

©Copyright 2013
Adam Greenhall

Wind Scenarios for Stochastic Energy Scheduling

Adam Greenhall

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2013

Reading Committee:

Richard Christie, Chair

Daniel Kirschen

Miguel Ortega-Vazquez

John Kramlich

Program Authorized to Offer Degree:
Electrical Engineering

University of Washington

Abstract

Wind Scenarios for Stochastic Energy Scheduling

Adam Greenhall

Chair of the Supervisory Committee:
Professor Richard Christie
Electrical Engineering

Large amounts of wind generation have been added to the power system in recent years. However, wind breaks many of the core assumptions in the process used to schedule energy and is particularly difficult to forecast accurately. Rather than scheduling based on a single forecast, stochastic Unit Commitment (UC) minimizes the expected cost over several wind scenarios for the next day. Stochastic UC is often held up as a solution to help alleviate the high costs related to uncertain renewables. Yet there is no widely accepted method for creating high quality stochastic scenarios.

In this dissertation, we examine two wind power scenario creation methods – *moment matching* and *analogs*. Moment matching is a general technique where scenarios are synthesized to match a set of statistics or moments. We propose a method for estimating these desired moments based on historical wind data. The analogs method looks back in time to find similar forecasts and uses the matching observations from those analogous dates directly as scenarios. This work proposes and tests a *simple* analogs method based solely on aggregate wind power forecasts.

The performance of these methods is tested on a realistic model of the Electric Reliability Council Of Texas (ERCOT) power system based on actual data from 2012. UC and dispatch simulations showed modest stochastic savings for the relatively flexible ERCOT model at 25% wind energy penetration. The scenario creation method and number of scenarios had a significant impact on these stochastic savings. Contrary to our hypothesis and the increase in perfect forecast savings, stochastic savings decreased as wind penetration increased to 30%. Stochastic savings are often largely due to a few high cost events during peak load periods; stochastic UC costs may be higher than deterministic UC for extended periods – generally when demand and marginal prices are low. Together these results paint a more nuanced picture of stochastic UC and provide a roadmap for future scenario creation research.

TABLE OF CONTENTS

	Page
Part I: Introduction	1
Part II: Literature Review	6
Chapter 1: Stochastic Energy Scheduling	8
1.1 Deterministic UC	8
1.1.1 Deterministic Thermal UC	8
1.1.2 Hydro scheduling	10
1.2 Decision Making Structure	14
1.2.1 Rolling Planning Scenario Trees	14
1.3 Stochastic Objective Functions	16
1.3.1 Conditional Value at Risk (CVaR)	16
1.3.2 Robust optimization	16
1.4 Solution Methods	18
1.4.1 Progressive Hedging (PH)	18
1.4.2 Stochastic Dual Dynamic Programming (SDDP)	18
1.5 Results	19
Chapter 2: Wind Modeling and Forecasting	22
2.1 Persistence	23
2.2 Normal Error	24
2.3 ARMA Based Wind Forecast	25
2.3.1 Statistical Preprocessing	26
2.3.2 Aggregate Modeling	27
2.3.3 Model Order Selection	27
2.3.4 ARMA in the UC Literature	27
2.4 Numerical Weather Prediction	28
2.5 Probabilistic and Scenario Forecasting	29
2.6 Analogs in Forecasting	29
2.6.1 Analog Techniques	29
2.6.2 Distance metrics	31
2.6.3 Performance	32
2.6.4 Wind forecasting	32

2.6.5	Summary	32
Chapter 3:	Scenario Creation	33
3.1	Monte Carlo Sampling Method	34
3.1.1	Scenario Reduction	34
3.2	Moment Matching Method	35
3.3	Numerical Weather Prediction (NWP) based scenarios	36
3.4	Hydro scenarios	36
Part III:	Methods	37
Chapter 4:	Scenario Creation	38
4.1	Moment Matching	38
4.1.1	Scenario Creation Problem Formulation	38
4.1.2	Scenario Reduction	40
4.1.3	Desired Moment Estimation	41
4.2	Analogs	46
4.2.1	Forecast data	46
4.2.2	Analog distance	48
4.2.3	Selection	48
4.2.4	Assigning probability	49
4.2.5	Caveats	49
4.3	Scenario Metrics	52
4.4	Summary	53
Chapter 5:	Stochastic Scheduling	54
5.1	Decision Making Structure	54
5.1.1	The horizon effect and non-anticipatory constraints	55
5.2	Evaluation	55
5.2.1	Margins of Error	55
Part IV:	Experimental Results	56
Chapter 6:	Experimental Design	57
6.1	Wind Modeling	57
6.1.1	Wind Energy Penetration Scaling	59
6.1.2	Scenario creation	60
6.1.3	Caveats	60
6.2	Thermal System Modeling	61
6.2.1	ERCOT system thermal model	61
6.3	Power System Simulation	61

6.3.1	Unit Commitment Simulation	62
6.3.2	Dispatch Simulation	63
6.3.3	Implementation	63
6.4	Summary	64
Chapter 7:	Results	65
7.1	Magnitude of Savings	65
7.1.1	Potential Savings from Perfect Forecasting	65
7.1.2	Stochastic Savings	65
7.2	Scenario Creation Method Effectiveness	70
7.2.1	Analog Limitations	70
7.2.2	Moment Matching Limitations	73
7.3	Scenario Metrics	77
7.3.1	Scenario Metric Distributions	77
7.3.2	Scenario Metrics with N_s	77
7.3.3	Scenario Metrics Explain Performance	79
7.4	Effect of Wind Shedding Cost	85
7.5	Computational Time	85
7.6	Potential Effects of Power System Modeling Assumptions	87
Part V:	Conclusions	90
Part VI:	Bibliography and Appendices	93
Appendix A:	List of Acronyms	103
Appendix B:	UC Formulation and Implementation	106
B.1	Day-Ahead Deterministic UC Formulation	106
B.1.1	Generator Limits	108
B.1.2	Curve linearization	110
B.1.3	Startup and shutdown costs	111
B.1.4	Initial and final constraints	111
B.1.5	Non-controllable generation	112
B.1.6	Wind Curtailment	112
B.2	The Real-Time Dispatch Problem Formulation	113
B.3	Handling Infeasible Cases	113
B.4	Testing	114
B.4.1	Constraints	114
B.4.2	Objective Components	115
B.5	Day-Ahead Stochastic UC Formulation	115

Appendix C:	Thermal generation modeling	117
C.1	Non-controllable generation	117
C.2	Modeled Parameters	118
C.2.1	Power limits	118
C.2.2	Ramp rate limits	118
C.2.3	Up and down time limits	118
C.2.4	Heat rate curves	118
C.3	Estimated Parameters	119
Appendix D:	Hydro System Modeling	121
D.1	Objectives	121
D.2	Constraints	122
D.3	Power Exports and Prices	123
D.3.1	Mid Columbia Price Model	123
D.4	Test Cases	123
D.5	Model Performance	124
Appendix E:	Short-Term Hydro Scheduling Formulation	126
E.1	Network model	128
E.2	Piecewise Linear (PWL) models	129
E.3	Hydro-thermal model integration	129
E.3.1	Decision Making Stages and Horizons	129

ACKNOWLEDGMENTS

Sincere thanks to all of those who have helped in the creation process, directly and indirectly.

Thank you to my committee: to Prof. Ortega-Vazquez and Prof. Kirschen for guidance and excellent suggestions, to Prof. Kirschen for the use of Hyak (without which the experimental portion of this work would not have been possible to its current extent), to Prof. Kramlich for kindly agreeing to be part of the committee, and to Prof. Christie for being an excellent advisor.

I'm grateful to have been part of the community of graduate students at the University of Washington. Thanks especially to Karen Studarus for being a sounding board, a sharp-eyed reader, and an excellent source of good ideas. Thanks also to Bryan Palmintier for bringing a new, deep perspective on power systems planning and lots of energy. I'm also lucky to have gotten encouragement and inspiration from the many mad scientist and mathematician friends I've known at UW.

This code for this research rests on the shoulders of many open source project giants. Thanks to the teams behind Python, Coopr, Pandas, GitHub, COIN OR, and many others.

This research was funded by in part by the Grainger Fellowship and the Hydro Research Fellowship in conjunction with the US Department of Energy. I'm grateful for your support.

I've learned a lot from conversations and work with the teams at Alstom Grid, Sandia National Labs, and 3Tier. Thanks to Xing Wang and Qianli Su at Alstom for tutoring me in the world of power systems software. Thanks to Jean-Paul Watson for the opportunity to visit the Coopr team at Sandia. Thanks also to all the folks at 3Tier - you made it fun to show up to work and taught me so much.

Finally thanks to my friends and family, for support, perspective, adventures, and love. Especially Amelia, I love you and can not thank you enough.

Part I

INTRODUCTION

One of the greatest challenges humanity faces today is minimizing energy production's environmental cost while generating the increasing amounts of energy that society requires. Wind energy is currently the low hanging fruit of renewable generation, representing the highest yield for the least cost. Over the past decade wind energy has become a significant part of the generation mix and today the sheer amount of wind energy coming online is forcing a re-imagination and re-invention of the power system. A major challenge is the scheduling of energy for a system with significant wind. Each day, operators at utilities and Independent System Operators (ISOs) schedule their generators to match the next day's predicted demand plus some margin for reliability, while trying to minimize cost. This process is called a Unit Commitment (UC). When hydroelectric projects provide energy to the power system there is an additional complex set of components to schedule. Day ahead wind forecast errors for today's power systems can be on the order of 1000 MW, and these errors will grow as wind penetrations increase. This wind uncertainty results in annual costs on the order of millions of dollars for a power system.

The Challenge of Uncertainty

Uncertainty is a relatively new problem in the power system, which was designed to operate with controllable generation and predictable demand. The system is scheduled on a day to day basis by the UC — a mathematical optimization problem where the objective is to meet the demand for power at lowest cost. Power demand as well as generation availability, costs, and limits are assumed to be known ahead of time. Cheaper generation is typically larger and slower to respond and often requires several hours notice to start up or shut down. Smaller and faster but more expensive natural gas “peakers” are used only when necessary. Hydro units can provide power at negligible operating cost and plants with reservoirs have some ability to store energy in the form of water. However, non-power constraints (e.g., fish, navigation, irrigation) often govern much of the hydro system's operation. Additionally hydropower systems are energy constrained – the amount of power that can be generated over a period of time is limited by the amount of water in the reservoir, most of which is allocated for the future.

Wind is currently represented in the UC using a single forecast. However, wind energy is difficult to forecast 24 h out, resulting in the addition of a great deal of uncertainty to the scheduling process. The day ahead point forecast error for wind in systems with significant penetrations can be on the order of a large conventional generator (i.e., 1000 MW)¹. Uncertainty at this scale calls for a level of preparedness and analysis which is not matched by current practice. Deterministic wind forecasts which ignore uncertainty are the industry standard for UC. Some operators struggling to deal with the uncertainties in tomorrow’s schedule will run the UC several times with different projections for load and wind, and then choose one commitment pattern that avoids likely problems [3]. This process highlights the operational need for incorporating more information about uncertainty into the scheduling process.

The result of our current deterministic approach to scheduling has been more expensive system operations than expected. Hydro units and peakers are regularly used to make up for wind forecast error and expensive emergency measures (e.g., deviations from reservoir targets and spill of water and wind) are taken more frequently. Wind also introduces significant short-term variations in the price of power and thus the value of water. The resulting non-optimal scheduling decisions are very expensive for power systems, adding up to millions of dollars annually. The joint short-term scheduling of hydro and wind is an unmet operational need in hydro systems like the Bonneville Power Administration (BPA), which had a 18.8% wind energy penetration in 2012. This issue will increase in severity as more wind is added to the system.

Potential Solutions

Power engineers have responded to these challenges by re-evaluating the assumptions on which the system is operated today, specifically by examining how uncertainty is handled in the system. There are several options currently under consideration [4] – continue to use reserve to offset uncertainty, build energy storage on a large scale, use demand response to reduce costs, hedge losses with financial instruments, dramatically improve forecasting, or schedule the system stochastically.

Reserve is the system’s current method for dealing with uncertainty by keeping significant extra fast-moving generating power available on short notice (usually in the form of peakers or hydro units). Many pragmatic system operators have simply “beefed up” their system’s reserve requirements in the face of wind. These operators state that the system simply needs to carry additional reserves proportional to the wind power being produced or the wind error expected (e.g., normal reserve levels plus 25% of current wind power output). This solution generally keeps the system running in

¹The ERCOT ISO has an average day ahead forecast error of about 10% of installed wind capacity [1]. ERCOT wind capacity is 9838 MW as of early 2012 [2].

all but the worst cases, but is extremely expensive. The trade off between reliability and expense is a classic compromise made in power systems. However, in a system with a large amount of wind generation, requiring additional reserves to match 25% of wind power is prohibitively expensive.

Energy storage would allow the power system to store wind energy and then later release that energy to match periods of high demand. But, the capital costs of energy storage technologies remain high, especially for storage with large power outputs and long time constants. Demand response is a solution which envisions electricity consumers reacting to the real time price of power, shaping the demand to the availability of low cost supply. This solution would require both large societal changes and infrastructure. Financial instruments to hedge against the cost of wind uncertainty are currently used in the more structured energy markets in the United States. This allows market participants to reduce their personal financial risk, but does not eliminate the overall cost. Better forecasting is currently a scientific and business endeavor for organizations including European Centre for Medium-Range Weather Forecasts (ECMWF), National Oceanic and Atmospheric Administration (NOAA), AWS Truepower, and 3Tier. However weather forecasts will never be perfect, especially at the day ahead time frame required for power systems scheduling. This leaves the option of stochastically scheduling the system.

Many researchers argue that stochastic optimization is a better method for handling large amounts of wind, i.e., one that results in cheaper and more reliable system operations. A stochastic UC formulation takes into account many possible scenarios for wind power and finds the optimal solution (e.g., lowest cost generator schedule) for all of those scenarios, weighted by their likelihood. Furthermore, stochastic scheduling is primarily a software solution and does not have the capital costs associated with storage or the operating costs of higher reserve. In 2009, Tuohy, Meibom, Denny, *et al.* [5] tested stochastic UC on a wind integration study of the Irish system with 34% wind energy penetration and reported reductions in cost between 0.25–0.9% of the deterministic UC solution. For large power systems, like the BPA which has annual revenues of around \$3 billion, a 1% savings would be considerable.

Challenges in Stochastic Scheduling

However, stochastic UC has yet to be used in a real control center [6]. There are three main challenges – unknown magnitude of savings, lack of a method for creating scenarios, and computational speed.

Magnitude of Savings The power systems research literature on stochastic UC for wind often is based on simple test systems and often fails to report stochastic savings relative to deterministic UC costs. There is no literature quantifying how stochastic savings may differ between systems, or

how they may vary as the system’s wind penetration increases.

Scenario Creation Methods There is an acute lack of an accepted method for creating scenarios in the weather forecasting, power systems, and operations research literatures. A scenario creation method must be able accurately represent the continuous distribution of wind uncertainty in a reasonable number of scenarios which are informative to the decision making process.

Computational Issues Long solution times, a well known characteristic of stochastic problems stemming from the increase in the number of variables needed to represent multiple scenarios, are a major concern. Utilities and ISOs require an optimal UC solution within a few hours, a requirement which is currently met for deterministic UC thanks to advances in computing hardware and optimization software. For a deterministic mixed integer programming (MIP) problem with V variables, stochastic optimization means that the problem size will be multiplied by the number of scenarios, becoming $N_s \cdot V$. As problem size increases, MIP problems are generally characterized by a non-linear increase in solution time. Stochastic optimization problems do not lend themselves to parallelization, although there is active research in this area on a technique known as Progressive Hedging (PH).

Underlying this concern about solution times is another inherent trade-off, this one from the stochastic optimization domain: lower cost solutions via more scenarios must be weighed against longer solution times. The research literature provides little guidance on the choice of N_s , often a study will use only value.

Methods and Experiments

This work compares two methods for scenario creation – the *analogs* method and the *moment matching* method. The stochastic scheduling costs of these methods are compared for a large scale test system based on the Electric Reliability Council Of Texas (ERCOT) ISO.

The *analogs* method creates scenarios based directly on historical data. The basic concept is also used in similar day forecasting, where a forecaster takes the current forecast data (e.g., temperature, pressure, etc.) and finds similar forecasts, or analogs, from the past. Then the observed values (e.g., wind power) from the most similar date are used as a prediction for the next day. When creating scenarios with the analogs method, every date in the historical archive represents a possible scenario. In this work only *simple analogs* based on aggregate power are tested, but an extension to analogs of weather data is discussed. The analogs method is compared with the *moment matching* method [7–9], which synthesizes scenarios to match a set of desired statistics.

Experiments were conducted on a test system representing ERCOT using modeled parameters and actual system data from 2012. Results are based on UC simulations running nine months in length. Stochastic UC costs are compared to the cost of a deterministic UC with the public aggregate forecast. Different numbers of scenarios are tested to examine the relationship between stochastic savings and computational costs. Different wind energy penetrations were tested to evaluate the performance.

Contributions

This dissertation contributes to our knowledge of the intersection between power systems operations, renewables forecasting, and stochastic optimization. We propose the *simple* analogs method, a scenario creation method based on historical aggregate power data. We also propose a method for estimating the desired moments, a key component of the moment matching method. In our experimental design, we lay out a framework for testing stochastic UC relative to both deterministic and perfect forecasts. We evaluate stochastic UC performance on a large scale, realistic power system model. Finally, we compare performance for two scenario creation methods, several values for the number of scenarios N_s , and several wind penetrations.

A Reader's Guide

The current state of the art for stochastic UC and scenario creation is reviewed in Part II. We detail our methods for creating scenarios and solving stochastic unit commitments in Part III. A set of experiments to test these methods are described in Chapter 6, with a detailed discussion² of the results in Chapter 7. A summary of the results, along with remaining research questions, is given in Part V. The appendices provide details on the problem formulations and methods used to model power systems.

²A digital copy of this dissertation, complete with color figures is available at <http://adamgreenhall.com/research/dissertation.pdf>

Part II

LITERATURE REVIEW

Stochastic unit commitment has been discussed in the power systems and operations research literature since the mid 1990s. The first papers on the subject in 1996 [10, 11] were interested in load forecast errors and the modeling of generator outages, a low-probability, binary random phenomenon that is quite different from the continuous, time-varying stochastic models of wind. The first major work on stochastic unit commitment with wind was the Wind Power Integration in Liberalized Electricity Markets (WILMAR) project [12] in 2006. The first truly modern formulation with a realistic power system was [5], which tested stochastic UC on a wind integration study of the Irish system with 34% wind energy penetration. Recent work has studied the effects of wind penetration levels [13], incorporated the CVaR risk metric [14], focused on curtailment [15], and attempted to incorporate NWP models [16].

All of these works that include wind share a common structure. Wind is modeled as a stochastic process and represented by a collection of discrete scenarios. Each scenario is a potential realization of the stochastic process through the entire time horizon (i.e., a time series) and has a given probability. A stochastic unit commitment which incorporates some subset of these scenarios is formulated as a discrete optimization problem and solved using an optimization solver.

The hydro scheduling literature has also been a significant source of large-scale stochastic optimization research. In most of these problems, hydro inflows are modeled as stochastic processes and decisions are made over longer time scales, with horizons from one week to a year. Hydro scheduling problems have different decision making structures and thus often employ different solution algorithms than those found in thermal stochastic UCs.

Scenario creation methods for these problems have varied widely. The most used method has been a sampling or Monte Carlo technique [5, 13, 17] based on a forecast model which includes an estimate of variance and assumes Gaussian forecast error. Other methods include drawing scenarios directly from a NWP model [16] and moment matching [8, 9].

There have also been significant contributions from areas outside the power systems literature. The statistics and atmospheric sciences communities have done significant work on modeling and

forecasting wind. On the other end of the process, the operations research literature has been developing models, algorithms, and solution tools for large stochastic optimization problems. The next few sections detail the best practices in the literature to date.

Chapter 1

STOCHASTIC ENERGY SCHEDULING

The core of the stochastic UC is driven by the functional requirements of the classic (i.e., deterministic) UC problem. In the classical thermal UC problem a utility needs a schedule of optimal generator statuses for the next operating period, e.g., 24 h. The UC result must be available well in advance, e.g., at noon the day before, as some thermal units require several hours of notice to change status and UC solution algorithms have significant run times. Most modern UC solution algorithms use mixed integer programming (MIP) formulations; there are many well specified MIP formulations of the deterministic UC problem available in the literature, including [18–20].

Hydroelectric systems add additional complexity to the problem, with additional constraints, different time scales, and the uncertainties of a water system. While traditional thermal systems required little scheduling beyond the day ahead and were relatively deterministic in nature, hydro systems are generally scheduled on several timescales, the longest of which have horizons of a few years. These long timescales require coverage of uncertainty, and much of the early stochastic power systems literature was concerned with long term hydro scheduling with uncertain inflows. Short term hydro scheduling is generally carried out on with horizons of a few days. §1.1.2 reviews the short term hydro scheduling problem.

Stochastic UC has a long history in the power systems and operations research literatures. We begin by examining the traditional deterministic UC, which is the core of any stochastic UC. §1.2 reviews the decision making structures which govern the formulation of stochastic problems. Different approaches to creating a stochastic objective function are covered in §1.3. Three stochastic optimization solution techniques are covered in §1.4. The chapter concludes with a summary of stochastic UC results seen in the literature and the power systems on which these results are based.

1.1 Deterministic UC

1.1.1 Deterministic Thermal UC

While early stochastic UC papers wrangled over the details of various solution methods, more recent papers are in agreement in their use of a mixed integer formulation and a standard mixed integer programming (MIP) solver. MIP algorithms and software advanced dramatically during the 1990s;

solution times for large MIP problems using CPLEX fell by a factor of up to 350 between 1988–1999 [21]. In 2006, Carrión and Arroyo [18] compared solution times and objective values for the solution of the same UC problem using several methods, including Lagrange Relaxation (LR), MIP, and several heuristic methods. In their study, the MIP method results in the lowest objective values and fastest solution times. While the level of detail covered varies, formulations are converging toward forms similar to those in [18–20].

The objective function for thermal UC is to minimize the sum of generator costs¹, including both the costs of starting up ($C_{\text{su}(g)}[t]$), the costs of shutting down ($C_{\text{sd}(g)}[t]$), and the operating costs of expending fuel ($C_{\text{op}(g)}[t]$).

$$\min \sum_t \sum_g C_{\text{op}(g)}[t] + C_{\text{su}(g)}[t] + C_{\text{sd}(g)}[t] \quad (1.1)$$

Operating costs are modeled as a polynomial function of power ($P_g[t]$) and are generally linearized (a simple process due to the convex nature of many of these curves).

$$C_{\text{op}(g)}[t] \approx a + b \cdot P_g[t] + c \cdot P_g^2[t] + d \cdot P_g^3[t] \quad (1.2)$$

Operating reserves required for deterministic UC in a system with uncertain renewables has been the subject of much literature and technical study. Several papers propose using stochastic security constrained UC including an Expected Energy Not Served (EENS) constraint to avoid fixed reserve requirements [17, 22, 23]. However, this approach is generally very computationally intensive for large systems. Matos and Bessa [24] propose a probabilistic forecast based method for computing reserve. Ortega-Vazquez and Kirschen [25] propose an iterative Monte Carlo based method for optimizing operating and reserve costs of an uncertain system. In current operating practice, “reserve requirements with moderate amounts of renewables (10-15% of energy) are comparable to those already maintained by the power system for contingencies and load uncertainty [26]”.

The horizon effect is a concern for any optimization problem in which future operations continue beyond the scope of the problem. In the case of UC, the commitment at the end of the operating period must not result in an infeasible problem for the next operating period. This is generally accounted for by extending the UC horizon beyond the operating period (e.g., a two day commitment

¹In a market centered framework, this is equivalent to the maximization of profits, assuming generators bid their incremental cost.

horizon for a day-long operating period) and ignoring the commitment pattern beyond the operating period.

Formulation efficiency for MIP UCs has been the subject of some recent debate. Early UC MIP formulations [27, 28] used three variables to represent generator status, start up, and shutdown. Carrión and Arroyo [18] point out that it is possible to reduce this to just a single variable representing status, with more complicated constraints to represent startup and shutdown limits. With fewer binary variables, the MIP solver is presented with a smaller problem, which should intuitively be easier to solve. However, Ostrowski, Anjos, and Vannelli [19] and Morales-Espana, Latorre, and Ramos [20] have recently and convincingly argued that a carefully crafted three variable formulation can actually produce a “tighter” and “more compact” MIP formulation which is actually faster to solve than the one status variable formulation.

1.1.2 Hydro scheduling

Hydrothermal systems are generally scheduled on three time frames [29–32]:

long-term several years horizon, considers aggregate hydro system and aggregate thermal system, climate based stochastic modeling, solved using multi-stage SDDP

mid-term about one year horizon, weekly or monthly intervals, inflows can be deterministic or stochastic. The results are an expected cost, shortage risk, and monthly release targets.

short-term one week horizon, hourly intervals, inflows are deterministic (an example hydro forecast is shown in Figure 1.1), power system constraints and water delays are considered

In this section we will focus on the short-term problem, as wind and thermal plants interact with the hydro system on this time scale. We will also describe the unique structure of the short-term hydro scheduling problem: a networked structure and non-convex production curves.

In 1978, [33] presented an optimization model of the BPA hydro system. The day-ahead schedule is solved using non-linear programming with soft constraints (penalty functions). The model is deterministic and does not include transmission constraints or a thermal system. [34] propose a Benders’ decomposition method for the hydro-thermal problem. [35] present a hydro-thermal model of the Swedish power system that includes zonal transmission constraints. The problem is decomposed into the hydro problem (solved by a specialized gradient method) and the thermal problem (which is reduced to an approximate aggregate cost function based on hydro production). [36] proposes a

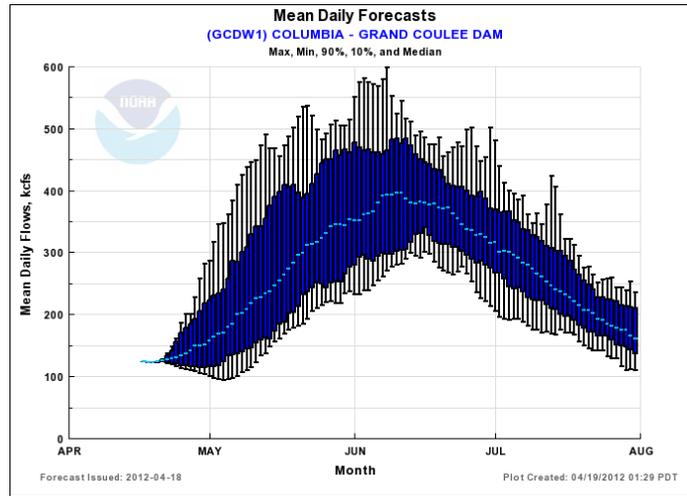


Figure 1.1: Mean daily flow forecast for the Columbia River at Grand Coulee, from NOAA. Note that the confidence intervals are very small for the first several days of the forecast.

method for calculating reserve (primary, secondary, and tertiary) for hydro-thermal-wind systems based on the predicted forecast error. However, normal wind error (see §2.2) is assumed and the proposed method is simply used to calculate reserve requirements, rather than to schedule the system. [37] formulate the hydrothermal problem using MIP and a mixed integer simplification of the non-convex hydro production curves for plants with multiple units. [38] propose a joint optimization strategy for wind and a hydro facility with storage. However, normal error is assumed and the focus is on a single facility bidding into a larger market.

In general the hydro literature which does include wind has used very simple “normal error” forecasting methods and/or focused on joint bidding of wind with storage. Stochastic scheduling of hydro-thermal-wind systems remains untested.

Network Structure

Cascaded reservoir systems have a networked flow structure, as illustrated in Figure 1.2. The total water available over the planning period is generally set by defined initial and final reservoir targets (levels) and predicted natural inflows. The short term hydro scheduling process defines the volume $v[t]$, outflow $o[t]$, and spill $s[t]$ for each time period within the planning period.

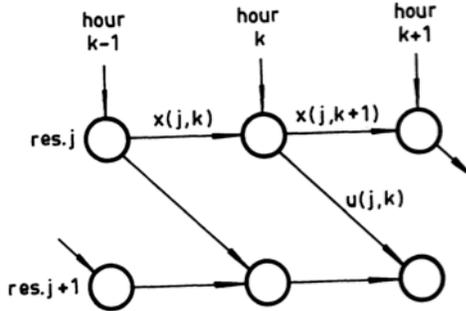


Figure 1.2: Reservoir models have a networked structure, from [34]

Production characterization

A hydro plant (indexed by i) has a non-linear power production curve that is dependent on the outflow $o_i[t]$ and the head $h_i[t]$:

$$P_i[t] = f_i(h_i[t], o_i[t])$$

Thus, the modeling of hydro production usually involves simplification and linearization. Several example curves are shown in Figure 1.3.

Brannlund, Bubenko, Sjelvgren, *et al.* [35] proposes a production function which is a polynomial function \tilde{f} of flow through the turbine plus a correction term z for head variation:

$$P_i[t] = f_i(h_i[t], o_i[t]) \approx \tilde{f}_i(o_i[t]) + z(h_i[t], o_i[t])$$

However this equation is not compatible with most MIP formulations as the correction term involves the product of two variables (a quadratic term which is generally not positive semi-definite). Many MIP formulations [34, 39] have simply used a fixed head curve and assumed that the short term effects of this correction are small.

Garcia-Gonzalez, Parrilla, Barquin, *et al.* [40] present an iterative method for finding the correct head. In each iteration, a simplified MIP problem with fixed head is solved, and then the values for head are updated based on the solution. The authors also give an excellent review of techniques for production curve linearization.

Borghetti, D'Ambrosio, Lodi, *et al.* [41] propose modeling production as a PWL function involving both the head and outflow terms, with a general formulation laid out in [42]. The basic idea of the linearization over two input variables is shown in Figure 1.4.

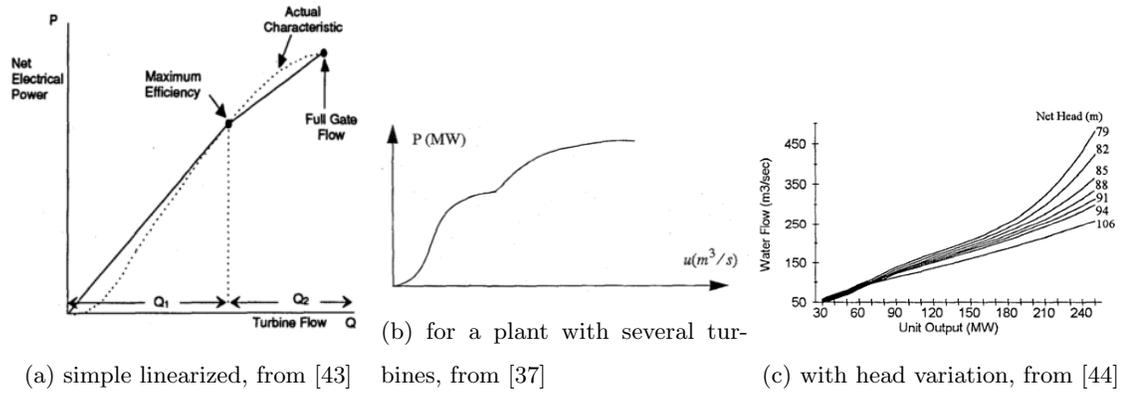


Figure 1.3: Hydro power production curve are a function of head and outflow. Several models for linearization exist.

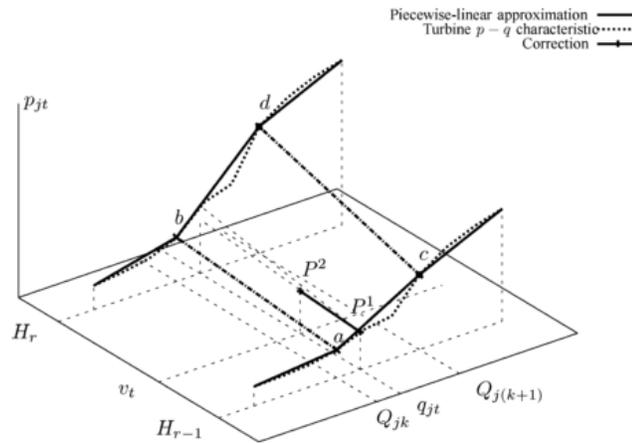


Figure 1.4: Two variable piecewise linear approximation of production curve, from Borghetti, D'Ambrosio, Lodi, *et al.* [41]. Piecewise linear segments of a flow Q to power P curve are defined for several head H values. Interpolation is used to get power values for head values other than those defined.

1.2 Decision Making Structure

Adding stochastic wind to the deterministic UC problem usually is done with a simple two stage stochastic formulation. The stochastic optimization literature refers to these stages as “here and now” and “wait and see” decisions [45]. In the first stage, i.e., when the UC is run, a single commitment decision is made for the set of generators over the schedule period. The exact amount of wind power is uncertain, as represented by the different wind scenarios. The second stage will be the whole operations period. During this operations period, the wind power, and thus the generator powers and the total system cost, will be realized. The logic of this decision making structure, where a single commitment pattern is required in the first stage, must be enforced explicitly in the formulation by “non-anticipatory” constraints. For the UC problem, these constraints ensure that the status variables are equal over all the scenarios.

1.2.1 Rolling Planning Scenario Trees

The “rolling planning” formulation proposed by the WILMAR tool [12] is a notable extension of the two stage stochastic UC formulation. Rolling planning is an attempt to account for the fact that wind prediction is not very accurate on a long time horizon. Additionally, new pieces of wind information (i.e., wind power measurements and better forecasts than the original day ahead) arrive on a continuous basis and can be incorporated into the decision making. The first intuition is that UCs should be run more frequently, every 3 h in the WILMAR case, to incorporate the new information (see the vertical axis of Figure 1.5). The second intuition is that uncertainty increases with the forecast horizon and therefore more uncertain periods should be modeled with more scenarios. The three stages of the “rolling planning” formulation and their corresponding number of scenarios are specified in Table 1.1. Barth, Brand, Meibom, *et al.* [12] note that a 24 stage (i.e., a day long operating period with hourly intervals) formulation would be an ideal representation (assuming the system was rescheduled every hour), but note that such a model quickly becomes computationally intractable as the total number of scenarios increases exponentially with the number of stages².

²If a UC is modeled with 24 stages and 100 wind scenarios per stage, the result would be 100^{23} total scenarios.

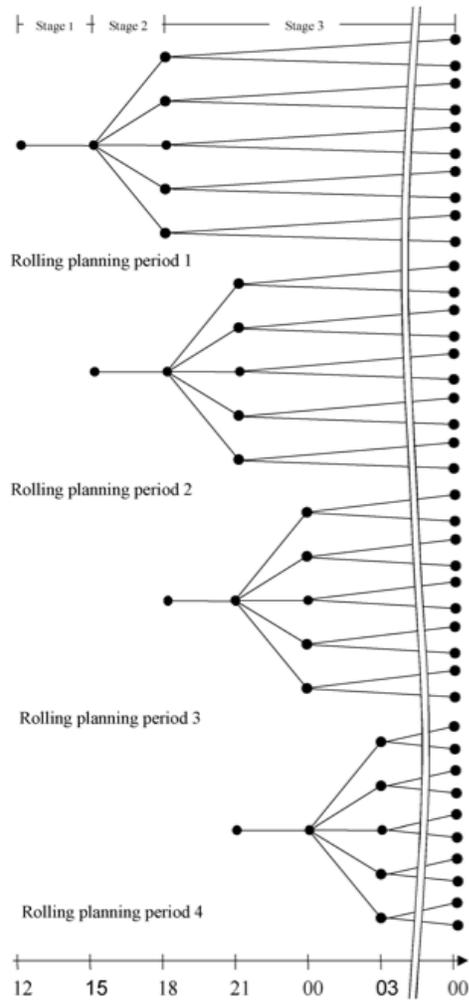


Figure 1.5: “Rolling planning” scenario tree structure proposed by the WILMAR project [12]

Stage	Length	Scenarios
First	3 h	1 (deterministic)
Second	3 h	5
Third	to end of period	10

Table 1.1: The WILMAR “rolling planning” formulation [12] has three stages for each UC run.

1.3 Stochastic Objective Functions

The classic stochastic optimization objective function is to minimize the expected cost over all scenarios:

$$E[C(x)] = \sum_s \rho_s \cdot C(x_s) \quad (1.3)$$

where $C(x_s)$ is the scenario cost and ρ_s is the scenario probability.

1.3.1 CVaR

Another objective formulation minimizes a measure of risk, such as value at risk or Conditional Value at Risk (CVaR) [4, 14, 46, 47]. This is equivalent to minimizing the expected cost under a region of the output distribution tails in Figure 1.6.

Formulating a CVaR objective involves adding a risk-adjusting term to the expected value. The variables involved are summarized in Table 1.2. The CVaR term has with two parameters: the weighting term β and the risk confidence level α . The CVaR weighting term β can take on values between $[0, 1)$ and determines the degree of risk averseness, with $\beta = 0$ being risk-neutral. An approximate way of thinking about the objective is to consider minimizing the expected cost over only the $(1 - \alpha) \cdot 100\%$ most expensive scenarios. The objective function becomes:

$$E[C(x)] + \beta \left(\zeta - \frac{1}{1 - \alpha} \sum_s \rho_s \eta_s \right) \quad (1.4)$$

and there are additional constraints:

$$\text{s.t. } \eta_s \geq 0 \quad \forall s \in S \quad (1.5)$$

$$\text{s.t. } \eta_s - C(x_s) = \zeta \quad \forall s \in S \quad (1.6)$$

where η_s (a per-scenario variable) can be thought of as an excess term which is constrained by the per-scenario cost $C(x_s)$ and the single variable ζ , in order to produce an approximation of the area under the tail of the distribution. Note that a large number of scenarios are required to accurately evaluate the costs at the tail of the distribution.

1.3.2 Robust optimization

A third objective formulation is to minimize the maximum possible cost over all of the scenarios, which is known as minimax or robust optimization. Robust optimization is the most risk-averse of the stochastic optimization objectives and thus results in a higher expected value for the objective function.

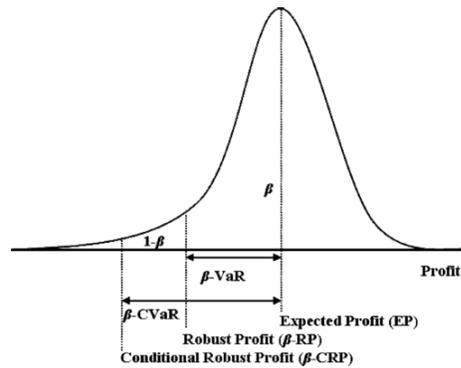


Figure 1.6: Conditional Value at Risk (CVaR) is represented as an area under the tail of an output distribution. In this case the problem is a maximization of profit, but the curve is equivalent for minimization of cost.

Variable	Description
α	risk confidence level
β	weighting term for degree of risk
$C(x_s)$	total cost of scenario s
ρ_s	scenario probability
η_s	scenario excess term
ζ	tail area approximation term

Table 1.2: CVaR objective function formulation variables and descriptions

1.4 Solution Methods

A majority of the more recent papers have solved the extensive form (EF) of the stochastic UC problem. In an EF problem, each variables and constraint in the problem is replicated per scenario and the entire problem is passed to the MIP solver as a whole. Other techniques exist for solving large stochastic optimization problems iteratively, including Benders decomposition, Progressive Hedging (PH), and Stochastic Dual Dynamic Programming (SDDP). These methods sample from the set of scenarios to create subsets, solve the subsets (often in parallel), and then make adjustments to the global problem based on the result. These methods are often used when optimization problems are computationally infeasible due to their size. A summary of the methods used in the stochastic UC with wind literature can be found in Table 1.3.

Cerisola, Baillo, Fernandez-Lopez, *et al.* [48] study the effectiveness of several algorithms applied to the thermal UC problem. LR, MIP, and two forms of Benders decomposition are compared. The Benders decomposition methods arrive at a similar solution in 25 minutes compared to the to a MIP solver runtime of 2 h. Cerisola *et al.* note that the drawback is that these decomposition methods are very difficult to implement.

1.4.1 Progressive Hedging (PH)

The PH algorithm was introduced by [49]. The algorithm decomposes a large stochastic problem into scenarios, allowing these subproblems to be solved in parallel. Woodruff, Watson, and Hart [50] claim that despite convergence not being guaranteed for mixed integer problems, PH has provided good quality solutions to problems from many disciplines.

PH was used in its Lagrange relaxation form by [11], one of the first stochastic UC papers. The algorithm has seen no subsequent use in the power systems UC literature. The PH algorithm has a MIP implementation in the Python Stochastic Programming (PySP) [50] software package, although tuning the parameters to a given problem is a non-trivial task.

1.4.2 Stochastic Dual Dynamic Programming (SDDP)

The SDDP algorithm originated with [51] and has been used widely by the hydro scheduling community for multi-stage problems. The algorithm consists of a forward and backward solution procedure, with the key advantage being a sampling of scenarios in the forward stage. The result is a solution time that is proportional to the number of scenarios sampled instead of the full number of scenarios (the full number of scenarios is exponential with number of stages). Thus, SDDP is used almost exclusively in problems with many stages.

1.5 Results

Performance metrics for stochastic optimization are lacking in power systems research. In most papers a single (small) number of scenarios is arbitrarily chosen and the resulting expected value of the objective function (e.g., total cost) is reported, sometimes with the time taken to solve the problem. In the power systems literature, only [52] reported testing more than one number of scenarios³. Few authors report the cost of the deterministic UC solution or the cost of the perfect solution. This makes it difficult to assess the potential savings that stochastic UC might provide for the system under consideration. A summary of the results reported in the stochastic UC with wind literature can be found in Table 1.4.

³The authors compared 4, 10, and 16 scenario solutions. However, scenarios – for load in this paper – are picked by assertion and cannot be said to represent the underlying process with any certainty.

Table 1.3: Comparison of methods for stochastic unit commitment with wind

Work	Wind Forecast Method	Scenario Creation	Solution Method
Takriti, Birge, and Long [11]		for load, by assertion	LR with PH
Castronuovo and Lopes [38]	normal error	Monte Carlo (MC) $N_s = 150$	
Barth, Brand, Meibom, <i>et al.</i> [12]	WILMAR ^a	MC $N_r = 10$	not reported
Wang, Shahidehpour, and Li [17]	normal error	MC $N_s = 3000$, $N_r = 10$	MIP
Tuohy, Meibom, Denny, <i>et al.</i> [5]	WILMAR ^a	MC	MIP CPLEX
Al-Awami and El-Sharkawi [14]	adjusted normal error	own method $N_s = 4000$, $N_r = 5$	MIP CPLEX
Constantinescu, Zavala, Rocklin, <i>et al.</i> [16]	NWP	own method $N_s = 100$, $N_r = 30$	MIP CBC
Lowery and O'Malley [9]	WILMAR ^a	moment matching, up to $N_s = 6$	MIP CPLEX

^a The WILMAR model uses ARMA(1,1) models of wind speeds at individual sites. A power curve is applied to these speeds and geographic correlation is adjusted for using a Gaussian covariance matrix. MC simulation is used to create a large, ad hoc number of scenarios. Scenario reduction is done using the method of [53] and $N_r = 10$.

Table 1.4: Comparison of results for work on stochastic unit commitment with wind

Work	System	Stochastic Savings*
Takriti, Birge, and Long [11]	Michigan system ¹	0.57% ^a
Castronuovo and Lopes [38]	wind-hydro plant profit maximization	not reported ^b
Barth, Brand, Meibom, <i>et al.</i> [12]	Nordpool model ²	not reported ^c
Wang, Shahidehpour, and Li [17]	IEEE 118 bus test system ³	not reported ^d
Tuohy, Meibom, Denny, <i>et al.</i> [5]	Irish 2020 System ⁴	0.25–0.9% ^e
Al-Awami and El-Sharkawi [14]	5 unit coordinated bidding	not reported ^f
Constantinescu, Zavala, Rocklin, <i>et al.</i> [16]	10 unit system ⁵	not reported ^g
Lowery and O’Malley [9]	Irish 2020 System ⁴	not reported ^h

* Stochastic savings are defined as the relative cost difference between the stochastic and deterministic UC solution costs, expressed as a percentage. For example, if the deterministic cost is \$100 and the stochastic cost is \$90, then stochastic savings would be 10%.

¹ The Michigan system was modeled with data for about 100 thermal units and six pumped hydro units. Historical load profiles from Detroit Edison were used to create scenarios.

² The model of Nordpool includes Germany, Denmark, Sweden, Norway and Finland. Units are aggregated into similar bundles, for a total of 350 aggregated units in the system. Wind is scaled for three cases to arrive at annual wind energy penetrations of about 6%, 12%, and 16%.

³ The modified IEEE 118 bus system has 76 units and is available on the author’s website. Instantaneous peak wind penetration is only 2.7%.

⁴ The Irish 2020 system is a wind integration study where the wind energy penetration is 34%. Thermal generation is modeled as a fairly inflexible mix, composed of 45 units totaling 8300 MW.

⁵ The 10 unit thermal system is based on [18]. Wind speed data from Illinois was scaled to create a 20% wind penetration test case.

^a See [11, Table III].

^b [38] calculate only the savings from a joint hydro-wind plant compared to an only wind plant. No deterministic solution is mentioned.

^c [12] evaluate the cost difference between a base case (6% wind energy penetration) and 12%, 16% cases. Deterministic cost is not mentioned.

^d [17] evaluate the cost difference between the system without any wind and with stochastically scheduled wind.

^e [5] test several rolling commitment strategies. Stochastic savings were calculated as low as 0.25% with 1 h rolling, up to 0.9% with 3 h rolling. Perfect savings were about 1.85%.

^f [14] compare uncoordinated thermal bidding (assuming no wind) and coordinated wind-thermal bidding.

^g [16] report cost information in evaluating the NWP sampling procedure, but does not report a deterministic cost.

^h [9] compare differences in stochastic solution costs for different moment selections. However, no comparisons are made to deterministic cost or to the results of the previous work of [5] on the same system.

Chapter 2

WIND MODELING AND FORECASTING

An accurate model of the underlying process is a critical for many scientific fields and the study of wind power production is no different. The modeling and forecast of wind has been the subject of much study by researchers at Argonne National Laboratories [16, 54, 55], who conclude that while “the main focus in the [wind power forecasting] industry so far has been on improving the accuracy of point forecasts . . . [but] probabilistic forecasts will become increasingly important in operational decisions” and that the “generation of representative wind power scenarios” will also be of importance. Also of special concern to power system operators are the ramp rate and duration of the ramp, as large decreases in wind power which happen very quickly can be especially difficult for systems to handle [56, 57].

Today, North American ISOs generally use some form of Numerical Weather Prediction (NWP) modeling to forecast wind power at longer lead times. Methods vary widely for forecasting at shorter lead times. The following sections describe the three main branches of wind power modeling found in the stochastic UC literature: persistence, Auto Regressive Moving Average (ARMA) based, and NWP. An additional section reviews the technique of analogs forecasting based on historical data. We will leave aside discussion of the large field of heuristic and machine learning methods for very short horizon forecasting, following the conclusions of the ANEMOS project [57] that the improvements these techniques provide are “usually deemed not enough to warrant the extra effort in training [the neural networks].”

Before discussing forecasting it is important to note that wind power is a difficult phenomenon to model from a statistical perspective due to the following characteristics:

non-linearity of power curve A typical wind speed to power curve is shown in Figure 2.1. Non-linearities in the mid-range and near the cut-out speed¹ are sources of significant power forecast error, even if wind speed is predicted with relative accuracy.

non-Gaussian Useful wind speeds are bounded between zero and cut-out speed. Over time, hourly average wind speed more closely represents a Weibull distribution [58]. Statistical models which traditionally assume a Gaussian distribution must apply a transformation. This applies to the ARMA models described in §2.3.

alternating atmospheric regimes Wind speed is driven by the weather, which can be thought of as having many different states - e.g., westerly wind vs. easterly wind, low pressure vs. high pressure. These states are not generally accounted for in time series models, although they can be explicitly modeled, as in [59, 60].

non-stationary Wind speed generally has strong diurnal and seasonal trends. When models that assume stationarity are used (e.g., ARMA), de-trending must be applied.

temporally and spatially autocorrelated Neighboring wind turbines and even nearby wind farms encounter wind and weather at predictable intervals. Models which assume independence between sites must be adjusted [59, 61, 62].

heteroscedastic Wind speed does not have a variance that is constant over its whole range. Higher wind speeds generally have significantly higher variance [63].

2.1 Persistence

Persistence forecasting, e.g., predicting that wind power in the the next hour will be equal to the current wind power, has been a common and surprisingly effective technique for short-term wind forecasting. This method is clearly less suitable for forecasting at longer time horizons. See Figure 2.4

¹The cut-out speed for a wind turbine is the maximum safe operating wind speed. When higher wind speeds are encountered the blades are “feathered” to be perpendicular to the wind and the machine is shut down. This can be a problematic phenomenon for the power system when many machines at a site all go from maximum power output to zero in a short period.

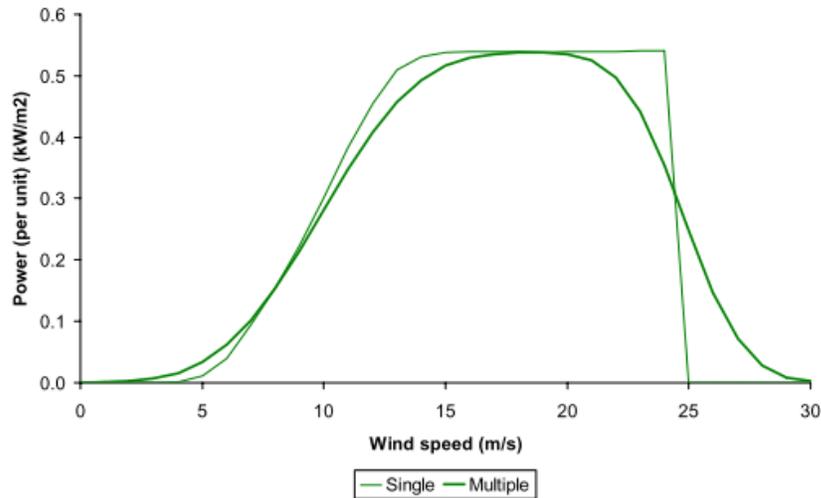


Figure 2.1: Typical wind speed to power conversion curve, from [64]. Note that the curve is non-linear.

for an example of the forecast performance of the persistence method for different forecast lead times. Persistence is generally included in academic discussions and industry software as a benchmark metric that must be outperformed by the proposed method.

2.2 Normal Error

Another widely used simple forecast method in the power systems literature has been to assume that a point forecast is available and that the forecast errors are normally distributed with zero mean and a standard deviation that increases with the forecast horizon. This “normal error” method is consistent with the power system’s existing theoretical treatment of load forecast error [36, 65], is easy to understand, and has been used by several researchers [17, 22, 36, 38]. However, given publicly available data it is also very easy to show that wind forecast error is not normally distributed [66–68]. The distribution of forecast errors for the BPA [69] and Eirgrid [70] balancing areas for the year 2010 are shown in Figure 2.2. Both forecast error distributions are clearly non-normal. [14] also pointed out that because these “normal error” methods don’t take into account inter-temporal correlation, it is easy to create scenarios where the change in power outputs (i.e., ramps) do not match reality.

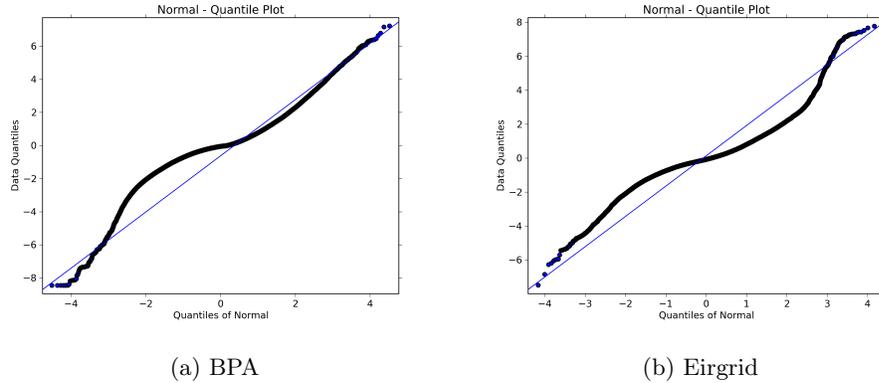


Figure 2.2: Q-Q plots of wind power forecast error show significant deviations from a Gaussian model, even when power is aggregated over an entire ISO. Q-Q plots are a method for comparing a distribution to normal; a perfectly normal distribution would fall along the diagonal line.

2.3 ARMA Based Wind Forecast

Wind speed can be modeled as a time-varying stochastic process. There is a rich statistical literature for modeling techniques for these processes, but one common technique is the Auto Regressive (AR) model [71]. An AR(p) model assumes that the future of the process X_t can be predicted solely based on a weighted sum of the past p time steps, plus a small white noise error component (ϵ_t), see Equation (2.1). The persistence forecast is in fact AR(1), the most basic class of AR models. A common extension is an ARMA(p, q), model, which adds an weighted sum of the most recent q error terms, see Equation (2.2).

$$X_t = \mu + \sum_{n=1}^p a_n (X_{t-n} - \mu) + \epsilon_t \quad (2.1)$$

$$X_t = \mu + \sum_{n=1}^p a_n (X_{t-n} - \mu) + \sum_{m=1}^q b_m (\epsilon_{t-m}) + \epsilon_t \quad (2.2)$$

One of the first papers on modeling wind power was Brown, Katz, and Murphy [58], sponsored by the BPA in 1984. The Brown et al. model is statistically quite complete – wind speed is de-trended and then transformed to take on a Gaussian distribution, forecast with an AR model, converted to power using a power curve. This set of preprocessing steps is detailed in the next section, as it is very similar to what is used in most ARMA based wind models.

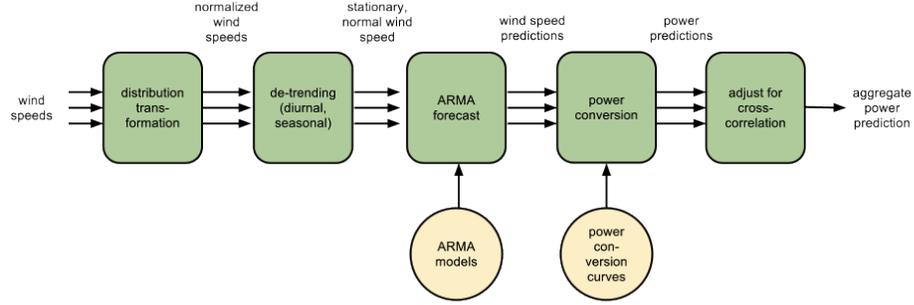


Figure 2.3: Wind power is forecast after undergoing several steps of statistical preprocessing. These steps are also used in model creation.

2.3.1 Statistical Preprocessing

There are many preprocessing steps that must be considered before wind speed data is ready to be modeled as an ARMA process. The overview of these processes for wind is shown in Figure 2.3. The following is an overview – the detail and order of these processes may vary. First each wind speed series is transformed to a more Gaussian distribution, often using a power or log ($c = 0$) transformation:

$$\hat{w}(t) = (w(t))^c \quad \forall t \in T \quad (2.3)$$

Then the series is de-trended to remove the diurnal and seasonal means and normalized by each variance. Given hourly means μ_h and variances σ_h for $h \in [0, 23]$ and full-length periodic series constructed from them:

$$\mu(t) = [\mu_{h0}, \mu_{h1}, \dots, \mu_{h23}, \mu_{h0}, \mu_{h1}, \dots]$$

$$\sigma(t) = [\sigma_{h0}, \sigma_{h1}, \dots, \sigma_{h23}, \sigma_{h0}, \sigma_{h1}, \dots]$$

the hourly de-trended (standardized) series becomes:

$$\hat{w}_s(t) = (\hat{w}(t) - \mu(t)) / \sigma(t) \quad (2.4)$$

The statistics for each period h are simple to estimate, given sufficient historical data. Seasonal de-trending follows a similar pattern.

At this point the normalized, detrended wind speeds $\hat{w}_s(t)$ can be fit to ARMA models using one of many algorithms, as supplied in most statistical and signal processing packages. Subsequent

ARMA forecasts for wind speed should then be converted to power using a power curve such as Figure 2.1 or by one of the more detailed methods reviewed in [57].

2.3.2 Aggregate Modeling

Given several wind farms there are a few methods for aggregation. Several works use covariance models for geographically separated farms [59, 64, 72]. It is common for large regions using NWP models to predict a few representative farms and follow some upscaling procedure to get the aggregate power output of a large area [57]. Sturt and Stbac [73] propose a time series method intended for modeling aggregate wind power output over a large geographic area. A second order AR model is used to model the short term variation in wind power, while a Gaussian process is used to model the long term variation. However, the model is not tested in the UC context.

2.3.3 Model Order Selection

A key question in time series modeling is what model order to choose - which values of (p, q) . [74] provide a detailed review of the statistical tests available for deciding on model order and conclude that ARMA(3, 2) is a statistically valid choice for modeling the wind speed at the Canadian site in their study. The Bayesian Information Criterion (BIC) and the Akaike Information Criterion (AIC) are the two most commonly used criterion for model order selection in the statistical literature.

2.3.4 ARMA in the UC Literature

Within the UC literature several researchers have used ARMA-based models of wind. The most influential of these has been WILMAR [64], the purpose-built tool for probabilistic wind forecast and scenario generation. This tool has been used as the foundation for several stochastic UCs, including [5, 12, 13]. The tool chooses an ARMA model with order (1, 1) based on the work of [75], for synthesizing wind speed forecast errors at a single location, perhaps on the basis of its simplicity. Note that [74] had earlier proposed a method for statistically validating order choice, but the WILMAR tool's selection of model order is not discussed.

An ARMA(1, 1) scenario creation method method simulates a series of forecast errors E_t for a single site:

$$E_t = a_0 E_{t-1} + b_0 \epsilon_{t-1} + \epsilon_t$$

where a_0 and b_0 are the ARMA weighting parameters and ϵ_t is a zero mean Gaussian random process with standard deviation σ_ϵ . The error series for $t \in [0, T]$ can be constructed this way by assuming that the initial error is zero, $E_t = 0$. The scenario values are equal to the sum of the forecast and

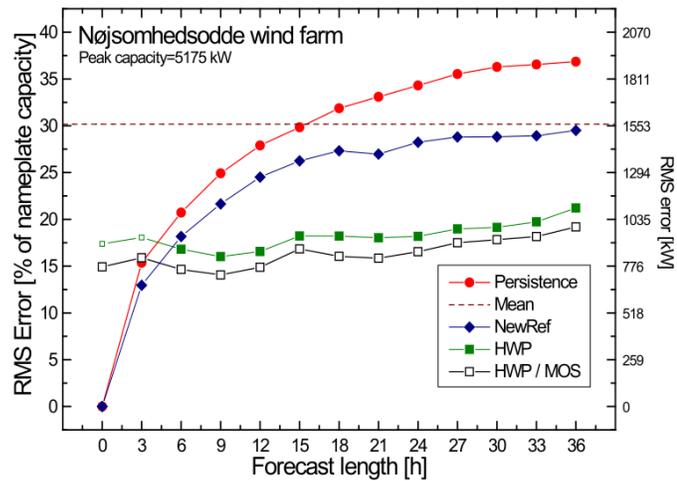


Figure 2.4: Forecast errors compared for several methods over several forecast horizons, from the ANEMOS project [57]. Both forecasts labeled HWP use NWP. The NewRef model is a weighted mix between a persistence model and the mean.

the error. The WILMAR tool then accounts for the spatial covariance of the wind speeds using a Gaussian model of covariance. A basic power conversion curve is used to translate a vector of wind speeds into an aggregate power output. However, as noted by researchers at Pacific Northwest National Laboratory (PNNL) [68] the WILMAR forecasting methodology “does not consider the cross-correlation of wind forecasts [between sites] for different time frames.”

2.4 Numerical Weather Prediction

The challenges of forecasting wind more than several hours in the future have been taken on by increasingly complex Numerical Weather Prediction (NWP) models. NWP is a method that takes mathematical models of the atmosphere and uses computational simulation to generate forecasts. Commonly used models include Global Forecast System (GFS), Weather Research and Forecasting (WRF), and ECMWF [76–78].

Figure 2.4 compares the performance of NWP and statistical models for various horizons. NWP forecasting is significantly more accurate for longer term predictions (i.e., horizons more than 3 h–6 h) than statistical models [57].

Additionally, most of the large ISOs in the US today forecast wind based on an ensemble of NWP models [1]. The forecasts from these ensemble models are often adjusted using statistical methods

(known as Model Output Statistics (MOS)) and blended with short term statistical forecasts. For a thorough review of work in NWP, see the reports from Argonne National Labs and the European ANEMOS project [54, 57].

2.5 Probabilistic and Scenario Forecasting

The forecasting method used depends on the model and on whether the intent is to generate point forecasts, probabilistic (interval) forecasts (Figure 2.5a), or scenarios (Figure 2.5b). The statistical work to date has largely focused on improving the accuracy of point forecasts. Some work has been done on probabilistic wind forecasts by Pinson [60, 79, 80] and on evaluation of these forecasts by [59]. Pinson and Madsen [72] propose a method to draw scenarios directly from these probabilistic forecasts and assert that this method takes into account the inter-temporal dependencies. However, their underlying statistical model is somewhat unclear, a Gaussian covariance structure is assumed, and the efficacy of their model was not tested in a UC context.

At the time of forecasting it is often beneficial to incorporate information on the current weather conditions. Both [59] and [60] propose choosing models which are fit to conditions similar to the current ones, e.g., a day with west-blowing wind should use the corresponding model for wind direction.

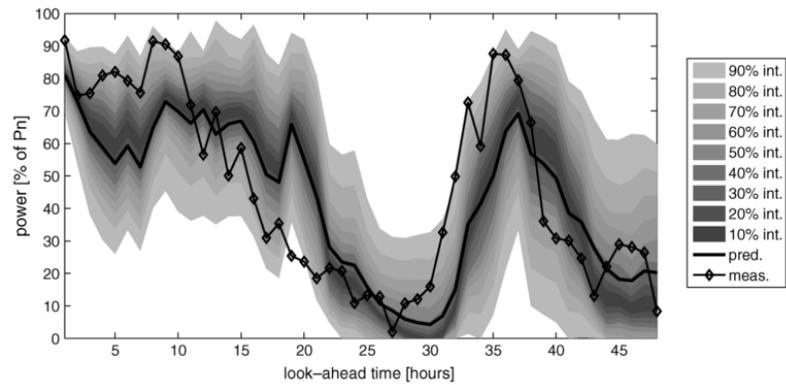
2.6 Analogs in Forecasting

Analog forecasting mines the past for similar events in order to predict the future. To search for similar events, forecast data (temperature, pressure, etc.) for the current day is compared to a historical archive of forecasts. Once a similar historical forecast, or analog, is found, the observed values (e.g., observed wind power) from that date are used to improve the current forecast.

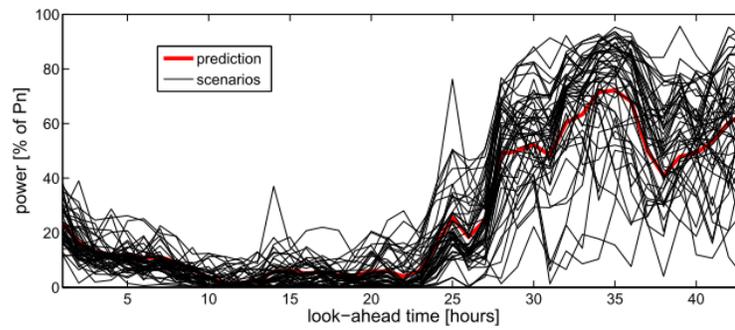
Analogs have been used in the forecasting community to create point forecasts [81–83], to correct raw NWP forecasts [84], and to create probabilistic forecasts [85]. Analogs have also been used with some success in the prediction of precipitation [81–83], stream flow [86], and wind speeds [84]. This section reviews some of the theory and practice typical in this literature.

2.6.1 Analog Techniques

There are three basic categories of analog techniques. Messner and Mayr [85] use a simple weather test system to compare the results of several probabilistic forecast methods. They compare performance to direct NWP model output and logistic regression. Several analog methods are tested including:



(a) Intervals for wind power output are forecast with the most likely regions shaded darker, from Pinson, Chevallier, and Kariniotakis [79]



(b) Discrete scenarios for wind power output are drawn from the underlying random process, from Pinson and Madsen [72]

Figure 2.5: Wind uncertainty can be represented as confidence intervals around a forecast or as discrete scenarios.

- **Analogs of deterministic forecasts (ADF)** is the most basic analogs method. The N_a smallest distances are selected and then assumed to be equally likely. A probability distribution is formed using ensemble relative frequencies over the N_a analog members.
- **Kernel estimation** methods involve weighting analogs by the closeness of their analogy (distance to the current forecast). Messner and Mayr [85] use a simple Gaussian kernel, with different standard deviation parameterizations.
- **Analog dressing** is used to improve an ensemble forecast. N_a analogs are found for each ensemble member (based on that member’s historical archive), potentially multiplying the size of the ensemble by N_a (if no dates are duplicated over ensemble members).

2.6.2 Distance metrics

There are two categories of distance metrics are in common use for analog selection:

RMS

The Root Mean Squared (RMS) distance is the standard distance metric in the analogs literature. The distances between two forecast data vectors (\vec{x}, \vec{y}) is calculated as:

$$D_{\text{RMS}}(x, y) = \sum_i \sqrt{(x_i - y_i)^2}$$

In [81], this metric is extended using feature scaling.

Rank distance

Hamill and Whitaker [82] proposed the rank distance as an alternative metric to RMS. They found that the using an RMS distance metric resulted in analog forecasts which are too close to the climatological mean, as there are more analogs to be found near the long-term average. To avoid this bias, the rank of the the current forecast data is compared against the rank of the entire archive (X_H):

$$\text{bias}_{\text{rank}}(x) = \sum_i \text{rank}(x_i) - \text{rank}(X_{H,i})$$

They conclude that using a rank corrected distance improved the analog method’s forecast skill for short forecast leads. Messner and Mayr [85] confirm that rank correction outperforms RMS distance.

2.6.3 Performance

Hamill and Whitaker [82] found that analogs methods produce deterministic forecasts with similar forecast skill to logistic regression techniques. The advantage of analogs is that they are much faster to compute than logistic regression. When compared to the effects of improving NWP models, [82] equate the improvements in skill made by analogs to “many years of sustained model development by a large staff of scientists.”

Messner and Mayr [85] compare the results of several probabilistic forecast methods on a simple weather test system and make several conclusions of interest. They find that analogs outperform traditional forecasting methods for forecast lead times more than 1 h , but do less well for a 1 h forecast lead time. Analog methods perform less well than traditional forecasting techniques at predicting very rare events. Finally, larger training sets improve short lead performance of analog methods but have little effect on long lead forecasting.

2.6.4 Wind forecasting

Delle Monache, Nipen, Liu, *et al.* [84] test the ability of post-processing methods to reduce raw NWP model errors. Analog methods are compared and combined with Kalman filtering methods. The tests are carried out on 10 m wind speeds for many Colorado locations using a Weather Research and Forecasting (WRF) modeling system with training periods between 2 and 12 months. They found that analogs perform better (in overall and spatial distribution of mean squared error, as well in as bias) than Kalman filtering or a combination of Kalman filtering and analogs. They also found that skill score improves markedly as the training period is increased from 4 to 8 months and then begins to level out.

2.6.5 Summary

The analogs method has been shown promise as a computationally easy way to improve the forecasting of many different weather phenomena. The past is mined for similar events and these analogs then become an ensemble forecast which can be averaged to get a deterministic forecast. Both RMS and rank are used to compare the distance between two forecasts. For good performance, the length of the historical archive should be relatively long. To date, analogs have been a forecast improvement method in the literature; analogs remain untested in the stochastic unit commitment literature.

Chapter 3

SCENARIO CREATION

Modeling random processes for stochastic optimization problems can be done in several ways. The most common is a recourse model which uses a discretization of the random process to quantify the potential future costs. Alternately, the random processes are represented as distributions and chance constraints are formulated to stipulate the probability of avoiding infeasibility [87]. Due to problem size and the corresponding speed of solution techniques, stochastic optimization for energy problems is handled by recourse models with discrete scenarios. A scenario is a potential realization of the stochastic process through the entire time horizon (i.e., a time series) and has a given probability. An example of a number of wind scenarios is shown in Figure 2.5b. In the stochastic optimization process scenarios are organized into trees; structures for scenario trees are discussed in §1.2. This section discusses several methods for scenario creation¹ and objective function formulation.

Accurately representing continuous random processes using discrete scenarios is a key problem in stochastic optimization. [88] give an excellent overview of methods which is summarized in this section, along with the methods used in the Unit Commitment and hydro scheduling literatures. A few techniques for the evaluation of scenario creation methods are reviewed.

Conditional sampling This method, which includes Monte Carlo (MC) sampling, is by far the most common. Extensions include variance reduction techniques (e.g., Latin Hypercube Sampling and Random Quasi Monte Carlo) and covariance corrections.

Moment matching Høyland and Wallace [7] propose a method to generate scenarios that match selected statistical properties (i.e., moments) of the reference distribution. This is achieved through a non-linear optimization which minimizes the difference between scenario and distribution moments.

Scenario reduction Reduction is often combined with MC sampling and attempts to reduce the size of a scenario tree while maintaining stochastic information.

¹In the literature the process of forming a finite set of discrete scenarios is generally referred to as scenario generation. In this work this process is referred to as scenario creation to avoid semantic confusion.

Internal sampling This method samples the distribution within the solution procedure by using Benders decomposition methods [89, 90].

Within the UC literature there have been many methods used to create scenarios. In some early work scenarios are based directly on previous years observations, as in [11], or are simply declared, as in [52]. However, the dominant method for scenario tree construction has been a MC sampling method combined with scenario reduction. More recently, a few researchers have tried a statistically based approach called moment matching. Little attention has been paid to the effect of the number of scenarios used on the solution cost.

3.1 Monte Carlo Sampling Method

Monte Carlo (MC) simulation based on a normal error model has been a common scenario creation technique in the UC literature [5, 13, 17]. The intuition behind MC sampling is simple: given a large number of samples², a Gaussian distribution can be adequately represented as a set of discrete values. Gaussian distributions have often been used in power systems literature, either as basic per-time models of wind forecast error [17, 22, 36, 38] or as part of an ARMA forecast error model [5, 12, 13]. Given a Gaussian model, MC sampling randomly draws N_s samples from the distribution for each time. These samples are joined together to become scenarios.

Al-Awami and El-Sharkawi [14] attempt to incorporate information about the statistical properties of change in wind power through a “shuffling” forecast procedure, assuming that statistics on change in wind power are known and normal. However, this method underestimates the likelihood of large down ramps in wind when the wind is blowing highest because the mean down ramp magnitude is reduced by the inclusion of lower wind events.

Homem-de-Mello, Matos, and Finardi [91] propose two alternate variance reduction techniques (Latin Hypercube Sampling and Random Quasi Monte Carlo) and suggests that they may produce better scenarios than MC. A long term hydro planning experiment is presented, but the results are inconclusive.

3.1.1 Scenario Reduction

Scenario reduction algorithms are often used to shrink the scenario set to a size, N_r , which a MIP solver can handle in a reasonable amount of time. Scenario reduction algorithms like those of [53] result in an optimal or near-optimal reduction given the number of final scenarios, but there is

²In the power systems literature, the number of scenarios needed to ensure accurate representation of a Gaussian is simply assumed to be very large; e.g., in [17] the number of 3000 was chosen by fiat.

information lost. In the data-dependent case described in [53], a 50% reduction in the number of scenarios resulted in a 10% reduction in scenario tree representation accuracy and a 98% scenario reduction corresponded to a 50% accuracy reduction.

A convincing rationale for the choice of the reduced number of scenarios, N_r , has yet to be presented in the power systems literature, either by statistical means or empirical comparison. Most studies, including those using the WILMAR model [5, 12, 13] as well as [14, 16], simply state the small number (N_r between 5-30) of scenarios used along with the result.

3.2 Moment Matching Method

Recently, Lowery and O'Malley [9] modified the WILMAR Scenario Tree Tool to incorporate moment matching. To create wind error scenarios, moments representing the per-time mean, standard deviation, skew, and kurtosis over all scenarios are used. Additionally, the autocorrelation between stages³ was included to produce smooth scenarios across the stages. A small number of scenarios (less than 6 scenarios for all stages) was used, with a similar experimental setup to Tuohy, Meibom, Denny, *et al.* [5]. Lowery and O'Malley [9] conclude that the consideration of higher order moments (skew and kurtosis) does result in a lower solution cost and that the moment estimation accuracy can have a significant impact on cost. However, [9] makes no quantitative comparison with the results of [5].

Moment matching has also been used to create scenarios for natural gas prices and electricity demand for generation expansion planning [8]. In [8] both demand and prices are modeled as correlated Geometric Brownian Motion processes, with moments of mean, variance, and skew for each variable and a correlation between variables. The main focus of [8] is a scenario reduction method which solves individual scenario optimization problems and then groups scenarios by their first stage variable values.

In [92], bid curves are modeled stochastically using the first four central moments are used to model electricity spot prices, regulation market spot prices, and demand, along with the correlation between variables. Statistics for regulation prices were found to be highly conditional on the spot price, thus the desired moments were constructed conditionally. Between one hundred and six hundred scenarios were generated.

Mello, Lu, and Makarov [68] propose a technique similar to moment matching but based on an ARMA process model. The objective function is a weighted sum of the absolute differences of the model statistics and the generated scenario statistics, including the mean, standard deviation, cross

³ In both [5] and [9] the WILMAR Scenario Tree Tool is a four stage decision model with each progressive stage covering a greater timespan and having a greater number of scenarios. See [9, Table IV] for details.

correlation (for multiple time series), and auto correlation. This scenario creation procedure was used to generate 150 scenarios, with correlated day ahead, hour ahead, and real-time forecasts, but was not tested in the UC context.

3.3 NWP based scenarios

Constantinescu, Zavala, Rocklin, *et al.* [16] implement a method that draws scenarios directly from the ensemble members of many NWP models (the study used a 30-member WRF ensemble) and feeds these to a UC. This method holds promise for the future of utility or ISO forecasting, but requires access to large, location specific meteorological datasets. Also, due to the complexity of the NWP models, the work in [16] was computationally limited in the number of scenarios that could be generated. About 500 CPUs would be required in order to generate NWP forecasts on an hourly basis. The effect of this limitation on the UC solution is not discussed.

3.4 Hydro scenarios

In the long-term hydro scheduling literature, stochastic scenarios are used to represent inflows. Scenarios for this problem are structured in multistage scenario trees with a specified number of branches per stage. Often very little is said about how these scenario trees are created, as is the case in [51, 93, 94]. AR models have been used to generate inflow scenarios, as in [95]. Stochastic Dual Dynamic Programming (SDDP) is often used to solve these problems. As this algorithm samples the scenario tree in the forward step, Monte Carlo based methods are common.

Part III
METHODS

A major consideration in any stochastic UC — and the major focus of this work — is the quality and quantity of the scenarios. This section details two statistically based methods for scenario creation — the moment matching method and the analogs method. Both of these methods allow a scheduler to take a deterministic forecast for wind power and leverage historical data to create many scenarios. The moment matching method synthesizes scenarios to match a set of moments or statistics. This work proposes a method for conditionally estimating the desired moments based on historical data making these moments conditional on the current forecast. The analogs method searches the past for forecasts similar to the current forecast and uses the observations from those analogous dates directly as scenarios. This work proposes a *simple* analogs method based only on the aggregate power forecast. Along with these methods, four scenario error criteria are proposed to help explain the performance of stochastic UC.

We also describe the decision making structure and MIP implementation of stochastic UC for power systems with significant penetrations of wind. A framework is given for the evaluation of stochastic UC results relative to the costs of deterministic UC.

Chapter 4

SCENARIO CREATION

4.1 *Moment Matching*

This section presents a wind scenario producing implementation of the moment matching method of Høyland and Wallace [7]. The basic idea of moment matching is that scenarios can be synthesized to match a set of desired statistics. In this implementation, the objective is based on moments representing the mean, standard deviation, skew, and kurtosis. An additional term representing the mean of the first difference is added to reflect the importance of the wind ramping. We propose that desired moments can be estimated conditionally on the forecast, based on a twenty quantile binning of the forecast values in the training set. The distance between each desired moment and the scenarios' moment is calculated using the RMS metric. Feature scaling is used to weight the importance of the moments evenly.

4.1.1 *Scenario Creation Problem Formulation*

The goal of the moment matching method is to produce N_s scenarios that have moments M_i which are similar to the desired moments \hat{M}_i . The variables in the optimization problem are the scenario values $x_s[t]$ (one value for each scenario s and time t) and probabilities ρ_s .

Objective

$$\min \sum_t \sum_i w_i \cdot D(M_i[t], \hat{M}_i[t]) \quad (4.1)$$

The RMS distance metric was used as the moment deviation function $D()$. The weighted (by w_i) set of moments for this implementation includes:

- M_1 : first order moment (the mean)
- M_2 : second order moment (representing the variance)
- M_3 : third order moment (representing the skew)
- M_4 : fourth order moment (representing the kurtosis)

- M_5 : mean of the first difference

The per-time scenario mean M_1 is calculated as the sum of the product of the scenario value $x_s[t]$ and the scenario probability ρ_s :

$$M_1[t] = \sum_s x_s[t] \cdot \rho_s \quad (4.2)$$

The second, third, and fourth order moments can be calculated as the sum of the difference between the scenario value and the first order moment $M_1[t]$ to the i^{th} power, where i is the order of the moment:

$$M_i[t] = \sum_s \rho_s \cdot (x_s[t] - M_1[t])^i \quad (4.3)$$

However, within the optimization this was found to be computationally difficult due to the recursion created by the $M_1[t]$ term. In the optimization implementation, the moments are estimated using the forecast value $\hat{M}_1[t]$ instead of the variable $M_1[t]$:

$$M_i[t] \approx \sum_s \rho_s \cdot (x_s[t] - \hat{M}_1[t])^i \quad (4.4)$$

The mean of the first difference is defined as:

$$M_5[t] = \sum_s (x_s[t+1] - x_s[t]) \cdot \rho_s \quad (4.5)$$

Weighting the moments' contribution to the objective function is a sensitive process in the moment matching algorithm. This implementation uses feature scaling (a method of standardizing the ranges of different variables) to set the weights such that:

$$w_i = e_i \cdot 1 / \sum_t |M_i[t]| \quad (4.6)$$

As the estimation of each of the higher order moments is dependent on having an accurate mean, the mean weighting factor is set to $e_1 = 10$ to emphasize its importance. The first difference mean weight is deemphasized by a factor of $e_5 = 0.1$, to reflect the wide variation seen in estimating this moment¹. All other moments are evenly emphasized with $e_i = 1$.

Required Constraints

The probability must be constrained to be between a low limit (to prevent the creation of scenarios with trivially small probabilities) and one:

$$\text{s.t. } \frac{1}{N_s}/100 \leq \rho_s \leq 1 \quad \forall s \in S \quad (4.7)$$

¹The observed first difference mean relative to the forecast first difference is shown in Table 4.4. The forecast generally overestimates the magnitude of the ramp. Due to this lack of certainty, this moment is weighted to be less influential than others.

The sum of the probabilities must equal one:

$$\text{s.t. } \sum_s \rho_s = 1 \quad (4.8)$$

Additional Constraints

Some constraints were added to the problem to improve the quality of the scenarios. The optimization results without these constraints often resulted in wind power scenario values which fluctuated between extremely low and extremely high values from one time interval to the next. To prevent this unrealistic behavior, additional constraints were added to limit scenario values and first differences. If the problem is not feasible with these additional constraints, they are removed and the problem is resolved with only the required constraints.

The limiting constraints ensure that a scenario does not deviate from the deterministic forecast by more than a pre-determined amount:

$$\text{s.t. } \underline{L}[t] \leq x_s[t] \leq \overline{L}[t] \quad \forall t \in T \quad \forall s \in S \quad (4.9)$$

The limits $\underline{L}[t]$ and $\overline{L}[t]$ are estimated by taking the min/max observed value for the appropriate forecast bin. For example, if the forecast $f[t]$ is 8500 MW, the bin might contain forecasts between 8400–9000 MW. The smallest/largest historical observed values for that bin might be $\underline{L}[t] = 6000$ MW and $\overline{L}[t] = 8800$ MW (this example bin is near capacity, so over-forecasting is more severe than under-forecasting). All scenario values for this time will fall within this range.

The first difference limiting constraints ensure that each scenario has a reasonable first difference:

$$\text{s.t. } \underline{\delta L}[t] \leq x_s[t+1] - x_s[t] \leq \overline{\delta L}[t] \quad \forall t \in T \quad \forall s \in S \quad (4.10)$$

The difference limits $\underline{\delta L}[t]$, $\overline{\delta L}[t]$ are estimated as the min/max of the first difference of the observed values for the appropriate forecast bin.

4.1.2 Scenario Reduction

The moment matching process as specified in the last section tends to produce scenarios which are very similar to each other. To avoid sending duplicate scenarios (which provide no additional information to the solution algorithm) to the solver, scenarios are evaluated on a similarity criteria and are combined if similar. For a discussion on the effects of the duplicate scenarios produced by moment matching, see §7.2.2.

Similarity criteria

Two scenarios i and j for a given day are labeled as being “similar” if:

1. They have a high correlation: $C[i, j] \geq 0.999$.
2. Their total absolute difference is small: $\sum_t |x_i[t] - x_j[t]| < \epsilon$. In this work ϵ is set to 10 MWh over the course of the day.

Scenario aggregation

Two similar scenarios i and j are aggregated into a combined scenario c by summing their probabilities:

$$\rho_c = \rho_i + \rho_j$$

and averaging their values (weighted by probability):

$$x_c[t] = (x_i[t] \cdot \rho_i + x_j[t] \cdot \rho_j) / \rho_c$$

Implementation

The non-linear problem formulation defined above was implemented in Python using the Pyomo package. Desired moments for a single day (changing hourly based on the forecast bin) were loaded into the optimization model and the solver IPOPT was used to solve the day’s problem.

4.1.3 Desired Moment Estimation

A key input to the moment matching process is the desired moments, $\hat{M}_i[t]$. These are the statistics which the scenarios are synthesized around. The moment matching method is a general theory — the process of determining which moments to use and how to estimate their values is application dependent.

In this work, we propose a method for estimating the desired moments conditional upon the forecast. When creating moments for wind power production, the historical forecasts and observations can be used. Low and high wind periods have significantly different statistics. The strategy used to account for these differences is to make the desired moments conditional on the capacity normalized² forecast values.

²All calculations are done in per-unit terms to account for any wind capacity additions that occur during the time period covering the training data.

The historical data set is grouped by normalized forecast quantile³. Then for each forecast quantile (and the corresponding set of training set times, T_b), the conditional statistics $m_{i,b}$ are calculated based on the observed wind $o[t]$. The bin mean is calculated as:

$$m_{1,b} = \frac{1}{N_b} \cdot \sum_{t \in T_b} o[t] \quad (4.11)$$

The second through fourth order moments are calculated as:

$$m_{i,b} = \frac{1}{N_b} \cdot \sum_{t \in T_b} (o[t] - m_{1,b})^i \quad (4.12)$$

and the first difference mean is calculated as:

$$m_{5,b} = \frac{1}{N_b} \cdot \sum_{t \in T_b} (o[t] - o[t-1]) \quad (4.13)$$

Choosing desired moments for testing data

First the testing data is binned using the same bin limits as the training data. Then higher order moment values for a given time $\hat{M}_i[t]$ in the testing set can be calculated by looking up the bin for the current forecast $f[t]$ — selecting b where $\underline{f}[t] < f_b \leq \bar{f}_b$ and using the bin moment as the desired moment value:

$$\hat{M}_i[t] = m_{i,b} \quad \forall i \in [2, 4] \quad (4.14)$$

The first moment is simply the actual forecast (rather than the bin mean):

$$\hat{M}_1[t] = f[t] \quad (4.15)$$

The first difference moment is calculated based on the first difference bin — selecting b where $\underline{\Delta f}[t] < \Delta f_b \leq \overline{\Delta f}_b$ — and then setting:

$$\hat{M}_5[t] = m_{5,b} \quad (4.16)$$

Table 4.2 gives an example of these binned desired moments and limits which are binned by the forecast. Table 4.4 shows the first difference moment $m_{5,b}$ and limits which are based on the first difference of the forecast.

³Histograms have bins of constant width, while quantiles have bins with constant count. Quantiles were chosen for this application, as wind is most often in a low capacity factor regime and data for the upper histogram bins would be estimated on a low number of samples.

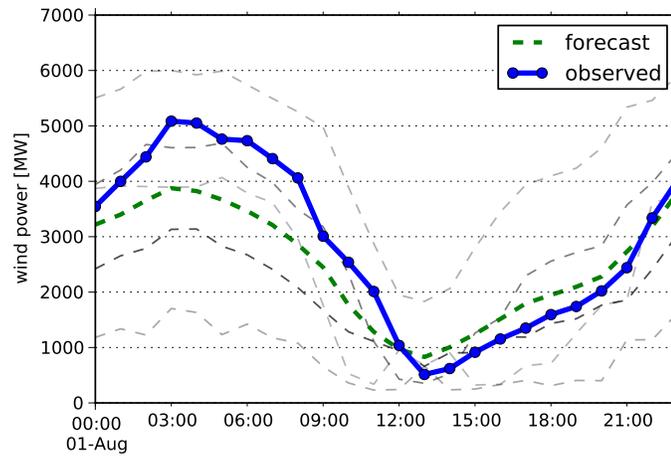


Figure 4.1: Five scenarios for ERCOT wind on August 1, 2012 were created using the moment matching method. The forecast and observed wind are also shown for comparison. The moment matching scenarios cover a wider range of possibilities than the forecast.

Figure 4.1 shows an example day of scenarios created for August 1, 2012 using the moment matching method.

\underline{f}_b	\bar{f}_b	bin count	$m_{1,b}$	$m_{2,b}$	$m_{3,b}$	$m_{4,b}$	$\underline{L}[t]$	$\bar{L}[t]$
0.0000	0.0501	162	$4.12 \cdot 10^{-2}$	$1.02 \cdot 10^{-3}$	$5.00 \cdot 10^{-5}$	$1.00 \cdot 10^{-5}$	0.1618	$4.5700 \cdot 10^{-3}$
0.0501	0.0760	161	$6.40 \cdot 10^{-2}$	$1.45 \cdot 10^{-3}$	$1.40 \cdot 10^{-4}$	$3.00 \cdot 10^{-5}$	0.2988	0.0129
0.0760	0.0987	161	$9.65 \cdot 10^{-2}$	$2.02 \cdot 10^{-3}$	$1.00 \cdot 10^{-4}$	$3.00 \cdot 10^{-5}$	0.3326	0.0240
0.0987	0.1213	161	$1.22 \cdot 10^{-1}$	$3.98 \cdot 10^{-3}$	$3.80 \cdot 10^{-4}$	$1.10 \cdot 10^{-4}$	0.4289	0.0233
0.1213	0.1454	161	$1.46 \cdot 10^{-1}$	$5.88 \cdot 10^{-3}$	$6.40 \cdot 10^{-4}$	$2.00 \cdot 10^{-4}$	0.4631	0.0336
0.1454	0.1702	161	$1.78 \cdot 10^{-1}$	$9.14 \cdot 10^{-3}$	$8.90 \cdot 10^{-4}$	$3.00 \cdot 10^{-4}$	0.4706	0.0324
0.1702	0.1957	161	$2.05 \cdot 10^{-1}$	$1.03 \cdot 10^{-2}$	$7.80 \cdot 10^{-4}$	$3.40 \cdot 10^{-4}$	0.5189	0.0282
0.1957	0.2229	161	$2.08 \cdot 10^{-1}$	$1.09 \cdot 10^{-2}$	$8.10 \cdot 10^{-4}$	$3.70 \cdot 10^{-4}$	0.4900	0.0342
0.2229	0.2482	161	$2.39 \cdot 10^{-1}$	$1.37 \cdot 10^{-2}$	$1.35 \cdot 10^{-3}$	$6.00 \cdot 10^{-4}$	0.5672	0.0539
0.2482	0.2719	161	$2.54 \cdot 10^{-1}$	$1.48 \cdot 10^{-2}$	$1.74 \cdot 10^{-3}$	$7.40 \cdot 10^{-4}$	0.5757	0.0663
0.2719	0.3036	161	$2.80 \cdot 10^{-1}$	$1.25 \cdot 10^{-2}$	$1.05 \cdot 10^{-3}$	$5.00 \cdot 10^{-4}$	0.6112	0.0953
0.3036	0.3376	161	$3.26 \cdot 10^{-1}$	$1.40 \cdot 10^{-2}$	$6.80 \cdot 10^{-4}$	$5.10 \cdot 10^{-4}$	0.6177	0.1161
0.3376	0.3635	161	$3.73 \cdot 10^{-1}$	$1.80 \cdot 10^{-2}$	$3.40 \cdot 10^{-4}$	$7.50 \cdot 10^{-4}$	0.6803	0.1073
0.3635	0.4012	161	$3.69 \cdot 10^{-1}$	$1.33 \cdot 10^{-2}$	$2.50 \cdot 10^{-4}$	$4.80 \cdot 10^{-4}$	0.6553	0.0991
0.4012	0.4459	161	$4.28 \cdot 10^{-1}$	$1.56 \cdot 10^{-2}$	$-5.00 \cdot 10^{-4}$	$5.70 \cdot 10^{-4}$	0.6755	0.1179
0.4459	0.4886	161	$5.06 \cdot 10^{-1}$	$9.72 \cdot 10^{-3}$	$-1.40 \cdot 10^{-4}$	$3.30 \cdot 10^{-4}$	0.8325	0.2306
0.4886	0.5360	161	$5.22 \cdot 10^{-1}$	$9.20 \cdot 10^{-3}$	$-6.40 \cdot 10^{-4}$	$3.10 \cdot 10^{-4}$	0.7425	0.2329
0.5360	0.6035	161	$5.76 \cdot 10^{-1}$	$9.71 \cdot 10^{-3}$	$-1.11 \cdot 10^{-3}$	$4.30 \cdot 10^{-4}$	0.7641	0.2457
0.6035	0.6821	161	$6.05 \cdot 10^{-1}$	$7.26 \cdot 10^{-3}$	$-7.50 \cdot 10^{-4}$	$2.70 \cdot 10^{-4}$	0.7782	0.2753
0.6821	1.0000	162	$6.30 \cdot 10^{-1}$	$4.44 \cdot 10^{-3}$	$-6.90 \cdot 10^{-4}$	$2.70 \cdot 10^{-4}$	0.7283	0.2264

Table 4.2: Quantile-conditional moment values and limits for ERCOT 2011–2012, per-unit normalized. The values \underline{f}_b and \bar{f}_b are the min/max forecast values for the quantile. Bin count is the number of training samples (hours) the moment is estimated on. The statistics $m_{i,b}$ are based the per-unit observed wind. $\underline{L}[t]$, $\bar{L}[t]$ limit the range of scenario values.

$\underline{\Delta f}_b$	$\overline{\Delta f}_b$	$m_{5,b}$	$\underline{\delta L}[t]$	$\overline{\delta L}[t]$
-1.00000	-0.06439	-0.04536	-0.43822	0.07881
-0.06439	-0.04646	-0.02660	-0.21289	0.36520
-0.04646	-0.03583	-0.02677	-0.15790	0.04776
-0.03583	-0.02894	-0.01850	-0.17283	0.14917
-0.02894	-0.02264	-0.01788	-0.11795	0.06333
-0.02264	-0.01720	-0.01095	-0.10608	0.10551
-0.01720	-0.01323	$-6.51000 \cdot 10^{-3}$	-0.10101	0.07318
-0.01323	$-8.94000 \cdot 10^{-3}$	$-8.35000 \cdot 10^{-3}$	-0.11265	0.11377
$-8.94000 \cdot 10^{-3}$	$-5.16000 \cdot 10^{-3}$	$-8.80000 \cdot 10^{-4}$	-0.09603	0.15996
$-5.16000 \cdot 10^{-3}$	$-1.60000 \cdot 10^{-3}$	$-1.70000 \cdot 10^{-3}$	-0.10685	0.12583
$-1.60000 \cdot 10^{-3}$	$1.59000 \cdot 10^{-3}$	$-3.90000 \cdot 10^{-4}$	-0.08062	0.09683
$1.59000 \cdot 10^{-3}$	$4.79000 \cdot 10^{-3}$	$3.69000 \cdot 10^{-3}$	-0.07474	0.09428
$4.79000 \cdot 10^{-3}$	$9.27000 \cdot 10^{-3}$	$6.34000 \cdot 10^{-3}$	-0.09146	0.16877
$9.27000 \cdot 10^{-3}$	0.01451	$8.75000 \cdot 10^{-3}$	-0.07938	0.16187
0.01451	0.01945	$3.14000 \cdot 10^{-3}$	-0.11027	0.07512
0.01945	0.02681	0.01412	-0.13567	0.17266
0.02681	0.03708	0.02039	-0.10483	0.16827
0.03708	0.05162	0.02210	-0.15152	0.18509
0.05162	0.07607	0.03187	-0.07442	0.21086
0.07607	1.00000	0.05568	-0.11201	0.25884

Table 4.4: Quantile-conditional first difference moments and limits for ERCOT 2011–2012, per-unit normalized. The moment $m_{5,b}$ is the mean of the first difference. $\underline{\delta L}[t]$, $\overline{\delta L}[t]$ limit the range of the first difference of the scenarios.

4.2 Analogs

Analog forecasting mines the past for similar events in order to predict the future. To search for similar events, forecast data (temperature, pressure, etc.) for the current day is compared to a historical archive of forecasts. Once a similar historical forecast, or analog, is found, the observed values (e.g., observed wind power) from that date are used to improve the current forecast.

To create scenarios using the analogs method, the observations from each date in the historical archive are treated as a possible scenario. The closer an analog's forecast data is to the current forecast, the more likely the scenario is. An outline of the analogs algorithm can be found in Algorithm 1. First the current forecast is compared with all historical forecasts, resulting in a set of analog distances. These distances are sorted and used to create an estimate of the probability of each analog. The observed values from the dates corresponding to the lowest N_s distances become scenarios. These steps are examined in more detail in the following sections.

4.2.1 Forecast data

In general forecast data consists of a set of predictions for many weather variables (often on the order of 50 variables, including 100 m wind speed, pressure, temperature, etc.) at several geographical points. However, in this work we test the performance of *simple* analogs, using only the aggregate power forecast data. While we expect that some of the benefits of using richer weather data will be lost when using *simple* analogs, this approach is conceptually simpler and requires significantly less data. *Simple* analogs allow scenario creation with only a historical record of forecasts and observed wind.

The analogs formulation laid out in this section can be generalized to multivariate weather data. The multivariate forecast matrix is simply reshaped to become a vector. It may become computationally expensive to compute the analog distances over a large archive with large numbers of variables. A dimensionality reduction technique (e.g., principal components analysis) may be required to make the problem tractable.

While not tested in this work, using weather data as the analog set predictor does allow for a simple, theoretically sound modeling of many wind farms. When creating scenarios for many wind farms using statistical methods, a wind speed correlation correction strategy [61, 62] can be used to create a single set of scenarios. With the analogs method, an aggregate wind scenario can be created by simply summing the observed wind farm power timeseries for the analog date.

Algorithm 1: An Analogs Method for Scenario Creation

Data: current forecast vector \vec{f}^*

Data: historical archive over $t \in T$ of :

observation vectors $\vec{o}[t]$

and forecast vectors $\vec{f}[t]$

Input: number of scenarios N_s

Result: a set of scenarios over $s \in [0, N_s)$, including

scenario values $\vec{x}[s]$

and probabilities $\rho[s]$

initialize distances, $\vec{x}[s], \rho[s]$

compute distances between the current forecast

and each forecast in the archive

for $t \in T$:

 | distances[t] $\leftarrow D(\vec{f}^*, \vec{f}[t])$

distanceValuesSorted $\leftarrow \text{sort}(\text{distances})$

convert distances to probabilities (see §4.2.4)

allProbabilities $\leftarrow \text{KDEprob}(\text{distanceValuesSorted})$

for $s \in [0, N_s)$:

 | # lookup the time corresponding to the s smallest forecast distance value

 | $t \leftarrow \text{distances.positionOf}(\text{distanceValuesSorted}[s])$

 | # set the scenario values to the observed values at t

 | $\vec{x}[s] \leftarrow \vec{o}[t]$

 | $\rho[s] \leftarrow \text{allProbabilities}[s]$

4.2.2 Analog distance

To find analogous days, the forecast for each day t in the historical archive \vec{f}_t is compared to the current forecast \vec{f}^* :

$$D[t] = D(\vec{f}^*, \vec{f}_t) \quad (4.17)$$

Consistent with much of the forecasting literature on analogs, we use a squared difference (L2 norm) as the comparison metric:

$$D(\vec{a}, \vec{b}) = \sum_i (a_i - b_i)^2 \quad (4.18)$$

Instead of computing distances per day it is possible to compare forecasts at each hour in the historical archive. This approach is advantageous as it allows for closer matches from a “larger” archive. The hourly distances for an example “current” date are shown as a timeseries in Figure 4.2.

4.2.3 Selection

Usually, the entire archive contains too many analogs for each analog to become a scenario. Also, analogs with large distances are very unlikely to be representative of the likely weather conditions for the date in consideration. To select a size N_s subset of the available analogs as scenarios, this implementation simply selects the dates for the N_s smallest distances, with the constraint that the selected distances must be at least 24 h apart. This constraint avoids the selection of many similar distance analogs, e.g., two analogs which are one hour shifted. The sorted order distances for an example date are shown in Figure 4.3. Scenarios values are set to the observed values \vec{o}_t for each of the times t belonging to the smallest N_s analog distances.

If forecasting was near perfect, the smallest distance analogs would be very similar. However, in practice the observed values for the lowest distance analog dates are generally fairly different due to forecasting error. See Figure 4.5, which shows the five lowest distance analogs for ERCOT on Aug. 1, 2012. The distribution of analog scenario spreads (max. minus min. values) for the ERCOT test case is shown in Figure 7.6. This histogram has a peak at about 15% of wind capacity. If forecasting error is low and more coverage of low likelihood events is desired it is possible to aggregate or sample analogs to form the scenario set. An investigation of these scenario selection methods is left for future research.

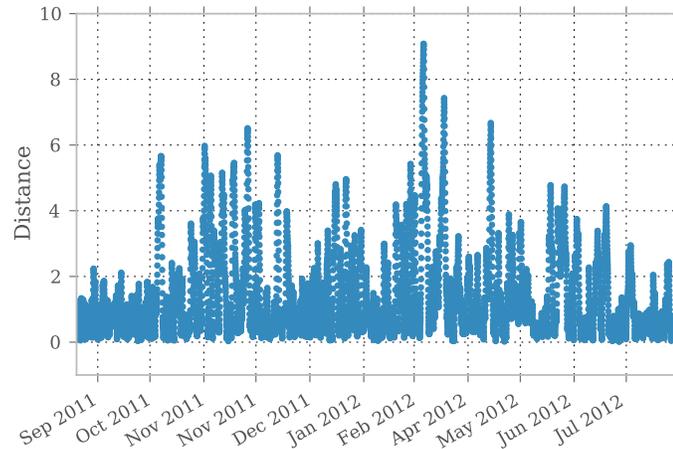


Figure 4.2: Analog distances between the forecast for 2012-08-01 and the forecast starting each hour in the historical archive (about a year worth of data prior to that August 2012). Distance is calculated as the sum of squares difference between the wind power forecasts in per-unit terms.

4.2.4 Assigning probability

Probabilities are assigned to the analogs based on their distances using Gaussian Kernel Density Estimation (KDE). A Gaussian mixture model is used to model the probability distribution of the distances over the entire historical archive⁴. Each analog is then assigned a probability equal to the mixture model’s PDF value at the analog’s distance. See Figure 4.4 for an example of the distance to probability transformation.

4.2.5 Caveats

There is an implicit assumption that the forecasting method does not change significantly over the entire historical period or that historical re-forecasts are made upon adopting a new method. In the experiment that follows we use public aggregate forecast data, and have no information to support or reject a constant forecast method hypothesis. An additional caveat is the potential effect that the length of the historical archive has on the quality of analog scenarios: short archives may not include rare events.

⁴Analog distances are bounded on $[0, \infty)$, while Gaussian models expect a two sided distribution. Ideally a one-sided kernel (i.e., a Weibull distribution) would be used, but non-Gaussian KDEs are not common. To correct for the one-sidedness of the distances distribution, KDE is performed on a vector of the original distance values (d_i) and their additive inverses ($-d_i$).

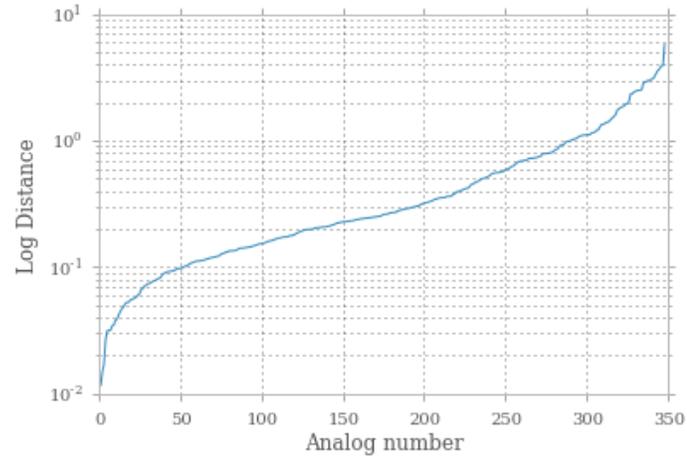


Figure 4.3: Sorted analog distances for 2012-08-01 are shown on a log scale. There are about 50 low-distance analogs, followed by a slow increase in distance value, followed by a few large distance analogs.

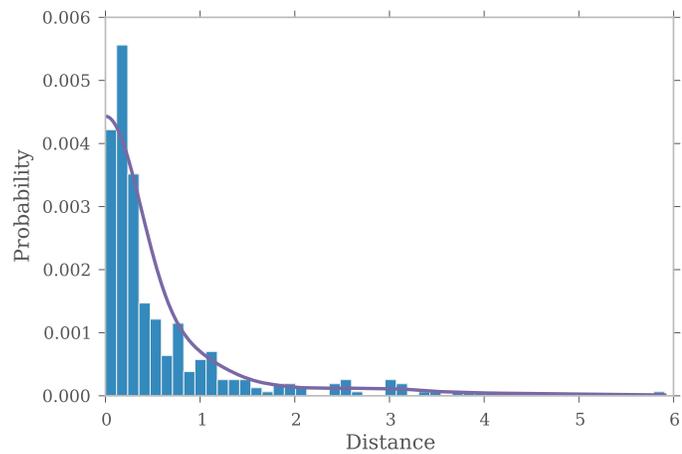


Figure 4.4: A normalized histogram of analog distances for 2012-08-01 is shown. A Gaussian KDE is used to model the distribution of the distances and is shown as line. Each analog in the historical archive is assigned a probability equal to the PDF value at the analog's distance.

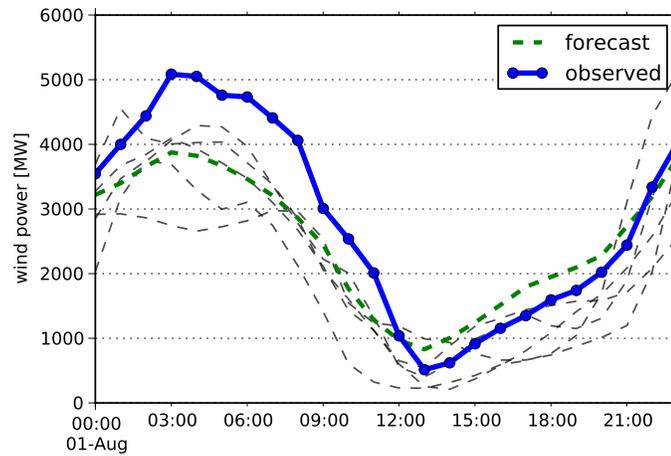


Figure 4.5: Five scenarios for ERCOT wind on August 1, 2012 were created using the analogs method. The forecast and observed wind are also shown for comparison. The analogs scenarios do cover a slightly wider range of possibilities than the forecast, but do not always anticipate the full range of the observed values for this date.

The recent, rapid installation of the wind power may also complicate the use of the analogs method for some systems. If the geographical distribution of the wind installations has not changed significantly over the historical archive, one solution (used here) is to handle all historical data in capacity normalized terms and then scale analogs to current capacity. Alternately, the historical power values for new wind farms could be simulated, using historical site-specific wind speed forecasts and a power curve. This observed power estimation process would better account for the changing geographic distribution of wind farms.

4.3 Scenario Metrics

The quality of a set of scenarios is determined by the observed system cost for the commitment which the scenarios produce. However, it is useful to be able to explain these cost results in terms of the scenario values relative to the observed values, with metrics akin to the forecast error for a deterministic UC. Stochastic UC is time-consuming and as researchers we would like to be able to interpret the results it produces. In this work four timeseries metrics are used to discuss scenario quality: scenario mean error, spread, and under/over prediction.

Mean error is defined as the difference between the per-time mean of the scenario values and the observed values $o[t]$:

$$\tilde{e}[t] = \bar{x}[t] - o[t] \quad (4.19)$$

where the scenario mean is the sum of the product of probability ρ_s and scenario value $x_s[t]$:

$$\bar{x}[t] = \sum_s \rho_s \cdot x_s[t]$$

Scenario mean error is closely related to deterministic forecast error: $e[t] = x[t] - o[t]$. In general we would expect a good scenario set to have a mean error less than or equal to forecast error. In practice this may be difficult to achieve with only a few scenarios.

Spread is defined as the difference between the scenario maximum and the scenario minimum:

$$\tilde{s}[t] = \bar{x}[t] - \underline{x}[t] \quad (4.20)$$

Deterministic forecasts have zero spread. Good scenarios should have low spreads for more common cases and higher spreads for less certain cases.

Under and Over Prediction are defined as the amount that all scenarios in the set are less than or greater than the observed wind. Under/over prediction can be thought of as how poorly the scenarios anticipate the extreme values. Over-forecast prediction equals the amount the minimum scenario overshoots the observed wind by:

$$\tilde{p}_u[t] = \begin{cases} \min(x_s[t]) - o[t], & \text{if } o[t] < \min(x_s[t]) \\ 0, & \text{else} \end{cases} \quad (4.21)$$

Under-forecast spread equals the amount the maximum spread undershoots the observed wind by:

$$\tilde{p}_o[t] = \begin{cases} \max(x_s[t]) - o[t], & \text{if } o[t] > \max(x_s[t]) \\ 0, & \text{else} \end{cases} \quad (4.22)$$

4.4 Summary

This chapter details the methods used in this work to create wind scenarios based on historical aggregate power data. Both the moment matching method and the *simple* analogs method are described. Scenario metrics akin to deterministic forecast error are also discussed.

The moment matching method for scenario creation presented here is similar to that of [9], but has a few key differences. The major improvement is the method for estimating the desired moments conditional on the current forecast. There are also limits placed on the scenario values, first differences, and probabilities to aid in the creation of more realistic scenarios.

The *simple* analogs method proposed in this work searches the historical archive for similar aggregate power forecasts for wind §4.2 also describes a method for creating scenarios using a set of analogs and their distances.

Similar to the concept of deterministic forecast error, this work proposes that scenarios can be described by four metrics: mean error, spread, and under/over prediction. These give us a framework for discussing stochastic UC costs in terms of the quality of scenarios.

Chapter 5

STOCHASTIC SCHEDULING

The core of any stochastic UC is the deterministic formulation. The deterministic thermal UC formulation used in this work is based on the MIP formulation of [18] and is described in detail in §B.1. A MIP formulation of the short term hydro scheduling problem is described in detail in Appendix §E.

The stochastic formulation can be thought of as a replication of the deterministic problem’s variables and constraints for each scenario, with a few modifications. These major differences are discussed in the following sections. The stochastic problems in this work are formulated in the extensive form, as described in detail in §B.5. This allows the problems to be solved by commercial MIP solvers, although in practice the number of scenarios which can be used are currently limited. The chapter closes with the discussion of a method for evaluating stochastic UC results.

5.1 Decision Making Structure

The stochastic UC problems in this work are formulated using two stages. Multiple stage formulations [5, 9] have certain advantages at longer time horizons; periods further from the present can be covered by more scenarios with decreased time resolution. However, the two stage formulation is a more natural fit to the existing power systems decision making model; units are committed in the day-ahead decision and unit power is decided in the real-time dispatch.

In a two-stage stochastic optimization formulation each variable must be either “here and now” (first stage) or “wait and see”. First stage variables must be decided at the outset, i.e., prior to the realization of the stochastic process, while second stage variables can vary depending on what happens.

For the thermal UC problem, unit statuses must be chosen in the day-ahead, without knowing the wind power, and in general cannot be changed. Power is adjusted in the realtime dispatch depending on the wind. To implement the logical requirement of “here and now” decisions, stochastic problems add “non-anticipatory” constraints to fix each first stage variable across all stages. The non-anticipatory constraint formulation for the thermal UC problem is described in Equation (B.35).

5.1.1 *The horizon effect and non-anticipatory constraints*

Power system operations also require avoiding the horizon effect; the start of the next operating period must be reachable given the solution for the current period. This is generally accomplished by running a UC for more hours than its result will be used for – in this implementation for 36 h. Beyond that horizon a new commitment will be in effect. Thus for times beyond the commitment horizon (e.g., hours 24–36), the normal non-anticipatory constraints on the generator statuses (i.e., the first stage decision variables) are not enforced. The commitment patterns may vary by scenario after the horizon, but all must be feasible.

5.2 *Evaluation*

Stochastic UCs must be evaluated against the benchmark of deterministic UC. In this work results are reported as *stochastic savings* in percentage terms, and defined as:

$$100 \cdot (C_{\text{deterministic}} - C_{\text{stochastic}}) / C_{\text{deterministic}}$$

As a lower bound on the solution, the perfect information case can also be considered — over the long term stochastic solutions can do no better than committing units while knowing the true value of the wind power for the next day. The cost of a deterministic or stochastic solution minus the perfect solution cost also serves as a measure of the potential for stochastic savings or the *cost of uncertainty*:

$$C_{\text{stochastic}} - C_{\text{perfect}}$$

5.2.1 *Margins of Error*

MIP solutions have a small margin of error due to the MIP gap, a solver stopping tolerance which is set at a small percentage of the objective function. Unless otherwise reported, this work uses a MIP gap of 0.01%. Solutions which differ by less percentage than this cannot be compared with certainty.

When comparing stochastic UC cost savings relative to the deterministic UC cost, the MIP gap creates margins of error in both directions. The deterministic solution cost may be lower by as much as the MIP gap, resulting in a lower stochastic savings. Alternately, the stochastic solution cost may be lower by as much as the MIP gap, resulting in a higher stochastic savings.

Part IV

EXPERIMENTAL RESULTS

What is the magnitude of savings that can be expected from implementing stochastic UC? Does stochastic UC become more worthwhile at high wind penetrations? Or, put another way, do savings increase as the amount of wind on the system increases?

Does the method used to create scenarios or the number of scenarios used influence stochastic savings? Does the same number of scenarios (and thus the same approximate computational cost) produce the same result across methods? Does a method's effectiveness increase dramatically with the more scenarios or are the improvements small?

This section examines these questions setting up a series of experiments based on stochastic UC and dispatch simulations of the ERCOT model system. The modeling methodology used to create the ERCOT test system is discussed in §6.1–6.2. The design of the rolling commitment and dispatch simulations is discussed in §6.3. In chapter 7, the results of the experiments are presented. Stochastic costs are measured against the cost of using a deterministic forecast. We close the chapter on results with an examination of computational times.

Chapter 6

EXPERIMENTAL DESIGN

A series of experiments were run to examine the computational costs and potential savings from stochastic scheduling using two different scenario creation methods. The experiments were carried out on a model of the ERCOT ISO. For some experiments wind forecasts and observations are scaled to simulate a system with a high wind penetration, in the range of 25–30% of energy produced by wind annually. To capture the longer term performance of stochastic scheduling, experiments are run for nine months¹, using real system load and wind data from 2012.

The focus of the experiments is on the effects of wind uncertainty. By using a historical period for simulation, we can examine how both deterministic wind forecast errors and stochastic scenario errors affect the costs of system operation. However, we do not attempt to estimate the true potential for stochastic UC cost savings in ERCOT. Nor do we attempt to implement an ISO-scale, production-ready stochastic UC implementation. The transmission system is neglected in the experiments; all generators and loads are assumed to be perfectly connected. There are no load forecast errors or generator outages in the experiments. For a discussion of these simplifications and their potential effects on the results, see §7.6.

As an introduction to the power system being modeled, Figure 6.1 shows a map of the generation in the ERCOT system based on the eGrid 2009 dataset, with plants sized by capacity. An overview of the total capacities for the ERCOT test system is shown in Table 6.1. The sections that follow describe the modeling of wind and thermal generation in ERCOT. The final section describes the commitment and dispatch processes used in the experiments.

6.1 Wind Modeling

Wind is modeled as an aggregate power source across an entire system, based on publicly available historical observations and forecasts. Forecasts were bias corrected based on historical training data, but otherwise assumed to be of good quality. Historical wind data for ERCOT was provided by 3Tier. Data was divided into training and testing sets, as described in Table 6.3. The distribution of observed wind for the testing set is shown in Figure 6.2.

¹Data availability for the wind forecasts limited the scope of the experiments to nine months. A full year or more of simulations would better capture the full extent of seasonal effects.

Table 6.1: ERCOT generator capacities and energy penetrations.

kind	capacity [MW]	energy penetration [%]
thermal	68,406	-
coal	20,414	-
natural gas	42,854	-
nuclear	5,138	-
wind	9,805	8.6
CHP	5,095	8.9

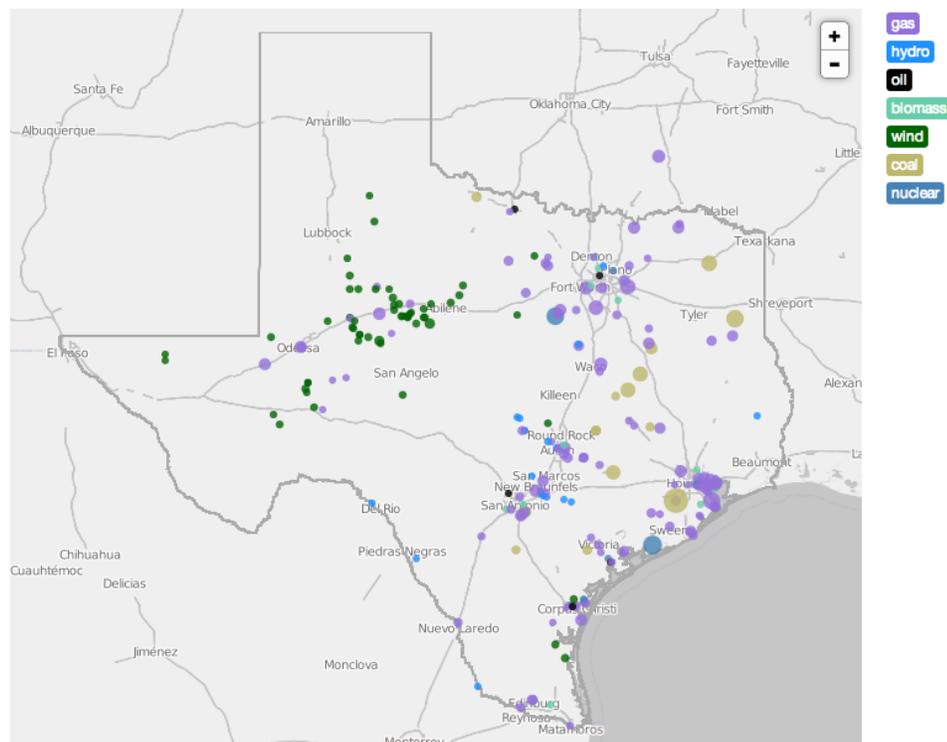


Figure 6.1: This map of ERCOT represents generating plants as dots which are sized by their nameplate capacities and colored by their fuel type. Many of the numerous wind farms in West Texas have a capacity of around 100 MW. The biggest plant is a 4000 MW coal plant just southwest of Houston. An interactive version is available at adamgreenhall.com/research/texas-map.

Set	Start	End
Training	2011-08-20	2011-12-31
Testing	2012-01-01	2012-09-30

Table 6.3: Training and testing sets for ERCOT wind data

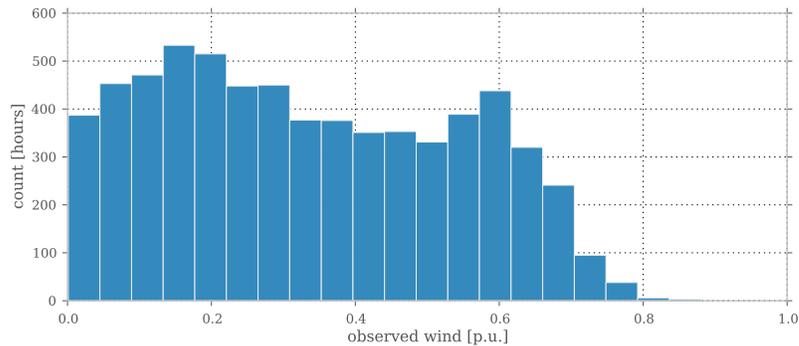


Figure 6.2: Capacity normalized wind distribution for the ERCOT testing set

6.1.1 Wind Energy Penetration Scaling

The ERCOT wind energy penetration for the testing period in 2012, at 8.6%, is comparatively low when set alongside most stochastic UC studies or even other control areas. The wind integration study on the Irish system [5, 9] had a wind energy penetration of about 34%. BPA in 2012 had a wind energy penetration of 18.8%².

For some of the following experiments, wind is scaled to create a higher wind energy penetration system. Scaling was done by multiplying both the wind forecast timeseries and the observed wind timeseries by the same factor: desired wind energy penetration divided by the 2012 wind energy penetration of 8.6%.

This simple scaling method produces forecast errors which have means and standard deviations reasonably consistent with those found in the 2008 ERCOT wind integration study [96, Table 4.1]. Simple scaling does result in larger minimum and maximum forecast errors than those in [96]. The different geographic distribution of future wind sites used in [96] produce a slight smoothing effect

²The BPA control area exports more energy than it uses internally. Wind energy penetration is only 8.6% (the same as ERCOT in 2012) if calculated based on internal demand plus interchange, rather than just internal demand.

on wind output and in turn reduce the forecast error extrema.

6.1.2 Scenario creation

Scenarios were constructed using both the moment matching method (§4.1) and analogs method (§4.2). Both methods use the capacity normalized forecasts and observations from the training data set (to create desired moments and find analogs, respectively). The number of scenarios N_s was varied, from as low as three scenarios up to thirty.

6.1.3 Caveats

For simplicity and data availability issues, these experiments use only aggregate wind power. Ideally the stochastic UC would take into account the constraints of the transmission system, especially in geographically large systems like ERCOT. The analogs scenario creation method could also benefit from the richer multi-variate weather forecast data (see §4.2.1) rather than the simple aggregate power analogs method used here.

There is also an implicit assumption in the analogs method that the forecasting method does not change significantly over the entire historical period or that historical re-forecasts are made upon adopting a new method. In the experiment that follows we use public aggregate forecast data from a major ISO, where forecasting methods are likely to be updated with some frequency.

An additional caveat is the potential effect that the length of the historical archive has on the quality of analog scenarios: short archives may not include rare events. The ERCOT dataset in these experiments has a particularly short training archive, less than one year. This short archive is likely to negatively affect analogs performance, particularly for the early spring and summer months, which have no similar season in the training dataset. This effect is somewhat alleviated by using an increasingly large archive, i.e., when selecting analogs for the date 2012-04-01, the historical archive consists of all prior dates – both the training set and the portion of the testing set up to 2012-03-31.

Finally, we note that the observed wind values in the dataset likely include some instances of wind shedding or curtailment. Curtailment events in the dataset reduce the accuracy of both scenario creation methods and mask wind events that were difficult for the power system to handle. This problem also affects forecasting algorithm training. If site-level wind speed observations are available, the real power output without curtailment can be modeled with reasonable accuracy. However, in this case data was not available to recreate a wind power dataset with the effects of curtailment removed.

6.2 Thermal System Modeling

Generator parameters are modeled from two data sets from the Environmental Protection Agency (EPA) — the eGrid dataset of generator information [97] and the Clean Air Markets dataset of observed heat and power timeseries [98]. This modeling is described in detail in Appendix C and also as an online demonstration, complete with a small amount of real data.

There are a few parameters, e.g., nuclear unit heat rates, which could not be estimated from data and had to be taken from the literature, see §C.3. Fuel costs are based on Energy Information Agency (EIA) average prices for 2012, see Table C.3.

Fast start units are defined as units which can be utilized in a realtime dispatch without being committed. A unit is labeled as fast starting if it is a gas turbine, or a combined cycle plant capable of ramping at least 50% of its capacity in one hour.

6.2.1 ERCOT system thermal model

The ERCOT power system model consists of 162 thermal generation units and a single aggregated Combined Heat and Power (CHP) unit which is treated as an uncontrollable fixed schedule. A summary of the thermal unit capacities is shown in Table 6.1. The overall system is fairly flexible due to a large proportion of natural gas units. The ERCOT model system has a capacity weighted flexibility index[99] of $F = 0.489$. There are 42 gas turbines and combined cycle units which are classified as fast-start units in the ERCOT model, a total fast-start capacity of about 8140 MW. Hydro units are few in number and in aggregate make up less than 500 MW in capacity in ERCOT³. Thus, in the model hydropower production is not modeled.

6.3 Power System Simulation

Power system scheduling simulation over the time period of many months is done by a rolling UC. That is, the solution for one day becomes the initial conditions for the next. The final cost of a commitment is evaluated after an additional step: real time dispatch. The goal of these modeling steps is to approximate the long term results of the day ahead and real time business and engineering processes used by system operators.

³All of the hydropower facilities in ERCOT are operated by municipal or public utilities. Many of these hydropower facilities are operated by the Lower Colorado River Authority (LCRA). The LCRA is a public utility which operates under a water management plan that allows it to generate electricity only as it releases water for downstream use, as part of its flood-management operations, or for grid emergencies.

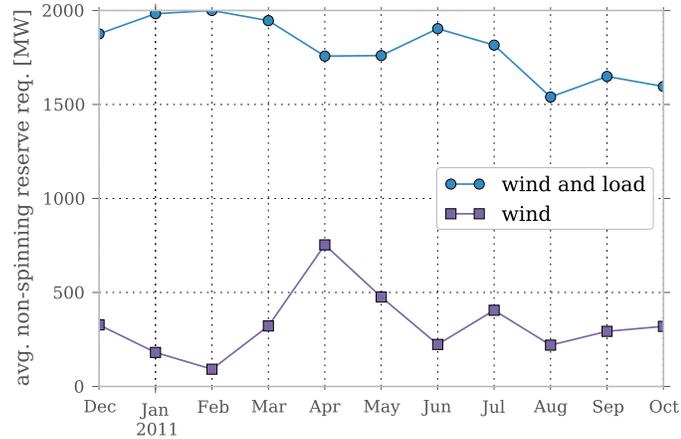


Figure 6.3: ERCOT separates contributions to its non-spinning reserve requirement into wind and load forecast error components. Non-spinning reserve in ERCOT is used to handle forecast uncertainty and to replenish spinning reserve. Data was taken from [101, Table 9].

6.3.1 Unit Commitment Simulation

The basic UC formulation and system parameters are described in §6.2 and in more detail in Appendix C. Commitments are based on 24 h periods, each of which is solved with a 36 h commitment, to avoid horizon effects. The stochastic UC formulation is described in §B.5.

Reserves

To focus on the effects of wind, there are no reserve requirements in the experiment. In operational practice, ERCOT handles day-ahead and hourly forecast error by setting a non-spinning reserve requirement based on the 95th percentile of historical load and wind forecast errors [24, 100, 101]. Non-spinning reserve must be able to within respond within 30 min. However, non-spinning reserve in the 2011 ERCOT system was in large part driven by load forecast error, as shown in Figure 6.3. Additionally the dispatch of non-spinning reserve to cover forecast errors happens on a timescale between day-ahead and real-time, with updated forecasts and decision making by the system operator. Reserves are not used in these simulations to handle wind uncertainty at the day ahead timescale to make the differences between deterministic and stochastic UC clear. In future work we hope to more completely represent the different time scales and decision making processes used to handle uncertainty.

Computational Parameters

The MIP solver stopping tolerance, called the MIP gap, was set to 0.01% of the linear programming (LP) optimal solution. Unless otherwise noted in the description of the experiment, solver time was limited to 1 h per day-long UC problem. This time limit was found to be long enough to allow most solutions to converge below the MIP gap, while allowing for reasonable experimental run times with large numbers of scenarios. If the solver failed to find a MIP solution that met the MIP gap criteria within this time limit, then the lowest cost feasible solution found by that time was recorded along with the resulting gap.

6.3.2 Dispatch Simulation

To estimate the actual system cost of implementing the day-ahead UC solution, a simulation of the real-time market (i.e., economic dispatch) was run for each day. The real-time dispatch problem is a bit different than the day-ahead UC problem. Wind generation is set to the observed value in place of the predicted. Generator statuses are fixed to the UC commitment solution pattern for all but the most flexible class of fast-start units. The dispatch problem is 24 h in length, as opposed to the 36 h commitment. The solution cost for this dispatch problem is the final system cost.

In the basic dispatch problem, all demand must be met at any cost. For the few days where the basic dispatch problem is infeasible due to wind high forecast error, the problem is rerun with the cost of unmet demand set to \$10,000 per MWh. This high cost representation in the objective function ensures that as little demand is shed as possible. To avoid a few occurrences having a disproportionate effect on the solution, the cost of unmet demand is not included in the solution costs. Instead, the counts and amounts of unmet demand events are reported alongside the generator costs in the results. This approach follows [5] and allows analysis of the results under many valuations of unmet demand.

Wind was modeled as having zero operational costs. Wind shedding is allowed, at no cost to the system unless otherwise noted. Wind shedding costs are examined in an independent experiment described in §7.4.

6.3.3 Implementation

Analog scenario creation and ERCOT generator modeling were implemented in Python, using the Pandas data analysis library. The moment matching scenarios were created using a non-linear optimization, solved using Coopr and the IPOPT solver.

The rolling commitment and dispatch problem is implemented using Minpower [102], a power systems optimization toolkit developed in the course of this work and described further in Appendix B. The core UC formulation in Minpower is thoroughly tested, a process described in §B.4. For these experiments Minpower uses the MIP solver CPLEX version 12.4. Each rolling UC experiment was run on a single node of the University of Washington’s Hyak cluster. One Hyak node is approximately as powerful as a desktop computer, with 16 GB of memory and eight 2 GHz dual-processor cores. Using Hyak allows multiple time-intensive experiments to be run in parallel.

6.4 Summary

This chapter describes the design of a series of stochastic UC experiments on a data-driven model of the ERCOT power system. Thermal generation in ERCOT is modeled based on publicly available data from the EPA. Wind power in ERCOT is modeled based on aggregate wind data from 2012. Having both forecasts and observations makes it possible to evaluate the dispatch cost of commitment decisions made based on real forecast errors. All experiments are based on nine months of rolling UC simulations, where the results of one day become the initial conditions for the next. The experiments are implemented using Minpower, an open-source power systems optimization toolkit created by the author.

Chapter 7

RESULTS

This chapter begins by examining the potential for wind forecast uncertainty savings on the ERCOT test system by comparing deterministic and perfect forecast UC costs. We then examine stochastic savings for the 25% and 30% wind penetration cases. The effectiveness of the moment matching and analogs methods are discussed in §7.2. Scenario error metrics are used to interpret the results in §7.3. The effects of pricing wind shedding are examined in §7.4. The chapter closes with a discussion of how computational times increase with the number of scenarios.

7.1 *Magnitude of Savings*

7.1.1 *Potential Savings from Perfect Forecasting*

Table 7.1 shows the UC cost of a perfect forecast relative to the UC cost of a deterministic forecast for different wind energy penetrations. Each cost difference is the result of two independent UC runs on the ERCOT 2012 test system. The deterministic forecast UC represents current practice. The perfect forecast provides a theoretical lower bound on the cost of UC under uncertainty. It is clear that at current wind penetration, the ERCOT test system would need a near perfect forecast to see any UC cost reductions. For higher wind penetrations, the cost of wind uncertainty — and potential savings — increase to moderate levels, about a 1% savings for the 30% penetration case. The relatively low savings potential can be partially attributed to the high flexibility of the ERCOT thermal generation mixture.

7.1.2 *Stochastic Savings*

Figure 7.1 presents stochastic savings for different scenario creation methods for the ERCOT test system with wind energy penetrations of 25% and 30%. Table 7.3 summarizes these results in tabular form. For the ERCOT model stochastic savings are small in percentage, but significant in dollars — in the 25% penetration case, a 0.2% savings represents about \$6M dollars over the nine month testing period. Stochastic savings of 0.2% are close to the lower end of the reported savings of 0.25–0.9% for the Irish wind integration system [5]¹. However, the level of perfect savings was higher in the Irish

¹See Table 1.4 for additional information on the results of previous stochastic UC studies.

Table 7.1: UC cost difference between perfect and deterministic forecasts for ERCOT model.

wind scaling	wind energy penetration [%]	perfect cost [\$B]	deterministic cost [\$B]	cost diff. [%]	demand not met [MWh]
1.0	8.6	4.57	4.57	0.05	0
1.8	15.0	4.20	4.21	0.26	0
2.9	25.0	3.68	3.71	0.85	0
3.5	30.0	3.46	3.50	1.09	2,090
4.7	40.0	3.12	3.18	1.75	8,523

system, 1.85% compared to 0.85% for ERCOT. This suggests that the moment matching method is capturing a larger portion of the available savings than the WILMAR scenario creation method used in [5].

However, stochastic cost savings decrease dramatically when wind energy penetration is increased from 25% to 30% for both moment matching and analogs methods. This is surprising, given the increase in perfect cost savings with increased penetration in Table 7.1. Looked at another way, a 0.2% stochastic savings out of the 0.85% perfect savings is a decent result for the 25% penetration case, capturing nearly a full quarter of the possible savings. But a 0.06% stochastic savings out of a more than 1% perfect savings is not a very good result for the 30% case.

The amount of demand not met does play a small role in the lower stochastic savings for the 30% penetration case. The deterministic UC fails to meet about 2GW of demand, while the moment matching method and ten analog cases meet all the demand. Table 7.5 shows the stochastic savings for the $N_s = 10$ scenario moment matching case with varying costs for unmet demand. If we assume that ERCOT would be able to engage demand response or emergency generation for cheaper than \$1000 per MWh, the stochastic savings for the 30% case remain smaller than for the 25% case.

To explain this unexpected result, Figure 7.2 shows the cumulative cost savings of stochastic solutions for a range of wind energy penetrations between 25–30%. This figure shows that the stochastic solutions can be significantly more expensive than the deterministic solutions through the winter and spring months, when demand is low. For the lower penetration cases, a few events in the peak demand summer season result in significant cost savings and an overall savings for the test period. However, in higher penetration cases many of these high savings events do not occur. For the analogs method, savings in the 25% case become increased costs in the 30% case.

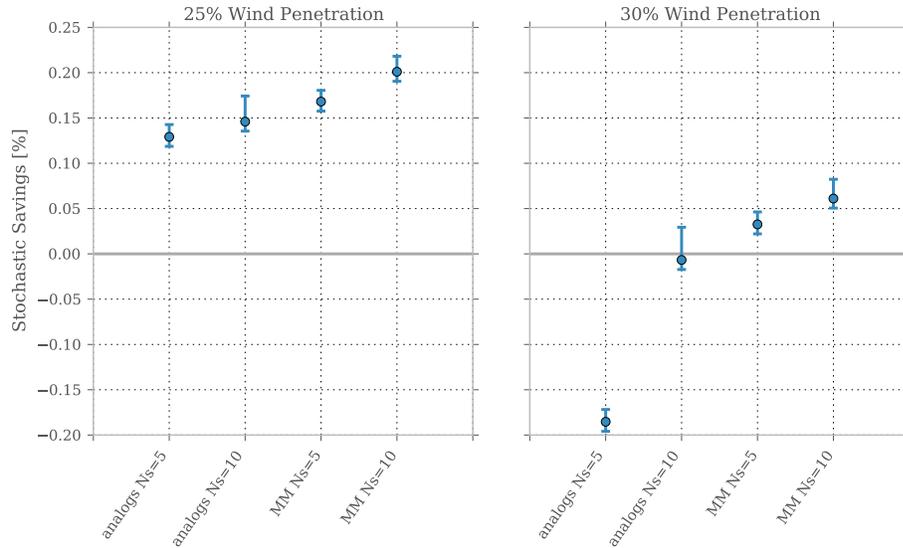
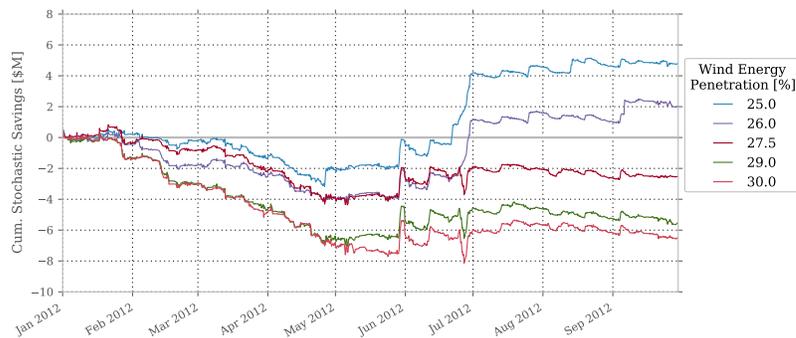


Figure 7.1: UC cost differences between deterministic and stochastic UC solution costs are compared. Stochastic savings is defined as the percentage difference between the total testing period stochastic and deterministic costs: $100 \cdot (C_{\text{deterministic}} - C_{\text{stochastic}}) / C_{\text{deterministic}}$. Error bars are based on the total MIP gap of the UC solutions. All solutions are for the ERCOT test system with wind energy penetrations of 25% and 30%.

Table 7.3: Summary of stochastic results for the ERCOT 25% and 30% wind penetration test cases. The stochastic savings are also shown in Figure 7.1.

WEP	method	N_s	cost [\$B]	uncertainty cost [\$M]	stochastic savings [%]	demand not met [MW]	realized WEP [%]
25	analog	5	3.704	26.7	0.13	671	24.6
		10	3.704	26.1	0.15	0	24.6
	MM	5	3.703	25.2	0.17	0	24.6
		10	3.702	24.0	0.20	0	24.6
30	analog	5	3.507	44.6	-0.19	2,654	28.7
		10	3.500	38.4	-0.01	0	28.7
	MM	5	3.499	37.0	0.03	0	28.8
		10	3.498	36.0	0.06	0	28.8



(a) Analogs



(b) Moment Matching

Figure 7.2: The cumulative stochastic savings using $N_s = 5$ scenarios is shown in millions of dollars, for several wind energy penetrations. For the analogs method, the net savings is only positive for the two lowest penetration cases.

As we will show in §7.3, stochastic UC saves the most compared to deterministic UC during periods of high demand and high forecast error. The basic explanation for lower stochastic savings at higher penetrations is that as wind penetration increases the lower net demand results in fewer opportunities for stochastic savings.

The seasonality of the savings in Figure 7.2 suggests that it may be possible to increase savings by selectively switching between the use of deterministic and stochastic UC, based on the predicted demand or marginal prices for the day. Deterministic UC would be used if demand is predicted to be low, while stochastic UC would be used only if high predicted demand indicated a potential for cost savings. Better quantification of this strategy is left for future research.

Table 7.5: Varying the cost of unmet demand met alters the stochastic savings for the 30% wind penetration simulation. The results shown are for the $N_s = 10$ scenario moment matching case.

cost of unmet demand [\$/MWh]	stochastic savings [%]
0	0.06
100	0.07
1,000	0.12
10,000	0.65

7.2 Scenario Creation Method Effectiveness

Stochastic savings differ significantly depending on the number of scenarios N_s . Figure 7.3 illustrates this effect over time for the ERCOT model 30% penetration case. It is clear that the number of scenarios can make a significant difference in the net savings. Also, individual events vary significantly in cost for different numbers of analogs. In contrast, different moment matching scenario sets produce much more similar costs, varying in the 30% case essentially only for one event near the end of July. This indicates that additional scenarios tend to be more similar to the existing set for the moment matching method than for the analogs method.

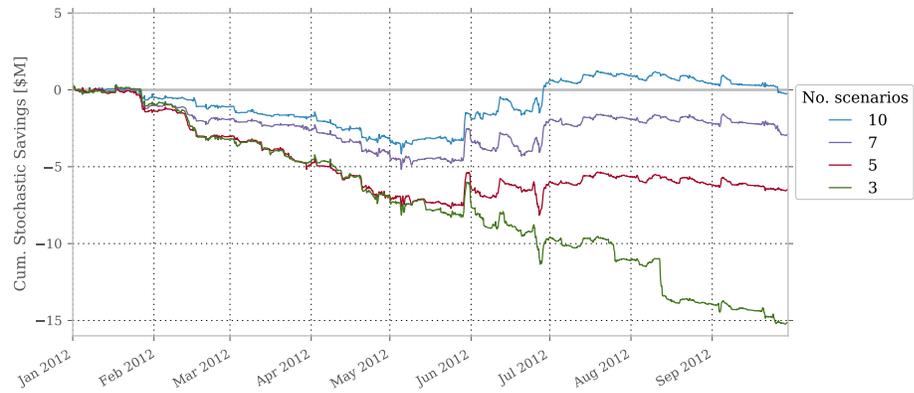
7.2.1 Analog Limitations

The simple analogs method is less effective at reducing costs than the moment matching method with a comparable number of scenarios. The $N_s = 5$ analogs method is even more costly than a deterministic UC for the 30% penetration case, suggesting that the choice of N_s is important and energy penetration dependent. In general these lackluster results show that simple analogs with a short historical archive are not a very effective scenario creation technique. A more ideal analogs method would have both a longer archive and a number of weather variables over which to search.

The difference in performance of the analogs method between the winter and spring months is also indicative of an archive that is too short. In Figure 7.3, all analogs results are significantly more expensive than deterministic for this period, while moment matching results are about equal. A spring season (where high winds are typical) is not present in the analogs historical archive; this limitation partially explains the high costs during this period.

One way to test the effect of these shortcomings is to use *perfect* analogs – where the observed values for the current day are used to search for similar observed values from the past. Contrast this to the simple analogs method of using the forecast to search for like historical forecasts and then taking the observed values from those dates as scenarios. The perfect analogs method requires perfect foresight, but is useful as a measure of sufficient archive length. If the archive is sufficiently long, there should be several very low distance past analogs for any possible observation. In such a case the perfect analogs method would produce scenarios which all closely matched the observed value and its performance will match perfect forecasting. However, if the archive is not sufficiently long the analogs will be less close, scenario spread will increase, and savings not meet the perfect forecast case.

The results for this experiment are tabulated in Table 7.7. As predicted for such a short historical archive, the perfect analogs method falls short of matching the performance of a perfect forecast,



(a) Analogs



(b) Moment Matching

Figure 7.3: The cumulative stochastic savings differs significantly for different numbers of scenarios. Results shown are for the ERCOT model with a wind energy penetration of 30%.

Table 7.7: Stochastic savings for the *perfect* analogs method are compared to simple analogs and the perfect forecast for the ERCOT test cases at 25% and 30% wind energy penetration. For the analogs methods $N_s = 5$ scenarios are used. Perfect analogs capture most but not all of the potential savings.

method	savings [%]	
	25%	30%
analogs	0.13	-0.19
perfect analogs	0.65	0.74
perfect	0.85	1.09

while significantly outperforming the simple analogs method. The perfect analogs method captures about 77% and 68% of the potential savings for the 25% and 30% penetration cases respectively. A longer historical archive would raise perfect analogs savings to nearly 100% of the potential savings and would likely raise simple analogs performance as well.

7.2.2 Moment Matching Limitations

Wind scenario creation using a moment matching method was observed to have the somewhat severe limitation of creating many duplicate scenarios. As only uniquely valued scenarios provide valuable information to the stochastic UC, these duplicate scenarios were aggregated. The result is an effective limit on the number of useful scenarios the method can create.

An experiment was run to test the effect of increasing N_s on the number of unique scenarios produced. The test was run on three randomly selected days, each from a different season. The number of scenarios was increased from $N_s = 2$ to $N_s = 42$; the moment matching algorithm was run once per number of scenarios per day. The number of unique scenarios, as well as the objective function and its components (the moment differences $D(M_i - \hat{M}_i)$) for this experiment are shown in Figure 7.4.

One might expect that with more scenarios a tighter fit to the higher order moments could be obtained, resulting in a lower objective function value. However, the plot shows that the objective function quickly reaches a plateau at about $N_s = 5$. The objective then is near constant, indicating that more scenarios do not create a better fit to the moments. The result of this is that some of the additional scenarios are duplicates, as shown in the top figure. The number of unique scenarios does not grow smoothly and varies by date, but it is clear that the average date will have less than N_s unique scenarios. In fact, when moment matching was run for the test case with $N_s = 30$, the average was just under 11 unique scenarios. A full table of unique scenario statistics is in Table 7.9.

The original proposal of moment matching discusses this phenomenon [7, see Section 4.1], labeling it as “underspecification” of the problem, and recommending adding additional constraints (i.e., additional moments). Another possible explanation comes from the mean and standard deviation’s quick convergence to their final values, and their dominance of the objective function. This suggests that a heavier weighting of the higher order moments might improve the solution behavior. However, when the heavier weighting for the higher order moments was attempted the solutions produced were not stable (i.e., either infeasible or feasible with highly unlikely scenario values). This instability is most likely due to the decreased relative weight of the mean, which the higher order moment calculations are dependent on. Better solution behavior may be possible, but only with significant effort in searching for the optimal selection of moments and tuning of the weights. The selection of additional, statistically meaningful, optimally weighted moments is left as future research.

In future stochastic UC implementations, improvements in computational power and solution methods will likely result in the ability to use large numbers of scenarios to further reduce costs.

In this environment, the moment matching method's limit of only ten to fifteen unique scenarios puts it at a disadvantage when compared with methods which can produce tens, hundreds, or even thousands of scenarios.

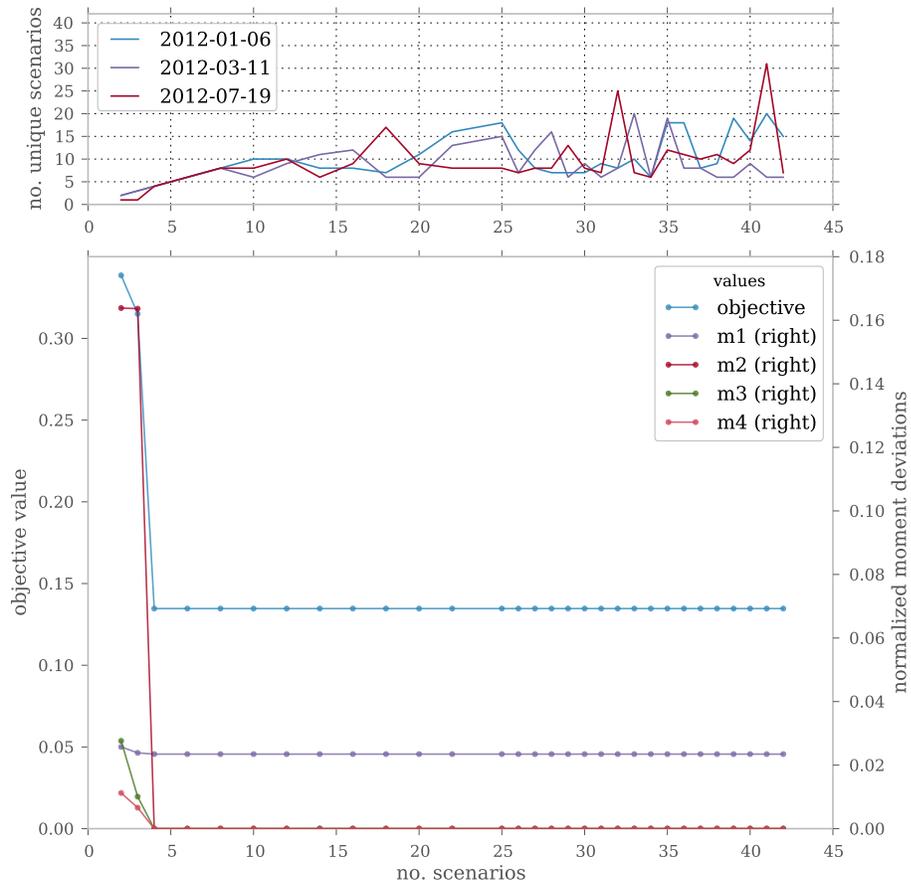


Figure 7.4: The moment matching method was used to produce scenarios for three days, with increasing number of scenarios. The upper plot shows the number of unique scenarios produced for each day. The bottom plot shows the objective function value (left axis) and the objective component (scaled moment deviation) values, averaged over the three days.

Table 7.9: The number of unique scenarios created by the moment matching method for a given day is generally less than the number of scenarios specified, N_s . Statistics on the number of unique scenarios for the ERCOT test case are given below.

N_s	mean	std. dev.	min.	max.
3	2.5	$8.9 \cdot 10^{-1}$	1	3
5	4.2	1.6	1	5
10	8.3	1.9	1	10
20	9.6	3.9	1	20
30	10.7	5.6	3	29

7.3 Scenario Metrics

Scenario metrics, introduced in §4.3, are a similar concept to deterministic forecast error; both provide a way of explaining UC results. Distributions of the metric values for the ERCOT test case are compared for the two scenario creation methods. Then we examine how metrics change with increasing N_s . Finally stochastic UC results are interpreted with their corresponding metrics.

7.3.1 Scenario Metric Distributions

Scenario mean error distributions are shown in Figure 7.5 for the ERCOT test case. The distribution of deterministic forecast error is shown for comparison and has a histogram that is more tightly clustered around zero than either the analogs or moment matching methods with $N_s = 5$. There are only small differences in the distributions of scenario mean error for moment matching and analogs.

Scenario spread is shown in Figure 7.6. The moment matching method produces a scenario spread distribution that is very different from that of the analogs method. The spread distribution of the analogs method roughly resembles the wind power distribution of the entire test case shown in Figure 6.2, but a bias toward lower magnitudes. This serves as a small validation of the analogs procedure – if two analog forecasts are good, the observed values for those analogs should be very close. For larger archives, scenario spread from taking the top N_s analogs should be smaller. For the moment matching method there is a distinct peak around 40% of capacity, which is likely the result of the algorithm trying to match the variance moment. The other peak at zero represents times when all scenarios have the same value. One hypothesis is that during these times other moments (e.g., the first difference) temporarily exerts greater influence on the objective than the variance moment.

Under and over prediction distributions are shown in Figure 7.7. Both scenario creation methods have a number of large under prediction values where the entire scenario set is at least 60% of capacity under the observed wind. Analog has more instances of small magnitude under/over predictions.

7.3.2 Scenario Metrics with N_s

One consideration for scenario creation methods is how much their performance improves with an increasing the number of scenarios. A summary of how the four metrics change with N_s

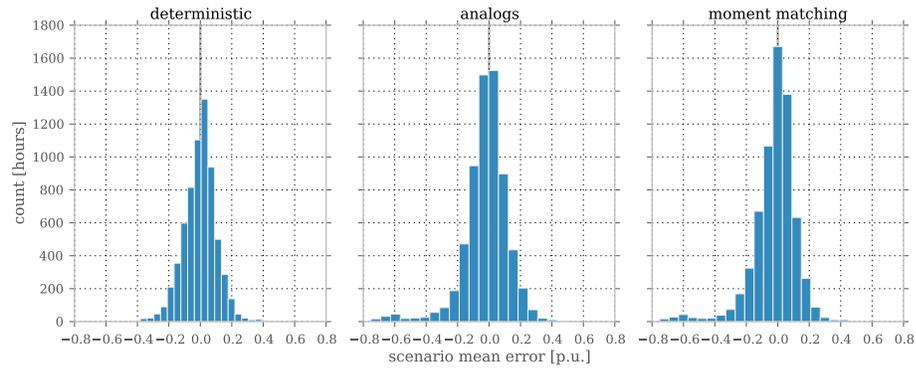


Figure 7.5: Hourly per unit scenario mean error for the ERCOT test case is shown as a set of histograms. Scenario mean error is similar to the deterministic forecast error, which is shown for comparison. For the analog and moment matching methods, error is based on $N_s = 5$ scenarios.

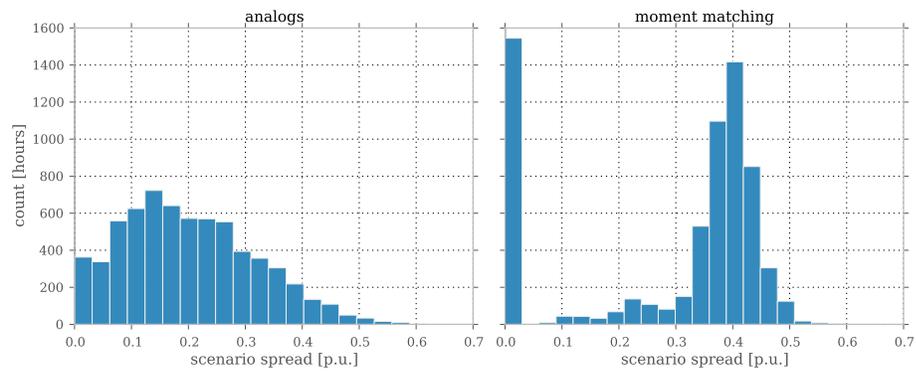


Figure 7.6: The distribution of hourly scenario spread is shown for $N_s = 5$. The values at zero for the moment matching method represent times when all scenarios have the same value.

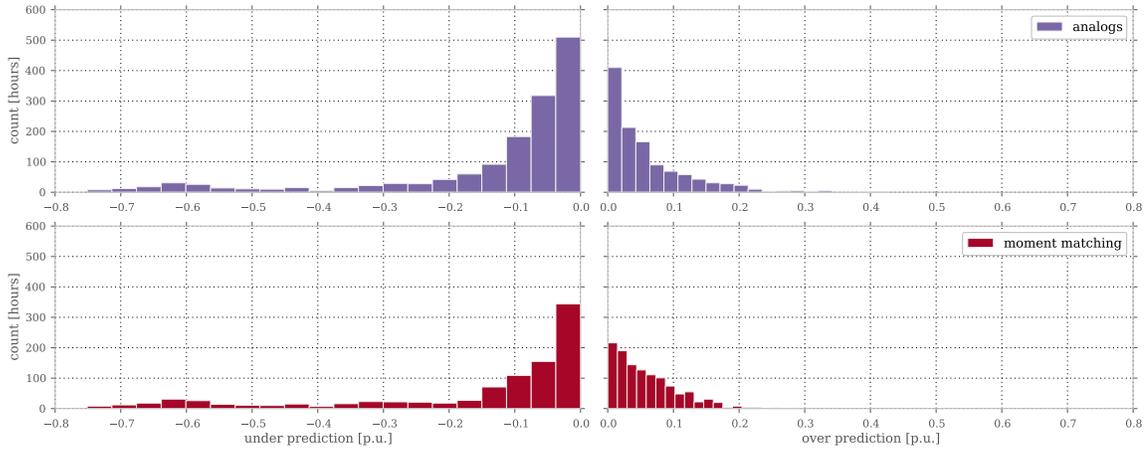


Figure 7.7: Under and over prediction distributions are shown for $N_s = 5$ scenarios. Both methods produce a number of cases where all scenarios are about 60% under the observed wind.

is shown in Figure 7.8. The moment matching method displays asymptotic behavior, reaching a final value around $N_s = 10$ scenarios. This behavior suggests that using more scenarios produced by the moment matching method as implemented will not reduce costs any further. An explanation for this behavior is discussed in §7.2.2. The analogs method improves metrics rapidly at first, and continues to improve most metrics more slowly for $N_s > 10$. The exception is the increase in mean error for $N_s = 30$, which suggests that analogs ranked greater than about thirty are too distant to be representative of the forecast for an archive of this size. This number of scenarios would likely be larger with a larger archive. An investigation of the limitations of the analogs method is discussed in §7.2.1.

7.3.3 Scenario Metrics Explain Performance

To see how scenario metrics can explain UC performance it is necessary to compare the metrics and the system cost. Demand is the major driver of costs, along with the costs of the units on the system – the cost effects of forecast and scenario errors are generally small compared to the effect of demand. As a starting point, Figure 7.9 shows the hourly costs of a perfect forecast UC solution relative to the demand. The shape of the system’s aggregate cost curve is roughly visible, spanning from about \$ 5–24 per MWh².

²These are average costs over the set of committed units – the marginal prices are in most cases much higher. The apparent cost floor at about \$200k in the figure is due to the fixed costs of the base load units.

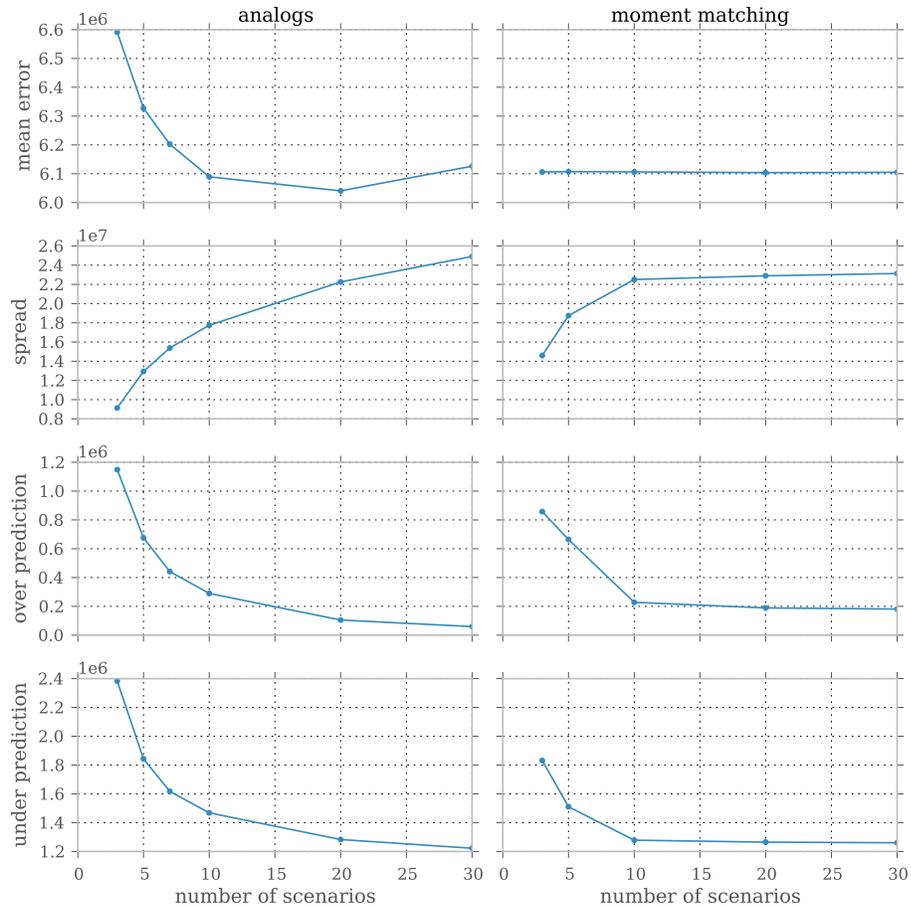


Figure 7.8: The absolute sum of scenario metrics (in MWh) is computed for different numbers of scenarios over the ERCOT test case at 2012 wind penetration levels. The scenario mean error and over/under prediction metrics decrease as more scenarios are added, while the scenario spread increases.

Adding wind uncertainty error increases the scatter around the perfect forecast cost curve. To see this effect, Figure 7.10 shows the difference in costs between the deterministic and perfect forecast UCs, or the *cost of uncertainty*, plotted against demand. Wind forecast error is denoted by color. Figure 7.12b shows an aggregated view of this plot, made by grouping the total uncertainty costs by demand and then by scenario mean error.

Moderate forecast errors of about $\pm 1\text{GW}$ contribute little to uncertainty costs. A number of higher cost are incurred during periods of moderately large under forecasts (less than -4GW and shown in light blue) and low to moderate loads – $P_D \leq 40\text{GW}$. This is likely due to the most expensive units being limited by a minimum power constraint. In order to respond to the higher than expected wind, the system must ramp down less expensive, but sufficiently flexible units before the higher cost limited units.

In the low demand periods there are also a few outliers where the deterministic cost is about \$200k cheaper than in the perfect case. These hours are often followed by a near \$200k positive outlier in the next hour (e.g., the optimization algorithm defers a unit startup for 1 h). A rolling mean of the cost difference over two hours greatly reduces the magnitude and number of negative outliers.

Periods of large over forecast error do not generally result in higher costs for periods of low demand ($P_D \leq 30\text{GW}$). But for the peak demand hours ($P_D \geq 40\text{GW}$), even moderate over forecasting generally results in significantly higher costs as high marginal cost peakers must be ramped up to cover the forecast error. Interestingly, there are no cases of large wind forecast errors during the peak hours where demand is above 60 GW in the test case. This is due to the seasonal properties of ERCOT demand and wind; demand peaks on summer afternoons when wind is generally low and forecast errors are small.

The same comparison is made for scenario mean error in Figure 7.11. Note that the leftmost chart is the same as Figure 7.10, but with a different color scale to account for the larger range of scenario mean errors. Similar to the deterministic forecast, analogs have scattering of high demand, high cost, high mean error hours. Moment matching lacks these cases, partially explaining its net savings for the 30% penetration ERCOT test case.

Scenario Metrics and Energy Penetration

Figure 7.12 breaks down the costs of uncertainty by demand and by scenario mean error. Low, medium and high demand groups are formed by cutting the hourly load into three quantiles, each with an equal number of hours. Similarly scenario mean errors are cut into three quantiles to form

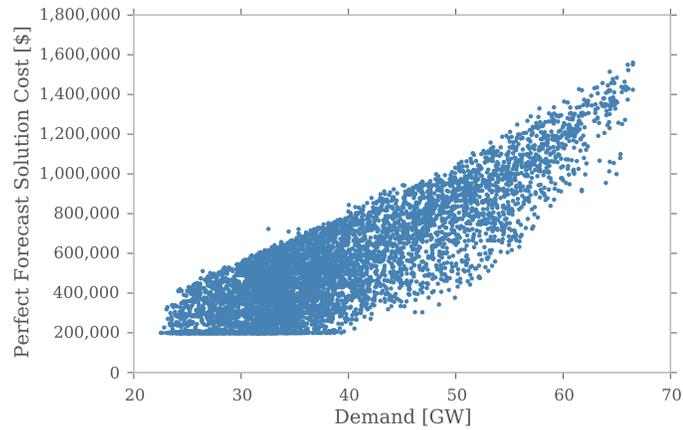


Figure 7.9: The total hourly cost of a perfect forecast UC is compared to system demand for the ERCOT 30% wind penetration test case. The shape of the aggregate system cost curve is roughly visible.

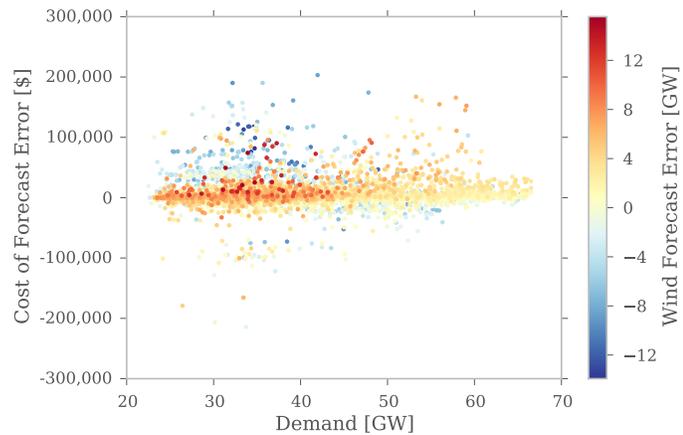


Figure 7.10: The hourly UC cost difference between a deterministic forecast and perfect forecast is shown against system demand for the ERCOT 30% wind penetration test case. Forecast errors are denoted by color – large under forecasts are colored blue, while large over forecasts are colored red.

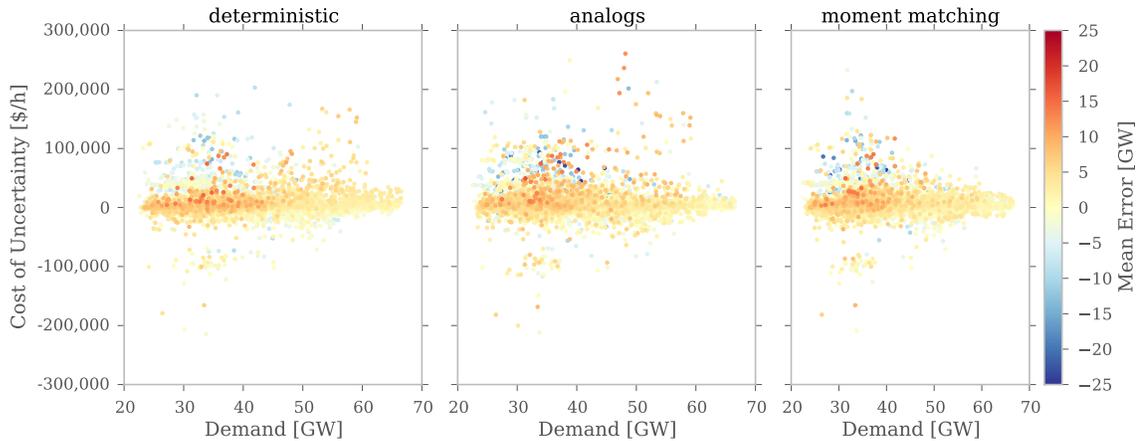
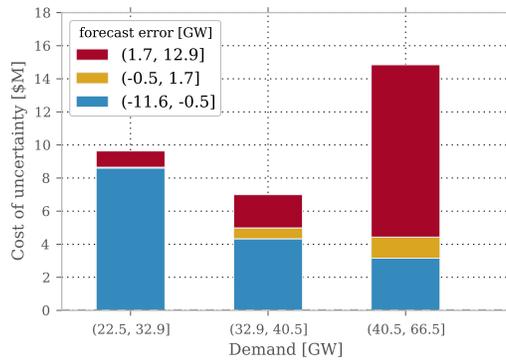


Figure 7.11: The hourly UC cost difference between a deterministic or stochastic case and the perfect forecast is shown against system demand for for the ERCOT 30% wind penetration test case. Scenario mean errors and deterministic forecast error are placed on the same scale and are denoted by color. The stochastic runs used $N_s = 5$ scenarios.

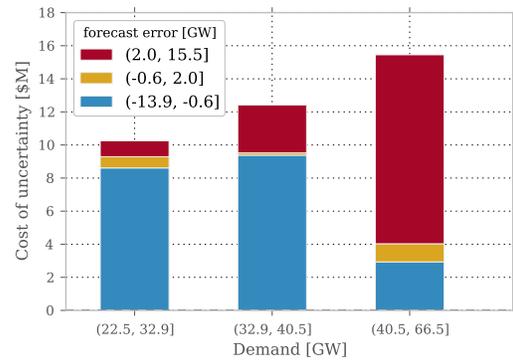
groups of where the mean error either under forecasts, has only a small error, or over forecasts. The resulting uncertainty cost breakdown is shown for both for both 25% and 30% penetration levels.

The largest uncertainty cost increases from 25–30% penetration levels are for the lowest (under forecast) mean error segment. For the analogs method the low and medium demand segments increase much more than the deterministic method, while the increase for the moment matching method is similar. For the largest demand segment, both stochastic methods have smaller uncertainty cost increases than the deterministic case.

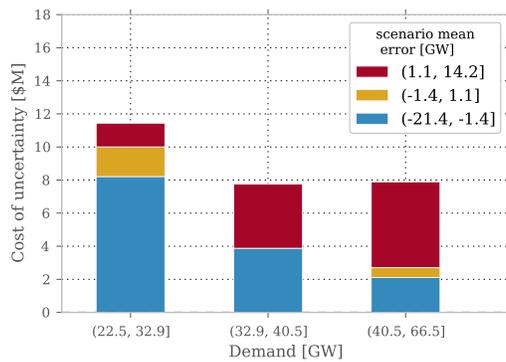
Periods of high mean error and high demand result in sharply increased costs, as additional supply in the real time market is very high cost. Stochastic savings are highest in these periods. Large negative mean errors during periods of low to medium demand also produce high uncertainty costs. These costs come in part from ramping down base load units which are higher cost but are not limited by their minimum power constraints to accommodate the unexpectedly high wind. However, these low-demand negative error periods are often just as expensive for stochastic UC as for deterministic UC. Stochastic savings are largely the result of high demand, high forecast error periods. When wind penetration is increased the net demand decreases, slightly shifting downward each point on the system cost curve and diminishing the best cases for stochastic savings. The net result is that stochastic UC is able to capture less of the potential savings for the 30% penetration case than for the 25% case.



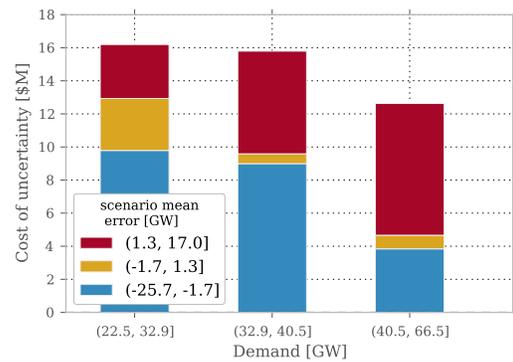
(a) Deterministic, 25% penetration



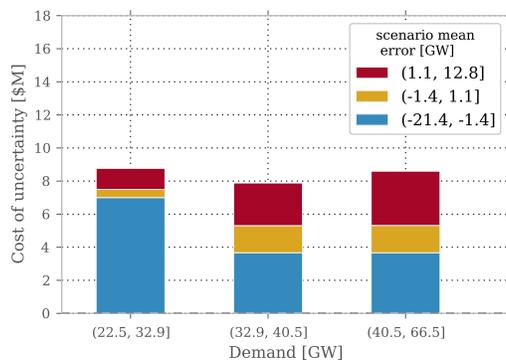
(b) Deterministic, 30% penetration



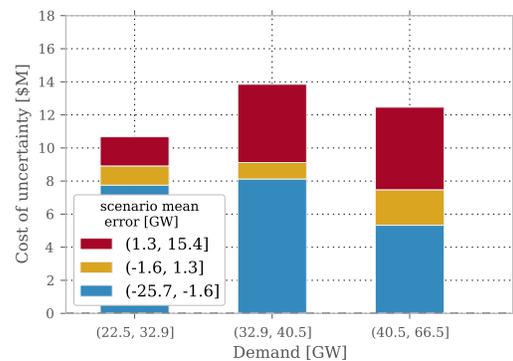
(c) Analogs, 25% penetration



(d) Analogs, 30% penetration



(e) Moment matching, 25% penetration



(f) Moment matching, 30% penetration

Figure 7.12: The cost of uncertainty is divided into quantiles of demand and scenario mean error. Quantiles are evenly sized by count rather than by bin width. Results shown are using the deterministic forecast UC and stochastic UC with $N_s = 5$, for the ERCOT model 25% and 30% wind penetration cases. In general negative mean errors (shown in blue) produce disproportionately large costs during periods low and medium demand. Small magnitude mean errors (shown in yellow) result in small costs. Large positive mean errors (shown in red) produce large costs in high demand periods, particularly for the deterministic method.

7.4 *Effect of Wind Shedding Cost*

One assumption in the experiments above is that the ISO can shed wind at zero cost. But depending on the regulation structure of the region, shed wind may have a cost. In the US the Production Tax Credit (PTC) often makes up a good portion of a wind producer's revenue stream. However, to be credited the \$23 per MWh PTC, that MWh must be used by the power system; shed power is not eligible for the PTC. Thus wind producers have fiscal incentive to avoid shedding and this pressure is often passed on to system operators. The most favorable policy to wind producers would be for the system operator to pay the PTC for any shed MWh. The costs of this policy are shown in Table 7.13 and compared to the zero cost case in Table 7.11.

As expected for an increase in the cost of shedding, less wind is shed but total system costs are higher, by about 0.7% for the deterministic case. With perfect foresight, this cost increase could be halved to 0.35%. Thus, the cost of uncertainty or potential savings increases to 1.21%. Stochastic savings also increase, up to 0.29% for the moment matching method. Also interesting is the reduction in demand not met for the analogs method from 672 MWh to zero. In this case the more flexible commitment pattern required to avoid shedding wind also avoids this instance of insufficient generation. In general it can be expected that pricing wind shedding will increase system uncertainty costs and thus increase stochastic savings.

7.5 *Computational Time*

Long solution times are often cited as a drawback of stochastic optimization in general and stochastic UC in particular. The box and whiskers plot in Figure 7.13 shows statistics for daily stochastic UC solution times over the 2012 test set for deterministic UC and stochastic UC with five, ten, and twenty scenario cases. For these experiments the MIP gap was set to 0.05% to reduce the solution times of the $N_s = 20$ tests. As a measure of solution quality, MIP gap statistics are shown in Table 7.15 for the days where the solver failed to reach the MIP gap within the 1 h time limit.

As Figure 7.13 shows, the difference in the median amount of time it takes for the solver to find a near optimal solution is significant. Medians range from 9 s for one scenario (deterministic UC), up to 507 s for ten scenarios, up to nearly the 1 h time limit for twenty scenarios.

From Table 7.15, we can see that while the deterministic case never exceeds the 1 h time limit, as the number of scenarios is increased the time limit is exceeded more often. The twenty scenario case exceeds the limit for nearly half of the 273 days in the test set. For these days which exceeded the time limit, the quality of the best solutions found, as measured by the MIP gap, is generally good for 5 and 10 scenarios, but degrades significantly for the 20 scenario case.

Table 7.11: Results with wind shedding priced at \$0 per MWh for the ERCOT 25% penetration test case, using $N_s = 5$ scenarios for the stochastic methods.

method	stochastic savings [%]	total cost [\$]	wind shed [MWh]	demand not met [MWh]
deterministic	0.00	$3.7091 \cdot 10^9$	1,105,930	0
analogs	0.13	$3.7043 \cdot 10^9$	1,119,082	672
moment matching	0.17	$3.7028 \cdot 10^9$	1,009,223	0
perfect	0.85	$3.6776 \cdot 10^9$	791,590	0

Table 7.13: Results with wind shedding priced at \$23 per MWh and an otherwise identical experimental setup to the results in Table 7.11.

method	stochastic savings [%]	total cost [\$]	wind shed [MWh]
deterministic	0.00	$3.7357 \cdot 10^9$	820,057
analogs	0.19	$3.7285 \cdot 10^9$	856,900
moment matching	0.29	$3.7249 \cdot 10^9$	668,370
perfect	1.21	$3.6905 \cdot 10^9$	475,273

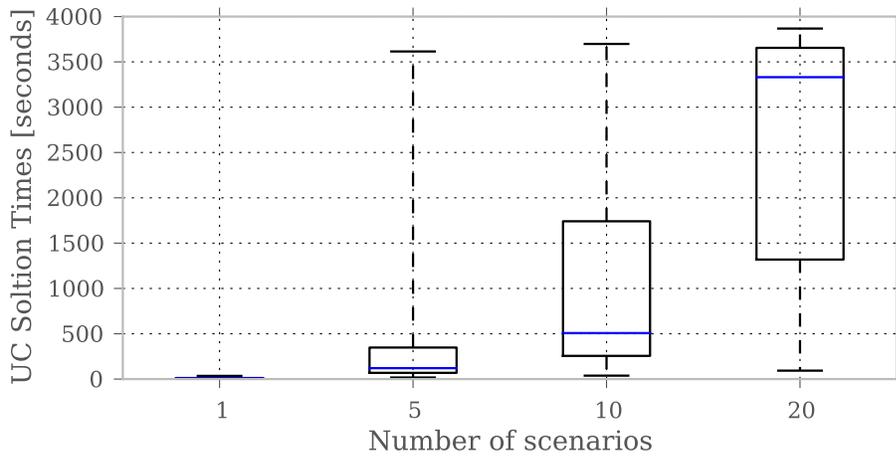


Figure 7.13: UC solution times are compared for different numbers of scenarios with a wind energy penetration of 30%. The one scenario set represents the deterministic forecast UC, while the five, ten, and twenty scenario sets are stochastic UC using analogs. The dashed whiskers represent the minimum and maximum solution times over the 2012 test set. Each box delimits the 25th and 75th percentiles, with the median solution time for the set shown in blue.

7.6 Potential Effects of Power System Modeling Assumptions

As discussed in Chapter 6, this model of the operations of the ERCOT system simplifies several components which are likely to affect the cost of operations. The experiments and results described in this chapter focus solely on the effects of day-ahead wind forecast uncertainty. These results should not be taken as a perfect estimation of total stochastic savings for the actual operations of the ERCOT ISO. In particular there are four factors which may alter solution costs: wind penetration scaling, load forecast errors, a constrained transmission system, and generation/transmission outages.

Wind scaling to increase system wind penetration is done in this work using a simple model, as discussed in §6.1.1. In particular this model results in larger forecast error extrema than a model which includes new wind installations which are more geographically distributed [96]. The exact effect this assumption has on costs depends on the particulars of the system and wind installations. On first pass, smaller magnitude forecast error extrema would likely reduce stochastic savings. However, for the ERCOT system we have seen that most stochastic savings come from over-forecast errors

Table 7.15: MIP gap statistics for solutions exceeding solver time limit.

N_s	# days	mean [%]	std. dev. [%]	max. [%]
1	0			
5	1	0.12		0.12
10	32	0.18	0.12	0.56
20	131	1.93	9.16	69.33

during high demand periods. Wind speeds and thus forecast errors are generally lower during the high-demand summer period in ERCOT — this period has a disproportionately low number of over-forecast error extrema. Thus, for ERCOT the effects of larger error extrema on stochastic savings may be less than expected.

Load forecast error was the focus of much of the early work in stochastic UC. This remains an important area today given the scale of modern ISOs (e.g., ERCOT’s peak hourly load in 2012 was 66 507 MW) and the widening scope of the electrical grid (e.g., electric vehicles and behind-the-meter renewable generation). Currently ERCOT allocates more reserves to handle load forecast error than wind forecast error (see Figure 6.3). This suggests that additional stochastic savings could be realized by modeling load forecast error stochastically.

Load forecast errors have traditionally been modeled as normal error processes [36, 65]. If load error is assumed to be independent of wind error, this would mean that the number of scenarios would become $N_{s,\text{load}} \cdot N_{s,\text{wind}}$. This increase in problem size would present a computational challenge to the current implementation. Alternately, the analogs method makes it possible to create scenarios for both load and wind. Given temperature observations from the analog dates, load forecast error could be modeled as a function of temperature forecast error. This joint scenario modeling would reduce the number of scenarios required, thus reducing solution times.

Transmission constraints play a significant role in the cost of energy in some power systems. A constrained transmission system increases the locational marginal price (LMP) for the areas affected. From an optimization perspective, transmission constraints create additional discontinuities

in the UC problem. Transmission constraints may require more expensive generation to compensate for forecast errors, raising the uncertainty costs and potentially increasing stochastic savings. In ERCOT, transmission constraints are most likely to become a factor in the peak-demand summer period, when stochastic savings are already high. However, the effects of transmission constraints on stochastic savings are likely to be highly system dependent. Incorporating the transmission system into the UC problem increases the problem size, potentially creating a significant computational challenge.

Generator and transmission outages are currently handled by spinning reserve requirements based on the $N - 1$ criteria. Wind and load forecast error are, in theory, handled separately by non-spinning reserve. Using spinning reserve to compensate for forecast errors would reduce the system's ability to handle an outage; mixing the two mechanisms for handling uncertainty should be discouraged unless other measures are put in place.

Stochastic UC has been used in the literature to model generator outages as a random phenomenon, instead of using a spinning reserve requirement [10, 11]. However, modeling outages introduces a huge number of scenarios. Currently a decomposition algorithm is necessary to solve the stochastic outages problem for any reasonably large system. Barring a unified method for handling power system uncertainty, it seems likely that stochastic UCs using MIP formulations will continue to deal with outages using a reserve criteria.

Part V

CONCLUSIONS

This research examines two wind scenario creation methods, moment matching and *simple* analogs. Both methods allow a scheduler to take a deterministic forecast for wind power and leverage historical data to create many scenarios. The moment matching method does this by synthesizing scenarios to match a set of statistics. The analogs method looks back in time to find similar forecasts and uses the observations from those analogous dates directly as scenarios.

These scenario creation methods are tested on a model of the ERCOT ISO based on historical data from 2012. Stochastic unit commitment, using the two methods and different numbers of scenarios, is compared to deterministic forecast UC and also to a perfect forecast. Wind is scaled to estimate the future effectiveness of stochastic UC on high wind penetration systems.

Findings

The costs of uncertainty in the ERCOT test system are relatively small due to the proportion of flexible natural gas units on the system. Even at 30% wind penetration, 3.5 times the current amount, a perfect forecast saved only about 1% or \$34 million over the nine month test period. Stochastic UC is able to capture a modest portion of these savings at 25% penetration.

However, stochastic savings may decline as wind penetration increases. Despite an increase in potential savings, stochastic savings declined significantly for all methods on the ERCOT test system as wind penetration moved from 25% to 30%. This unexpected result is due in part to the reduction in net demand, which diminishes the cases with the highest potential for stochastic savings – high demand periods with large over-forecast errors.

The number of scenarios is an important parameter and can strongly affect savings. A “good” value for N_s is likely to be system and wind penetration dependent.

Similarly, the method of scenario creation strongly affected savings. The moment matching method resulted in higher savings than the analogs method for this dataset. This result may not hold if more historical data is available, weather data is used to create analogs, or more scenarios are used and the moment matching method continues to produce duplicate scenarios.

Stochastic UC may cost more than deterministic UC for extended periods – generally when

demand and marginal prices are low. Net savings are often largely the result of a few high cost events during peak load periods.

A large proportion of the total uncertainty costs also comes from under forecasts during low to medium demand periods. During these periods many baseload units may be unable to ramp down due to minimum power constraints, resulting in high costs from under forecasts. However, stochastic UC does not perform better than deterministic UC for these cases.

Computational times for stochastic UC are higher than for deterministic UC, but are still reasonable with 2012-era desktop computing hardware and software. For the ERCOT model with about 160 units (but without transmission or security constraints) MIP solver run times of under one hour were easily attainable for up to $N_s = 10$.

Contributions

This research advances our knowledge of the intersection between power systems operations, renewables forecasting, and stochastic optimization. Specifically this work:

1. proposes the *simple* analogs method for scenario creation;
2. proposes a method for estimating the desired moments, a key component of the moment matching method;
3. lays out a framework for testing stochastic UC relative to both deterministic and perfect forecasts; and
4. evaluates stochastic UC performance on a large scale, realistic power system model with two scenario creation methods, several values of N_s , and several wind penetrations.

The experiments in this work involved the implementation of several pieces of software. In particular two of these components were designed to be open source and are available for use and improvement by future researchers:

1. a methodology for modeling US power systems based on publicly available historical data; and
2. Minpower, a thoroughly-tested toolkit for power systems optimization.

Future Research

Scenario creation and stochastic UC remains an open area for research. Many of the methods presented in this dissertation could be improved upon to produce better scenarios and to attain stochastic UC solutions with lower costs in less time. This section outlines areas for future work that may lead to such improvements.

Improvements in the moment matching method may be possible with better tuned moment weights (w_i), additional moments, or a penalization of similar scenarios. A moment matching method which produced more unique scenarios would enable lower cost solutions for an increased number of scenarios N_s .

Analog scenario creation has even more potential for improvement. Weather data could be used to match analogs instead of simple aggregate power data. This would likely result in analogs which are much closer to the observed value. The analogs distance metric could be tuned to wind ramping as well as energy differences. The distance metric might also deemphasize differences at longer time horizons. A better method for selecting scenarios from the set of all historical analogs could improve inclusion of rare events. Research to determine the necessary historical archive length would enable higher quality analog scenarios.

Both scenario creation methods have the potential to be applied to other sources of uncertainty in the power system. Hydro inflows in particular might benefit from a synthesis of long-term forecasting models and the analogs method.

The seasonality of the stochastic UC results suggest that in some power systems a hybrid deterministic-stochastic model may yield higher savings than either method alone. A switching strategy could be based on the predicted demand or marginal prices for the day. Deterministic UC would be used if demand is predicted to be low, while stochastic UC would be used only if high predicted demand indicated a potential for cost savings.

The formulation of the stochastic UC can affect the run-time and limit the number of scenarios, potentially decreasing the savings from using stochastic UC. Work done to improve UC formulations [19, 20] could be applied to stochastic UC. A rolling method UC method [5] could be tuned to the particular system to better capture the stochastic time frame and lower costs. A CVaR based objective function [14, 46] could be used to improve the risk balance between rare and frequent events and potentially decrease overall stochastic solution costs.

Finally, savings due to stochastic UC are likely to vary from system to system and with wind energy penetration. The number of scenarios used may play a large role in the commitment costs. Evaluations of the potential for stochastic savings should be made for other systems, especially those with unique flexibility requirements.

Part VI

BIBLIOGRAPHY AND APPENDICES

BIBLIOGRAPHY

- [1] J. Rogers and K. Porter, “Central Wind Power Forecasting Programs in North America by Regional Transmission Organizations and Electric Utilities : Revised Edition 2008 to 2010 Central Wind Power Forecasting Programs in North America by Regional Transmission Organizations and Elect”, NREL, Tech. Rep. March, 2011.
- [2] ERCOT. (Mar. 2012), [Online]. Available: http://www.ercot.com/news/press_releases/show/495.
- [3] X. Wang, personal communication, 2012.
- [4] J. Morales, A. Conejo, and J. Perez-Ruiz, “Short-term trading for a wind power producer”, *IEEE Transactions on Power Systems*, vol. 25, no. 1, pp. 554–564, Feb. 2010.
- [5] A. Tuohy, P. Meibom, E. Denny, and M. O’Malley, “Unit Commitment for Systems With Significant Wind Penetration”, *IEEE Transactions on Power Systems*, vol. 24, no. 2, pp. 592–601, May 2009.
- [6] L. Jones, “Strategies and Decision Support Systems for Integrating Variable Energy Resources in Control Centers for Reliable Grid Operations”, Alstom Grid, Tech. Rep., 2012.
- [7] K. Høyland and S. Wallace, “Generating scenario trees for multistage decision problems”, *Management Science*, vol. 47, no. 2, pp. 295–307, Feb. 2001.
- [8] Y. Feng and S. M. Ryan, “Scenario construction and reduction applied to stochastic power generation expansion planning”, *Computers and Operations Research*, vol. 40, no. 1, pp. 9–23, Jan. 2013.
- [9] C. Lowery and M. O’Malley, “Impact of Wind Forecast Error Statistics Upon Unit Commitment”, *IEEE Transactions on Sustainable Energy*, vol. 3, no. 4, pp. 760–768, Oct. 2012.
- [10] P. Carpentier, G. Gohen, J.-C. J. Culioli, and A. Renaud, “Stochastic optimization of unit commitment: a new decomposition framework”, *IEEE Transactions on Power Systems*, vol. 11, no. 2, pp. 1067–1073, May 1996.
- [11] S. Takriti, J. Birge, and E. Long, “A stochastic model for the unit commitment problem”, *IEEE Transactions on Power Systems*, vol. 11, no. 3, pp. 1497–1508, 1996.
- [12] R. Barth, H. Brand, P. Meibom, and C. Weber, “A Stochastic Unit-commitment Model for the Evaluation of the Impacts of Integration of Large Amounts of Intermittent Wind Power”, *2006 International Conference on Probabilistic Methods Applied to Power Systems*, pp. 1–8, Jun. 2006.
- [13] P. Meibom, R. Barth, B. Hasche, H. Brand, C. Weber, and M. O’Malley, “Stochastic optimization model to study the operational impacts of high wind penetrations in Ireland”, *IEEE Transactions on Power Systems*, vol. 26, no. 3, 2011.

- [14] A. Al-Awami and M. El-Sharkawi, “Coordinated Trading of Wind and Thermal Energy”, *IEEE Transactions on Sustainable Energy*, 2011.
- [15] J. Restrepo and F. Galiana, “Assessing the yearly impact of wind power through a new hybrid deterministic/stochastic unit commitment”, *IEEE Transactions on Power Systems*, vol. 26, no. 1, 2011.
- [16] E. M. Constantinescu, V. M. Zavala, M. Rocklin, S. Lee, and M. Anitescu, “A Computational Framework for Uncertainty Quantification and Stochastic Optimization in Unit Commitment With Wind Power Generation”, *IEEE Transactions on Power Systems*, vol. 26, no. 1, pp. 431–441, Feb. 2011.
- [17] J. Wang, M. Shahidehpour, and Z. Li, “Security-constrained unit commitment with volatile wind power generation”, *IEEE Transactions on Power Systems*, vol. 23, no. 3, pp. 1319–1327, 2008.
- [18] M. Carrión and J. Arroyo, “A computationally efficient mixed-integer linear formulation for the thermal unit commitment problem”, *IEEE Transactions on Power Systems*, vol. 21, no. 3, Aug. 2006.
- [19] J. Ostrowski, M. F. Anjos, and A. Vannelli, “Tight Mixed Integer Linear Programming Formulations for the Unit Commitment Problem”, *IEEE Transactions on Power Systems*, vol. 27, no. 1, pp. 39–46, Feb. 2012.
- [20] G. Morales-Espana, J. M. Latorre, and A. Ramos, “Tight and Compact MILP Formulation for the Thermal Unit Commitment Problem”, *IEEE Transactions on Power Systems*, pp. 1–12, 2013.
- [21] R. E. Bixby, M. Fenelon, Z. Gu, E. Rothberg, and R. Wunderling, “Mip: theory and practice—closing the gap”, *System Modelling and Optimization: Methods, Theory and Applications*, vol. 174, pp. 19–49, 2000.
- [22] F. Bouffard, “Stochastic security for operations planning with significant wind power generation”, *IEEE Transactions on Power Systems*, vol. 23, no. 2, 2008.
- [23] G. Liu and K. Tomsovic, “Quantifying Spinning Reserve in Systems With Significant Wind Power Penetration”, *IEEE Transactions on Power Systems*, vol. 27, no. 4, pp. 2385–2393, Nov. 2012.
- [24] M. a. Matos and R. J. Bessa, “Setting the Operating Reserve Using Probabilistic Wind Power Forecasts”, *IEEE Transactions on Power Systems*, vol. 26, no. 2, pp. 594–603, May 2011.
- [25] M. A. Ortega-Vazquez and D. S. Kirschen, “Estimating the spinning reserve requirements in systems with significant wind power generation penetration”, in *Power Energy Society General Meeting*, Jul. 2009.
- [26] B. Palmintier, “Incorporating operational flexibility into electric generation planning”, PhD thesis, Massachusetts Institute of Technology, 2013.

- [27] L. Garver, "Power generation scheduling by integer programming-development of theory", *Power Apparatus and Systems, Part III. Transactions of the American Institute of Electrical Engineers*, vol. 81, no. 3, pp. 730–734, 1962.
- [28] T. Dillon, K. W. Edwin, H.-D. Kochs, and R. J. Taud, "Integer programming approach to the problem of optimal unit commitment with probabilistic reserve determination", *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-97, no. 6, pp. 2154–2166, 1978.
- [29] M. Pereira and L. Pinto, "Application of Decomposition Techniques to the Mid - and Short - Term Scheduling of Hydrothermal Systems", *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-102, no. 11, pp. 3611–3618, Nov. 1983.
- [30] M. Pereira, "Optimal scheduling of hydrothermal systems-an overview", in *IFAC Symposium on Planning and Operation of Electric Energy Systems*, 1985.
- [31] S. M. Amado and C. C. Ribeiro, "Short-Term Generation Scheduling of Hydraulic Multi-Reservoir Multi-Area Interconnected Systems", *IEEE Transactions on Power Systems*, vol. 2, no. 3, pp. 758–763, 1987.
- [32] T. Ohishi, S. Soares, and M. de Carvalho, "A short term hydrothermal scheduling approach for dominantly hydro systems", *IEEE Transactions on Power Systems*, vol. 6, no. 2, pp. 637–643, May 1991.
- [33] C. R. Gagnon and J. F. Bolton, "Optimal Hydro Scheduling at the Bonneville Power Administration", *Transactions on Power Apparatus and Systems*, vol. PAS-97, no. 3, pp. 772–776, May 1978.
- [34] H. Habibollahzadeh and J. A. Bubenko, "Application of Decomposition Techniques to Short-Term Operation Planning of Hydrothermal Power System", *IEEE Transactions on Power Systems*, vol. 1, no. 1, pp. 41–47, 1986.
- [35] H. Brannlund, J. A. Bubenko, D. Sjelvgren, and N. Andersson, "Optimal Short Term Operation Planning of a Large Hydrothermal Power System Based on a Nonlinear Network Flow Concept", *IEEE Transactions on Power Systems*, vol. 1, no. 4, pp. 75–81, 1986.
- [36] L. Soder, "Reserve margin planning in a wind-hydro-thermal power system", *IEEE Transactions on Power Systems*, vol. 8, no. 2, pp. 564–571, May 1993.
- [37] O. Nilsson and D. Sjelvgren, "Mixed-integer programming applied to short-term planning of a hydrothermal system", in *Proceedings of Power Industry Computer Applications Conference*, IEEE, 1995, pp. 158–163.
- [38] E. Castronuovo and J. Lopes, "On the optimization of the daily operation of a wind-hydro power plant", *IEEE Transactions on Power Systems*, vol. 19, no. 3, pp. 1599–1606, 2004.

- [39] G. Chang and C. Su, “A practical mixed integer linear programming-based short-term hydro scheduling”, in *IEEE/PES Transmission and Distribution Conference and Exhibition*, vol. 3, IEEE, 2002, pp. 1606–1610.
- [40] J. Garcia-Gonzalez, E. Parrilla, J. Barquin, J. Alonso, a. Saiz-Chicharro, and a. Gonzalez, “Under-relaxed iterative procedure for feasible short-term scheduling of a hydro chain”, *2003 IEEE Bologna Power Tech Conference Proceedings*, vol. 2, pp. 934–939, 2003.
- [41] A. Borghetti, C. D’Ambrosio, A. Lodi, and S. Martello, “An MILP Approach for Short-Term Hydro Scheduling and Unit Commitment With Head-Dependent Reservoir”, *IEEE Transactions on Power Systems*, vol. 23, no. 3, pp. 1115–1124, Aug. 2008.
- [42] C. D’Ambrosio, A. Lodi, and S. Martello, “Piecewise linear approximation of functions of two variables in MILP models”, *Operations Research Letters*, vol. 38, no. 1, pp. 39–46, Jan. 2010.
- [43] M. Piekutowski, T. Litwinowicz, and R. Frowd, “Optimal short-term scheduling for a large-scale cascaded hydro system”, *IEEE Transactions on Power Systems*, vol. 9, no. 2, pp. 805–811, May 1994.
- [44] a.T. Dasigenis and a.R. Garcia-San Pedro, “Real-time hydro coordination and economic hydro optimization”, in *Power Industry Computer Applications*, IEEE, 1994, pp. 150–157.
- [45] R. J. Wets, “Stochastic programming models: wait-and-see versus here-and-now”, in *Decision Making Under Uncertainty*, Springer, 2002, pp. 1–15.
- [46] R. Rockafellar and S. Uryasev, “Optimization of conditional value-at-risk”, *Journal of risk*, vol. 2, pp. 21–42, 2000.
- [47] R. Jabr, “Robust Self-Scheduling Under Price Uncertainty Using Conditional Value-at-Risk”, *IEEE Transactions on Power Systems*, vol. 20, no. 4, pp. 1852–1858, Nov. 2005.
- [48] S. Cerisola, a. Baillo, J. M. Fernandez-Lopez, a. Ramos, and R. Gollmer, “Stochastic Power Generation Unit Commitment in Electricity Markets: A Novel Formulation and a Comparison of Solution Methods”, *Operations Research*, vol. 57, no. 1, pp. 32–46, Jan. 2009.
- [49] R. Rockafellar and R. Wets, “Scenarios and policy aggregation in optimization under uncertainty”, *Mathematics of operations research*, vol. 16, no. 1, pp. 119–147, 1991.
- [50] W. E. Woodruff, J.-P. Watson, and D. L. Hart, “PySP : Modeling and Solving Stochastic Programs in Python”, *Mathematical Programming Computation*, 2011.
- [51] M. V. F. Pereira and L. M. V. G. Pinto, “Multi-stage stochastic optimization applied to energy planning”, *Mathematical Programming*, vol. 52, no. 1-3, pp. 359–375, May 1991.
- [52] C. C. Carøe and R. Schultz, “A Two-Stage Stochastic Program for Unit Commitment Under Uncertainty in a Hydro-Thermal Power System”, in *Konrad-Zuse-Zentrum fur Informationstechnik*, 1998.
- [53] J. Dupačová, N. Gröwe-Kuska, and W. Römisch, “Scenario reduction in stochastic programming: An approach using probability metrics”, *Mathematical Programming*, 2003.

- [54] Argonne National Laboratory, “Wind Power Forecasting: State-of-the-Art 2009”, Argonne National Laboratory, Tech. Rep., 2009.
- [55] A. Botterud, Z. Zhou, J. Wang, R. Bessa, H. Keko, J. Sumaili, and V. Miranda, “Wind power forecasting, unit commitment, and electricity market operations”, in *Power and Energy Society General Meeting*, IEEE, 2011.
- [56] Y. V. Makarov, P. V. Etingov, Z. Huang, J. Ma, B. B. Chakrabarti, K. Subbarao, C. Loutan, and R. T. Guttromson, “Incorporating wind generation and load forecast uncertainties into power grid operations”, PNNL, Tech. Rep. January, Apr. 2010.
- [57] G. Giebel, R. Brownsword, G. Kariniotakis, M. Denhard, and C. Draxl, “The State of the Art in Short-Term Prediction of Wind Power”, ANEMOS Project, Tech. Rep., 2011.
- [58] B. Brown, R. Katz, and A. H. Murphy, “Time series models to simulate and forecast wind speed and wind power”, *Journal of Climate and Applied Meteorology*, vol. 23, 1984.
- [59] T. Gneiting, K. Larson, K. Westrick, M. G. Genton, E. Aldrich, and T. Report, “Calibrated Probabilistic Forecasting at the Stateline Wind Energy Center : The Regime-Switching Space-Time (RST) Method”, University of Washington, Tech. Rep., 2004, pp. 1–27.
- [60] P. Pinson and G. Kariniotakis, “Conditional prediction intervals of wind power generation”, *IEEE Transactions on Power Systems*, vol. 25, no. 4, pp. 1845–1856, 2010.
- [61] U. Focken, M. Lange, K. Mönnich, H. Waldl, H. Beyer, and A. Luig, “A statistical analysis of the reduction of the wind power prediction error by spatial smoothing effects”, *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 90, no. July 2001, pp. 231–246, 2002.
- [62] J. Morales, R. Míguez, and a.J. Conejo, “A methodology to generate statistically dependent wind speed scenarios”, *Applied Energy*, vol. 87, no. 3, pp. 843–855, Mar. 2010.
- [63] T. Thorarinsdottir and T. Gneiting, “Probabilistic forecasts of wind speed: ensemble model output statistics by using heteroscedastic censored regression”, *Journal of the Royal Statistical Society*, vol. 173, no. 2, p. 1, 2010.
- [64] Rüdiger Barth, L. Söder, C. Weber, H. Brand, and D. J. Swider, “Documentation Methodology of the WILMAR Scenario Tree Tool”, WILMAR, Tech. Rep. January 2006, 2006.
- [65] G. Gross and F. Galiana, “Short-term load forecasting”, *Proceedings of the IEEE*, vol. 75, no. 12, pp. 1558–1573, 1987.
- [66] H. Bludszweit, J. A. Dominguez-Navarro, and A. Llombart, “Statistical Analysis of Wind Power Forecast Error”, *IEEE Transactions on Power Systems*, vol. 23, no. 3, pp. 983–991, Aug. 2008.
- [67] S. Tewari, C. J. Geyer, and N. Mohan, “A Statistical Model for Wind Power Forecast Error and its Application to the Estimation of Penalties in Liberalized Markets”, *IEEE Transactions on Power Systems*, vol. 26, no. 4, pp. 2031–2039, Nov. 2011.

- [68] P. E. Mello, N. Lu, and Y. Makarov, “An optimized autoregressive forecast error generator for wind and load uncertainty study”, *Wind Energy*, vol. 14, no. 8, pp. 967–976, Nov. 2011.
- [69] Bonneville Power Administration. (2013). Wind Generation Data, [Online]. Available: <http://transmission.bpa.gov/Business/Operations/Wind>.
- [70] Eirgrid. (2012). System Performance Data, [Online]. Available: <http://www.eirgrid.com/operations/systemperformancedata/>.
- [71] G. E. P. Box and G. M. Jenkins, *Time series analysis : forecasting and control*. Holden-Day, 1976.
- [72] P. Pinson and H. Madsen, “From Probabilistic Forecasts to Statistical Scenarios of Short-term Wind Power Production”, *Wind Energy*, 2008.
- [73] A. Sturt and G. Stbac, “Time series modelling of power output for large-scale wind fleets”, *Wind Energy*, vol. 14, pp. 953–966, 2011.
- [74] R. Billinton, H. Chen, and R. Ghajar, “Time-series models for reliability in evaluation of power systems including wind energy”, *Microelectronics Reliability*, vol. 36, no. 9, 1996.
- [75] L. Soder, “Simulation of wind speed forecast errors for operation planning of multiarea power systems”, in *Probabilistic Methods Applied to Power Systems, 2004 International Conference on*, IEEE, 2004, pp. 723–728.
- [76] N. C. for Weather and C. Prediction. (2012). Global forecast system, [Online]. Available: <http://www.emc.ncep.noaa.gov/>.
- [77] N. C. for Atmospheric Research. (2012). About the weather research and forecasting model, [Online]. Available: <http://www.wrf-model.org/>.
- [78] E. C. for Medium-Range Weather Forecasts. (2012). European centre for medium-range weather forecasts, [Online]. Available: <http://www.ecmwf.int/>.
- [79] P. Pinson, C. Chevallier, and G. N. Kariniotakis, “Trading Wind Generation From Short-Term Probabilistic Forecasts of Wind Power”, *IEEE Transactions on Power Systems*, vol. 22, no. 3, pp. 1148–1156, Aug. 2007.
- [80] P. Pinson, H. Nielsen, J. Møller, H. Madsen, and G. Kariniotakis, “Non-parametric probabilistic forecasts of wind power: required properties and evaluation”, *Wind Energy*, vol. 10, no. 6, pp. 497–516, 2007.
- [81] B. Rajagopalan and U. Lall, “A k-nearest-neighbor simulator for daily precipitation and other weather variables”, *Water Resources Research*, vol. 35, no. 10, p. 3089, 1999.
- [82] T. M. Hamill and J. S. Whitaker, “Probabilistic Quantitative Precipitation Forecasts Based on Re-forecast Analogs: Theory and Application”, *Monthly Weather Review*, vol. 134, no. 11, pp. 3209–3229, Nov. 2006.

- [83] L. Panziera, U. Germann, M. Gabella, and P. V. Mandapaka, “NORA-Nowcasting of Orographic Rainfall by means of Analogues”, *Quarterly Journal of the Royal Meteorological Society*, vol. 137, no. 661, pp. 2106–2123, Oct. 2011.
- [84] L. Delle Monache, T. Nipen, Y. Liu, G. Roux, and R. Stull, “Kalman Filter and Analog Schemes to Postprocess Numerical Weather Predictions”, *Monthly Weather Review*, vol. 139, no. 11, pp. 3554–3570, Nov. 2011.
- [85] J. W. Messner and G. J. Mayr, “Probabilistic Forecasts Using Analogs in the Idealized Lorenz96 Setting”, *Monthly Weather Review*, vol. 139, no. 6, pp. 1960–1971, Jun. 2011.
- [86] S. Gangopadhyay, B. L. Harding, B. Rajagopalan, J. J. Lukas, and T. J. Fulp, “A nonparametric approach for paleohydrologic reconstruction of annual streamflow ensembles”, *Water Resources Research*, vol. 45, no. 6, W06417, Jun. 2009.
- [87] S. Sen and J. L. Hige, “An introductory tutorial on stochastic linear programming models”, *Interfaces*, vol. 29, no. 2, pp. 33–61, 1999.
- [88] M. Kaut and S. Wallace, “Evaluation of scenario-generation methods for stochastic programming”, Tech. Rep. 1, 2003, pp. 1–16.
- [89] G. Dantzig and G. Infanger, “Large-scale stochastic linear programs: Importance sampling and Benders decomposition”, *Operations Research*, 1991.
- [90] G. Infanger, “Monte Carlo (importance) sampling within a Benders decomposition algorithm for stochastic linear programs”, *Annals of Operations Research*, vol. 39, pp. 69–95, 1992.
- [91] T. Homem-de-Mello, V. L. Matos, and E. C. Finardi, “Sampling strategies and stopping criteria for stochastic dual dynamic programming: a case study in long-term hydrothermal scheduling”, *Energy Systems*, vol. 2, no. 1, pp. 1–31, Jan. 2011.
- [92] S.-E. Fleten and E. Pettersen, “Constructing Bidding Curves for a Price-Taking Retailer in the Norwegian Electricity Market”, *IEEE Transactions on Power Systems*, vol. 20, no. 2, pp. 701–708, May 2005.
- [93] B. G. Gorenstin, N. Campodonico, J. P. da Costa, and M. Pereira, “Stochastic optimization of a hydro-thermal system including network constraints”, *IEEE Transactions on Power Systems*, vol. 7, no. 2, pp. 791–797, May 1992.
- [94] R. Kelman, L. Barroso, and M. Pereira, “Market power assessment and mitigation in hydrothermal systems”, *IEEE Transactions on Power Systems*, vol. 16, no. 3, pp. 354–359, 2001.
- [95] S. Rebennack, B. Flach, M. V. F. Pereira, and P. M. Pardalos, “Stochastic Hydro-Thermal Scheduling Under CO₂ Emissions Constraints”, *IEEE Transactions on Power Systems*, vol. 27, no. 1, pp. 58–68, Feb. 2012.

- [96] GE Energy, “Analysis of Wind Generation Impact on ERCOT Ancillary Services Requirements”, Tech. Rep., 2008.
- [97] EPA. (2013). eGrid, [Online]. Available: www.epa.gov/cleanenergy/energy-resources/egrid/.
- [98] —, (2013). Clean Air Markets Data, [Online]. Available: <http://camddataandmaps.epa.gov/gdm/>.
- [99] J. Ma, V. Silva, R. Belhomme, D. S. Kirschen, and L. F. Ochoa, “Evaluating and Planning Flexibility in Sustainable Power Systems”, *IEEE Transactions on Sustainable Energy*, vol. 4, no. 1, pp. 200–209, Jan. 2013.
- [100] D. Maggio. (2013). Methodology for calculating reserves in the ercot market, [Online]. Available: http://www.uwig.org/San_Diego2012/Maggio-Reserve_Calculation_Methodology_Discussion.pdf.
- [101] G. Energy. (2013). Pjm renewable integration study: review of industry practice and experience in the integration of wind and solar generation, [Online]. Available: <http://pjm.com/~media/committees-groups/task-forces/irtf/postings/pris-task3b-best-practices-from-other-markets-final-report.ashx>.
- [102] A. Greenhall, R. Christie, and J.-P. Watson, “Minpower: A power systems optimization toolkit”, in *2012 IEEE Power and Energy Society General Meeting*, IEEE, Jul. 2012, pp. 1–6.
- [103] A. Greenhall. (2013). Minpower, [Online]. Available: <http://adamgreenhall.github.io/minpower>.
- [104] Sandia National Labs. (2013). Coopr Solvers, [Online]. Available: <https://software.sandia.gov/trac/coopr/wiki/GettingStarted/Solvers>.
- [105] W. E. Hart, J.-P. Watson, and D. L. Woodruff, “Pyomo: modeling and solving mathematical programs in Python”, *Mathematical Programming Computation*, vol. 3, no. 3, Aug. 2011.
- [106] J. P. Vielma, S. Ahmed, and G. Nemhauser, “Mixed-integer models for nonseparable piecewise-linear optimization: unifying framework and extensions”, *Operations research*, vol. 58, no. 2, pp. 303–315, 2010.
- [107] A. J. Wood and B. F. Wollenberg, *Power Generation, Operation and Control*, 2nd. John Wiley & Sons, 1996.
- [108] Northwest Power and Conservation Council, “Sixth Northwest Conservation and Electric Power Plan Appendices”, Northwest Power and Conservation Council, Tech. Rep. February, 2010.
- [109] WECC. (2012). WECC Transmission Expansion Planning Policy Committee 2020 Base Case Dataset, [Online]. Available: <http://www.wecc.biz/committees/BOD/TEPPC/>.

Appendices

Appendix A
LIST OF ACRONYMS

ISO	Independent System Operator
BPA	Bonneville Power Administration
USACE	United States Army Corps of Engineers
EPA	Environmental Protection Agency
EIA	Energy Information Agency
FCRPS	Federal Columbia River Power System
ERCOT	Electric Reliability Council Of Texas
PNNL	Pacific Northwest National Laboratory
PTC	Production Tax Credit
CAISO	California ISO
UC	Unit Commitment
ED	economic dispatch
OPF	optimal power flow
LP	linear programming
RMS	Root Mean Squared
NG	natural gas

CHP	Combined Heat and Power
LR	Lagrange Relaxation
LP	linear programming
MIP	mixed integer programming
SDDP	Stochastic Dual Dynamic Programming
NWP	Numerical Weather Prediction
CVaR	Conditional Value at Risk
WILMAR	Wind Power Integration in Liberalized Electricity Markets
PySP	Python Stochastic Programming
PH	Progressive Hedging
SOS	Special Ordered Sets
EF	extensive form
ARMA	Auto Regressive Moving Average
AR	Auto Regressive
MC	Monte Carlo
KDE	Kernel Density Estimation
LMP	locational marginal price
BIC	Bayesian Information Criterion
AIC	Akaike Information Criterion

MOS	Model Output Statistics
GFS	Global Forecast System
WRF	Weather Research and Forecasting
ECMWF	European Centre for Medium-Range Weather Forecasts
NOAA	National Oceanic and Atmospheric Administration
PWL	Piecewise Linear
OASIS	Open Access Same-Time Information System
EENS	Expected Energy Not Served
LCRA	Lower Colorado River Authority

Appendix B

UC FORMULATION AND IMPLEMENTATION

Minpower is an open source power systems optimization toolkit designed by the author [102, 103]. The toolkit is designed to make working with the classical problems of economic dispatch (ED), optimal power flow (OPF), and UC simple and intuitive. Minpower is also built for flexibility and is well positioned to be a platform for research on stochastic resources. Powerful generic optimization solvers (e.g., CPLEX) can be used by Minpower, allowing for reasonable solution times on even large-scale problems. The Minpower code is open source and is set up for collaborative authorship and maintenance. Tutorials and thorough documentation are available at <http://adamgreenhall.github.io/minpower/>. For an overview, see [102].

Minpower uses Coopr, an optimization package which can use any major solver (e.g., CPLEX, Gurobi, GLPK) [104]. Coopr provides both an optimization modeling language, called Pyomo, and a stochastic modeling and solver engine, called PySP [50, 105]. Thus, Minpower can handle stochastic problems as well as deterministic ones.

B.1 Day-Ahead Deterministic UC Formulation

Minpower uses a formulation based on the deterministic UC formulation of Carrión and Arroyo [18]. This mixed integer formulation is simple due to its use of only a single binary variable¹ per generator per time period² – the status $u_g[t]$. The basic configuration of the optimization model formulation is described in this section. For the sake of clarity, variations on this basic formulation required for infeasible cases (i.e., load and wind shedding) are described separately following this section.

¹Only one binary variable is used if the cost curve is convex. If the curve is non-convex, the linearization method in §B.1.2 must be used, resulting in additional binary variables.

²The single status variable formulation of [18] may not be the fastest-solving UC formulation. Recent work suggests that a three variable formulation may be more efficient.

Table B.1: Indices for UC formulation

index	set	description
g	G	generators
t	T	times

Table B.2: Parameters for UC formulation

parameter	description	units
\underline{p}_g	min. power	MW
\bar{p}_g	max. power (capacity)	MW
\underline{r}_g	min. ramp rate	MW/hr
\bar{r}_g	max. ramp rate	MW/hr
\underline{m}_g	min. down time	hours
\bar{m}_g	min. up time	hours
$c_{su}(g)$	start up cost	\$
$c_{sd}(g)$	shut down cost	\$
$P_D[t]$	system load at time t	MW
$R[t]$	reserve required at time t	MW
$P_g[t_i]$	initial power	MW
$u_g[t_i]$	initial status	MW
η_g	hours in status at initial time	hours
$\bar{\eta}_g$	hours that generator must remain on at start of commitment	hours
$\underline{\eta}_g$	hours that generator must remain off at start of commitment	hours
C_{wc}	cost of wind curtailment	\$/MW
$P_w[t]$	wind power available	MW

Table B.3: Variables for UC formulation

variable	description	units
$u_g[t]$	unit status (on/off)	
$P_g[t]$	power output	MW
$\bar{P}_g[t]$	power output capacity	MW
$C_{\text{op}(g)}[t]$	operating cost	\$
$C_{\text{su}(g)}[t]$	start up cost	\$
$C_{\text{sd}(g)}[t]$	shut down cost	\$
$P_{\text{wu}}[t]$	wind power used	MW

The objective of the UC problem is to minimize the total generation cost:

$$\min \sum_t \sum_g C_{\text{op}(g)}[t] + C_{\text{su}(g)}[t] + C_{\text{sd}(g)}[t] \quad (\text{B.1})$$

while meeting the load:

$$\text{s.t. } P_D[t] = P_w[t] + \sum_g P_g[t] \quad \forall t \in T \quad (\text{B.2})$$

and the spinning reserve requirement:

$$\text{s.t. } P_D[t] + R[t] = \sum_g \bar{P}_g[t] \quad \forall t \in T \quad (\text{B.3})$$

B.1.1 Generator Limits

Additionally each generator has several constraints on its own operation. Where not denoted by an explicit set, these constraints apply for each time for each generator; for clarity we leave out the $\forall t \in T \quad \forall g \in G$ notation. Also for clarity, constraints around the initial and final time periods are not shown here, but are fully described in §B.1.4.

The power must be within machine limits:

$$\text{s.t. } u_g[t] \cdot \underline{p}_g \leq P_g[t] \leq u_g[t] \cdot \bar{p}_g \quad (\text{B.4})$$

and within the generator's ramping limits:

$$\text{s.t. } u_g[t] \cdot \underline{r}_g \leq P_g[t] - P_g[t-1] \leq u_g[t-1] \cdot \bar{r}_g \quad (\text{B.5})$$

Startup/shutdown ramping limits

Some units have minimum power limits which exceed their maximum/minimum ramp rates. For these units special startup/shutdown ramping limits \underline{rs}_g and \overline{rs}_g can be specified.

$$\text{s.t. } \underline{r}_g[t] \leq P_g[t] - P_g[t-1] \leq \overline{r}_g[t] \quad (\text{B.6})$$

where:

$$\underline{r}_g[t] = \underline{r}_g \cdot u_g[t] + \underline{rs}_g \cdot (u_g[t] - u_g[t-1]) \quad (\text{B.7})$$

$$\overline{r}_g[t] = \overline{r}_g \cdot u_g[t-1] + \overline{rs}_g \cdot (u_g[t-1] - u_g[t]) \quad (\text{B.8})$$

Defaults are used if values for \underline{rs}_g and \overline{rs}_g are not specified and the unit's minimum power limit exceeds its maximum/minimum ramp rates. The defaults are set such that the unit is able to startup and shutdown:

$$\underline{rs}_g = -\underline{P}_g \quad (\text{B.9})$$

$$\overline{rs}_g = \underline{P}_g \quad (\text{B.10})$$

Minimum up/down times

There are also inter-temporal constraints on when a generator is able to change its status, including a constraint set that limits the generator to be on for at least \overline{m}_g hours:

$$\begin{aligned} \text{s.t. } \hat{m}_g (u_g[t] - u_g[t-1]) &\leq \sum_{\tau=t}^{t+\hat{m}_g-1} u_g[\tau] \\ \forall g \in G \quad \forall t \in [\overline{\eta}_g + 1, \dots, T - \overline{m}_g + 1] \end{aligned} \quad (\text{B.11})$$

and a constraint set that limits the generator to be off for at least \underline{m}_g hours:

$$\begin{aligned} \text{s.t. } \underline{m}_g (u_g[t-1] - u_g[t]) &\leq \sum_{\tau=t}^{t+\underline{m}_g-1} 1 - u_g[\tau] \\ \forall g \in G \quad \forall t \in [\underline{\eta}_g + 1, \dots, T - \underline{m}_g + 1] \end{aligned} \quad (\text{B.12})$$

Available capacity

The generator's power output capacity $\overline{P}_g[t]$ is useful for calculating reserve³. Capacity is bounded by several constraints. By definition, the output capacity is bounded by the maximum output:

$$\text{s.t. } 0 \leq \overline{P}_g[t] \leq u_g[t] \cdot \overline{p}_g \quad \forall t \in T \quad \forall g \in G \quad (\text{B.13})$$

³If reserve is not required in the formulation, the capacity variable and the associated constraints can be dropped from the formulation.

and is also bounded by the ramp up limits:

$$\text{s.t. } \bar{P}_g[t] \leq P[t-1] + \bar{r}_g u_g[t-1] \quad (\text{B.14})$$

and by the ramp down limits:

$$\text{s.t. } \bar{P}_g[t] \geq P[t-1] - \underline{r}_g u_g[t-1] \quad (\text{B.15})$$

B.1.2 Curve linearization

Curve linearization is handled in two different ways. If the curve is convex, then constraint lines are added to enforce the cost being over each line. This approach does not add additional variables to the formulation, making it more efficient than the non-convex formulation. For non-convex curves, the curve is characterized by line segments and a binary and real valued variable must be added to the formulation. These type of constraints are known as Special Ordered Sets (SOS) constraints and may increase solution time considerably. Both convex and non-convex curve formulations are implemented in Minpower using Pyomo's `Piecewise` class, which is based on the formulations of [106].

Convex curve modeling

To linearize a convex curve $C = f(x)$ for a number of segments L , first take $L-1$ evenly spaced points on the curve, $[x_1, x_2, \dots, x_{L-1}]$ and their corresponding outputs $[y_1, y_2, \dots, y_{L-1}]$. Then calculate the line parameters – slope m_l and intercept b_l for each segment. Finally, add a constraint for each segment in the form:

$$\text{s.t. } C \geq m_l \cdot x + b_l \quad (\text{B.16})$$

Non-convex curve modeling

The linearization of a non-convex curve $C = f(x)$ for a number of segments L (and $B = L - 1$ breakpoints) is a more complicated process. First a binary variable z_l is added to the problem for each segment and a real valued variable for each breakpoint f_b . The variable f_b represents how far along a particular segment line the value lies as a fraction of the segment's two endpoints and is bounded on the interval $[0, 1]$. Thus the cost is represented as the sum of the product of fraction variables and their breakpoint outputs:

$$C(x) = \sum_b f_b \cdot y_b$$

A constraint is added to ensure that there is only one active segment:

$$\text{s.t. } \sum_l z_l = 1$$

and that the fractions add up to one:

$$\text{s.t. } \sum_l f_b = 1$$

and that the fractions are non-zero only on the active segment:

$$\text{s.t. } f_b \leq z_b + z_{b-1} \quad \forall b \in [1, \dots, B]$$

$$\text{s.t. } f_{b=0} \leq z_{l=0}$$

$$\text{s.t. } f_{b=B} \leq z_{l=L}$$

and that the input variable x (e.g., a generator's power) is equal to the sum of the product of breakpoint inputs and fractions:

$$\text{s.t. } x = \sum_b f_b \cdot x_b$$

Constant cost terms

Constant cost terms may be added to generator curves, after being multiplied by the status variable. The total operating cost for a generator at a given time is then the sum of the constant term a_g and the output of the curve model function C_g :

$$C_{\text{op}(g)}[t] = u_g[t] \cdot a_g + C_g(P_g[t]) \quad (\text{B.17})$$

B.1.3 Startup and shutdown costs

To implement startup and shutdown cost, two constraint sets are added:

$$\text{s.t. } C_{\text{su}(g)}[t] \geq C_{\text{su}(g)} \cdot (u_g[t] - u_g[t-1]) \quad (\text{B.18})$$

$$\text{s.t. } C_{\text{sd}(g)}[t] \geq C_{\text{sd}(g)} \cdot (u_g[t-1] - u_g[t]) \quad (\text{B.19})$$

B.1.4 Initial and final constraints

For clarity the constraints around the initial and final time periods are not shown above, but are fully described here.

For minimum up time the generator cannot turn off until the initial up time period $\bar{\eta}_g$ has passed:

$$\text{s.t. } \sum_{\tau=1}^{\bar{\eta}_g} 1 - u_g[\tau] = 0 \quad \forall g \in G \quad (\text{B.20})$$

Similarly, until the initial down time period $\underline{\eta}_g$ has passed, the generator cannot turn on:

$$\text{s.t. } \sum_{\tau=1}^{\underline{\eta}_g} u_g[\tau] = 0 \quad \forall g \in G \quad (\text{B.21})$$

In the final time periods:

$$\begin{aligned} \text{s.t. } 0 &\leq \sum_{\tau=t}^T u_g[\tau] - (u_g[t] - u_g[t-1]) \\ &\forall g \in G \quad \forall t \in [T - \bar{\eta}_g + 2, \dots, T] \end{aligned} \quad (\text{B.22})$$

$$\begin{aligned} \text{s.t. } 0 &\leq \sum_{\tau=t}^T (1 - u_g[\tau]) - (u_g[t-1] - u_g[t]) \\ &\forall g \in G \quad \forall t \in [T - \underline{\eta}_g + 2, \dots, T] \end{aligned} \quad (\text{B.23})$$

B.1.5 Non-controllable generation

Non-controllable generation, including wind, is handled in the current model by simply creating a fixed schedule of power for each non-controllable generator. The total generation is then a summation of the power variables for controllable generators and the fixed parameters for non-controllable generation.

For stochastic problems, non-controllable stochastic generators have one schedule (i.e., set of power parameters) per scenario. Non-controllable but non-stochastic (i.e., CHP) units still have only one schedule.

B.1.6 Wind Curtailment

Wind curtailment can be allowed in the Minpower formulation of UC using the `economic_wind_shed` configuration option. Minpower can also assign a cost to wind curtailment, C_{wc} . By default wind can be curtailed at no cost to the system ($C_{wc} = 0$), but this can be changed by using the `cost_wind_shedding` option, e.g., to \$26 MWh (the value of the Production Tax Credit (PTC)).

If wind shedding is allowed, the formulation is altered by adding a set of variables to represent wind power used $P_{wu}[t]$, which is constrained to be less than the scheduled wind $P_w[t]$:

$$\text{s.t. } P_{wu}[t] \leq P_w[t] \quad \forall t \in T \quad (\text{B.24})$$

A curtailment cost term is added to the objective function, which becomes:

$$\min \sum_t \sum_g C_{\text{op}(g)}[t] + C_{\text{su}(g)}[t] + C_{\text{sd}(g)}[t] + C_{\text{wc}} \cdot (P_w[t] - P_{\text{wu}}[t]) \quad (\text{B.25})$$

and the power balance equation becomes:

$$\text{s.t. } P_D[t] = P_{\text{wu}}[t] + \sum_g P_g[t] \quad \forall t \in T \quad (\text{B.26})$$

B.2 The Real-Time Dispatch Problem Formulation

The real-time problem involves the observed wind power, rather than the forecast wind used in the day ahead. The generators have already been committed, all that remains is to calculate the optimal economic dispatch for each hour in the commitment horizon (the times beyond the horizon are discarded and left to be calculated by the next day's commitment and dispatch). This is implemented by fixing the values of all of the generator status variables and changing the value of the wind power parameters to match the observed values. This creates an LP problem, which is generally quickly solved and is the basis for the observed cost of the solution.

B.3 Handling Infeasible Cases

When a UC is infeasible there is a hierarchy of less strict problems that are tried. First, if the infeasible problem is a real-time dispatch and there are “fast-start” units (i.e., peakers, which can be called upon with little notice) which have not been committed in the day-ahead, the problem is resolved with the fast-starters allowed to start up and meet the load (i.e., the fast-start status variables are left unfixed). If this resolve fails or if the infeasible problem is a day-ahead problem, the problem is re-formulated to allow the shedding of load and non-controllable generation. For shed-able load, a “demand fulfilled” variable P_{Df} is added along with a corresponding constraint:

$$\text{s.t. } 0 \leq P_{Df}[t] \leq P_D[t] \quad \forall t \in T \quad (\text{B.27})$$

the power balance constraint is modified:

$$\text{s.t. } P_{Df}[t] = \sum_g P_g[t] \quad \forall t \in T \quad (\text{B.28})$$

and a cost of load shedding (C_{ls}) term is added to the objective function:

$$\min \sum_t \sum_g C_{\text{op}(g)}[t] + C_{\text{su}(g)}[t] + C_{\text{sd}(g)}[t] + C_{ls} \cdot (P_D[t] - P_{Df}[t]) \quad (\text{B.29})$$

Generally, load shedding is assigned a very high cost (e.g., $C_{ls}=\$10\,000$ MWh). The result is that in general, load is shed only if there is no other option.

B.4 Testing

Minpower is extensively tested. Almost all constraints and objective components are programmatically tested by a UC (or ED, where appropriate) problem test case. All tests are automatically run for each Minpower release using the continuous integration testing platform Travis CI. The results of the tests are available at: <http://travis-ci.org/adamgreenhall/minpower>. Additional documentation of the testing code is available at: <http://minpowertoolkit.com/api/minpower.tests.html>. The following tables summarize the tests in place related to the UC problem formulation:

B.4.1 Constraints

Description	Equation	Test case
min power	(B.4)	power_minimum
max power	(B.4)	power_maximum
min ramp limit	(B.5)	ramp_down
max ramp limit	(B.5)	ramp_up
min shutdown ramp limit	(B.7)	hot_shut_down
max startup ramp limit	(B.8)	cold_ramp_up
default shutdown ramp limit	(B.9)	pmin_startup_limit
default startup ramp limit	(B.10)	pmin_shutdown_limit
min up time	(B.11)	min_up_time
min down time	(B.12)	min_down_time
start up cost	(B.18)	start_up_cost
shut down cost	(B.19)	shut_down_cost
reserve limited by power max	(B.13)	reserve_fixed_amount
power balance with load shedding	(B.28)	load_shedding
initial min ramp limit		ramp_down_initial
initial max ramp limit		ramp_up_initial
initial min up time	(B.20)	initial_min_up_time
initial min down time	(B.21)	initial_min_down_time
final min up time	(B.22)	final_min_up_time
final min down time	(B.23)	final_min_down_time

B.4.2 Objective Components

Description	Term	Equation	Test case
Constant cost term	a	(B.17)	fixed_costs_when_off
Convex cost curve	$C_{\text{op}(g)}[t]$	B.1.2	cubic_convex
Concave cost curve	$C_{\text{op}(g)}[t]$	B.1.2	cubic_non_convex
Startup cost	$C_{\text{su}(g)}[t]$	(B.18)	start_up_cost
Shutdown cost	$C_{\text{sd}(g)}[t]$	(B.19)	shut_down_cost

B.5 Day-Ahead Stochastic UC Formulation

The stochastic formulation can be thought of as a replication of the deterministic problem’s variables and constraints for each scenario. We use a two stage formation of the stochastic UC problem. Generator statuses $u_{g,s}[t]$ are first stage variables; they must be chosen in the day-ahead, without knowing the wind power, and in general cannot be changed. Generator power $P_{g,s}[t]$ is adjusted in the realtime dispatch depending on the wind. The addition of non-anticipatory constraints enforce the decision making structure.

Duplication of Variables and Constraints

To set up the stochastic UC problem, we start with the deterministic formulation and duplicate all variables — generator power $P_{g,s}[t]$, status $u_{g,s}[t]$, and cost curve linearization variables where necessary — and all constraints for each scenario. This process is handled automatically in Minpower which uses PySP to create the extensive form problem. As an example of constraint replication, the power balance equation (B.2), is replaced by its per-scenario form⁴:

$$\text{s.t. } P_D[t] = P_{w,s}[t] + \sum_g P_{g,s}[t] \quad \forall t \in T \quad \forall s \in S \quad (\text{B.30})$$

This replacement and replication is repeated with scenario-specific variables for each constraint, (B.2) to (B.23).

Expected Cost Objective Function

This work uses an expected cost formulation of the objective function:

$$\min E[C] = \sum_s \rho_s \cdot C_s = C_1 + \sum_s \rho_s \cdot C_{2,s} \quad (\text{B.31})$$

⁴assuming that wind is the only stochastically modeled variable in the system

where ρ_s is the scenario probability, C_s is the scenario cost. C_1 is the cost of the first stage variables, which do not differ between scenarios. First stage costs for the UC problem come from starting up or shutting down generation:

$$C_1 = \sum_t \sum_g C_{\text{su}(g)}[t] + C_{\text{sd}(g)}[t] \quad (\text{B.32})$$

$C_{2,s}$ is the per-scenario second stage costs. In the unit commitment these are generator operating or fuel costs:

$$C_{2,s} = \sum_t \sum_g C_{\text{op}(g,s)}[t] \quad (\text{B.33})$$

where a generator's operating cost depends on the scenario power $P_{g,s}[t]$. For a quadratic representation of the cost curve, a generator's scenario cost for an hour would look like:

$$C_{\text{op}(g,s)}[t] = a + b \cdot P_{g,s}[t] + c \cdot P_{g,s}^2[t] \quad (\text{B.34})$$

Non-anticipatory constraints

Non-anticipatory constraints are added to enforce the decision making structure. For the thermal UC problem, statuses $u_{g,s}[t]$ are the only first stage variables. Thus for each time in the commitment horizon T_{na} (i.e., $T_{\text{na}} = 24$ h, see §5.1.1 for additional discussion), a constraint is added to keep each status variable from anticipating any particular scenario. In PySP, the non-anticipatory constraint formulation uses a “root” variable, $u_g[t]$, which is constrained to be equal the first stage variable for each scenario:

$$u_g[t] = u_{g,s}[t] \quad \forall s \in S \quad \forall t \in T_{\text{na}} \quad (\text{B.35})$$

Appendix C

THERMAL GENERATION MODELING

A significant amount of work was put into modeling the parameters of actual power systems. The primary advantage of modeling real systems is a more realistic estimate of solution times and costs. Additionally real data tends to force the resolution of problems that toy problems do not expose. A key component of modeling electric power systems is generator parameters. However these parameters are very difficult to find in the United States due to generator owner concerns (justified or not) about how their competitors would use data on their plants.

This chapter explains the process behind the modeling of the generator data, using two datasets from the EPA [97, 98]. Static data on generators (e.g., fuel type) comes from the EPA’s eGrid 2010 dataset [97].

Hourly time-series data (specifically power output, heat input, and fraction of the hour in operation) for emissions-producing thermal units is available from the EPA’s Clean Air Markets program [98]. This timeseries data can be processed to infer parameters — heat rates, power limits, ramping limits, and up/down time limits¹. A demonstration of this process (including example code) is available at <http://adamgreenhall.com/research/power-data>.

C.1 Non-controllable generation

Wind farms and CHP plants are aggregated and taken as non-controllable units. Hourly observed wind and load data are available upon request from ERCOT. BPA provides historical load and aggregate wind data (both observed and forecast) at 15 min increments on its website. The historical ERCOT public forecast data was provided by special arrangement with 3Tier. Hourly aggregate CHP data was calculated from the Clean Air Markets hourly dataset.

¹Plants were consistent between datasets, but unit names and numbers of units per plant sometimes varied between the EPA’s two datasets, often in combined cycle plants. When mismatches were irreconcilable, data was aggregated into a single-unit combined plant.

C.2 Modeled Parameters

C.2.1 Power limits

Minimum and maximum power limits are estimated on the subset of the time-series data where the unit is operational for the full hour. Outliers are discarded by binning the data and discarding low frequency bins. Combined plant maximums are set to the sum of the units' maximum powers and minimums are assumed to be the minimum of the units' minimum powers.

C.2.2 Ramp rate limits

Minimum and maximum (i.e., max. down and max. up) ramp rates are determined by calculating the difference in power output for sequential hours where the unit is on for both entire hours. Outliers are discarded as described above. If the unit was never on for two hours, it is assumed that the unit is not significantly ramp limited. Combined plant ramp maximums are calculated based on the sum of the units' powers.

C.2.3 Up and down time limits

A histogram of number of contiguous hours on and number of contiguous hours off was created. Up/down limits were set to the minimum hours up/down which occurred more than once. Minimum contiguous periods longer than 48 h were discarded. Combined plant time limits are determined based on a "on" criteria where any unit in the plant is on. This method is currently unable to distinguish between emergency shutdowns/startups or outages and regular operations. The result is an underestimation of real minimum limits for some plants.

C.2.4 Heat rate curves

Heat rates curves were approximated by fitting a curve to binned historical values of heat input and power output. Hours where the unit was on for only part of the hour were discarded, to reduce noise. For each unit a histogram is formed, binning the heat input values into 20 bins. The mean of corresponding power output values for the bin is recorded. The bin means of power and heat are used to fit a series of polynomials – linear, quadratic, and cubic. The monotonically increasing polynomial with the best fit is chosen to represent the unit. If no polynomial is monotonically increasing, the linear polynomial is chosen. Combined plant heat rates are based on the sum of the units' heats and sum of units' powers.

C.3 Estimated Parameters

Nuclear units (which are not subject to Clean Air Market reporting requirements) are modeled according to parameters taken from [108], shown in Table C.2. Start up costs are based on data from [109] and are shown in Table C.1.

Startup Costs

Table C.1: Start up costs used in simulation of the ERCOT system. Based on [109]

Fuel	Turbine	Startup Cost (\$)
Coal		3581
NG	steam	3581
NG	combined cycle	3581
NG	gas turbine	7811
Nuclear		100000

Nuclear Generation Parameters

Table C.2: Parameters used to model the nuclear generation of ERCOT. Based on [108]

Parameter	Value
power minimum	70% capacity
ramp rate limits	$\pm 10\%$ capacity
down time minimum	18 h
up time minimum	120 h
startup cost	\$100000
heat rate	$10.4P$ [MMbtu/MW]
uranium price	\$0.65/MMbtu

Fuel Costs

Cost curves for thermal generators are created by multiplying the appropriate fuel cost by the heat rate curve. A simple constant approximation for 2012 fuel prices from the Department of Energy's Energy Information Agency is made as follows:

Table C.3: Constant fuel price assumptions used in thermal generation cost model, in \$/MMbtu

Fuel	Price
uranium	0.65
coal	2.40
natural gas	3.45

Appendix D

HYDRO SYSTEM MODELING

Power generation in the Pacific Northwest is heavily dominated by hydro. Most of the energy produced comes from the cascade of dams that make up the Federal Columbia River Power System (FCRPS). Hydroelectric systems are even more complex to model and to schedule than thermal power systems – flows of water, rainfall and snowmelt, and often fish passage concerns all dictate hydropower system operations. This appendix describes the modeling of the BPA system, focusing on the core of the FCRPS - ten primary BPA controlled projects which are most relevant to the short-term scheduling problem. A mixed integer programming (MIP) formulation of the deterministic short term hydro scheduling problem is given in Appendix §E.

D.1 Objectives

BPA operates the FCRPS with many objectives, including various non-economic objectives. In this complex system it is often not possible to meet all of the constraints and the objective may include “soft constraint” terms, i.e., penalizations of constraint violations. In general hydro systems can have many possible objectives for short-term scheduling, including:

- economic dispatch (i.e., maximize the value of water)

- minimize spill

- minimize spot market purchases (or maximize sales)

- minimize hour-to-hour variation in production

- minimize the amount of constraint violation

- minimize the number of times constraints are violated

- maximize flexibility (e.g., maximize the distance between operations and constraint bounds)

Table D.1: The ten main federally owned projects of the FCRPS.

Name	Capacity [MW]
Grand Coulee	6,735
Cheif Joseph	2,607
Lower Granite	930
Little Goose	930
Lower Monumental	930
Ice Harbor	690
McNary	1,120
John Day	2,480
The Dalles	2,060
Bonneville	1,211

However, many of these objectives are difficult to formulate precisely or are heavily situation dependent (e.g., how much should deviations in spill be penalized relative to deviations in outflow ramp rate?). For the purposes of this experiment, the objective function is dramatically simplified to an economic dispatch - profits are maximized given the demands and constraints on the system.

D.2 Constraints

BPA operates with a large number of complex, frequently changing constraints which are negotiated to satisfy many stakeholders. This work uses a subset of this constraint set, capturing the major hydrological constraints on the system during the test periods.

The first-order constraints on the system are physical project limits – maximum and minimum limits on power, elevation, outflow, and outflow ramp rate (i.e., rate of change). Some projects also have limits on elevation and tailwater ramp rates. There is also a set of constraints that enforces the cascaded hydraulic structure of the projects. Then there are constraints that vary with time (e.g., seasonal or hour-based constraints). These include limits on spill and potentially more restrictive limits on some of the first-order project constraints.

Basic project constraint data was provided by BPA, along with the curve characterizations for

the outflow-tailwater, elevation-volume, and production. The historical behavior of the river system was modeled based on the publicly available United States Army Corps of Engineers (USACE) dataset. This dataset provides project-level hourly timeseries, including forebay elevation, tailwater elevation, volume, power generated, outflow, spill, and inflow.

D.3 Power Exports and Prices

BPA sells a large proportion of the power it generates – about 54% in 2012 – outside its narrowly defined control area. Many of these sales are governed by long term contracts, for which records are not readily available to the public. Some energy is also sold in the day ahead, through the Open Access Same-Time Information System (OASIS). Dow Jones tracks and records the quantity and price of these day ahead sales at various points throughout the United States. The primary price index for the Northwest is the Mid-Columbia or “Mid-C” price.

This work uses a simplifying Mid-Columbia price assumption to model the relationship between the BPA control area and the rest of the Western Interconnect. It is assumed that the BPA system is perfectly connected internally¹, with a single energy-limited connection to the rest of the Western interconnect. Any power which is scheduled to be sold in the day ahead schedule is priced at the Dow Jones Mid Columbia price. Export limits are based on BPA’s publicly available historical transmission limits dataset.

D.3.1 Mid Columbia Price Model

Exports prices are modeled by the Mid Columbia price index. The actual index is a commercial product sold by Dow Jones Inc. and is not available to non-subscribers. However, a summary of the index consisting of the high and low prices for each day is available through the EIA and from the energy market information company ICE. We assume for this model that the hourly price varied between the high and low price in proportion to the combined demands of the BPA and California ISO (CAISO) systems. Figure D.1 shows the Mid-Columbia price as modeled for 2012, along with the net demand for power for CAISO and BPA.

D.4 Test Cases

Two week-long scheduling cases curated by BPA and representing the dates in Table D.3 were used to test the scheduling algorithm. The cases represent both winter and summer regimes. A full year

¹BPA is a relatively well connected system internally, so much so that engineers sometimes refer to the “Federal bus”.

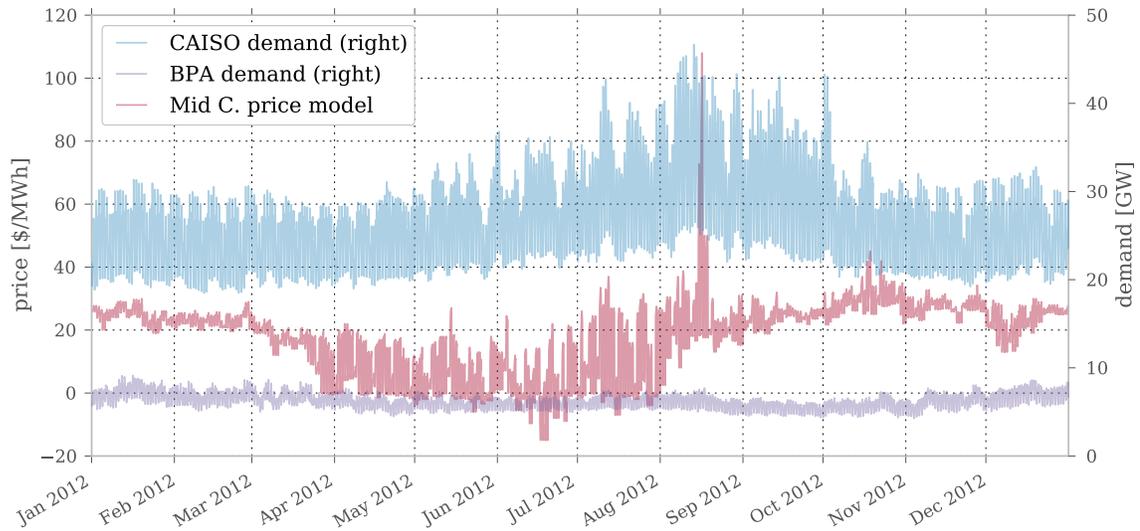


Figure D.1: The Mid Columbia day ahead price index is shown as modeled. Context for this price is given by the system demands for CAISO and BPA on the right axis.

Table D.3: Scheduling test cases for the BPA model based on historical data.

	time period	inflows	export prices
case 1	Nov 29 to Dec 13, 2010	low	low
case 2	Aug 2-16, 2010	high	high

case based on 2012 was also prepared using the USACE dataset. In this case, inflow schedules as well as daily minimum and maximum reservoir targets were created based on the historical record of operations.

D.5 Model Performance

The short-term hydro scheduling problem proved to be very computationally challenging. In the MIP formulation, the reservoir curve models and production curve models for the hydro projects all increase the number of binary variables, which generally makes MIP problems more difficult to solve. The hydro problem is also characterized by a “flat optimum”, i.e., many solutions perform very similarly due to the problem’s networked structure. Extensive tuning of the curve formulation

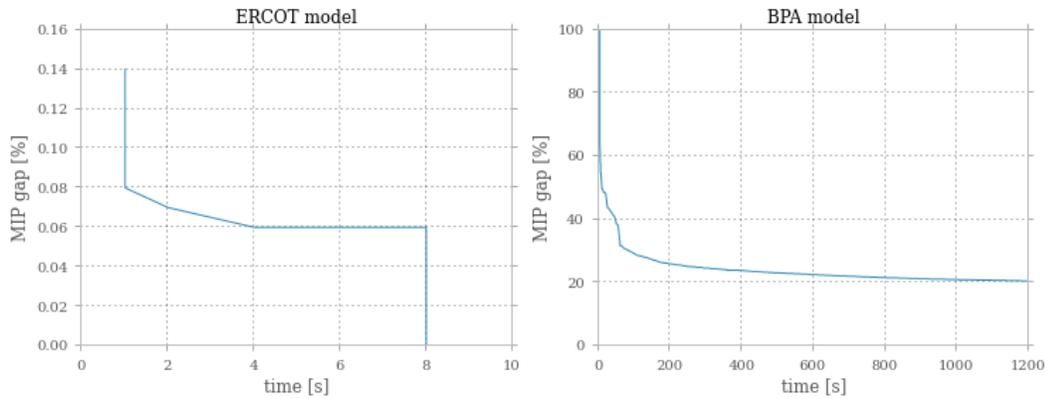


Figure D.2: In deterministic scheduling experiments, the rate of MIP convergence for the BPA model was much slower than convergence of the ERCOT model (note the different scales).

was carried out to minimize the number of binary variables, which resulted in significant convergence speed ups. However even after tuning, the problem remained very difficult for the solver. Figure D.2 compares the convergence of the ERCOT model (see §6.2) to the BPA model described above. A day-long ERCOT model UC problem converges with in a few seconds, while the BPA model scheduling problem slowly converges to a MIP gap of about 20% – far above the normal stopping criteria (generally in the range of 0.01–1%). The model as formulated is not likely to be able to produce meaningful results with reasonable computational times.

The MIP formulation and the economic objective function are the two most likely causes for these long computational times suggest. Alternative objective functions may make the optimum less flat by penalizing certain solutions. It may be necessary to use an altogether different optimization algorithm, e.g., Lagrange Relaxation, heuristic methods, or non-linear mixed integer programming. The disadvantage of many of these methods compared to using a standard MIP solver tends to be a complex problem-dependent implementation. Alternately, it may be possible to seed the MIP problem with a good feasible solution using one of these algorithms, thus speeding up the branch and bound search. The BPA’s short-term hydropower scheduling problem is highly complex and an successful implementation of the solution requires further work.

Appendix E

SHORT-TERM HYDRO SCHEDULING FORMULATION

This section details a deterministic optimization model for a multi-reservoir hydroelectric system. The first component of this model is the network model §E.1 that governs the flows between reservoirs and their volumes. The second component of the model is the hydro production curves. The final section discusses how this model can be integrated with the thermal UC optimization model.

The cascaded hydro system described by this model is based on the Federal Columbia River Power System (FCRPS) as described by the BPA. The resulting short-term hydro-thermal scheduling problem was implemented in Minpower-Hydro.

Table E.1: Indices for short term hydro scheduling formulation

index	set	description
i	I	hydro reservoirs
t	T	times

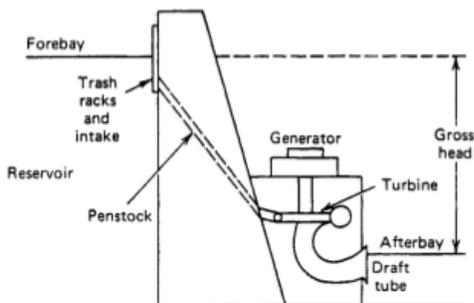


Figure E.1: Overview of hydroplant components. From Wood and Wollenberg [107].

Table E.2: Variables for short term hydro scheduling formulation. Variables below the line are derived variables.

variable	description	units
$e_i[t]$	forebay elevation	ft
$o_i[t]$	outflow	cfs
$s_i[t]$	spill	cfs
$o_{\text{net},i}[t]$	net outflow	cfs
$n_i[t]$	net inflow	cfs
$P_i[t]$	power production	MWh
$v_i[t]$	reservoir volume	ft ³
$t_i[t]$	tailwater elevation	ft
$h_i[t]$	head	ft

Table E.3: Parameters for short term hydro scheduling formulation. Parameters below the line represent final conditions.

parameter	description	units
$r_i[t]$	net incremental upstream inflows	cfs
$\underline{e}_i[t], \bar{e}_i[t]$	min, max. elevation	ft
$\underline{v}_i[t], \bar{v}_i[t]$	min, max. volume	ft ³
$\underline{o}_i[t], \bar{o}_i[t]$	min, max. outflow	cfs
$\underline{s}_i[t], \bar{s}_i[t]$	min, max. spill	cfs
$\underline{P}_i[t], \bar{P}_i[t]$	min, max. power	MW
$\underline{\delta e}_i[t], \bar{\delta e}_i[t]$	min, max. rate of change in elevation	ft/h
$\underline{\delta o}_i[t], \bar{\delta o}_i[t]$	min, max. rate of change in outflow	cfs/h
$\underline{\delta P}_i[t], \bar{\delta P}_i[t]$	min, max. rate of change in power	MWh/h
$\underline{v}_{\text{final}(i)}, \bar{v}_{\text{final}(i)}$	min, max. final volume	ft ³
$\underline{e}_{\text{final}(i)}, \bar{e}_{\text{final}(i)}$	min, max. final elevation	ft

E.1 Network model

Following [34], the change in reservoir volume is defined as the net inflow minus the net outflow.

$$\text{s.t. } v_i[t+1] - v_i[t] = n_i[t] - (o_i[t] + s_i[t]) \quad (\text{E.1})$$

The net inflow is defined as the incremental inflow plus the sum of the delayed (by τ_{ij} hours) outflows of any immediately upstream reservoirs (j):

$$n_i[t] = r_i[t] + \sum_j (o_j[t - \tau_{ij}] + s_j[t - \tau_{ij}]) \quad (\text{E.2})$$

Head is defined as the forebay elevation minus the tailwater elevation. Tail water elevation is approximated as a PWL function of net outflow.

$$h_i[t] = e_i[t] - \text{TW}_i(o_i[t] + s_i[t]) \quad (\text{E.3})$$

Forebay elevation is defined as a PWL function of reservoir volume:

$$e_i[t] = \text{VE}_i(v_i[t])$$

Each reservoir variable can have a set of time varying constraints:

$$\text{s.t. } \underline{e}_i[t] \leq e_i[t] \leq \bar{e}_i[t]$$

$$\text{s.t. } \underline{o}_i[t] \leq o_i[t] \leq \bar{o}_i[t]$$

$$\text{s.t. } \underline{s}_i[t] \leq s_i[t] \leq \bar{s}_i[t]$$

$$\text{s.t. } \underline{P}_i[t] \leq P_i[t] \leq \bar{P}_i[t]$$

$$\text{s.t. } \underline{v}_i[t] \leq v_i[t] \leq \bar{v}_i[t]$$

Forebay elevations, outflows, and power outputs can have limits on the rate of change:

$$\text{s.t. } \underline{\delta e}_i[t] \leq e_i[t] - e_i[t-1] \leq \bar{\delta e}_i[t]$$

$$\text{s.t. } \underline{\delta o}_i[t] \leq o_i[t] - o_i[t-1] \leq \bar{\delta o}_i[t]$$

$$\text{s.t. } \underline{\delta P}_i[t] \leq P_i[t] - P_i[t-1] \leq \bar{\delta P}_i[t]$$

Additionally, the reservoirs are constrained by reservoir targets at the final time T . This can be expressed in terms of volume:

$$\text{s.t. } \underline{v}_{\text{final}(i)} \leq v_i[T] \leq \bar{v}_{\text{final}(i)}$$

or in terms of elevation:

$$\text{s.t. } \underline{e}_{\text{final}(i)} \leq e_i[T] \leq \bar{e}_{\text{final}(i)}$$

E.2 PWL models

The relationship between outflows and tailwater elevation $TW_i()$ and the relationship between forebay elevation and volume $VE_i()$ can both be modeled as piecewise linear functions. This is easily done by one of the methods outlined in §B.1.2.

Power production for a hydro project depends on both head and outflow. This work follows the two variable PWL formulation of [41, 42]. The basic idea of the linearization over two input variables is shown in Figure 1.4.

In this formulation’s MIP implementation, two sets of binary variables are added to denote membership in the segments of flow and one set is added for the segments of head. Two sets of per segment continuous variables are added to denote position within the segment. Several constraints are added to enforce the segment logic and the linearized power value [42, see Section 2.3].

E.3 Hydro-thermal model integration

Hydro units have such small incremental operating costs that in practice they are not considered in the calculation of total system cost. The hydro variables and constraints are simply added to those of the thermal optimization problem. The coupling between the two problems occurs in the power balance and reserve constraints, where hydro generation must be included in the summation of generation power.

E.3.1 Decision Making Stages and Horizons

Mid- to long term hydro scheduling is generally done with multiple decision stages, but the short term, multi-reservoir problem can be simplified to a two stage framework. Due to the time it takes water to flow between reservoirs, net outflows are “here and now” decisions or first stage variables. In the realtime dispatch or second stage, each project can use the arriving water how it chooses, within its constraints. A summary of first and second stage variables is given in Table E.4.

Hydro scheduling problems handle the horizon effect by using reservoir targets. These targets are created by the longer term scheduling process to ensure that the reservoirs retain enough water to meet future predicted demands.

Problem	First Stage	Second Stage
Thermal	status, $u_g[t]$	power, $P_g[t]$
Hydro	net outflow, $o_{\text{net},i}[t]$	power, $P_i[t]$ outflow, $o_i[t]$ spill, $s_i[t]$

Table E.4: First and second stage variables for thermal UC and short term hydro scheduling