, \quad \delta_i \doteq \begin{cases} [\hat{B}^{-1} \alpha]_i, & i < n, \\ 0, & \text{otherwise.} \end{cases} \quad [20]$$

Consider a line (i, j) . Since the flow \mathbf{f}_{ij} on line (i, j) equals $\beta_{ij}(\boldsymbol{\theta}_i - \boldsymbol{\theta}_j)$ we have that

$$\mathbf{f}_{ij} = \bar{\theta}_i - \bar{\theta}_j - (\delta_i - \delta_j) \mathbf{S} + \sum_{k=1}^{n-1} (\pi_{ik} - \pi_{jk}) \mathbf{w}_k, \quad [21]$$

and therefore, $E(\mathbf{f}_{ij}) = \beta_{ij}(\bar{\theta}_i - \bar{\theta}_j)$ and, since the \mathbf{w}_i are pairwise independent,

$$\text{var}(\mathbf{f}_{ij}) = \beta_{ij}^2 \sum_{k \in \mathcal{W}} \sigma_k^2 (\pi_{ik} - \pi_{jk} - \delta_i + \delta_j)^2. \quad [22]$$

Likewise, denoting by \mathbf{P}_g the power produced by generator g , we have $E(\mathbf{P}_g) = \bar{p}_g$, $\text{var}(\mathbf{P}_g) = \alpha_g^2 \sum_{j \in \mathcal{W}} \sigma_j^2$. As a result of the above we have:

Lemma 0.1. For a given choice of vectors $\bar{\mathbf{p}}$ and α , each quantity $\boldsymbol{\theta}_i$ or \mathbf{f}_{ij} is an affine function of the random variables \mathbf{w}_i .

Proof. Follows from eqs. (19) and (20). ■

Formulation. Let the vectors $\bar{\mathbf{p}}$ and α (both in \mathbb{R}^{n-1}) be fixed, and consider the following system of inequalities, on variables $\delta_i, \bar{\theta}_i$ (all for $1 \leq i \leq n-1$) and $\mathbf{f}_{i,j}, s_{i,j}$ (for each line (i, j)) [we also have the additional quantities δ_n and θ_n as variables fixed at zero, for convenience of notation]:

$$\text{for } 1 \leq i \leq n-1, \quad \sum_{j=1}^{n-1} \hat{B}_{ij} \delta_j = \alpha_i; \quad \delta_n = 0, \quad [23]$$

$$\forall (i, j): \quad \beta_{ij}^2 \sum_{k \in \mathcal{W}} \sigma_k^2 (\pi_{ik} - \pi_{jk} - \delta_i + \delta_j)^2 \leq s_{ij}^2, \quad [24]$$

$$\text{for } 1 \leq i \leq n-1, \quad \sum_{j=1}^{n-1} \hat{B}_{ij} \bar{\theta}_j - \bar{p}_i = \bar{b}_i, \quad \bar{\theta}_n = 0, \quad [25]$$

$$\forall (i, j): \quad \bar{f}_{i,j} - \beta_{ij}(\bar{\theta}_i - \bar{\theta}_j) = 0. \quad [26]$$

Theorem 0.2. Consider the affine control given by a pair of vectors $\bar{\mathbf{p}}$ and α satisfying (17). Then $(\delta, \bar{\theta}, \bar{\mathbf{f}}, s)$ is feasible for (23)-(26) if and only if: (a) for each bus i , $\theta_i = E(\boldsymbol{\theta}_i)$; (b) for each line (i, j) , $\bar{f}_{i,j} = E(\mathbf{f}_{i,j})$; (c) for each line (i, j) , $s_{i,j}^2 \geq \text{var}(\mathbf{f}_{i,j})$; (d) for each generator g , $E(\mathbf{P}_g) = \bar{p}_g$ and $\text{var}(\mathbf{P}_g) = \alpha_g^2 \sum_{j \in \mathcal{W}} \sigma_j^2$.

Proof. Suppose $(\delta, \bar{\theta}, \bar{\mathbf{f}}, s)$ is feasible for Eqs. (23)-(26). By Eq. (25) we have that at $1 \leq i \leq n-1$, $\bar{\theta}_i = [\hat{B}^{-1}(\bar{\mathbf{b}} + \bar{\mathbf{p}})]_i = E(\boldsymbol{\theta}_i)$, and by Eqs. (26), (25), $\bar{f}_{i,j} = E(\mathbf{f}_{i,j})$, for each line (i, j) . Similarly, Eq. (20),(24),(22) imply $s_{i,j}^2 \geq \text{var}(\mathbf{f}_{i,j})$ as desired. (d) Holds by construction. The converse is similar. ■

Note: it is easily seen that under the conventions for bus n , constraints (11-14) of the main text are equivalent to Eqs. (23)-(26),(17).

We let $0 < \epsilon_{ij} < 1$ (and, $0 < \epsilon_g < 1$) denote the probabilistic tolerance for line (i, j) (resp., for generator g), and denote by ϵ the $(m + |\mathcal{G}|)$ -vector of tolerances.

Definition 0.3. Given a vector ϵ of tolerances, a control $(\bar{\mathbf{p}}, \alpha)$ is $(1 - \epsilon)$ -strong, if $\text{Prob}(|\mathbf{f}_{i,j}| > f_{ij}^{max}) < \epsilon_{ij}$, for each line (i, j) . and $\text{Prob}(p_g^{\min} \leq \mathbf{P}_g \leq p_g^{max}) \geq 1 - \epsilon_g$, for each generator g .

For the next result, suppose we have a fixed control $(\bar{\mathbf{p}}, \alpha)$, and consider Eq. (19). By expanding \mathbf{S} as $\sum_k \mathbf{w}_k$, we see that for any line

(i, j) , $f_{i,j}$ equals a constant plus a linear combination of the w_k . It follows that if the w_k are normally distributed, then so is each flow value $f_{i,j}$.

Additional notation: For real $0 < r < 1$ we write $\eta(r) = \phi^{-1}(1 - r)$.

Lemma 0.4. Let (\bar{p}, α) satisfy Eq. (17). Suppose (\bar{p}, α) is $(1 - \epsilon)$ -strong. Then there exists a vector $(\delta, \theta, \bar{f}, s)$ feasible for Eqs. (23)-(26) such that, in addition, for all lines (i, j) ,

$$|\bar{f}_{i,j}| \leq f_{ij}^{max} - \eta(\epsilon_{ij}) s_{ij}, \quad [27]$$

and for all generators g ,

$$p_g^{min} + \eta(\epsilon_g) \alpha_g \sqrt{\sum_{j \in \mathcal{W}} \sigma_j^2} \leq \bar{p}_g \leq p_g^{max} - \eta(\epsilon_g) \alpha_g \sqrt{\sum_{j \in \mathcal{W}} \sigma_j^2}. \quad [28]$$

Conversely, if there exists a vector $(\delta, \theta, \bar{f}, s)$ satisfying Eqs. (23)-(26) and Eqs. (27)-(28) then (\bar{p}, α) is $(1 - 2\epsilon)$ -strong.

Proof. If $Prob(|f_{i,j}| > f_{ij}^{max}) < \epsilon$ then using Theorem 0.2 we have that Eqs. (27) hold. For the converse, we have that Eq. (27) implies $Prob(f_{i,j} > f_{ij}^{max}) < \epsilon$ and $Prob(f_{i,j} < -f_{ij}^{max}) < \epsilon$. Likewise with Eqs. (28). ■

Comment: Eqs. (27), together with Eq. (26), are equivalent to the constraint Eq. (15) of the main text.

Cutting-Plane Algorithm

In the case of a grid with thousands of buses and lines, the optimization problem given by Eqs. (17),(23)-(26),(27)-(28) (or, equivalently, Eqs. (10-15) of the main text), amounts to a large-scale convex conic programming problem. Experience with realistic examples with thousands of lines shows that commercial optimization packages are unable to solve the resulting problems “out of the box.” Here we outline a simple algorithm that proves effective and fast, a so-called “cutting-plane” algorithm [27], [50]. From a theoretical standpoint the algorithm is motivated and justified by the “efficient separation is equivalent to efficient optimization” paradigm that underlies the ellipsoid method [50].

Without constraints (24),(28), the optimization problem is a linearly constrained convex quadratic programming problem, still somewhat large (in the case of large grids) but within the reach of commercial solvers. Our algorithm iteratively replaces these constraints with a number of linear approximations which are algorithmically discovered. For a given line (i, j) write $C_{ij}(\delta) \doteq \beta_{ij} \sqrt{\sum_{k \in \mathcal{W}} \sigma_k^2 (\pi_{ik} - \pi_{jk} - \delta_i + \delta_j)^2}$; then constraint (24) can be written as $C_{ij}(\delta) \leq s_{ij}$. Given a vector δ^* , this constraint can be relaxed by the (outer) approximation $C_{ij}(\delta^*) + \frac{\partial C_{ij}(\delta^*)}{\partial \delta_i} (\delta_i - \delta_i^*) + \frac{\partial C_{ij}(\delta^*)}{\partial \delta_j} (\delta_j - \delta_j^*) \leq s_{ij}$, which is valid for all δ . A similar linearization applies to Eqs. (28). This technique gives rise to the following iterative algorithm. The algorithm maintains a *linear* system of inequalities $A(\bar{p}, \alpha, \delta, \theta, \bar{f}, s)^T \geq b$, henceforth referred to as the *master system*, which is initialized to the set of all our constraints but with the conics removed; i.e. Eqs. (17),(23),(25)-(26),(27). Denoting by $F(\bar{p}, \alpha)$ the objective function in Eq. (10) of the main, the algorithm iterates through the following steps:

1. Solve $\min\{F(\bar{p}, \alpha) : A(\bar{p}, \alpha, \delta, \theta, \bar{f}, s)^T \geq b\}$. Let $(\bar{p}^*, \alpha^*, \delta^*, \theta^*, \bar{f}^*, s^*)$ be an optimal solution.
2. If all conic constraints are satisfied up to numerical tolerance by $(\bar{p}^*, \alpha^*, \delta^*, \theta^*, \bar{f}^*, s^*)$, **Stop**.
3. If all chance constraints are satisfied up to numerical tolerance by (\bar{p}^*, α^*) , **Stop**.
4. Otherwise, add the outer inequality arising from that constraint (24) which is most violated to the master system.

As the algorithm iterates the master system represents a valid relaxation of the constraints Eqs. (10-15) of the main text; thus the objec-

tive value of the solution computed in Step 1 is a valid lower bound on the value of problem. Each problem solved in Step 1 is a linearly constrained, convex quadratic program. Computational experiments involving large-scale realistic cases show that the algorithm is robust and rapidly converges to an optimum.

Note that Step 3 is not redundant. The stopping condition in Step 2 may fail because the variance estimates are incorrect (too small), nevertheless the pair (\bar{p}^*, α^*) may already satisfy the chance constraints. Checking that it does, for a given line (i, j) , is straightforward since the flow $f_{i,j}$ is normally distributed (Lemma 0.1) and its mean and variance can be directly computed from (\bar{p}^*, α^*) .

Table 1 displays typical performance of the cutting-plane algorithm on (comparatively more difficult) large problem instances. In the Table, ‘Polish2’ is the case described in the main text (loads increased 10% as in the text). Polish1 and Polish3 are two others included in MATPOWER [29]. All Polish cases have uniform random costs on [0.5, 2.0] for each generator and ten arbitrarily chosen wind sources. The average wind power penetration for the four cases is 8.8%, 3.0%, 1.9%, and 1.5%. ‘Iterations’ is the number of linearly-constrained subproblems solved before the algorithm converges. ‘Barrier iterations’ is the total number of iterations of the barrier algorithm in CPLEX over all subproblems, and ‘Time’ is the total (wallclock) time required by the algorithm. For each case, line tolerances are set to two standard deviations and generator tolerances three standard deviations. These instances all prove unsolvable if directly tackled by CPLEX or Gurobi.

Case	Buses	Generators	Lines	Time (s)	Iterations	Barrier iterations
BPA	2209	176	2866	5.51	2	256
Polish1	2383	327	2896	13.64	13	535
Polish2	2746	388	3514	30.16	25	1431
Polish3	3120	349	3693	25.45	23	508

Example

As an additional experiment we studied the Polish national grid (obtained from MATPOWER as explained above) under simulated 20 % renewable penetration spread over 18 wind farms, co-located with the 18 largest generators. This co-location should lessen the risk associated with renewable fluctuation (which should be partially “absorbed” by the co-located generators). Figure 11 studies the resulting risk exposure under standard OPF. The chart shows the number of lines that attain several levels of overload probability. The situation in the chart is unacceptable: it would lead to frequent tripping of at least four lines.

In contrast, Figure 12 shows the performance attained by the chance-constrained OPF in the same setting as that of Figure 11. Notice the drastic reduction in overload probabilities – the system is stable. Moreover, this is attained with a minor increase in cost (less than one percent) while the computational time is on the order of 10 seconds.

Figures




Fig. 11. Figure shows number of lines that are overloaded with given probability values in simulation of 2746 bus Polish power grid using standard OPF with 20% wind penetration distributed over 18 wind farms. In particular, two lines are overloaded half of the time, and one line is overloaded one-third of the time, constituting a situation with unacceptable systemic risk.




Fig. 12. Same as Figure 11, but under chance-constrained OPF. Notice that the largest overload probability is 200 times smaller than in the case of standard OPF. Moreover, the cost increase is by less than one percent

ACKNOWLEDGMENTS. We are thankful to S. Backhaus and R. Bent for their comments, and to K. Dvijotham for help with BPA data. The work at LANL was carried out under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy at Los Alamos National Laboratory under Contract No. DE-AC52-06NA25396. MC and SH also acknowledge partial support of the DTRA Basic Research grant BRCALL08-Per3-D-2-0022. DB was partially supported by DOE award DE-SC0002676.




Fig. 13. Standard OPF solution on Polish grid; full rendering

1. BPA-WIND. (2012) Wind generation & total load in the bpa balancing authority (<http://transmission.bpa.gov/business/operations/wind/>).
2. EERE. (2008) 20% of wind energy by 2030: Increasing wind energy's contribution to us electricity supply. *Department of Energy (DOE) DOE/GO-102008-2567*.
3. CIGRE. (2009) Technical brochure on grid integration of wind generation. *International Conference on Large High Voltage Electric Systems*.
4. DENA. (2005) Energy management planning for the integration of wind energy into the grid in germany, onshore and offshore by 2020.
5. Gonzalez, G & et al. (2006) Experience integrating and operating wind power in the peninsular spanish power system. point of view of the transmission system operator and a wind power producer. *CIGRE*.
6. Chertkov, M, Pan, F, & Stepanov, M. (2010) Predicting failures in power grids: The case of static overloads. *IEEE Transactions on Smart Grids* 2, 150.
7. Chertkov, M, Stepanov, M, Pan, F, & Baldick, R. (2011) Exact and efficient algorithm to discover extreme stochastic events in wind generation over transmission power grids. *CDC-ECC* pp. 2174–2180.
8. Bagsorkhi, S & Hiskens, I. (2012) Analysis tools for assessing the impact of wind power on weak grids. *SysCon* pp. 1–8.
9. Nemirovski, A & Shapiro, A. (2006) Convex approximations of chance constrained programs. *SIAM Journal on Optimization* 17, 969–996.
10. Bergen, A & Vittal, V. (2000) *Power System Analysis*. (Prentice Hall. Inc).
11. Huneault, M. (1991) A survey of the optimal power flow literature. *Power Systems, IEEE Transactions*.
12. Kundur, P. (1994) *Power System Stability and Control*. (McGraw-Hill, New York, NY, USA).
13. CAISO. (2007) Integration of renewable resources: Transmission and operating issues and recommendations for integrating renewable resources on the california iso-controlled grid (<http://www.caiso.com/1ca5/1ca5a7a026270.pdf>).
14. Makarov, Y, Loutan, C, Ma, J, & de Mello, P. (2009) Operational impacts of wind generation on california power systems. *Power Systems, IEEE Transactions on* 24, 1039–1050.
15. Wang, G, Negrete-Pincetic, M, Kowli, A, Shafiepoorfar, E, Meyn, S, & Shanbhag, U. (2011) in *Control and Optimization Theory for Electric Smart Grids*, eds. Chakraborty, A & Illic, M. (Springer).
16. Force, U.-C. P. S. O. T. ((2004)) Report on the august 14, 2003 blackout in the united states and canada: Causes and recommendations. <https://reports.energy.gov>.
17. Prékopa, A. (1970) On probabilistic constrained programming. *Proc. Princeton Symposium on Math. Programming* pp. 113–138.
18. Charnes, A, Cooper, W, & Symonds, G. (1958) Cost horizons and certainty equivalents: an approach to stochastic programming of heating oil. *Management Science* 4, 235–263.
19. Miller, L & Wagner, H. (1965) Chance-constrained programming with joint constraints. *Operations Research* 13, 930–945.
20. Bertsimas, D & Sim, M. (2004) Price of robustness. *Operations Research* 52, 35–53.
21. Bludszuweit, H, Dominguez-Navarro, J, & Llombart, A. (2008) Statistical analysis of wind power forecast error. *Power Systems, IEEE Transactions on* 23, 983–991.

22. Hodge, B.-M & Milligan, M. (2011) Wind power forecasting error distributions over multiple timescales. *Detroit Power Engineering Society Meeting* pp. 1–8.
23. Zhang, H. (2011) Chance Constrained Programming for Optimal Power Flow Under Uncertainty. *Power Systems, IEEE Transactions on* 26, 2417–2424.
24. Ozturk, U & Mazumdar, M. (2004) A solution to the stochastic unit commitment problem using chance constrained programming. *Power Systems, IEEE* 19, 1589–1598.
25. Wang, J, Shahidehpour, M, & Li, Z. (2008) Security-constrained unit commitment with volatile wind power generation. *Power Systems, IEEE Transactions on* 23, 1319–1327.
26. Zhang, S, Song, Y, Hu, Z, & Yao, L. (2011) *Robust optimization method based on scenario analysis for unit commitment considering wind uncertainties*. pp. 1–7.
27. Boyd, S & Vandenberghe, L. (2004) *Convex Optimization*. (Cambridge University Press).
28. Goldfarb, D & Alizadeh, F. (2003) Second-order cone programming. *Mathematical Programming* 95, 3–51.
29. Zimmerman, R. D, Murillo-Sánchez, C. E, & Gan, D. (2011) Matpower, a matlab power system simulation package. *IEEE Trans. Power Sys.* 26, 12–19.
30. Wyld, M. (2008/08/27) Wind energy bumps into power grid limits. *New York Times*.
31. Zimmerman, R, Murillo-Sanchez, C, & Thomas, R. (2011) Matpower: Steady-state operations, planning, and analysis tools for power systems research and education. *Power Systems, IEEE Transactions on* 26, 12–19.
32. Lavaei, J & Low, S. (2012) Zero duality gap in optimal power flow problem. *IEEE Transactions on Power Systems* 27, 92–107.
33. Dörfler, F, Chertkov, M, & Bullo, F. (2012) Synchronization in Complex Oscillator Networks and Smart Grids. <http://arxiv.org/abs/1208.0045>.
34. Bent, R, Toole, G, & Berscheid, A. (2012) Transmission network expansion planning with complex power flow models. *Power Systems, IEEE Transactions on* 27, 904–912.
35. Chen, J, Thorp, J, & Dobson, I. (2005) Cascading dynamics and mitigation assessment in power system disturbances via a hidden failure model. *International Journal of Electrical Power; Energy Systems* 27, 318 – 326.
36. Nedic, D, Dobson, I, Kirschen, D, Carreras, B, & Lynch, V. E. (2006) Criticality in a cascading failure blackout model. *International Journal of Electrical Power; Energy Systems* 28, 627 – 633.
37. Pfitzner, R, Turitsyn, K, & Chertkov, M. (2011) *Statistical classification of cascading failures in power grids*. pp. 1–8.
38. Pfitzner, R, Turitsyn, K, & Chertkov, M. (2011) Controlled Tripping of Overheated Lines Mitigates Power Outages. *arxiv:1104.4558*.
39. Bienstock, D. (2011) Adaptive online control of cascading blackouts. *Proc. 2011 IEEE Power and Energy Society General Meeting* pp. 1–8.
40. Bienstock, D. (2011) Optimal control of cascading power grid failures. *Proc 2011 CDC-ECC* pp. 2166–2173.
41. Eppstein, M & Hines, P. (2012) A "random chemistry" algorithm for identifying collections of multiple contingencies that initiate cascading failure. *Power Systems, IEEE Transactions on* 27, 1698–1705.
42. Bernstein, A, Hay, D, Uzunoglu, M, & Zussman, G. (2012) Power grid vulnerability to geographically correlated failures - analysis and control implications (<http://arxiv.org/abs/1206.1099>).
43. Bienstock, D & Iyengar, G. (2006) Faster approximation algorithms for covering and packing problems. *SIAM Journal on Computing* 35, 825–854.
44. Bienstock, D. (2002) *Potential Function Methods for Approximately Solving Linear Programming Problems, Theory and Practice*. (Kluwer).
45. Kraning, M, Chu, E, Lavaei, J, & Boyd, S. (2012) Message passing for dynamic network energy management (http://www.stanford.edu/~boyd/papers/decen_dyn_opt.html).
46. CPLEX. (2012) ILOG CPLEX Optimizer (<http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/>).
47. GUROBI. (2012) Optimizer (<http://www.gurobi.com/>).
48. MOSEK. (2012) Optimizer (<http://www.mosek.com/>).
49. WECC. (2012) Western electricity coordination council (<http://transmission.bpa.gov/business/operations/wind/>).
50. Grötschel, M, Lovász, L, & Schrijver, A. (1998) *Geometric Algorithms and Combinatorial Optimization*. (Springer-Verlag).