

Commitment and Dispatch with Uncertain Wind Generation by Dynamic Programming

Jeremy J. Hargreaves, *Member, IEEE*, Benjamin F. Hobbs, *Fellow, IEEE*

Abstract— Fluctuating wind production over short time periods is balanced by adjusting generation from thermal plants to meet demand. Thermal ramp rates are limited, so increased variation in wind output as wind penetration increases can add to system operating costs because of the need for more thermal operating reserves. Traditional deterministic modeling techniques fail to fully capture these extra costs. We propose a stochastic dynamic programming (SDP) approach to unit commitment and dispatch, minimizing operating costs by making optimal unit commitment, dispatch, and storage decisions in the face of uncertain wind generation. The SDP solution is compared with two other solutions: (1) that of a deterministic dynamic program with perfect wind predictions to find the cost of imperfect information, and (2) that of a simulation model run under a decision rule, derived from Monte Carlo simulations of the deterministic model, to assess the cost of sub-optimal stochastic decision making. An example application to the Netherlands generation system shows that these two types of costs can amount to several percent of total production costs, depending on the amount of installed wind capacity. These are the conclusions of a single case study with simplified assumptions. Nonetheless, the results indicate that efforts to improve wind forecasting and to develop stochastic commitment models may be highly beneficial.

Index Terms—Stochastic Dynamic Programming, Markov Chains, Monte Carlo Simulation, Power Market Models, Renewable Energy Integration

I. NOMENCLATURE

A. Indices and Sets

- I Set of all aggregate generating units, indexed by i
- J Set of all component units within an aggregate unit, indexed by j
- Y_t Set of all possible states that meet the demand in t

B. Variables

- w_t Wind generation in t , discretized by Δ in $[0, \bar{w}]$
- $g_{i,t}$ Generation of i in t , discretized by Δ in $[\underline{G}_i(\underline{u}_i), \bar{G}_i(\bar{u}_i)]$
- s_t Storage in t , discretized by Δ in $[0, \bar{S}_t]$
- d_t Demand in t , discretized by Δ in $[\underline{D}_t, \bar{D}_t]$
- ws_t Wind spilled in t
- ll_t Loss of load in t
- rr_{jd} Ramping rate of component unit j meeting demand d in

aggregate unit

- n_{jd} Binary variable: 1 if unit is committed, 0 otherwise
- $u_{i,t}$ Integer variable: number of component units committed within an aggregate unit i in t
- mc_t Marginal cost in t
- x_{jd} Continuous generation level of unit j meeting demand d
- y_{jd} Binary variable: 1 if $n_{jd} - n_{jd-1} = 1$, 0 otherwise

C. Parameters

- A Constant from heat rate curve
- B First order component of heat rate curve
- C Second order component of heat rate curve
- \bar{S}_t Storage capacity in t
- \underline{RR}_i Rate of ramping down unit i (\overline{RRS} for storage)
- \overline{RR}_i Rate of ramping up unit i (\overline{RRS} for storage)
- $\underline{G}_i(u_i)$ Minimum generation of unit i given commitment u_i
- $\bar{G}_i(u_i)$ Maximum generation of unit i given commitment u_i
- Δ Increment between levels of discretized state variables
- UP Penalty for underproduction
- OP Penalty for overproduction
- CL Loss during charging of storage
- DL Loss during discharge of storage
- Dem_d Demand at level d of aggregate unit, discretized by Δ in $[\underline{G}_i(\underline{u}_i), \bar{G}_i(\bar{u}_i)]$
- T Number of stages in SDP, indexed by t
- D Number of demand levels in aggregate unit, index by d
- $Norm_d^R$ Ramp normalization factor for demand d
- $Norm_d^C$ Cost normalization factor for demand d
- W_1, W_2 Weights applied to aggregation model objectives
- $SU_{j,d}$ Start-up cost of unit j

D. Functions

- $C(u_{it}, u_{it+1})$ Cost of unit commitment for unit i moving from time t to $t+1$
- $C_X(X_t)$ Cost of unit dispatch, storage losses, and loss of load

II. INTRODUCTION

ENERGY security concerns and the desire to minimize greenhouse gas emissions has driven energy policy in many countries over the past decade. Renewable energy generating capacity rose steeply in response to the resulting initiatives. In 2008, the US became the world's largest wind energy producer after wind capacity increased by 50% in a single year, adding 8558 MW [1]. High wind capacity growth is likely to continue through the next decade. Many states have

J.J. Hargreaves is with E3, San Francisco, CA (jeremy@ethree.com) and B.F. Hobbs is with The Johns Hopkins University, Baltimore, MD, USA (bhobbs@jhu.edu). This work was supported by NSF Grant ECS 0621920 and by Energieonderzoek Centrum Nederlands.

ambitious portfolio standards [2], and Congress has been considering proposals for national renewable portfolio standards. However, reliable and economic operation of a generation system with large proportions of wind will be a challenge.

Wind generated electricity is variable and uncertain with limited control. Currently, it is often modeled as a negative load on the grid in most of the US because its operating costs are the lowest, though in Texas, for example, high penetration of wind generation has already led to curtailment in some circumstances [3]. Wind leads to increased variance and uncertainty in the net load met by the thermal part of the system, which must therefore be capable of higher production ramping rates. Higher operating costs are expected in comparison to an equivalent but constant reduction in demand. A power system should ideally be scheduled such that the selected unit commitment and dispatch will minimize expected system operating costs given uncertainty in the demand and supply.

Integration of wind generation has been a significant area of research over the last few years. The aims of such research and the different approaches taken are summarized in [4]. Smith et al. [5] describe the technical difficulties and benefits that arise from wind integration, and present results from past integration studies. Many studies have been commissioned or carried out by regional planning organizations to determine the benefits and challenges of wind integration within their area of operation. In the US, the biggest of these include the Eastern Wind Integration and Transmission Study (EWITS) [6], and the Western Wind and Solar Integration Study (WWSIS) [7]. The typical approach taken by such studies is summarized in [8]. They simulate the behavior of system operators to find the costs of wind integration, and the effects on system operation, as in [9, 10]. US and European studies have been compared in [11]. They use historical wind speed, demand, and wind and demand forecast data to simulate the response of power systems to increased wind penetration. Uncertainty in wind generation is dealt with as a forecast error on an hourly or greater time scale. Statistical analysis is used to determine a reserve requirement, ensuring that forecasting errors in wind generation can be dealt with without reliability problems. When sub-hourly periods are considered in [12], sub-hourly wind generation is predicted with persistence forecasting until the next hourly wind forecast.

These integration studies compare costs of integration for different levels of wind generation. However, commitment and dispatch are simulated rather than optimally chosen to minimize expected system operating cost. Optimal unit commitment models [e.g. 13,14] have traditionally focused on selecting commitment and dispatch to minimize cost while meeting a deterministic demand. However, a purely thermal system does face uncertain demand [15], and deviations from the forecasted demand have been met by system operating reserve requirements [16]. However, as growing wind penetrations increases uncertainty in the net demand facing thermal plants, using predetermined operating reserve margins to ensure feasibility may yield suboptimal solutions.

Some models that address optimal commitment and dispatch with uncertain wind generation have been developed.

Wang et al. [17] present a mixed integer programming based approach to commitment with uncertain wind generation that ensures reliability considering several Monte Carlo-generated scenarios. This method is capable of modeling large systems; however dispatch decisions are made without considering the statistical dependence of future wind generation on present wind status. Short-term uncertainty in wind generation can be better captured by a probabilistic definition of wind generation in the next model time step, based on the known current and previous wind generation, leading to lower operating cost solutions [18]. Modeling the German market, Swider and Weber [19] took a stochastic recombining tree approach to unit commitment with probabilistic wind transitions between time periods. However, their time step was too large to capture unit ramping constraints. Tuohy et al. [20] used the WILMAR scheduling tool [21] to capture the stochastic behavior of the wind variable with rolling planning horizons. Scenarios are created using Monte Carlo simulation. Stochastic commitment decisions are then made given updated wind forecast information and the probability of different scenarios occurring. This approach comes closer to capturing the uncertainty faced by system operators. However, perfect forecasts are still assumed within each scenario, and a spinning reserve requirement is used to deal with wind uncertainty. Further, a single hour is the shortest time period used in any of the stochastic models. The stochastic variability of wind generation over short time periods (time steps of 15 minutes or less), when ramping rate constraints strongly influence dispatch decisions, is not yet addressed in unit commitment modeling efforts. Optimal commitment and dispatch models that have smaller time steps while accounting for the dependence of wind output on previous periods may find significantly different solutions and costs than models using larger time steps or scenarios.

We propose a stochastic dynamic programming (SDP) approach that uses (1) a shorter time step to better capture ramp rate limits and (2) a stochastic process representation of wind output rather than scenarios. We then compare its results with those of two other unit commitment models to address two questions. First, what is an upper bound to the value of better wind forecasts? Second, what is the value of stochastic unit commitment and dispatch relative to decision making based on deterministic models?

To answer the first question, the results from the SDP approach are compared to those of a deterministic model solved multiple times using Monte Carlo-generated scenarios of wind generation. The latter is used to represent the widely used deterministic approach that accounts for variation but not uncertainty. We demonstrate the weaknesses in modeling short-term unit commitment and dispatch with deterministic models; because those models assume perfect forecasts of wind and demand, they understate the costs of wind integration. The difference in expected performance of the SDP and the Monte Carlo/deterministic approach is the economic value of perfect wind and demand forecasts, which is an upper bound to the value of improvements in such forecasts.

The second question is addressed by comparing the SDP results with those of a model that uses a heuristic rule for commitment and dispatch decisions. A fixed operating reserve

requirement is an example of a heuristic decision rule that might not be optimal in all situations. The heuristic we consider is derived by first analyzing the results of the above Monte Carlo/deterministic model to identify the most commonly taken commitment decision taken within each system state (across all the scenarios considered), and then use a single time step optimization to find the lowest cost dispatch for the next time interval, subject to that commitment. The results of this heuristic operation strategy are compared to the results from the SDP to quantify the cost of suboptimal decision making and the possible benefits of stochastic optimization of dispatch and unit commitment.

Depending on the magnitude of the cost of suboptimal decision making, this model comparison may show either that system operating costs can be estimated and satisfactory operating strategies found using deterministic modeling approaches combined with heuristic decision rules, or, alternatively, that large scale commitment and dispatch models would benefit from explicitly including uncertainty. The adequacy of a deterministic decision rule approach is likely to depend on the level of wind penetration.

The following section describes the models we apply. The SDP formulation is given in III.A followed by the stochastic model of wind production in III.B. Descriptions of the Monte Carlo/deterministic model and decision rule model are given in III.C and III.D respectively. To overcome the curse of dimensionality found in dynamic programming, the Netherlands generation units were aggregated into four groups, each with similar characteristics. The aggregation process is described in III.E. The results of unit aggregation are presented in IV.A and the model comparisons are presented in IV.B, including our estimates of the value of perfect information and the value of stochastic optimization for the Netherlands system under various level of wind investment.

III. MODELS

A SDP-based commitment and dispatch model is a stochastic optimization that identifies an optimal strategy, defined as the immediate decision, for each system operating state at each time stage, that minimizes expected future costs. This can be viewed as the search for the lowest cost probabilistic paths through a stochastic network. In our general model, each stage includes the following state variables: electricity demand, generation from wind, generation from each thermal unit in the system, commitment status of each unit, and how much electricity is being stored. Each of these variables is discretized through selection of some basic increment of energy, resulting in a finite number of levels that each variable can realize. Demand and wind are assumed to evolve over time according to a Markov process. The set of states within each stage contains all the combinations of state variables that meet demand in that stage.

A. SDP Formulation

The model can be described through a two stage optimization by the following recursive Bellman equation [22]:

For each $\{t, Y_t\}$:

$$F_t(Y_t) = \text{Min}_{\{u_{t+1}\}} C(u_{i,t}, u_{i,t+1}) + \sum_{w_{t+1}} P(w_{t+1}, d_{t+1} | w_t, d_t) * \left[\text{Min}\{g_{t+1}, s_{t+1}, ws_{t+1}, ll_{t+1}\} \sum_{d_{t+1}} (C(X_{w_t}) + F_{t+1}(Y_{t+1})) \right] \quad (1)$$

s.t.

$$\sum_i (g_{i,t+1}) + w_{t+1} - ws_{t+1} + ll_{t+1} + (s_{t+1} - s_t) - d_{t+1} = 0 \quad (2)$$

$$\underline{G}_i(u_{i,t+1}) \leq g_{i,t+1} \leq \overline{G}_i(u_{i,t+1}) \quad \forall i \in I \quad (3)$$

defined such that $\overline{G}_i(u_{i,t}) + \Delta = \underline{G}_i(u_{i,t} + 1)$

$$\underline{RR}_i(g_{i,t}) \leq g_{i,t+1} - g_{i,t} \leq \overline{RR}_i(g_{i,t}) \quad \forall i \in I \quad (4)$$

$$\underline{RRS} \leq s_{t+1} - s_t \leq \overline{RRS} \quad (5)$$

$$0 \leq s_t \leq \overline{S}_t \quad (6)$$

with $X_t = \{g_b, g_{t+1}, s_b, s_{t+1}, ll_{t+1}\}$ and $Y_t = \{w_b, d_b, u_b, g_t, s_t\}$ such that $Y_t \in \mathbf{Y}_t$.

The two minimizations in (1) correspond to the two decisions being made by the system operator within a time step. In the first, or outer, minimization, a unit commitment decision for each aggregate unit is made before the wind generation is realized. The decision made gives the best commitment decision based on a minimization of the expected cost of generation. The second, or inner, minimization occurs after wind generation is realized. This is the unit dispatch and electricity storage decision, which is subject to the constraints (2)-(6).

In the general form, the model can account for ramping costs [23] from g_t to g_{t+1} . However, to reduce model solution time, it was assumed that ramping occurs as an instantaneous step change at the beginning of each 15 minute period. The cost function $C_X(X_t)$ therefore calculates the fuel cost of continuous operation at g_{t+1} over the period, unit start-up and shut-down costs, plus the cost of any losses due to storage charging and discharging,¹ and the cost of any loss of load (7):

$$C_X(X_t) = \sum_{i \in I} (A_i + B_i g_{i,t+1} + C_i g_{i,t+1}^2) F_i + (s_{t+1} - s_t)^+ \left(\frac{CL}{1 - CL} \right) mc_t + (s_t - s_{t+1})^+ \left(\frac{DL}{1 - DL} \right) mc_t \quad (7) + (ll_{t+1})^+ UP + (ll_{t+1})^- OP$$

Some unit commitment characteristics were not considered including minimum unit up times and down times and start-up costs as a function of the time that a unit has been shut down.

¹ The costs of charging and discharging losses from storage are approximate. To ensure contributions to and from storage remained a multiple of Δ , the lost energy from charging and discharging was covered by incrementing the operation of the marginal unit in $t+1$ by an amount equal to the loss. Costs from charging losses are accrued at the marginal unit cost at time of charging and discharging costs occur at the time of discharging. Ramping rate and maximum generation constraints of the marginal unit are not accounted for during this calculation. This approach is justified since losses from storage are likely to be small in comparison to the ramping limits of the gas units, which have themselves been approximated to multiples of Δ . Moreover, the gas units are shown in the results to rarely operate at maximum production levels. Modeling large storage capacities might justify more sophisticated accounting of ramp and capacity constraints, though at the cost of longer model run times.

A heuristic approach to solving the SDP simulates market conditions where spilling wind or loss of load are last resort options, and if a feasible commitment and dispatch is available without utilizing them, it is chosen in preference. Unless (2)-(6) cannot be satisfied, the inner minimization is over $\{g_{t+1}, s_{t+1}\}$, setting $w_{s_{t+1}}$ and l_{t+1} equal to zero; if there is no feasible $\{g_{t+1}, s_{t+1}\}$ in that case, then nonzero spill and loss of load are considered by allowing wind to be spilled and penalties on loss of load to be imposed. Loss of load is penalized at a rate of \$1000/MWh and spill is penalized at \$30/MWh, and not allowed to exceed w_{t+1} .

The conditional probability in (1) captures the uncertainty a system operator faces when making decisions for the next period. This formulation assumes that the behavior of wind generation can be described by a first-order Markov process. Although higher order Markov processes could be represented by conditioning upon additional lagged wind state variables, the size of the commitment and dispatch problem precludes their use (the DP ‘curse of dimensionality’). However, based on our analysis of Netherlands wind generation time series in section III.B, we conclude that the daily evolution of wind production on a 15 minute time step is adequately estimated by a first order Markov chain.

The size of the commitment and dispatch problem also limits the number of generating units that can be modeled by the SDP approach. Each unit i in I is therefore an aggregate unit constructed from component units with similar characteristics. A component unit is a physical generating unit in the Netherlands. Section III.E describes the process of unit aggregation.

The model was solved for 24 hours (or 96 15-minute stages), given a known initial wind state. The initial generation state was not set, rather costs were found for all possible initial generation states. The results presented in this paper use the results from the lowest cost initial generation state.

B. Markov Chain Wind Model Estimation and Evaluation

To simplify the model formulation for the comparison presented in this paper, demand was assumed to be deterministic, leaving wind as the only uncertain variable. A year of aggregate Netherlands wind generation data, recorded at 15 minute intervals, and scaled to 35% onshore and 65% offshore with a total capacity of 4600 MW was provided by ECN [24]. Short-term wind speed simulation models that capture the statistical characteristics of a wind time series are often complex. However wind speed data can first be translated in to wind power production by truncating the tails of the wind speed distribution in which power production is zero. Papaefthymiou and Klockl [18] show that the resulting generation time series can be simulated using a Markov Chain.

Markov Chain simulation requires that the continuous stochastic variable, in this case wind, be discretized into a number of bins or states. For example, in the case of a maximum wind generation of 4500MW, nine evenly sized bins are shown in Fig. 2, where the median value of each bin represents its contents. In the SDP, the median values must be multiples of increment Δ .

The second step of Markov Chain simulation is to construct a Markov transition matrix. First order Markov chains have a two dimensional transition matrix that records the probability of moving from each discrete wind state in time t to any other

in time $t+1$, shown for the nine bin example in (8). This square matrix is populated by the probability of each state to state transition by recording the frequency of each transition occurring in the historic record.

$$P(w_{t+1} | w_t) = \begin{array}{c|cccc} w(MW)_t \setminus t+1 & 250 & 750 & \cdots & 4250 \\ \hline 250 & p_{11} & p_{12} & \cdots & p_{19} \\ 750 & p_{21} & p_{22} & \cdots & p_{29} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 4250 & p_{91} & p_{92} & \cdots & p_{99} \end{array} \quad (8)$$

To verify the fit of the simulated data to the real time series, and to generate random wind time series for use in the deterministic model, the following process is followed. A cumulative probability transition matrix is formed from the transition matrix where each entry $c_{mn} = \sum_{q=1, \dots, n-1} p_{mq}$, with m indexing the states in stage t and n indexing the states in $t+1$.

When simulating wind generation data, a uniform random number, U , on $[0,1]$ was generated to find the wind state in the next stage. If the system is in state m in stage t , the wind state in the next stage would be state n such that $c_{m,n-1} < U \leq c_{m,n}$.

This procedure is described in detail in [18]. Four year long time series were created, ensuring that they were representative of the stationary Markov chain distribution.

Discretizing a continuous variable results in loss of information. The fewer discrete states that are considered, the more information is lost. Conversely, if too many states are chosen, probabilities in the transition matrix (8) will be subject to more sample error, because certain state to state combinations will occur rarely [18]. In this analysis, the number of discrete states was chosen to give the best fit in terms of two statistical characteristics of the historical time series: the autocorrelation function (acf) and the unconditional probability density function (pdf). A comparison of the acfs evaluates how well the simulated data captures the degree of persistence in the real time series [18]. Meanwhile, comparing the pdfs of simulated wind output data to that of the real data determines how well the Markov process captures the long term frequency of wind generation levels.

Information is also lost if a single transition matrix is assumed for transitions for all hours of the day. This is certainly inappropriate in regions with strong diurnal wind patterns. The SDP is capable of using time-of-day as well as monthly differentiated transition matrices to reflect the diurnal and seasonal patterns often observed in wind time series. In this case, a separate transition matrix would be constructed for each period of the day. However, constructing separate transition matrices for different time periods reduces the data used to construct each transition matrix. Transition matrices at the 15 minute level were not possible due to the lack of data for each period. Dividing the historical record into hourly periods however did yield good results based on the Dutch data. For this reason, we consider a third characteristic to assess the performance of differentiated transition matrices. In particular, the average daily wind generation profile from differentiated transition matrices was compared to the historical average daily wind profile to assess the quality of fit.

The acf comparison for several different Markov models using a single transition matrix is shown in fig 1. These are

shown for 5, 8 and 9 bins. Shown in the same figure are the results from hourly differentiated transition matrices. Only the best hourly differentiated model is shown, which had 9 bins.

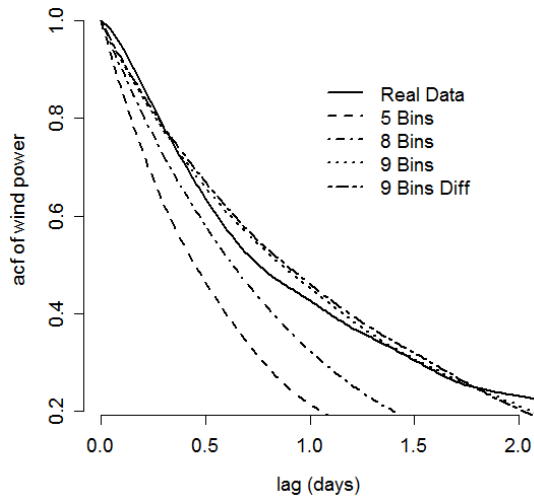


Fig. 1. The acf of wind production data simulated with varying bin numbers in the Markov model vs the acf of the real time series.

The results of the acf analysis show that as bin size is increased from 5 to 9 bins for the single transition matrix model, the fit to the real data improves. The fit at 9 bins shows close agreement to the persistence of the real data. Higher numbers of bins did not improve the fit further. The same trend was observed using hourly differentiated transition matrices (only the 9 bins results are shown in fig. 1).

The pdf comparison between the real data and the 9 bin single transition matrix model simulation is given in fig 2(a). The expected wind generation in the scaled real data provided by ECN is 1739MW. In the simulated data it is 1766MW, so close agreement to the real data is found. This together with the acf comparison shows that the characteristics of the Netherlands wind generation time series are well captured in the SDP when using a 9 bin first order Markov chain to represent the stochastic wind variable.




Fig. 2. The pdf of 9 bin Markov chain simulated data vs real data. *Left:* (a) single transition matrix. *Right:* (b) differentiated hourly transition matrices

The pdf of the differentiated 9 bin model is shown in fig 2(b). A small bias towards higher wind states is seen in the simulated wind data. This may be caused by a lack of data to form accurate differentiated hourly transition matrices.

In the third comparison of simulated and historical data, we compare the steady state expected wind generation for the differentiated 9 bin model and the actual data in fig 3. Though

the shape of the diurnal trend is well captured, a slight bias towards higher levels of wind generation is seen here too. This upward bias has two components, the larger being the lack of data. The fit could be expected to improve with increasing size of the historical dataset. With a larger dataset, differentiation by month as well as hour may further improve the fit.




Fig.3. Daily expected wind gen profile for real and simulated time series

The second component, a systematic upward bias of up to 1% at low wind generation levels appears to be introduced by using the median bin wind generation as the value of wind generation associated with each bin in the simulated data. That bin value is restricted by the SDP methodology to be a multiple of the increment used in the dynamic program. However, that upward bias can be minimized by selecting upper and lower limits for each bin that approximately set the average wind generation in each bin equal to the corresponding bin value. This could be accomplished by an optimization in which limits are chosen to minimize the difference between the bin value and average wind generation in each bin.

These slight upward shifts in average values are unlikely to materially affect the main conclusions of this paper concerning the impacts of wind variability and forecast uncertainty upon expected costs.

It is important to note that the Markov chain derived here is specific to a single year of wind generation data. Not only would more years of data improve the fit and allow greater chain differentiation, it would also result in a model more representative of annual wind energy output. Ideally the historical data should approximate the long term mean energy output of wind in the region being modeled.

The average daily demand profile for the Netherlands (based on year 2006) is shown below in fig. 4 for comparison. Wind generation increases in unison with increasing demand in the morning. However, as demand peaks in late afternoon, wind generation drops off.




Fig.4. Deterministic daily demand profile

C. Monte Carlo/Deterministic Dynamic Program (MCDP)

To capture the uncertainty in wind production using a scenario-based deterministic modeling approach, the SDP was simplified to a deterministic dynamic program that optimizes subject to perfect wind forecast. A set of wind generation scenarios, A , was created with Monte Carlo sampling of the transition matrix (8) using the same procedure used to simulate wind data in section III.B. The result was a vector of wind outputs, w_a , for each generation scenario $a \in A$. For each $a \in A$, $P(w_{a,t+1}|w_{a,t})$ was set to 1 and all other transition probabilities were set to 0, reducing the SDP to a deterministic DP.

For each level of wind capacity investigated, we ran 1000 scenarios. Averaging the costs from the 1000 MCDP runs gave the expected cost. This represented the total operating cost given perfect forecasts of wind generation over the 24 hour period examined with the model. This provides a lower bound (subject to sample error from the Monte Carlo sampling) for the expected cost from an optimal stochastic solution from the SDP, since the SDP optimizes for the more realistic case of imperfect knowledge of future wind output.

D. Suboptimal Decision Rule Model

The suboptimal decision rule is a heuristic method of making commitment and dispatch decisions for the next time period. It was designed to be a readily implementable by system operators using existing deterministic unit commitment and dispatch models.

Over a sufficient number of simulations of a deterministic model such as the MCDP of the previous section, many commitment and dispatch decisions can be observed for each state in each time period t , and the most common commitment decision made in each state can be noted. The most common dispatch decisions made once wind generation in $t+1$ is realized can also be recorded. This database of the most frequently made decisions in the deterministic commitment and dispatch model runs can then be used to derive an operating rule for system operators, prescribing what is to be done under each system state.

This approach to creating an operating rule has two problems. First, the number of possible combinations of state variables at any point in time is extremely large. Not all of those combinations will occur, even if thousands of scenario-based deterministic model runs are made. Thus there will be states for which there are few or no observed decisions. Second, the most commonly made commitment decisions from the MCDP will often only satisfy demand in $t+1$ for the wind generation state in $t+1$ with the highest probability. To meet other possible wind generation states, wind spill or loss of load penalties may be incurred, and the expected cost may be higher than for an optimal stochastic policy from a SDP.

To solve the first problem, if no commitment decision is available in the operator's database, then a single step SDP optimization is used to determine the next transition. That is, given the probability distribution of wind generation in the next period, what are the lowest cost commitment and dispatch decisions to take in t without regard to stages beyond $t+1$? If the database of decisions recorded from scenario modeling is sufficiently populated, this single state optimization

will rarely occur. Furthermore, because this is a single stage optimization of one period's decisions, an approach capable of modeling larger numbers of generating units could be substituted for the SDP, such as proposed in [4]. An alternative to the single step optimization would be to use a surrogate commitment decision from a closely related state that has an available decision in the database.

Turning to the second problem, commitment decisions are available for a given state from the database of simulation runs; however a dispatch decision is unavailable for some of the possible wind states. In this case the single step optimization is used to decide dispatch, spill, and loss of load while making the commitment that occurs most often for that state in the database of simulations.

The expected cost of using this decision rule will be higher than the SDP since decisions taken in scenario modeling with perfect foresight often leave the system inadequately insured against rare or unexpected changes in wind generation. This could result in higher fuel costs or loss of load. The SDP, in contrast, by design minimizes expected cost.

E. Unit Aggregation

Unit aggregation is necessary because of the curse of dimensionality; it is not possible to include generation of each individual unit as a state variable. Intractable energy model formulations are common with large numbers of units, and aggregation is often used to reduce problem sizes. An example is given in [25], where Langrene et al. solve a large scale dispatch model with dynamic constraints by making use of unit aggregation. An SDP containing aggregate units represents only a subset of the possible unit dispatch configurations. Thus operating costs may be different than when units can act independently; for instance, the additional flexibility may result in lower costs. Applications of this methodology that have fewer aggregations and more units will result in better approximations of the cost. In the example application to the Netherlands system, however, both the SDP and the DP contain the same aggregate units, so the comparative cost results are valid given the aggregate system modeled. Future work in this area will focus on reducing the aggregation required, and investigating its effect on expected cost estimates.

We undertake aggregation in two steps. First, we divided the generating units into subsets, each representing one aggregate unit. Second, we define the characteristics of each aggregate unit, including min and max capacity; ramp rates as a function of the level of aggregate unit output; and operating cost as a function of the level of aggregate unit output.

Addressing the first step, aggregate units were defined corresponding to generation type including coal, natural gas, and combustion turbines. The natural gas aggregate was split further into two subgroups of units with similar marginal costs.

Once aggregate units were chosen, their characteristics were found in the second step. Aggregation of units for use in the SDP cannot follow simple stacking by marginal or average cost because of the short time step used in the SDP. Such stacking would be appropriate only if component units in the aggregate were flexible enough to transition from stacked cost dispatch to stacked cost dispatch during all transitions encoun-

tered during system operation within the 15 minute time step; however, ramping rates make this impossible. As such, the ramping rates and operating costs of an aggregate unit, when meeting a certain demand in $t+1$, are influenced by the state of each component unit meeting demand in t ; as a result, ramp rate constraints prevent an aggregate unit from moving to the minimum cost operating point in every stage. The ideal definition of an aggregate unit will meet some objective, for example maximizing ramping rates across all levels of output of the aggregate unit. We used a mixed integer linear program (MILP) to aggregate units according to two competing objectives: minimizing operating costs and maximizing ramp rates, each normalized over the ranges of possible output. In that aggregation process, it was necessary to assume linear variable costs for each component unit. However, the more accurate quadratic cost functions for component units were the basis of the final calculation of the operating costs function for the aggregate units.

The formulation of the MILP aggregation model is below:

$$\text{Max}_{\{\bar{rr}, \underline{rr}, n, x\}} W_1 \sum_{d=1, \dots, D} \frac{\sum_{j \in J} (\bar{rr}_{j,d} + \underline{rr}_{j,d})}{\text{Norm}_d^R} - W_2 \sum_{d=1, \dots, D} \frac{\sum_{j \in J} F(n_{j,d}, x_{j,d}, y_{j,d})}{\text{Norm}_d^C} \quad (9)$$

$$\sum_{j \in J} x_{j,d} = \text{Dem}_d \quad \forall d = 1..D \quad (10)$$

$$\bar{rr}_{j,d} \leq \bar{G}_j - x_{j,d} \quad \forall j \in J, d = 1..D \quad (11)$$

$$\bar{rr}_{j,d} \leq n_{j,d} R_j \quad \forall j \in J, d = 1..D \quad (12)$$

$$\underline{rr}_{j,d} \leq x_{j,d} - n_{j,d} \underline{G}_j \quad \forall j \in J, d = 1..D \quad (13)$$

$$\underline{rr}_{j,d} \leq n_{j,d} R_j \quad \forall j \in J, d = 1..D \quad (14)$$

$$x_{j,d} \geq n_{j,d} \underline{G}_j \quad \forall j \in J, d = 1..D \quad (15)$$

$$x_{j,d} \leq n_{j,d} \bar{G}_j \quad \forall j \in J, d = 1..D \quad (16)$$

$$x_{j,d+1} \leq x_{j,d} + n_{j,d} R_j + (1 - n_{j,d}) \underline{G}_j \quad \forall j \in J, d = 1..D - 1 \quad (17)$$

$$y_{j,d} \geq n_{j,d} - n_{j,d-1} \quad \forall j \quad (18)$$

with $F(n_{j,d}, x_{j,d}) = A_j n_{j,d} + S U_j y_{j,d} + (B_j + C_j (\bar{G}_j - \underline{G}_j) / 2) x_{j,d}$.

The objective function (9) minimizes operating costs while maximizing ramping rates according to the weights placed on each objective. A solution that optimizes one objective will not optimize the other; for instance, maximizing rampability requires more units being operated between their lower and upper bounds than if units are committed to minimize cost; thus, the weighted objective will achieve a compromise between the two goals. The optimization is subject to an energy balance (10), ramping rate constraints (11)-(14), and unit commitment constraints (15) and (16). Constraint (17) ensures that no component unit violates its ramping rate constraints when the aggregate unit ramps generation up or down by Δ , the increment in the state variability discretization. Constraint

(18) imposes a start-up cost for units transitioning to committed status.²

IV. EXAMPLE APPLICATION: NETHERLANDS CASE STUDY

A. Aggregate Unit Summary

The characteristics of the four aggregate units are summarized in Table 1. These were generated from information on Dutch generating units, including fixed operating costs, quadratic heat rate curves, capacities, ramping rates and fuel costs. Start-up costs were estimated from industry data [e.g., 26]. The output of combined heat and power plants is deducted from the load, assuming that they are committed to meet heat demand and not in response to power prices. No storage was modeled in the example.

Component units within the gas fueled aggregates were also assumed to be capable of starting up to minimum generation and shutting down from minimum generation in a 15 minute time period. This reflects the assumption that a unit is required to reach minimum generation before being connected to the grid, and ramping down generation to minimum levels before being disconnected. However, the SDP makes commitment decisions 15 minutes prior to a unit coming online, biasing the model towards higher system flexibility than is actually the case, at least for gas steam units. Higher flexibility is expected to lead to lower total operating costs; this limitation of the model is to be addressed in future work. This is the opposite bias of most wind integration studies, which instead assume fixed commitment schedules. For example, [4] assumes day ahead commitment and [20] used 3 hour commitment blocks. Component units within the coal aggregate in the SDP were assumed to follow a fixed commitment schedule determined at the beginning of the day. For the example here, all coal units were committed to be on-line for the full 24 hours.

Each aggregate unit was created using the model (9)-(18) - presented in the previous section, using equal weights on each of the two objectives. The generation range of each unit was discretized using a value of 250MW for Δ , striking a balance between SDP processing time and accuracy. Component gas units with a minimum marginal cost less than \$46/MW were assigned to aggregate unit Gas 1, and those above were assigned to Gas 2. We chose this rule because with unconstrained ramping and units committed by marginal cost, those

² Significant improvements in model run time can be made by adding the following cuts to the feasible region.

$$\sum_{j \in J} n_{j,d} C_{\min} \leq \text{Dem}_d \quad \forall d = 1..D$$

$$\sum_{j \in J} n_{j,d} C_{\max} \geq \text{Dem}_d \quad \forall d = 1..D$$

These are necessarily satisfied by any feasible solution.

Our use of aggregate units locks in an ordering of unit commitment (de-commitment) that each component unit must follow as the aggregate unit increases (decreases) generation, removing some flexibility from the model compared to a unit commitment model that models each of the component units separately. The possible states of generation to meet demand in each stage in the SDP are therefore a subset of all possible combinations of component unit states. The solution to the SDP will therefore likely be suboptimal, having a higher expected cost than the true optimum. However, our use of consistent aggregate units for all three modeling approaches investigated in this paper allows for a consistent comparison, allowing us to answer our two questions about value of perfect forecasts and stochastic optimization.

in the lower marginal cost group would be committed before those in the higher group. Since each component coal unit is committed for all 24 hours, coal output ranges from the collective minimum generation of all coal units to their maximum.

TABLE I
CHARACTERISTICS OF AGGREGATE UNITS

Aggregate Unit	Generation		# Component Units	# Commitments to max
	Max	Min		
Coal	4750	2500	11	0
Gas 1	4000	0	11	9
Gas 2	5250	0	22	9
CT	250	0	9	1

The final column in Table 1 is the number of times one or more component units are committed during an increase in generation of the aggregate unit over the entire range of output. For example, the set of committed component units in Gas 1 changes 9 times when ramping from 0 to 4000 MW.

The total variable cost as well as its derivative (marginal cost) for each aggregate unit are shown below in Fig. 5. The marginal cost of Gas 1 is not monotonic. This is due to the balance struck by the aggregation model (9)-(18) between ramping flexibility and cost for each of the aggregate units. The ramping capabilities of each unit are shown in fig. 6. These have been rounded to the increment size of 250MW. The upward ramp rate is highest at zero generation for the gas units because there is the potential to turn on all generators at the same time. Rampability decreases as the maximum generation of the aggregate unit is approached because more of the component units are operating at or near their capacity. Likewise, maximum downward ramp rates decrease as generation decreases because minimum operating levels are approached for more component units.




Fig. 5. Aggregate unit characteristics. Left: (a) variable costs. Right: (b) marginal costs.




Fig. 6. Maximum ramp rates for aggregate units. Left: (a) ramp up. Right: (b) ramp down.

B. Model Comparison

Each of the three models was run under the same conditions for four different levels of installed wind capacity. The expected wind generation is shown in fig. 7. These are all above the 2009 Dutch installed capacity of 2220 MW. At the highest installed capacity, annual wind production would amount to about 27% of the country’s annual energy requirements. This is broadly comparable to 2020 renewable targets in Europe and California. The initial state of wind generation at hour zero was set at two wind states above minimum wind generation. The Markov wind model based on hourly differentiated 9 bin transition matrices ((as described in Section III.B) is used in each case. The evolution of expected wind generation is dependent on the initial wind state in the model.




Fig. 7. Expected wind generation at different levels of installed wind capacity

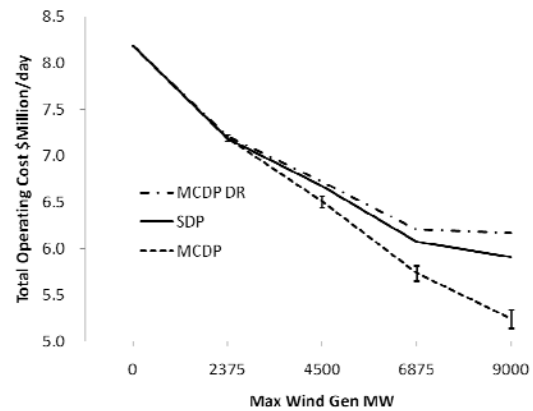


Fig. 8. Total predicted 1 day system operating costs using differentiated transition matrices

Figure 8 shows the results of the model comparisons. The average cost for the 1000 deterministic solutions using Monte Carlo-generated perfect forecasts (the MCDP model) is given along with its 95% confidence interval, accounting for sampling error in the Monte Carlo process. That cost is the lowest among the three models because commitment and dispatch decisions are made knowing exactly what the wind will do in the future, i.e., perfect forecasting. If, for instance, it is known that wind will drop precipitously exactly two hours from now, then units can be redispatched over those two hours to increase the amount of ramp available when needed. This level of foresight is unrealistic. The SDP is more realistic, as it represents uncertainty in wind forecasts. It yields the theoretically optimal feasible system operating decisions, given that uncertainty. Finally, the decision rule (MCDP DR) model

yields higher costs because it uses a heuristic rule to make commitment and dispatch decisions in each state, rather than minimizing expected cost.

Both the decision rule and MCDP costs closely approximate those of the SDP at the lowest level of wind integration. But as the amount of installed wind capacity increases, the MCDP increasingly underestimates the attainable system cost until, at a maximum wind generation of 9000 MW, the difference between the SDP and MCDP solutions is 11% (fig. 8). The difference between the SDP and MCDP costs gives the value of perfect wind forecasts. This shows that scenario-based modeling using a deterministic model and perfect wind forecasting, as used in some wind integration studies, can significantly underestimate system operating costs, thus overestimating the value of wind generation at high levels of wind integration.

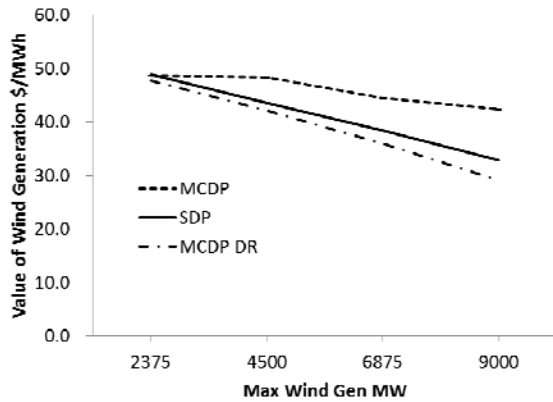


Fig. 9. Marginal value of wind generation found using differentiated transition matrices

Fig. 9 illustrates this overestimation, showing the incremental value of wind estimated by each model (i.e., the negative of the slopes in Fig. 8). At 9000 MW maximum wind generation, the values of wind found by the SDP and DR models are 77% and 69% respectively of that found with perfect forecasts (MCDP). Fig. 9 also shows that wind investment provides diminishing marginal returns, with its value decreasing from about 50 to 30 \$/MWh for the SDP and DR models.

To explore why these results occur the average dispatch by hour for the SDP and MCDP models are shown in Figs. 10 and 11, respectively. For each model, the dispatch under the lowest and highest levels of wind integration are presented.




Fig. 10. SDP expected dispatch chart. Left: (a) 2375 MW wind capacity. Right: (b) 9000 MW wind capacity




Fig. 11. MCDP expected dispatch chart. Left: (a) 2375 MW wind capacity. Right: (b) 9000 MW wind capacity.

At the lowest level of wind integration, the dispatch from the SDP and MCDP models are very similar. Coal operates at or very close to maximum capacity and, other than during the low demand period early in the morning, Gas 2 does most of the load following. Yet at the highest level of integration, dispatch from the two models is markedly different. In each stage in the MCDP model, the model chooses the optimal commitment and dispatch given the known wind generation for the rest of the day. In contrast, the SDP chooses the optimal decision given the full range of wind profiles that can occur. This difference is observable in the dispatch charts. The SDP operates units to maintain a higher level of system ramping ability than does the MCDP, resulting in the large difference in operating costs. The MCDP produces 11% more energy than the SDP from coal over the day; 2% more from cheap gas; 46% less from expensive gas; and 32% less from combustion turbines.

As installed wind capacity increases, the cost penalty from using the heuristic DR model rather than the optimal SDP also increases. This shows that basing an operating rule upon commitments from the MCDP model with perfect forecasts can yield significantly higher costs than stochastic optimization, although the differences are insignificant at low levels of wind penetration. At a wind capacity of 9000 MW, the largest difference in operating cost between the two (i.e., the benefit of stochastic optimization) is 4%. However, this is less than half the largest difference between MCDP and SDP costs (i.e., the value of perfect forecasts). Whether this is a general result that would occur for other systems cannot be determined, however, without applying these models to other cases.

The percentage of wind spilled at each level of wind integration is shown in figs. 12 and 13. All models show negligible wind spill at the lowest level of wind integration. As wind investment increases though, wind spill rises significantly. The SDP shows the highest wind spill out of all the modeling approaches at the higher levels of wind integration. Spilled wind may be highly volatile from stage to stage because the level of spill is not explicitly optimized in our SDP approach. Rather wind is spilled as a last resort to maintain feasibility during transitions from states with no other options. Meanwhile, as might be expected, the MCDP has lower spill because units can be committed ahead of time in anticipation of precisely what wind generation will occur in each hour. Finally, the MCDP decision rule model shows higher spill than the MCDP because unit commitment decisions are made with heuristics based on imperfect information.




Fig. 12. Wind spill. Left: (a) 2375 MW wind capacity. Right: (b) 4500MW wind capacity




Fig. 13. Wind spill. Left: (a) 6875 MW wind capacity. Right: (b) 9000 MW wind capacity

The deterioration in performance of the decision rule heuristic can be inferred from comparing figs. 12 and 13 with the corresponding unserved load shown in figs. 14 and 15. For a wind capacity of 6875 MW, the spill observed in the SDP solution early in the morning is replaced by unserved load in the DR solution. Similarly at 9000 MW maximum wind generation, the DR has significantly higher unserved load than the SDP at times when the SDP is spilling more wind than the decision rule model. This shows that the SDP prefers to commit thermal units that can ramp up to meet the maximum possible demand net of wind in the next period, at the expense of more spilled wind if instead net demand is low and the thermal units are against their minimum run or ramp down constraints. This makes sense because the penalty on unserved load is very high. The decision rule model however must use the most common commitment decision found in the MCDP for a given state in t . Commitment decisions made in the MCDP are likely to meet the most probable wind state in the next stage at lowest cost without allowance for less probable but large fluctuations in wind. The DR model is therefore predisposed to commitment too few units in meeting uncertain demand, resulting in higher levels of unserved demand than the SDP.




Fig. 14. Unserved load. Left: (a) 2375 MW wind capacity. Right: (b) 4500MW wind capacity




Fig. 15. Unserved load. Left: (a) 6875 MW wind capacity. Right: (b) 9000 MW wind capacity

V. CONCLUSION

As the penetration of wind power increases, the difference in total operating cost found by stochastic optimization (SDP) versus a scenario-based deterministic model (MCDP) increases significantly, resulting in a cost difference equal to 11% of the SDP total operating cost at the highest level of wind integration studied. This difference is the value of perfect forecasts, and also indicates the extent to which using such deterministic models for wind integration studies can understate the cost of wind integration. This difference is equivalent to the deterministic model overstating the marginal value of wind generation by almost 50%, relative to the SDP marginal value at the highest wind penetration. The implicit assumption of perfect information in deterministic models is the cause of this inflation of the value of wind.

As expected, the decision rule results in higher costs than the SDP model (and thus lower marginal value of wind). This is in part due to increased unserved demand. However, the probability of unserved demand occurring is low, which partially accounts for the cost of suboptimal decision making at the highest level of wind integration being just 4%. Thus, although the commitment decision made by a heuristic rule can leave the system ill prepared for improbable events, handling these events with a single stage optimization results in a total cost for the decision rule that is close to that of the SDP.

The results suggest that the benefit of using stochastic optimization as opposed to operating rules based on Monte Carlo analysis of a deterministic model are relatively modest for intermediate or small levels of wind integration. This conclusion is, of course, specific to the particular model parameters and input datasets used. The wind data tested was aggregated from turbine sites across the Netherlands, perhaps resulting in lower variance than wind generation serving more local markets. Transmission constraints may inhibit the interregional hedging of turbine sites and should be investigated in further work. Finally, the Netherlands system has a higher proportion of gas units than other markets, and these provide added flexibility and thus lower integration costs than in markets with larger proportions of coal or nuclear.

In addition, the computational limitations of using a SDP to model a large system necessitates significant approximations. The need to combine generators into a small number of aggregate units means that a commitment and dispatch path for components of each aggregate unit must be assumed. It is unclear how this assumption would affect the comparison between models. Discretizations of demand, generation and

time are also necessary in the SDP, as are assumptions of relatively short unit start-up and shut-down times for gas units. Those assumptions overstate the flexibility of the system, and may result in underestimation of the cost differences between models. Unit flexibility is comparatively more valuable in stochastic optimization than deterministic optimization because the latter can prepare for large thermal demand changes ahead of time.

These limitations on drawing general conclusions are the basis for further work on stochastic unit commitment and dispatch. Work on improving the applicability of this methodology by incorporating start-up and shut-down times and reducing the aggregation of units will help better quantify the value of perfect forecasting. This work is important to determining the benefits of better forecasts and whether the added complexity of stochastic unit commitment and dispatch models are a worthwhile investment.

Acknowledgments. The authors gratefully acknowledge the collaboration of colleagues at ECN, including S. Hers, J. van Stralen, and C. Kolokathis, A. Brand of ECN provided the wind data.

VI. REFERENCES

- [1] IEA. "IEA Wind Energy Annual Report 2008." 2009.
- [2] N.C. Solar Center. *State Renewable Incentives*. 2009. www.dsireusa.org (accessed October 22, 2009).
- [3] M.B. Lively, "Renewable Electric Power - Too much of a Good Thing: Looking At ERCOT." *US Association for Energy Economics Dialogue*, 2009: Vol.17, No.2, 21-27.
- [4] L. Soder, H. Holttinen, "On Methodology for Modeling Wind Power Impact on Power Systems", *Int. J. Global Energy Issues*, Vol. 29, No 1/2, 2008.
- [5] J. C. Smith, M.R. Milligan, E.A. DeMeo, B.Parsons, "Utility Wind Integration and Operating Impact State of the Art", *IEEE Trans. Power Systems*, 2008: Vol. 22, No. 3, 900-908.
- [6] NREL, "Eastern Wind Integration and Transmission Study", www.nrel.gov/wind/systemsintegration/ewits.html, 2010.
- [7] NREL, "Western Wind and Solar Integration Study", www.nrel.gov/wind/systemsintegration/wwsis.html, 2010.
- [8] E. Ela, M. Milligan, B. Parsons, D. Lew, D. Corbus, "The Evolution of Wind Power Integration Studies: Past, Present, and Future", *Power & Energy Society General Meeting*, 2009.
- [9] B.C. Ummels, M. Gibescu, E. Pelgrum, W. L. Kling, A. J. Brand, "Impacts of Wind Power on Thermal Generation Unit Commitment and Dispatch." *IEEE Trans. Energy Conv.*, 2007: Vol. 22, No.1, 44-51.
- [10] H. Holttinen, et al. "Impacts of large amounts of wind power on design and operation of power systems, results of IEA collaboration." *8th International Workshop on Large-Scale Integration of Wind Power into Power Systems*. Bremen, 2009.
- [11] H. Holttinen, P.Meibom, A. Orths, M. O'Malley, B.C. Ummels, J.O. Tande, A. Estanqueiro, E. Gomez, J. C. Smith, E. Ela, "Impacts of Large Amounts of Wind Power on Design and Operation of Power Systems; Results of IEA Collaboration", NREL Conference Paper, 2008.
- [12] Y.V. Makariv, C. Loutan, J. Ma, P. de Mello, "Operational Impacts of Wind Generation on California Power Systems", *IEEE Trans. Power Systems*, 2008: Vol. 24, No. 2, 1039-1050.
- [13] B.F. Hobbs, "Optimization methods for electric utility resource planning." *Euro. J. Operational Research*, 1995: Vol.83, 1-20.
- [14] M. Ventosa, A. Baillo, A. Ramos, M. Rivier, "Electricity Market Modeling Trends," *Energy Policy*, 33, 897-913, 2005.
- [15] R. Doherty, M. O'Malley. "New approach to quantify reserve demand in systems with significant installed wind capacity." *IEEE Trans. Power Systems*, 2005: Vol. 20, No.2, 587-595.
- [16] B.F. Hobbs, M. H. Rothkopf, R. P. O'Neill, H. Chao. *The Next Generation of Electric Power Unit Commitment Models*. Norwell, MA: Kluwer, 2001.
- [17] J. Wang, M. Shahidehpour, Z. Li. "Security Constrained Unit Commitment With Volatile Wind Power Generation." *IEEE Trans. Power Systems*, 2008: Vol. 23, No. 3, 1319-1327.
- [18] G. Papaefthymiou, B. Klockl. "MCMC for Wind Power Simulation." *IEEE Trans. Energy Conv.*, 2008: Vol. 23, No. 1, 234-240.
- [19] D.J. Swider, C. Weber. "The Costs of Wind Intermittency in Germany: Application of a Stochastic Electricity Market Model." *European Trans. Elect.Power*, 2007: Vol. 17, 151-172.
- [20] A. Tuohy, P. Meibom, E. Denny, M. O'Malley, "Unit Commitment for Systems with Significant Wind Penetration", *IEEE Trans. Power Systems*, 2009: Vol. 24, No. 2, 592-601.
- [21] P. Meibom, H.V. Larsen, R. Barth, H. Brand, C. Weber, O. Voll, "Wilmar Joint Market Model Documentation", Riso National Laboratory, Denmark, 2006.
- [22] R. Bellman, *Dynamic Programming*. Princeton University Press, 1957.
- [23] C. Wang, S. M. Shahidehpour, "Effects of Ramp-Rate Limits on Unit Commitment and Economic Dispatch", *IEEE Trans. Power Systems*, 1993: Vol.8, No.3,1341 - 1350.
- [24] ECN, Department of Policy Studies, Netherlands. www.ecn.nl/units/ps/
- [25] A.H. van der Weijde, B.F. Hobbs, "Locational-based Coupling of Electricity Markets: Benefits from Coordinating Unit Commitment and Balancing Markets", *J. Regulatory Econ.*, 2011: Vol. 39, No. 3, 223-251.
- [26] N. Langrené, W. van Ackooij, F. Bréant, "Dynamic Constraints for Aggregated Units: Formulation and Application", *IEEE Trans. Power Systems*, 2011: Vol. 26, No. 3, 1349-1356.

Jeremy J. Hargreaves (M '10) received the B.S.E. and M.S.E. degrees in chemical engineering from Imperial College London and Ph.D. degree in Environmental Engineering from the Johns Hopkins University. He is a consultant at Energy and Environmental Economics (E3), San Francisco, CA, where his projects address distributed renewable energy, transmission planning, and wind integration.

Benjamin F. Hobbs (F '08) received the Ph.D. degree from Cornell University, Ithaca, NY. He is Schad Professor of Environmental Management in the Departments of Geography & Environmental Engineering and Applied Mathematics & Statistics, Johns Hopkins University, Baltimore, MD. He also directs the JHU Environment, Energy, Sustainability & Health Institute. Dr. Hobbs chairs the California ISO Market Surveillance Committee.