

A STOCHASTIC PROGRAMMING APPROACH TO POWER PORTFOLIO OPTIMIZATION

Suvrajeet Sen
Lihua Yu
Talat Genc

Raptor Laboratory, SIE Department,
University of Arizona,
Tucson, AZ 85721, U.S.A.
September 20, 2004

ABSTRACT

We consider a power portfolio optimization model that is intended as a decision aid for scheduling and hedging (DASH) in the wholesale power market. Our multi-scale model integrates the unit commitment model with financial decision-making by including the forwards and spot market activity within the scheduling decision model. The methodology is based on a multi-scale stochastic programming model that selects portfolio positions that perform well on a variety of scenarios generated through statistical modeling and optimization. When compared with several commonly used fixed-mix policies, our experiments demonstrate that the DASH model provides significant advantages.

1 INTRODUCTION

Deregulation is an evolving process. In many states (including Arizona), the major electricity producers have the responsibility of meeting a certain “native load” that constitutes the regulated portion of the business. Beyond this regulated native load, a power producer may buy or sell power in the wholesale electricity market in a manner that the producer finds profitable. Prior to the emergence of electricity markets, profitability was determined simply by the ability of a power producer to convert fuel into electricity in a least-cost manner. Hence minimization of generation costs provided the appropriate strategy. With the emergence of wholesale electricity markets, a utility can manage its power production and revenue potential by trading within this market. A forward (contract) for power is a financial instrument that allows a power producer to buy or sell power for delivery on a future (maturity) date at a price that is agreed upon several months earlier. As weather patterns, economic activity, and market prices evolve, these power portfolios can be rebalanced so as to maximize expected profitability, while appropriately

balancing risk exposure. In this environment, judicious decision-making can mean the difference between survival and demise of a power company.

The DASH model for power portfolio optimization provides a tool that helps decision-makers coordinate production decisions with opportunities in the wholesale power market. Before providing the technical details of our approach, we provide a brief outline of some of the major determinants of profitability in electricity markets. Following this description, we describe statistical models that are used for developing scenarios used within the stochastic programming model. The latter model consists of a financial sub-model and a generation sub-model that are used to determine the profitability of any portfolio position. We also describe an alternative investment strategy based on a certain type of fixed-mix policy which is commonly used by electricity traders. This strategy provides a “base-case” against which we compare the results of a stochastic programming model. Our results are based on data obtained from Pinnacle West Capital, which is a holding company for Arizona Public Service, the largest investor owned electric utility in Arizona. In order to maintain confidentiality of their data, our results will be presented in terms of percentage gain. The backtesting experiment, which covers a five month operating period from January 2001 through May 2001, shows a monthly advantage of approximately 7% in favor of the stochastic programming approach. The DASH model has also been tested against a variety of synthetic scenarios. These experiments reveal the robustness of the forward decisions recommended by DASH.

1.1 Contributions of this Paper

Portfolio optimization models have been investigated using stochastic programming in many recent papers (e.g. Carino and Ziemba (1998), Wu and Sen (2000)), and the volume edited by Ziemba and Mulvey (1998) provides extensive coverage of asset/liability modeling. By the same token the electric utility industry has also applied stochastic programming for hydro-electric generation scheduling (Jacobs et al. (1995)), unit commitment under uncertainty (e.g. Carpentier et al. (1996), Escudero et al. (1996), Takriti, Krasenbrink and Wu, (2000), Nowak and Roemisch (2000), Baccard et al. (2001), Growe-Kuska et al. (2002)). Fleten, Wallace and Ziemba (2002) have discussed a model that combines hydro-electric systems scheduling as well as investments in electricity markets. Our paper is in the spirit of their work, although our modeling approach has several differences that we outline below. First we provide a comprehensive approach in which statistical models of the markets and decision models of the producer are integrated, and the methodology is evaluated through extensive simulation experiments. Moreover we

propose spot market and power generation models that operate on a fine enough time-scale to allow for modeling on-peak as well as off-peak electricity products and electricity generation. In addition we allow a variety of generators within a fine time-scale unit commitment model. On the financial side we allow contracts to be modeled on a larger time-scale (i.e. monthly) and moreover include both electricity and gas markets. Since gas is often the marginal fuel used by power producers facing peak-load, modeling the gas market provides much more realistic estimates of future marginal cost of electricity. Our model captures the impact of generation costs (which are typically obtained from short term unit commitment models) on investments in electricity commodities (typically on a monthly time-scale). This multi-scale approach allows much greater fidelity than has been attempted to date. Naturally, the resulting model is far too complex for solution using off-the-shelf MILP solvers. It turns out that the challenge of incorporating multiple scales, both in modeling uncertainty, and in the decision-making process, leads to the main contributions of this paper.

We design a new nested column generation approach that decomposes the model into smaller sub-problems that are coordinated within the new algorithm. In addition to providing an algorithmically tractable approach, this new algorithm maintains modularity by solving a fine time-scale (electricity generation) model and a coarse time-scale (financial investment) model separately. Another important advantage of the new algorithm is that it is relatively straightforward to study a sequence of instances, with alternative probability estimates for the scenarios. This is because the method is based on column generation, and a change in probability distribution only shows up in the objective function of the model. Thus, the new algorithm is amenable for day-to-day implementations in which probability estimates may evolve, and new instances may have to be resolved.

Finally, this paper also describes our statistical modeling effort for scenario generation in the decision phase, as well as the evaluation phase. The scenarios generated for the decision phase are used within the stochastic programming model, whereas scenarios used in the evaluation phase are meant to test the robustness of decisions provided by the stochastic programming (DASH) model. Thus, our paper provides a comprehensive treatment including statistical modeling, optimization, and simulation.

2 SCOPE OF THE DASH MODEL

To begin with, we outline the manner in which we expect the decision process to unfold. At the start of each month, financial analysts/traders for the producer wish to reevaluate/rebalance their power portfolio. At this point, they may invoke some decision model (e.g. DASH) which recommends the mix of

power products that the producer ought to hold. While the decision model itself may be dynamic (as in DASH), the trader only commits to a recommendation for the current month. After the appropriate re-balancing trades are executed, the traders wait and observe the market until the end of the month, at which point, they update the decision model by “rolling the horizon” forward, and providing up-to-date information to the decision model which then provides an updated recommendation for the next month. While it is possible to use the DASH model at decision-epochs that are less than a month long, the portfolios within DASH are represented at monthly intervals.

Market modeling is another feature incorporated within DASH. In some cases, power producers trade electricity in multiple markets. For example, a California utility may trade in Palo Verde (AZ) and the California-Oregon Border (COB). For the sake of this model however, we will consider only one market for electricity. In addition to electricity, the model also allows interactions with one natural gas market. On the generation-side, the unit commitment decisions are made on a monthly basis, and allow us to incorporate heat-rates, start-up costs, minimum downtimes, etc. The current model does not accommodate hydro generation, although this extension is currently under consideration. Finally, we also set aside modeling market power as such an extension would involve a game theoretic setting, a particular version of which is explored in Genc, Reynolds and Sen (2003).

2.1 Electricity Demand

In a completely deregulated market, the traditional notion of load takes a back-seat to demand curves relating prices and quantities. However the extent of deregulation is in a state of flux in most states in the U.S. For instance in Arizona, retail tariffs are regulated by the state Corporation Commission and are held constant over long periods of time. Electric utilities are required to serve the “native load” that arises from their customers at regulated retail rates. There are several different demand models that have been studied in conjunction with the current DASH model, including time-series that use temperature as one of the main factors. Retail prices, as they are fixed, are not considered.

2.2 The Wholesale Electricity Market

The model allows both electricity forward contracts and spot market activity. While the model does not accommodate options, these can be included without adding to the computational burden. While prices in the electricity market (especially the spot market) vary on an hourly basis, we have discretized time according to a sixteen hour on-peak period, and an eight hour off-peak period for each day.

2.2.1 Forward Contracts for Power

For the purposes of our model, forward contracts will be assumed to be “monthly”, so that planning for period t refers to some month t in the future. A power forward contract can be either a financial contract or physical contract. If a forward contract is settled before its maturity date, it is a financial forward contract since no physical transactions are necessary; otherwise, it is a physical forward contract for which the electricity is delivered physically. A physical forward contract can be settled by taking a position in financial forwards and letting the financial contract mature on the delivery date. Note that the megawatts committed (bought or sold) to the market in period j influences the total electricity generated during period t , $t > j$. To facilitate profit-making, trading decisions must consider future load projections and generation capacity, both of which are subject to uncertainty. If the decisions for the delivery month (t) could be treated independently of other months, then one could develop a model that could treat each delivery month independently. However, such an assumption might expose the firm to a greater risk level than might be acceptable. This is because the (financial) risk exposure of a firm depends on the mix of instruments in its portfolio at any point in time. Hence, it is not sufficient to simply consider profitability for a delivery month; the collection of forwards held at any point in time is an important determinant of risk exposure.

The current price of any forward contract is usually assumed to be known. However, forward prices for each delivery month will evolve over time until the delivery month. As one might expect, this evolution is uncertain on the decision-making date. In the current version of the DASH model, we use a non-parametric approach in which historical data is used to create a vision for the future (e.g., the next six months). This vision is based on creating a number of scenarios of “returns” (percentage change in prices) which may be revealed in the future. The actual process of developing these scenarios is discussed in the next section.

2.2.2 Spot Market for Power

As with forward contracts, on-peak and off-peak power have different price trajectories, and are modeled separately. However, there are two important observations in modeling the spot market. The time-scale for spot prices can be hourly. In the interest of computational tractability, we treat spot market on a daily basis, and allow it to fluctuate according to the sixteen-hour on-peak and eight-hour off-peak peri-

ods. Also, the spot prices for each day (d) during the month (t) must be correlated to the forward prices associated with the scenario (s) that unfolds.

2.3 Unit Commitment

The technological constraints of this socio-technical model arise in the unit commitment problem. Traditionally, unit commitment models are used to determine a short-term (weekly) power generation schedule. While they have also been used to estimate annual production costs, the deterministic nature of the original models (e.g. Bertsekas et al. (1983)) do not lend themselves to mid- and long-term analysis. More recently, these models have been extended to accommodate uncertainty in load forecasts, fuel prices, etc. (Takriti, Birge, and Long (1996), Takriti, Krasenbrink, and Wu (2000), and Nowak and Roemisch (2000)). Recent advances in unit commitment models are summarized in Hobbs et al. (2001). Escudero and Pereira (2000) provide an introduction to the impact of electricity markets on unit commitment. The models mentioned above are typically focused on a short-term scheduling issue (a week or two at most). Due to the medium-term nature (i.e., one year) of many financial instruments, it is difficult to measure their impact using short-term models. Our multi-scale approach integrates the unit commitment model with financial decision-making by including the forwards and spot market activity within the scheduling decision model.

3 STATISTICAL INPUT MODELS

With the exception of the unit commitment model, all features discussed in the previous section are represented by statistical models. The main purpose of these statistical models is to help generate a finite number of scenarios that are represented in the form of a scenario tree. A scenario models the evolution of information during the decision process (Birge and Louveaux (1997)). It is important to emphasize that our scenario generation procedure is intended to work in concert with the model proposed in Section 4, which in turn must be amenable to algorithmic treatment. Thus no single module should create bottlenecks for another. The notation used in this section is summarized in Appendix A.

We begin with an overview of the scenario generation process that will provide inputs into a large-scale stochastic programming decision model. As we shall see in Section 4, forward decisions are modeled as “here-and-now” decisions, and must therefore satisfy non-anticipativity requirements of stochastic decision processes. On the other hand, spot market and generation decisions are handled through “wait-and-see” models, and as a result they adapt to a sample path. The difference between “here-and-

now” decisions versus “wait-and-see” decisions call for different mechanisms for scenario generation, depending on the type of decision. Accordingly, there are two types of scenario generation processes; a coarse time-scale discretization for the forward (here-and-now) decisions, and fine time-scale models for generation and spot market (wait-and-see) decisions. The coarse time-scale discretization will be based on “matching” medians of forward returns data, similar in spirit to the idea of matching moments or other properties suggested by Hoyland and Wallace (2001). On the other hand, we use detailed econometric models for load and spot price trajectories. The latter approach is similar to models used by Growe-Kuska et al. (2002) in which load trajectories are incorporated within a stochastic unit commitment model. To be concrete, we illustrate a scenario tree in Figure 1.

Figure 1: Scenario Tree Structure

Columns in Figure 1 represent the vector of forward prices for on-peak power, off-peak power and gas, in each subsequent month. That is, for a scenario tree that covers a six-month period, nodes in the first month require eighteen forward prices, those in the second month require fifteen and so on. A forward scenario is one path through the tree as shown in Figure 1. Along each scenario (path), spot prices (on-peak and off-peak) as well as loads (on-peak and off-peak) will form a sample path (trajectory). Associated with each forward scenario, we will sample several spot prices, and load trajectories from econometric models developed below. This two-pronged approach maintains simplicity of the forward scenario tree and the resulting stochastic programming model, whereas the details of the spot prices (and loads) provide a better approximation of daily processes. Together the statistical input models and the stochastic programming decision model provide a computationally tractable approach to the forward decision making issue. In order to verify whether this approach does indeed suffice, we perform an experiment to study the quality of the decisions when they are used under scenarios that are generated using more complicated econometric models in Section 6.3.4. We should also mention that scenario generation for stochastic programming is currently an area of active research, and readers may refer to Hoyland, Kaut, and Wallace (2003), as well as Heitsch and Roemisch (2003) for recent approaches to scenario generation.

The sequence of models presented below are provided in subsections that parallel those in Section 2.

3.1 Modeling Electricity Demand

Our load data represents an eleven-year period (1990–2000) of hourly loads in an APS service area. Since each day is modeled by “on-peak” and “off-peak” segments, we begin by transforming the hourly data into averages for each segment. The hours 6 a.m. to 10 p.m. are considered on-peak, and the remaining hours are considered off-peak. In order to give the reader a sense of the load data, Figure 2 provides a three year sequence of on-peak loads. The off-peak loads also portray similar cyclical and seasonal trends, and these are confirmed by the Kendal-Tau and Turning Point tests (see Kendall, Stuart and Ord (1983)).

Figure 2: On-Peak Load Data for 3 Years

Based on seasonality of loads depicted in Figure 2, we partition the data for a year into four groups, each representing a season. The first has a decreasing trend, the next an increasing trend and so on. For each group/partition, we use d to denote a day, \tilde{L}_d denotes the load. Assuming an annual growth rate of g , we propose the following model:

$$\tilde{L}_d = L_d(1 + g) ,$$

where

$$L_d = \alpha_0 + \alpha_1 d + \varepsilon_d$$

$$\varepsilon_d = \sum_{i=1}^7 \beta_i \varepsilon_{d-i} + \eta_d + \beta_8 \eta_{d-1} , \quad \eta_d \sim N(0,1)$$

In order to create load scenarios from such a model, we generate pseudo random η_d .

For the data set we investigated, the de-trended load (for both on-peak and off-peak segments) follows ARIMA(7,0,1) for each partition. This is consistent with the study of Dupacova, Growe-Kuska and Roemisch (2000) who examined hourly loads (which can be viewed as high frequency data) and concluded that SARIMA(7,0,9) \times (0,1,0) was an appropriate model for hourly loads. Following Box and Jenkins (1976), we also performed several tests to investigate the quality of the model. These include plots of the remaining residuals, the autocorrelation function, partial autocorrelation function, p-values of Ljung-Box statistics, and qq-normal plot of residuals. These results can be obtained from the authors upon request. The diagnostic tests validated the sufficiency of the demand model.

3.2 Modeling Electricity Forward Prices

This forms the core of the DASH scenario generation procedures. The inputs we use are the forward prices for the preceding year, together with recent trends in the market. Let us first focus on the forward

prices for the preceding year. These are available as hourly quotes which we transform into on-peak and off-peak average prices. We have the following format for the prices: $\pi_{\tau\kappa e}$, where π is the price, and τ, κ and e denote the contract week, delivery week and segment, respectively. Here, the range of indices are: $\tau = 1, 2, \dots, 52$, $\kappa = \{1, 2, \dots, N\}$, $e \in \{on, off\}$, N denoting the last week in which delivery will happen. For example, $\pi_{1,8,on}$ is the price (\$/MWh) on January 7th (i.e., end of week 1) for on-peak power delivered starting on March 1st (for the entire month of March). However we use “returns” to predict prices; that is, $r_{\tau,\kappa,e} = (\pi_{\tau+4,\kappa,e} - \pi_{\tau,\kappa,e}) / \pi_{\tau,\kappa,e}$. Since the index τ reflects an index for weeks, the subscript $\tau + 4$ denotes a period that is four weeks removed from period τ . Assuming that there are four weeks in a month, the return $r_{\tau,\kappa,e}$ denotes the relative change in price during the month starting in week τ .

There are two important reasons behind our choice to model returns rather than prices. First, this approach allows us to treat different power contracts (associated with different months) with the same scenario tree, thus reducing the complexity of modeling the evolution of prices associated with each type of contract. We have empirically verified that it is the interval of time between contract and delivery that is important for modeling returns, and not the actual contract. Hence the same scenario tree remains a valid representation of returns for alternative contracts. Secondly, the econometrics literature recommends that “returns” are better for predictive purposes because empirical evidence suggests that they appear to have better properties (e.g. stationarity) from a computational point of view (Taylor (1986)).

A discrete scenario tree may now be formed by grouping returns into subsets for each period (i.e., month), and modeling the return process as one that allows probabilistic transitions from one subset to another over time. In order to maintain computational tractability, we consider only two subsets in each period: “High” and “Low” return states. Thus, the resulting scenario tree can be represented by a binary tree in which the returns can assume “High” or “Low” values over the course of the decision process.

To assign “High” and “Low” values for the return states, we adopt a sampling-based procedure that is guided by recent observations of the return series. The nominal value that we assign to each state (“High” or “Low”) is the median of the corresponding group for that period. However, without accommodating extreme values, the scenario tree (and consequently the decisions themselves) overlooks extreme events, thus opening up the possibility for catastrophic losses. We will of course, include some loss constraints within the decision model, but in the absence of extreme scenarios, such constraints can only have limited impact. Accordingly, we use a combination of medians and extreme values (“Min” and “Max”) to assign values to the “High” and “Low” states. The precise manner in which we choose

one or the other depends on a heuristic guided by market conditions prior to running the model. Finally, the formation of the scenario tree requires a specification of transition probabilities between nodes representing information states. Recall that our scenario tree is binary, and hence there are only two probabilities that need to be specified. In the event that our heuristic produced two nodes that are represented by medians (“High” and “Low” respectively), then we simply use equal conditional probabilities for these two transitions. On the other hand, if the heuristic produces an extreme value for one path, and a median for the other, then we associate a conditional probability of $\frac{1}{4}$ for the extreme value, and $\frac{3}{4}$ for the median value. These conditional probabilities reflect approximately the number of times the return either exceeded the high median or fell below the low-median within the data set. Our heuristic does not produce two extreme values from any node, and hence this possibility is not considered.

The above process creates a binary scenario tree for the return series, which is then used to create price scenarios used by the stochastic programming model described in the next section.

3.3 Modeling Gas Forward Prices

The process used to model gas forward prices is similar to the process described in the previous subsection (on electricity forward prices). We will also assume that the returns for gas and electricity are perfectly correlated so that a scenario obtained from the electricity forward return tree generates a similar scenario from the gas forward return tree (for justification, see Section 6.3.3.2).

3.4 Modeling Electricity Spot Prices

Recall that the forward price process is discretized on a monthly basis. However, spot prices must be modeled on a different time-scale. As discussed earlier, on-peak and off-peak spot prices will be modeled on a daily basis, with the understanding that they will be correlated with an appropriate forward price scenario. As with the forwards, we resort to modeling the return series of spot prices, and make the assumption that forward prices reflect the expected spot prices. This assumption avoids creating a “money machine.”

The spot prices during a delivery month are generated from the following formulation (of spot returns): $r_{e,d,t,\omega}^p = r_{e,t,\omega}^f + \sigma_{e,t} z_{d,t,\omega}$, where ω is the node number of the forward scenario tree, $\sigma_{e,t}$ is the standard deviation of spot returns which changes from delivery month to delivery month, and $r_{e,t,\omega}^f$ is the daily equivalent of the forward return ($r_{t,t+1,e}$) on node ω for month t . The quantity z represents a stan-

dard normal random variate. Here $\sigma_{e,t}$ may be interpreted as the volatility associated with on-peak and off-peak returns during month t and are estimated using a GARCH (Generalized Auto Regressive Conditional Heteroskedasticity) model (Bollerslev (1986)). Because the expectation of spot market prices may be assumed to equal the expected forward prices (Hull (1997)), the above relationship between spot and forward returns captures both the first as well as second moments of the spot price process. Assuming that forward return ($r_{e,t,\omega}^f$) is given and volatility of spot returns ($\sigma_{e,t}$) is computed, we draw as many outcomes of the standard normal variate z as there are days (d) in a period (t). This process yields a sample path of spot returns.

4 THE DECISION MODEL

The DASH model may be classified as a multi-stage stochastic integer program that recommends forward decisions on a here-and-now basis, whereas the operational decisions (generation, spot market activity, etc.) are used to evaluate the viability of the portfolio. In this sense, the generation and spot market decisions are adaptive (i.e., wait-and-see), and allow us to compute medium term (six months to a year) decisions without being mired in daily (here-and-now) details.

In formulating the stochastic program, all decision variables and many independent parameters depend on the scenario. However, in the interest of simplifying the notation we have suppressed this dependence below. We remind the reader that all forward variables will be required to satisfy the non-anticipativity requirements of stochastic programming (Birge and Louveaux (1997)). The formulation is presented in two parts: the financial problem and generation costing problem.

4.1 The Financial Problem

In the following formulation, all power decision variables are in MW.

Scenario Independent Parameters

α : Max liquidity limit coefficient;

T : Number of periods;

P : Regulated power price;

J_t : The number of segments in period t ; in a month consisting of 28 days, there are 56 segments;

ε : {on-peak, off-peak};

H_e : Hours of one on/off peak segment, $H_e=16\text{h}$ for $e=\text{on-peak}$, and 8h for $e=\text{off-peak}$;

$p(j)$: Peak status(on/off) of segment j ;

YL_{0te} : Power forward in long position for delivery period t , peak e held initially;

YS_{0te} : Power forward in short position for delivery period t , peak e held initially;

YG_{0t} : Gas forward for delivery period t held initially;

PT_t : Profit target for period t ;

Scenario Dependent Parameters

$PP_{\tau e}$: Price of energy forward (MWh) for delivery period t , peak e (on/off peak) at contract period τ ;

$PG_{\tau t}$: Price of gas forward for delivery period t , at contract period τ ;

PS_{ij} : Price of energy in spot market in period t , segment j ;

D_{ij} : Electricity demand in period t , segment j ;

Scenario Dependent Decision Variables

$FP_{\tau e}$: Power forward for delivery period t , peak e (on/off peak), signed at contract period τ (positive for long position, negative for short position);

$FP_{\tau e}^+$, $FP_{\tau e}^-$: Power forward in long and short position respectively for delivery period t , peak e (on/off peak), signed at contract period τ , an upper bound is imposed on this variable;

$FG_{\tau t}$: Gas forward in long position for delivery period t , signed at contract period τ ;

$YP_{\tau e}$: Total power forward for delivery period t , peak e held at contract period τ (positive for long position, negative for short position);

$YP_{\tau e}^+$, $YP_{\tau e}^-$: Total power forward in long and short position respectively for delivery period t , peak e held at contract period τ , an upper bound is imposed on this variable;

$YG_{\tau t}$: Total gas forward for delivery period t held at contract period τ ,

SP_{tj} : Power exchanged with spot market in period t , segment j (positive for purchase, negative for sale);

ZP_{te} : Total power forward cost for delivery period t , peak e ;

ZG_t : Total gas forward cost for delivery period t ;

C_{tj} : Total generation cost in period t , segment j ;

Scenario Dependent Constraints

$$FP_{\tau te} = FP_{\tau te}^+ - FP_{\tau te}^- \quad \tau \in \{1, \dots, T\}, t \in \{\tau, \dots, T\}, e \in \mathcal{E}; \quad (1)$$

$$YP_{\tau e} = YP_{\tau e}^+ - YP_{\tau e}^- \quad \tau \in \{1, \dots, T\}, t \in \{\tau, \dots, T\}, e \in \mathcal{E}; \quad (2)$$

$$YP_{\tau e}^+ = YP_{(\tau-1)te}^+ + FP_{\tau e}^+ \quad \tau \in \{1, \dots, T\}, t \in \{\tau, \dots, T\}, e \in \mathcal{E} \quad (3)$$

(Power forward balance in long position at period τ);

$$YP_{\tau e}^- = YP_{(\tau-1)te}^- + FP_{\tau e}^- \quad \tau \in \{1, \dots, T\}, t \in \{\tau, \dots, T\}, e \in \mathcal{E} \quad (4)$$

(Power forward balance in short position at period τ);

$$YG_{\tau t} = YG_{(\tau-1)t} + FG_{\tau t} \quad \tau \in \{1, \dots, T\}, t \in \{\tau, \dots, T\}, e \in \mathcal{E} \quad (5)$$

(Gas forward balance at period τ);

$$\sum_{t \in [\tau, T]} FP_{\tau e}^+ \leq \alpha \sum_{t \in [\tau, T]} YP_{(\tau-1)te}^+ \quad \tau \in \{1, \dots, T\}, e \in \mathcal{E} \quad (6)$$

(Max liquidity limit for long position);

$$\sum_{t \in [\tau, T]} FP_{\tau e}^- \leq \alpha \sum_{t \in [\tau, T]} YP_{(\tau-1)te}^- \quad \tau \in \{1, \dots, T\}, e \in \mathcal{E} \quad (7)$$

(Max liquidity limit for short position);

$$ZP_{te} = \frac{J_t}{2} \sum_{\tau \in [1, t]} PP_{\tau e} FP_{\tau e} H_e \quad t \in \{1, \dots, T\}, e \in \mathcal{E} \quad (8)$$

(Total power forward cost for delivery period t , peak e);

$$ZG_t = \sum_{\tau \in [1, t]} PG_{\tau} FG_{\tau} \quad t \in \{1, \dots, T\} \quad (9)$$

(Total gas cost for period t);

$$\sum_{j \in [1, J_t]} [(D_{ij}P - SP_{ij}PS_{ij})H_{p(j)} - C_{ij}] - \sum_{e \in \mathcal{E}} ZP_{te} \geq PT_t \quad t \in \{1, \dots, T\} \quad (10)$$

(Monthly profit target control).

Constraints (1-5) constitute balance constraints (dynamics). Constraints (6-7) provide a way to control the extent to which a portfolio is allowed to change from one period to the next period. These constraints help avoid speculation, thus limiting risk exposure. The lower the value of α (not necessarily less than 1), the tighter the control is on the forward trajectory allowed by the model. In addition to the above liquidity constraints, we also include an upper limit on forward positions which is described in the variable definition. The costs associated with the forward decisions are captured in (8-9). By constraint (10), the monthly profit target control is enforced. However, if this constraint cannot be satisfied, we include this target within the objective function via a penalty term. Finally, there are two important factors required to specify the financial problem.

- Non-anticipativity constraints require that scenarios which share the same history until period t should be associated with decisions which have the same values until period t . These linear constraints couple decisions from different scenarios, thus allowing a well hedged plan.
- The objective function for the financial problem maximizes expected profits associated with the portfolio. In calculating the profits, we accommodate the generation cost, which is computed via the model discussed next.

4.2 The Generation Problem

The electricity products model includes on-peak and off-peak power, during 16 and 8 hour segments of a day, respectively. Because of the non-storability of electricity, the utilization of generation assets and the associated costs must be captured during these on-peak and off-peak periods. Hence it is necessary for production-costing models to accommodate this time-scale. With each scenario we associate a series of generation problems, and each generation problem models a period of power production. Thus for any scenario, there will be the same number of generation problems as there are periods in the financial model. In this formulation, the generation and spot market variables are allowed to be adaptive. Although gas generators are included with the generation model, we do not model gas storage capacity, so that gas inventory from day-to-day is not included in this model. In essence, this is equivalent to assum-

ing that there is infinite capacity for gas storage and its inventory cost is negligible. This assumption together with the assumption that the forward price of a path and expected spot price are equal implies that decisions in the spot market can be subsumed by decisions in the forward market.

We now proceed to a description of the generation model. As before, the notation suppresses the dependence on scenarios.

Scenario Independent Parameters

I : The set of generators;

d : Index of days.

$j(d)$: Indices of the two segments associated with day d ;

$t(d)$: The period associated with day d ;

ML : Maximum acceptable daily loss;

Gas : The set of gas generators;

$Coal$: The set of coal generators;

Nuc : The set of nuclear generators;

CP_t : Coal price for period t ;

NP_t : Nuclear fuel price for period t ;

Q_i : Maximum generation capacity of generator i ;

q_i : Minimum generation capacity of generator i ;

L_i : Minimum up time requirement for generator i ;

l_i : Minimum down time requirement for generator i ;

$F_i(x)$: Consumption function of fuel for generation of x due to generator i ($F_i(x) = a_i + b_i x$ where a_i and b_i are parameters);

Scenario Dependent Parameters

W_{ij} : Scheduled outage ($W_{ij} = 0$, if outage is scheduled in period t , segment j for generator i ; 1, otherwise);

ϖ_{ij} : Forced outage ($\varpi_{ij} = 0$, if outage is forced in period t , segment j for generator i ; 1, otherwise);

Scenario Dependent Decision Variables

TG_{ij} : Total generated power in period t , segment j ;

G_{ij} : Power generated by generator i in period t , segment j ;

U_{ij} : Operation decisions for generator i in period t , segment j

(Binary decision variable; $U_{ij}=1$, if generator i is on in period t , segment j ; $U_{ij}=0$, otherwise);

SG_{ij} : Consumption of gas in period t , segment j ;

SC_{ij} : Consumption of coal in period t , segment j ;

SN_{ij} : Consumption of nuclear fuel in period t , segment j ;

Scenario Dependent Constraints

$$YP_{te} + SP_{ij} + TG_{ij} = D_{ij} \quad t \in \{1, \dots, T\}, j \in \{1, \dots, J_t\}, e = p(j) \quad (11)$$

(Demand-generation- forward - spot relationship);

$$YG_{tt} = \sum_{j \in [1, J_t]} SG_{ij} \quad t \in \{1, \dots, T\} \quad (12)$$

(Total gas consumption for period t);

$$SG_{ij} = \sum_{i \in Gas} F_i(G_{ij}) \quad t \in \{1, \dots, T\}, j \in \{1, \dots, J_t\} \quad (13)$$

(Gas consumption for period t , segment j);

$$SC_{ij} = \sum_{i \in Coal} F_i(G_{ij}) \quad t \in \{1, \dots, T\}, j \in \{1, \dots, J_t\} \quad (14)$$

(Coal consumption for period t , segment j);

$$SN_{ij} = \sum_{i \in Nuc} F_i(G_{ij}) \quad t \in \{1, \dots, T\}, j \in \{1, \dots, J_t\} \quad (15)$$

(Nuclear fuel consumption for period t , segment j);

$$C_{ij} = \frac{ZG_t}{J_t} + NP_t SN_{ij} + CP_t SC_{ij} \quad t \in \{1, \dots, T\}, j \in \{1, \dots, J_t\} \quad (16)$$

(Generation cost for period t , segment j)

$$TG_{ij} = \sum_{i \in I} G_{ij} \quad t \in \{1, \dots, T\}, j \in \{1, \dots, J_t\} \quad (17)$$

(Total generated power at period t , segment j);

$$q_i U_{ij} \leq G_{ij} \leq Q_i U_{ij} \quad i \in I, t \in \{1, \dots, T\}, j \in \{1, \dots, J_t\} \quad (18)$$

(Operating range for each generator);

$$U_{ij} - U_{it,j-1} \leq U_{it\tau} \quad i \in I, t \in \{1, \dots, T\}, j \in \{1, \dots, J_t\}, \tau \in \{j+1, \dots, \min(j+L_i-1, J_t)\} \quad (19)$$

(Minimum up-time requirement);

$$U_{it,j-1} - U_{ij} \leq 1 - U_{it\tau} \quad i \in I, t \in \{1, \dots, T\}, j \in \{1, \dots, J_t\}, \tau \in \{j+1, \dots, \min(j+L_i-1, J_t)\} \quad (20)$$

(Minimum down-time requirement);

$$U_{ij} \leq W_{ij} \quad i \in I, t \in \{1, \dots, T\}, j \in \{1, \dots, J_t\} \quad (21)$$

(Scheduled outage);

$$U_{ij} \leq \varpi_{ij} \quad i \in I, t \in \{1, \dots, T\}, j \in \{1, \dots, J_t\} \quad (22)$$

(Forced outage);

$$\sum_{j \in j(d)} [(D_{ij} P - SP_{ij} PS_{ij}) H_{p(j)} - C_{ij}] - \frac{\sum_{e \in \varepsilon} Z P_{te}}{J_t / 2} + ML \geq 0, \quad \forall d, t = t(d) \quad (23)$$

(Maximum daily loss constraint).

The alternative time indexes used in the model result in a multi-scale model with the financial decisions being made on a monthly time index, and the generation decisions being indexed by segments that are either eight or sixteen hours long.

In constraint (11), the forward contract position $Y P_{te}$ involves an exchange of electricity (i.e. the net physical amount of electricity exchanged) in month t . The max-daily loss constraint is imposed in order to provide a measure of risk control on the decisions. Note that we have two risk constraints: a daily “loss” constraint in the generation problem (23) and a monthly profit target constraint in the financial problem (10). These constraints are based on profit targets, and failure to meet these targets determines the extent of “loss.” Such targets are commonly set by management, and adopted by traders for the purposes of hedging.

The price of gas used in the model includes the cost of delivery from the gas market. Moreover, the estimated cost of gas forwards is prorated according to the number of segments in the period/month (see (16)). There are more accurate ways to allocate the cost of gas forwards to each segment, but variables introducing usage-based allocation for each segment result in many more coupling variables between the

financial and generation problems, and that would limit the ease with which these submodels may be decomposed within an algorithm. Accordingly, we have adopted the formulation of (16) and (23). Although (23) is a financial constraint, it is included within the generation model. Because spot market prices and demands are modeled on a daily basis, it is best to incorporate risk control on a daily basis, and hence this constraint appears in the generation model. However the model may become infeasible in instances in which the target ML is unattainable. In such instances it may be recommended that the user include such a measure within a penalized objective function for the generation problem.

Finally, we discuss the objective function for the generation problem. Since many of the decision variables have been expressed without explicitly indicating dependence on the scenario, we begin by stating a scenario dependent objective function as follows.

$$Max \quad \sum_{t,j} [(D_{ij}P - SP_{ij}PS_{ij})H_{p(j)} - C_{ij}] - \sum_{t,e} ZP_{te} \quad (24)$$

This expression is a scenario dependent objective function which reflects the profit/loss of spot market activity, as well as the cost of power generation under one scenario. The complete objective function maximizes the expected profit obtained by accumulating a weighted average in which the scenario objectives (24) are weighted by their respective probabilities. In the absence of monthly and daily targets as specified in (10, 23), and liquidity (6, 7), such an objective may be interpreted as a risk neutral preference. However, the inclusion of (10, 23) provide down-side risk control and should be viewed as a mechanism for monthly and daily hedging. In addition, (6, 7) accommodate the user's perception on liquidity in the market. Collectively, these constraints incorporate dynamic piecewise-linear risk measures within the DASH model. Moreover in contrast to most other hedging schemes (see Eydeland and Wolyniec (2003)) where decisions are evaluated using a terminal goal, the DASH model incorporates "path-dependent" goals. This extension however requires a new algorithmic strategy which we describe next.

5 A NESTED COLUMN GENERATION DECOMPOSITION STRATEGY

The stochastic programming model presented in the previous section is a very large-scale optimization problem. Fortunately, the model is amenable to solution using decomposition techniques. This discussion is best motivated by studying the structure of the DASH model.

5.1 Here-and-now Problem Embedded with Wait-and-see Problems

In section 4, the DASH model was presented in terms of its two main submodels: the financial model and the generation model. The financial decisions in this model are made monthly. In each decision period, forward positions are chosen for each of the future delivery months. As the market evolves, these positions will be rebalanced in order to react to changes in the market. Because the forward decisions are made before forward prices are realized, they should be treated as here-and-now decisions. Thus forward scenarios will provide monthly evolution of prices and the here-and-now (financial) decisions will be required to be non-anticipative with respect to the forward prices scenario tree. We should reiterate that the forward prices refer to multiple stochastic processes including on-peak/off-peak power and gas.

Note that the focus of the DASH model is forward decisions, with generation costs merely providing the basis for economic decisions. Unlike the financial decisions, the generation decisions are assumed to be made on a segment-by-segment basis (i.e., 16 hour on-peak, and 8 hour off-peak). That is, for each forward scenario, generation decisions follow the evolution of load and spot prices during each month of a given scenario. This suggests a wait-and-see (adaptive) approach for the generation decisions. Nevertheless, it should be noted that if there are two scenarios that have the same monthly data history (i.e. loads, forwards and spot prices) until month t , then the generation history associated with these scenarios should also be the same until month t . We refer to this property as non-anticipativity with respect to the monthly data process.

Proposition: Suppose that the forward decisions are non-anticipative (with respect to the monthly data process). Assume that any tie-breaking rule for alternative optima in the generation problem is applied in such a manner as to ensure that generators are dispatched in the same order, given the same history of circumstances (among scenarios sharing the same partial data-path). Then the generation decisions also satisfy non-anticipativity with respect to the monthly data process.

Proof: Constraints (19) and (20) in the generation problem restrict the range of the segment index τ to the range $\tau \in \{j+1, \dots, \min(j+L_i-1, J_t)\}$. Hence generation decisions of one month are decoupled from generation decisions of any other, except as required by the forward decisions. Since the forward decisions in (11) are non-anticipative with respect to the monthly data process, and any two scenarios having the same monthly data history have the same loads and spot prices, consistency of the dispatching order

ensures that generation decisions for one scenario are also optimal for the other (in the respective generation problems). Hence the generation decisions are non-anticipative with respect to the monthly data process, and this completes the proof.

The resulting structure is therefore one that involves a multi-stage here-and-now stochastic program that has a sequence of large wait-and-see MILPs embedded within it. This structure turns out to be amenable to decomposition because the unit commitment model, an MILP, is much easier when treated as a wait-and-see problem, than as a here-and-now problem (Takriti, Krasenbrink and Wu, (2000), Nowak and Roemisch, (2000)).

In order to give the reader a sense of the magnitude of each scenario problem, the financial decisions involve $T(T-1)/2$ for each of the following types of forwards: on-peak electricity, off-peak electricity and gas. This certainly seems manageable for reasonable values of T (e.g. $T=6$ or 12). For the generation problem, each day corresponds to two segments (on-peak and off-peak), each served by $|I|$ generators. Hence, a month-long unit commitment model involves 56 segments, and thus $56|I|$ binary variables. By aggregating some of the generators, it is possible to solve such problems with reasonable computational effort. However, if the number of generators is large, then, it may be more convenient to solve a weekly unit commitment problem instead.

5.2 Nested Column Generation Decomposition Strategy

Our approach decomposes the stochastic program into three interrelated optimization problems which are motivated by a nested column generation (i.e. Dantzig-Wolfe) decomposition strategy. The algorithm is best motivated by studying the structure of the model. Figure 3 illustrates the original DASH model which consists of non-anticipativity constraints (see Appendix B) and all scenario sub-problems. In order to model the non-anticipativity constraints, we introduce the notation $n(s,t)$ to denote the node of the scenario tree associated with scenario s in period t . The gas and power forward variables are referred to as node variables and designated as $FG_{\pi}^{n(s,t)}$ and $FP_{\pi e}^{n(s,t)}$. Non-anticipativity constraints are enforced by requiring that all scenario decisions for gas and power forwards equal the appropriate node variables as shown in Figure 3.

Figure 4 summarizes the structure of each scenario problem, which consists of complete forward dynamics and all generation problems. The forward positions that are ultimately realized (on a delivery date) appear in the generation model as shown in Figure 4. Given that both Figures 3 and 4 depict block-angular matrices, it is natural to consider an algorithm in which column generation is carried out in a

nested manner; that is, we develop a non-anticipativity master problem (see Figure 5) whose responsibility is to seek non-anticipative forward decisions by choosing convex combinations of columns that represent each scenario. Similarly, Figure 6 depicts a master program for any scenario, and the columns generated here represent forward positions for the scenario. These positions are proposed by the generation problem. Thus, the nested column generation approach adopted here involves interactions between three problems described below (refer to Figures 4-6). The precise formulations are provided in Appendix B.

1. We use a non-anticipativity master problem to enforce non-anticipativity restrictions. Each scenario is represented by a collection of columns in this problem, and its goal is to find a convex combination of columns of each scenario that also satisfy non-anticipativity restrictions. Initially, a Phase I problem is solved to obtain a feasible solution to this problem. We note that the objective function coefficient for each column in this problem represents the total profit under a particular scenario of forward prices, spot prices, and electricity demand.
2. Given the price for achieving non-anticipativity from the master problem described above, a mid-level coordinating problem (scenario master problem) is formulated to make the best forward decisions for each scenario. This is essentially the same formulation as the financial problem described in section 4.1. However, the summation of forward decisions for a certain delivery period is once again represented via a convex combination of forward columns that are generated by the corresponding generation problem where the summation of forward decisions appear in the demand constraint (11). Here, we decompose the whole embedded generation problem (wait-and-see) into several generation problems, one for each delivery period (month).
3. Finally, the lowest level problem (generation problem), which generates the aggregation (i.e. summation) of forward columns for the higher levels, consists of a series of unit commitment problems. As with the second level coordinator, this problem assumes that the scenario is given, and a series of deterministic instances of the unit commitment problem are solved. The prices of forwards in this problem are modified by the dual prices from forward balance constraints in the mid-level coordinator (scenario master problem).

The use of Dantzig-Wolfe decomposition for bounding within a branch-and-bound method for specially structured integer programs is sometimes referred to as a branch-and-price algorithm (Barnhart et al. (1998)). In these methods, columns are generated for each node (of a branch-and-bound tree) at which a linear programming relaxation is solved. Within the context of stochastic programming, Lulli

and Sen (2004) have shown how stochastic integer programs can be solved using branch-and-price algorithms. For our specific application, we find that solving the root node of this branch-and-bound tree provides a very reasonable approximation because the difference between upper and lower bounds are within a small fraction of a percentage. Here, the upper bound is calculated by using the forward decisions recommended by DASH, and then running each scenario problem using binary variables (not a convex combination). Of course, the lower bound is provided by the DASH model. For the data sets reported in Section 6, the percentage difference between the upper and lower bounds was about 0.15%. Consequently, there was not much to be gained by going beyond the root node relaxation (i.e., branch-and-price was deemed unnecessary here).

A few more remarks regarding the advantages of our algorithm are in order. First, the nested approach allows us to maintain modularity, so that generation costing and financial decision are performed by coordinated, yet independent models. Moreover, such modularity promotes the ability to use a distributed computing environment, which has its own advantages (scalability, reliability etc.). Finally, we observe that in cases where a problem instance is the result of changes in a previously solved data set (e.g. due to changes in the probability measure, or having a new plant come on-line), the non-anticipativity master problem can be warm-started using previously generated columns, thus allowing efficient re-solves. Indeed this is the feature that makes it possible to provide decision support to analysts who gain insights from answers to “what-if” questions.

Figure 3: Original Problem Structure

Figure 4: Scenario Sub-problem Structure

Figure 5: Non-anticipativity Master Structure

Figure 6: Scenario Master Structure

In order to solve DASH problem with the above decomposition strategy, we apply CPLEX 7.0 to implement the nested column generation decomposition algorithm. The CPLEX Barrier Optimizer is used to solve Non-anticipativity master problem, and the Simplex Optimizer to solve the Scenario Master Problem. It is well known (see Carpenter, Lustig, and Mulvey (1991)) that interior point methods are superior to the Simplex method for stochastic programs with the split variable formulation, as in our case. These are very large, sparse problems. The generation problem is a mixed-integer program, and we use CPLEX Mixed Integer Optimizer to solve it.

6 EXPERIMENTAL RESULTS

This section is subdivided into several subsections. In the first of these subsections (6.1), we provide a brief summary of the data sets we use for testing the structure and performance of the model. In section 6.2 we study the behavior of the model; in particular, we study whether the inclusion of integer variables, and hedging constraints are necessary. Section 6.3 is devoted to several experiments, which are intended to test the performance of the new model against a benchmark. We also investigate the performance of the model against “out-of-sample” scenarios (see sections 6.3.2, 6.3.3, and 6.3.4) that may be generated in a variety of ways. Finally, section 6.4 provides a summary of experienced computational times.

6.1 Summary of the Data Set

The experiments reported in the following subsections are based on data obtained from Pinnacle West Capital, the largest investor-owned utility in Arizona. Before presenting the experimental results in details, we briefly summarize the data set. For all runs reported here, we used an initial position of forwards amounting to 15% of the averaged electricity load for a certain period. The electricity market data for our study reflects prices at Palo Verde, AZ, whereas, the gas market prices reflect data from Henry Hub, LA plus the cost of delivery. For the sake of this study, transaction costs were not included, although such calculations are easily accommodated within a simulation. In order to obtain results within reasonable computational resources, we aggregated plants into groups based on the type of fuel used by the plant. This reduced the number of generation-types to three: nuclear, coal, and gas. The DASH decisions that we are using for the simulation experiments are made once a month starting from January 2001 through May 2001. These months reflect a period of turmoil in the California energy markets, although the worst part of the crisis preceded that period. Finally we note that although the decision model aggregates the generation plants into groups based on fuel-type, the entire set of generators (say 50 plants) can be incorporated within the simulation.

6.2 Analysis of Model Structure

In developing the DASH model, we have made several choices regarding its structure. In making these choices we have tried to balance the realism underlying model formulation and its algorithmic tractability. In this subsection we study the behavior of the forward decisions as a function of alternative model

structures. There are two complicating factors that arise within the DASH model: one arises from the inclusion of integer variables, and the other arises from the inclusion of profit/loss control and max liquidity limit constraints. The complications due to the inclusion of integer variables are well known in the operations research literature. The other complication is more specific to power portfolio optimization, and calls for a brief discussion. It is noted that the DASH model would be a lot simpler without the profit/loss control and max liquidity limit constraints. If there are no such constraints, the model can be decomposed into many simpler separable sub-models, each corresponding to a future delivery month. We investigate the impact of model structure on decisions resulting from simpler models that do not accommodate the features incorporated in DASH. In order to investigate the impact of such features, we made Jan. 2001 runs based on a scenario tree with 200 scenarios. We first discuss the impact of the integer variables (operation decision variables) in the generation model. Then we discuss the influence of the profit/loss control (10, 23) and max liquidity limit constraints (6-7).

In order to investigate the influence of the integer variables, we executed the model under two alternative assumptions: one that includes integer variables; and another that relaxes them. In order to maintain confidentiality of the data, we report the percentage change in decisions instead of the real decisions (percentage = (forward decision with integer variables - forward decision without integer variables)*100/decision without integer variables). From Figure 7, we observe that incorporation of the integer variables significantly changes the forward decisions.

Figure 7: Influence of Integer Variables in the Generation Model

Note that the difference in gas forward decisions is not particularly significant. This may be attributed to the fact that the start-up constraints for gas turbines are not significant. However, significant differences arise with the power forward decisions. The latter may be attributed to the fact that other generation technologies, such as nuclear power plants, have significant start-up constraints.

Next, we investigate the impact of the profit/loss control (10, 23) and max liquidity limit constraints (6, 7). As before, we only report the percentage change in decisions (percentage = (forward decision with constraints (6, 7, 10, 23) - decision without these constraints)*100/decision without these constraints). Once again, the gas forward decisions are not significantly different, but the difference in power forward decisions continues to be significant, as shown in Figure 8. This suggests that forward decisions (for hedging) are less dependent on the most expensive fuel (i.e., last dispatched technology), and more dependent on the remaining technologies. Moreover, the fact that most of the differences are

Figure 8: Influence of Profit/Loss Control and Max Liquidity Limits

positive imply that the magnitude of power forward decisions are greater in the presence of the hedging constraints than the magnitude resulting from models that do not have these constraints.

6.3 Results of Simulation Experiments

The experiments reported here are intended to investigate the quality of decisions provided by the DASH model. In order to do so, we study the performance of DASH decisions in three settings: a) a backtesting exercise in which the decisions are evaluated with respect to the data observed during the first half of 2001, b) a simulation exercise in which the DASH decisions were evaluated when scenarios were generated from an extended scenario tree which involved four branches (instead of two) at each node, and c) a simulation exercise in which the price models are completely independent of the scenario tree models used by DASH. These price models are described subsequently in this section. Our intent is to test the performance of DASH under various “stresses” so that weaknesses can be identified.

For each monthly DASH run, we generated a scenario tree with 200 scenarios, and new trees were generated for successive monthly decisions. For comparative purposes, our tests will be carried out against a benchmark known as the fixed-mix policy, which is relatively common in this industry. In the portfolio optimization, fixed-mix refers to rebalancing strategy in which a constant mix of holdings is maintained, depending on the fluctuations in the market. Fleten, Wallace and Ziemba (2002) compare the fixed-mix strategy for the financial portfolio optimization with corresponding stochastic linear programming models. We perform similar experiments to compare a fixed-mix approach to the one proposed in this paper. However, in addition to the obvious difference between the models, there are two distinctions between our experiments, and those reported in Fleten, Hoyland, and Wallace (2002). First, we compare the decisions from DASH with those of the best fixed-mix strategy (among a finite set), and second, our simulations are performed to emulate a succession of decisions on a rolling horizon basis. Hence the comparisons we report are dynamic, and based on realistic data.

The fixed-mix strategy for power portfolio optimization may be described as follows.

On any contract date, an appropriate hedging position for a future delivery date (month) is one that is determined according to the following strategy. Make a prediction of expected demand and expected capacity for the delivery month. If expected demand exceeds expected capacity, then assume a long position for forwards in that delivery month, and the quantity of this transaction should be a fraction “ f ” of the difference. On the other

hand, if expected capacity exceeds expected demand, then one should assume a short position for forwards in the delivery month being considered. Once again, the quantity associated with this transaction should be a fraction “ f ” of the difference.

One can devise several variations on this scheme. For instance, instead of using expected demands and capacities, one may use scenarios to determine scenario-dependent strategies, and then use some weighted average to determine the exact mix. For our experiments we only tested the basic scheme outlined in the previous paragraph. However, we ran our simulations using several values of the fixed-mix fraction f , including 0, 0.1, 0.2, 0.3 and 0.4. We should also note that the simulations used here incorporate greater details on generation capacities than that used within the DASH decision model. Moreover, each comparison experiment between DASH policy and fixed-mix policy was made based on the same set of scenarios, though we did generate different sets of scenarios for each comparison simulation experiment.

6.3.1 The Backtesting Experiment

As outlined in the introduction, this experiment covers a five-month operating period from January 2001 through May 2001, with hedging decisions being made once each month. The decisions at the beginning of each month are, of course, made prior to observing the markets. Once the transactions are carried out, no portfolio changes are allowed for the rest of the month. During this period, we run a generation costing simulation based on weekly unit commitment and calculate the actual weekly profit. There are two steps for this procedure. In the first step, we forecast power demands and spot prices for the coming week, and run a weekly generation problem based on forward decisions for current period. In step 2, we calculate the actual profit based on the scheduled generation, actual demands and spot prices. Following this procedure, we can simulate the actual profits week by week during the current month. At the start of the next month, we once again use the fixed-mix policy to obtain the newly rebalanced positions, and the process resumes again. Moreover, since all rules carry out the same number of transactions, the difference in transactions costs between the different policies can be ignored. Finally, costs/revenues are calculated using the unit commitment (generation) model which includes spot market and forward activity. Thus revenues are accounted for in a delivery month only.

Figure 9: Comparisons between Fixed-mix Strategies

To maintain confidentiality, we will report performance in terms of fractions, with the profit of the best policy assuming the value of 1. Figure 9 is based on outputs that showed that using $f = 0.1$ provided

the most profitable fixed-mix strategy. Note that although some other fractions appear to be competitive during certain months, using $f = 0.1$ provides the overall winner among the fixed-mix strategies.

Next we proceed to experiments with the Stochastic Programming approach. These experiments were run with the same data as above, except that the fixed-mix hedging rule was replaced by decisions from the DASH model. During each month (January 2001 through May 2001), we run the stochastic programming model once. As before, decisions are made before observing market prices at Palo Verde, AZ and Henry Hub, LA. The planning period used within the decision model was five months long (i.e. $T = 5$). Hence as in the previous experiments, delivery dates of six months in the future were permitted in the model. Thus, the experimental setup and data are exactly the same as in the fix-mix study, and this permits comparisons between hedging decisions from stochastic programming and the fixed-mix rule.

Figure 10: Comparing DASH with Fixed-mix Strategy
(Backtesting Data)

In Figure 10 we use the best fixed-mix strategy ($f = 0.1$) as the basis for our comparisons. We made two series of runs with the DASH decisions: one using $\alpha = 0.3$ (i.e., 30% change allowed in the portfolio from one month to the next), and another series of runs using $\alpha = 0.5$ (i.e., 50% change allowed from month to month). Both series of DASH runs perform significantly better than the best fixed-mix strategy. It turns out that the series of revenues for $\alpha = 0.5$ exceeds that for the best fixed-mix strategy by approximately 7% per month, on average. This is a significant advantage in favor of the DASH model.

Before closing this subsection, we should comment on a certain initialization bias that results from restrictions imposed by the initial portfolio. Recall that when we allow a 50% change allowed in the portfolio from month to month ($\alpha = 0.5$), it takes about 2 months for the effect of the initial portfolio to wear off. It is therefore appropriate to focus our attention on the performance of DASH (with $\alpha = 0.5$) for months 3, 4 and 5. Similarly, when $\alpha = 0.3$, the output for months 4 and 5 are critical.

6.3.2 Results of Experiments with Synthetic Scenarios from an Extended Scenario Tree

In order to test the robustness of the decisions provided by the DASH model, we created synthetic scenarios and tested the decisions provided by the model against these scenarios. In conducting this phase of our experiments, we did not re-optimize to allow DASH to adapt to the observed (synthetic) scenario; instead, we used the decisions obtained from the backtesting experiment, and applied those to the synthetic scenarios. Hence the gains reported here are lower bounds on potential improvements. In these experiments, we follow the same simulation procedure as the one in the backtesting experiment.

The synthetic scenarios of this section were created in two steps. First, we create a series of forward prices from a discrete-time stochastic process with each time step reflecting the passage of a month. During each month, we draw a random number representing a particular outcome of forward prices. We allow four such outcomes in any month: {Max, High-Median, Low-Median, Min}. The values for these quantities are obtained from historical data as described in Section 3.2, and the probability of these outcomes is assumed to be $\{1/8, 3/8, 3/8, 1/8\}$. Note that over a five month period, we can create a total of 1024 scenarios. For the purposes of our tests, we generate 30 scenarios, against which the model is tested. For each of these scenarios, we also generate spot market prices, and loads. The latter are created in the same manner as described in Section 3.

Due to the initialization bias in the first two months (see Section 6.3.1), the comparison we report pertains to months 3, 4 and 5. This comparison involves the DASH model ($\alpha = 0.5$) and the fixed-mix strategy using $f = 0.1$. Figure 11 depicts the fraction of differences (i.e. (DASH – Fixed-Mix/Fixed-Mix)) over all 30 scenarios, for months 3, 4 and 5. Upon examining this figure, it is clear that DASH is the winner over most scenarios, with the magnitude of wins being significantly higher than the magnitude of losses. A summary of Figure 11 in terms of win-loss statistics is provided in Table 1.

The win-loss advantages in favor of DASH are unmistakable. Moreover, these results may underestimate the gains because the DASH model was not re-optimized based on observations of the evolving (synthetic) scenario.

Figure 11: Comparing DASH ($\alpha = 0.5$) with Fixed-mix Strategy ($f=0.1$)

Table 1: Win-Loss Statistics with scenarios from extended tree

Month\Statistic	Wins-Losses for DASH	Average Size of Wins	Average Size of Losses
3	22-8	3.29%	4.70%
4	19-11	11.19%	6.82%
5	19-11	21.68%	12.91%

6.3.3 Generating Synthetic Scenarios using Alternate Price Models

While the results of the previous subsection are encouraging, it remains to be seen how the model might perform under scenarios that are based entirely on more detailed econometric models. In order to do so, we modeled the price processes (gas, electricity forwards and spot) directly, rather than modeling “returns” as in formulating the scenario tree (see Section 3). These price models are similar in spirit to stud-

ies by Eydeland and Geman (1999), but as one might expect, our prices create scenarios. The details are discussed below, and are significantly different from previous work.

6.3.3.1 Gas Spot Pricing

Because gas generators are usually the last ones to be dispatched for merit ordered electricity generation, the marginal cost of electricity reflects gas prices. Hence, it is natural to first model gas spot prices, followed by gas spot price scenarios, and then electricity spot price scenarios. The reader may recall that the DASH model does not allow activities in the gas spot market, and the entire reason for studying gas spot prices is to help generate electricity spot prices. This process of generating gas spot prices is based on the following simple algorithm.

Step 0. Calculate trend (μ), and standard deviation (σ) using the first m days of gas spot prices.

Step 1. (Calculate n sample paths for the remaining days). For $i = 1, \dots, M$ (samples) we generate price trajectories for $d = m + 1, \dots, D$ as follows: $\varphi_{i,d} - \varphi_{i,d-1} = \mu + \sigma \varepsilon_{i,d}$.

This process simulates a discrete Brownian motion provided that the trend and volatility remain constant. However, if the forecasting interval (i.e. D) is long, then it may be worth considering time dependent trend and volatility (e.g., Taylor (1986), Engle (1982)). This can be accomplished by dividing the forecasting interval into shorter segments, and then estimating the associated trend and volatility parameters. Our experimental studies were performed for only a five-month interval, and for such a small interval, the estimated parameters were stable over the interval.

6.3.3.2 Spot Pricing of Electricity

In the course of this study, we have observed high correlation between Henry Hub gas spot prices and Palo Verde electricity spot prices. For instance in year 2000 the correlation coefficient between on-peak electricity and gas spot prices is 0.58. The same figure (i.e. correlation coefficient for 2000) for off-peak electricity and gas prices is 0.84. In year 2001 these correlations are 0.62 and 0.7 respectively. As indicated earlier, the reason for this high correlation is that the marginal fuel is gas in the Southwest. Hence electricity spot prices are very sensitive to gas spot prices in these wholesale markets. Using this information we propose the following relation.

$$(PS)_d^e = \lambda_0 + \lambda_1 (PS)_{d-1}^e + \lambda_2 \varphi_d \xi_d^e$$

Here, $e \in \{on, off\}$ identifies on-peak and off-peak electricity and the notation $(PS)_d^e, \varphi_d$ denote spot price for electricity and gas respectively for day d ; moreover ξ_d^e represents a random variable whose distribution we infer from the data. The motivation for this model is as follows. Since electricity spot prices are persistent, the model includes the lagged spot price and moreover, any change in gas spot prices has a nonlinear effect on the electricity spot price due to the inclusion of the third term in the above equation. Note that while on-peak and off-peak electricity spot prices are correlated, we believe that this correlation is due to their dependence on gas prices, which is reflected in the model stated above. From historical data we have determined the distribution of the error term in the model. For $\lambda_0 = 0$ and $\lambda_1 = \lambda_2 = 1$ we used the “best fit” function of Arena (<http://www.software.rockwell.com>) to determine the distribution for the error term. For the on-peak prices the distribution turned out to be the normal with mean 0.1 and 5.5 standard deviation; for off-peak corresponding distribution was Normal with parameters 0.04 and 2. Given initial electricity spot price and gas spot prices generated as in the previous section, we generate electricity spot price trajectories by sampling randomly from the above distributions. Likewise, given other gas spot price trajectories we obtain M on-peak and off-peak spot price scenarios.

6.3.3.3 Gas Forward Pricing

We formulate gas forward prices in two steps. First given initial forward prices for five months from now, we determine the trend and standard deviation. Then we use Brownian motion with drift process to generate scenarios. In the following, the unit of time is one month. Specifically

$$GF_{\tau,t} - GF_{\tau,t-1} = \mu + \sigma \varepsilon_t, \quad \varepsilon_t \sim N(0,1), \text{ for } t = 1, 2, \dots, 5, \text{ and } \tau \leq t,$$

where $GF_{\tau,t}$ denotes the gas forward price at time τ delivered at period t , μ and σ are the drift (trend) term, and standard deviation respectively. We initialize $GF_{0,0}$ as gas spot price at time 0. Given $\tau = 1$ we perform M independent samples of the random variable ε_t , for $t = 1, 2, \dots, 5$. Then we increment τ and for $t = \tau, \dots, 5$ use M independent samples of the random variable $GF_{\tau,t}(1 + \delta(\text{rand}[-l, l]))$ to obtain prices in period τ ; here rand is a uniform random number between $-l$ and $+l$, and l is the standard deviation of the price data collected in period $\tau - 1$ for all future prices (subsequent to period $\tau - 1$). Here δ is the scaling parameter of the random noise $\text{rand}[-l, l]$. One can estimate this scaling parameter based on observed price series.

6.3.3.4 Electricity Forward Pricing

It is known that some special characteristics of electricity such as non-storability, weather conditions, unscheduled outages and transmission disruptions make this market different from all other commodity markets. Eydeland and Geman (1999) note that market for power derivatives is not complete because hedging portfolios do not exist or are at least very difficult to identify. This incompleteness implies the non-existence of a unique derivative price, hence the wide bid-ask spread observed on certain contracts. Vehvilainen (2002) concludes that no analytical connection has been established even between the spot price and electricity forward prices, and that standard financial models may not necessarily apply in electricity markets. Consequently, Eydeland and Geman (1999) propose to approximate power forward prices as $PF_{\tau,t} = p_0 + \phi(w_{\tau,t}, L_{\tau,t})$, where p_0 , $w_{\tau,t}$, $L_{\tau,t}$, and ϕ designate base load price, forward price of marginal fuel (gas, oil, etc.), expected load (or demand) for date t conditioned on the information available at time τ , and a “power stack” function (see Eydeland and Geman (1999)) which can either be actual or implied from option prices. Analogous to their formulation, we propose on and off-peak forward prices as

$$PF_{\tau,t}^e = GF_{\tau,t} \cdot \exp(a^e L_{\tau,t}^e + b^e), \quad (25)$$

where $e \in \{on, off\}$ and a^e , b^e are coefficients. Note that this form of ϕ is rather complex, because load ($L_{\tau,t}^e$) is modeled by either ARIMA or GARCH models, and moreover, the gas forward prices ($GF_{\tau,t}$) are modeled as Brownian motion with some volatility term structure. Note that the standard option pricing models would arise if one were to assume that the load is normally distributed and gas forward prices follow geometric Brownian motion.

In order to implement (25), we follow a sequential procedure that is similar to the one used for gas forwards. Since forward prices are monthly we first convert daily load into average monthly load. Using this load, and gas forward prices we estimate coefficients a^e , b^e by using non-linear least squares. Given $\tau = 1$ we use M previously generated values of $GF_{\tau,t}$ and $L_{\tau,t}^e$ (see Sections 6.3.3.3 and 3.1 respectively.) to obtain M values of $PF_{\tau,t}^e$ for $t = 1, 2, \dots, 5$. Then we increment τ and for $t = \tau, \dots, 5$ use M independent samples of the random variable $PF_{\tau,t}^e \cdot (1 + \delta(\text{rand}[-h, h]))$ to obtain prices in period τ ; here rand is a uniform random number between $-h$ and $+h$, in which h is the standard deviation of the price data collected in period $\tau - 1$ for all future prices (subsequent to period $\tau - 1$). δ is the scaling parameter.

6.3.4 Results of Experiments with Synthetic Scenarios from Alternate Scenario Models

As with experimental results reported in Section 6.3.2, we consider the performance of the DASH model relative to fixed-mix models for months 3, 4 and 5. With the scenarios obtained in Section 6.3.3, the fixed-mix strategy using $f = 0.4$ was the most profitable among all fixed-mix strategies that were tried. Hence we compare this strategy with the DASH model ($\alpha = 0.5$). Figure 12 depicts the fraction of differences (i.e. $(\text{DASH} - \text{Fixed-Mix})/\text{Fixed-Mix}$) over all 30 (i.e. $M=30$) scenarios, for months 3, 4 and 5. Upon examining this figure, it is once again clear that DASH is the winner over most scenarios, with the magnitude of wins being significantly higher than the magnitude of losses. A summary of Figure 12 in terms of win-loss statistics is provided in Table 2. As in Table 1, the win-loss advantages in favor of DASH are unmistakable.

Figure 12: Comparing DASH ($\alpha = 0.5$) with Fixed-mix Strategy ($f=0.4$)

Table 2: Win-Loss Statistics from alternative synthetic scenarios

Month\Statistic	Wins-Losses for DASH	Average Size of Wins	Average Size of Losses
3	19-11	9.4%	5.0%
4	19-11	13.4%	4.8%
5	22-8	28.0%	12.1%

6.4 Computational Performance

In this subsection, we report the computational performance associated with the DASH problem and the nested column generation algorithm. For a DASH problem with 200 scenarios, the deterministic equivalent LP contains over 1.5×10^6 constraints and 7.5×10^5 variables. We used CPLEX7.0 (<http://www.ilog.com/products/cplex>) as the LP engine to solve the problem. And our experiments were conducted on a workstation SUN Ultra80 with two processors and 1 GB RAM.

As noted in Section 5.2, we used the CPLEX Dual Simplex Optimizer and Mixed Integer Optimizer to solve the scenario master problems and generation subproblems respectively. However, we chose CPLEX Barrier Optimizer to solve the non-anticipativity master problem. This choice was made after our initial experimentation revealed that the Barrier Optimizer performed much better than the Primal

Simplex option. In Table 3, we report the associated CPU times from DASH runs for each of the five periods.

Table 3: Computational CPU times (mins.)

Month	CPU Time
1	533.75
2	527.73
3	507.78
4	529.25
5	509.72
Average	521.65

Considering the huge size of realistic instances, such as those solved in this paper, the above computational performance is reasonable. However, we should reiterate that there are two main advantages of our solution method: a) it allows efficient warm-starts, and b) it is amenable for distributed computation. Both have the potential to speed up the method considerably. For instance, we created a data set in which the prices were perturbed by 5% from the original data set. We executed this program using a warm-start option in which columns from a previous run are used to initialize the master problems. For such a run, the program stopped in about 30 minutes, which is certainly acceptable. Our results using a distributed computing framework will be reported in a subsequent paper.

7 CONCLUSIONS

In this paper we described a large-scale modeling effort which demonstrates the viability of stochastic programming as a modeling paradigm for decision-making under uncertainty. We described our input-modeling effort, the DASH decision model, the DASH algorithm, and several experimental tests to study the robustness of the approach. Our experimental evidence suggests that the stochastic programming approach provides a powerful and robust tool for scheduling and hedging in wholesale electricity markets. There are several additional features (e.g. options, swaps, forced outages, hydroelectric generators etc.) that are being incorporated into the DASH model, and future papers will report on these extensions. The integration of market and production data with statistical models, optimization models, and simulation within one software framework requires fairly heavy investments in modeling and simulation technology. However, as demonstrated by our experiments, such an investment is very likely to bear fruit.

In terms of future work, we envision several avenues, some that will enhance the DASH model, and others that will enhance the algorithmic implementation. In the former category, we suggest the inclusion of game-theoretic approaches that will allow us to model the both strategic and tactical market power. Another extension includes the possibility of accommodating gas inventory within the model. While such models will be more realistic, they will create computing challenges that may call for the introduction of supercomputing into this arena. This of course leads to the second category of algorithmic enhancements that will be required in the future.

Appendix A

List of Symbols within Statistical Models

y, t, τ, d, s : Index for year, month, week, day, scenario respectively.

\tilde{L}_d : Load on the day d .

L_d : Load estimate on the day d .

g : Annual growth rate in load.

ε, η : Normally distributed variables.

β_i, α_j : Coefficients.

τ, κ : Contract week and delivery week, where $\tau = 1, 2, \dots, 52$, $\kappa = \{1, 2, \dots, N\}$.

e : Type of load segment, where $e \in \{on, off\}$

ω : Node number.

z : Standard normal random variate.

$\pi_{\tau\kappa e}$: Forward price with contract week τ , delivery week κ , segment e .

$r_{\tau,\kappa,e}$: Rate of change in price during the month starting in week τ .

$\sigma_{e,t}$: Standard deviation of spot returns in segment e in month t .

$r_{e,t,\omega}^f$: Daily equivalent of the forward return on node ω for month t .

$r_{e,d,\tau,\omega}^p$: Daily equivalent of the spot return on node ω for month t .

μ, σ : Trend, and standard deviation respectively.

$\varphi_{i,d}$: Gas spot price on day d . from sample i .

D : Forecasting interval.

M : Replication number.

$(PS)_d^e$: Power spot price.

ξ_d^e : Random variable.

λ_i : Coefficients.

$GF_{\tau,t}$: Gas forward price.

rand: Uniform random number generator.

δ : Scaling parameter.

p_0 : Base load price.

$w_{\tau,t}$: Forward price of marginal fuel (gas, oil, etc.)

$L_{\tau,t}$: Expected load for date t conditioned on the information available at time τ .

ϕ : “Power stack” function.

a^e, b^e : Coefficients

l, h : Estimated standard deviation of prices.

Appendix B

Nested Column Generation Algorithm for DASH

The nested column generation procedure applies the Dantzig-Wolfe decomposition principle in a hierarchical manner. At the highest level of the hierarchy, prices associated with non-anticipativity constraints are generated, and passed to scenario problems. The scenario problems (which form the second level of the hierarchy) generate imputed prices for forward contracts (on-peak power, off-peak power and gas forwards), and spot market activity. At the lowest level in the hierarchy, generation problems provide plans based on the prices provided by higher layers, and these form columns in the higher layer decision problems. The specific formulations are provided below.

First level of the nested column generation procedure:

Let p^s is the probability for scenario s , λ_j^s is a variable associated with scenario s and iteration j (first level), $Scenarioprofit_j^s$ is the profit associated with scenario problem s and iteration j . S is the set of scenarios. The notation $n(s,t)$ to denote the node of the scenario tree associated with scenario s in pe-

riod t . The gas and power forward variables are referred to as node variables and designated as $FG_{\pi}^{n(s,t)}$ and $FP_{\pi e}^{n(s,t)}$.

0. While the stopping rule is not satisfied, continue.

Stopping rule: the optimal reduced cost for each scenario problem s is greater than or equal to 0.

1. Solve the non-anticipativity master problem to get new dual prices by solving the following problem.

$$\begin{aligned} \text{Max} \quad & \sum_s p^s \sum_j \lambda_j^s (\text{Scenario profit}_j^s) \\ \text{s.t.} \quad & \\ & \sum_j \lambda_j^s FG_{\pi}^{sj} = FG_{\pi}^{n(s,t)} \quad s \in S, \tau \in \{1, \dots, T\}, t \in \{\tau, \dots, T\} \\ & \sum_j \lambda_j^s FP_{\pi e}^{sj} = FP_{\pi e}^{n(s,t)} \quad s \in S, \tau \in \{1, \dots, T\}, t \in \{\tau, \dots, T\}, e \in \mathcal{E} \\ & \sum_j \lambda_j^s = 1 \quad s \in S \end{aligned}$$

In this problem, non-anticipativity restrictions are enforced.

2. Given the dual prices from step 1, solve the scenario problems associated with each scenario s and return to step 0.

Second level of the nested column generation procedure:

Solve the Scenario Problem Associated with Scenario s

0. While the stopping rule is not satisfied, continue.

Stopping rule: the optimal reduced cost for each generation subproblem associated with period t is greater than or equal to 0.

1. Solve the scenario master problem

$$\begin{aligned} \text{Max} \quad & \text{Scenario profit} - \sum_{\tau, t \geq \tau, e} (\pi_{FP_{\pi e}} FP_{\pi e} + \pi_{FG_{\pi}} FG_{\pi} + \mu_s) \\ \text{Scenario profit} = \quad & \sum_t (-(\sum_e ZP_{te} + ZG_t) + \sum_i \beta_t^i (\text{Genprofit}_t^i)) \end{aligned}$$

s.t.

Constraints (1), (3)-(4), (6)-(7)

$$YP_{\pi e} = YP_{\pi e}^+ - YP_{\pi e}^- \quad \tau \in \{1, \dots, T\}, t \in \{\tau + 1, \dots, T\}, e \in \mathcal{E}$$

$$YG_{\pi} = YG_{(\tau-1)t} + FG_{\pi} \quad \tau \in \{1, \dots, T\}, t \in \{\tau + 1, \dots, T\}, e \in \mathcal{E}$$

$$\sum_i \beta_t^i YP_{ite}^i = YP_{t-1,te}^+ - YP_{t-1,te}^- + FP_{ite} \quad t \in \{1, \dots, T\}, e \in \mathcal{E}$$

$$\sum_i \beta_t^i YG_{it}^i = YG_{(t-1)t} + FG_{it} \quad t \in \{1, \dots, T\}$$

$$\sum_i \beta_t^i ZP_{ite}^i = \frac{J_t}{2} \sum_{\tau \in \{1, t\}} PP_{\pi e} FP_{\pi e} H_e \quad t \in \{1, \dots, T\}, e \in \mathcal{E}$$

$$\sum_i \beta_t^i ZG_{it}^i = \sum_{\tau \in \{1, t\}} PG_{\pi} FG_{\pi} \quad t \in \{1, \dots, T\}$$

$$\sum_{j \in \{1, J_t\}} [(D_{ij} P - \sum_i \beta_t^i SP_{ij}^i PS_{ij}) H_{p(j)} - \sum_i \beta_t^i C_{ij}^i] - \sum_{i, e \in \mathcal{E}} \beta_t^i ZP_{ite}^i \geq PT_t \quad t \in \{1, \dots, T\}$$

$$\sum_i \beta_t^i = 1 \quad t \in \{1, \dots, T\}$$

where the last six constraints are coupling constraints connecting the scenario master problem with the generation subproblems, π and μ are dual prices associated with non-anticipativity constraints in the non-anticipativity master problem (duals from the first level), $Genprofit_t^i$ (defined in step 2 below) is the profit associated with period t generation subproblem and iteration i , β_t^i is a variable associated with period t generation subproblem and iteration i .

2. Given the dual prices from step 1, solve the generation subproblem for each period t .

$$\begin{aligned} \text{Max} \quad & Genprofit - \sum_e (\pi_{YP_{\pi e}} YP_{\pi e} + \pi_{ZP_{\pi e}} ZP_{\pi e}) - \pi_{YG_{\pi}} YG_{\pi} - \pi_{ZG_{\pi}} ZG_{\pi} \\ & + \pi_{PT_t} (\sum_j (SP_{ij} PS_{ij} H_{p(j)} + C_{ij}) + \sum_e ZP_{ite}) - \mu_t \\ Genprofit = & \sum_j ((D_{ij} P - SP_{ij} PS_{ij}) H_{p(j)} - NP_t SN_{ij} - CP_t SC_{ij}) \end{aligned}$$

s.t.

Constraints (11)-(23)

where π and μ are dual prices associated with coupling constraints in scenario master problem .

ACKNOWLEDGMENTS

This research has been partially supported by a grant from the National Science Foundation. A preliminary (and abridged) version of this paper was presented at the Winter Simulation Conference, San Diego (December 2002). We thank several of our colleagues at Pinnacle West Capital, the holding company for APS. In particular, we thank Ned Lesnick (formerly at Pinnacle West Capital), Wen Lin, Bill Maese, Steve Murphy, and Gary Zhu for their insights and support. Ben Hobbs' comments on the previous version greatly enhanced this paper, and we thank him for his efforts. We also thank the referees, and the Associate Editor for thorough and timely reviews.

REFERENCES

ARENA. <http://www.software.rockwell.com>.

BACAUD, L., C. LEMARECHAL, A. RENAUD, C. SAGASTIZABAL. 2001. Bundle Methods in Stochastic Optimal Power Management: A Disaggregated Approach Using Preconditioners. *Computational Optimization and Applications*. 20, 227-244.

BARNHART, C, E.L. JOHNSON, G.L. NEMHAUSER, M.W.P. SAVELSBERGH, AND P.H. VANCE. 1998. Branch-and-Price: Column Generation for Solving Huge Integer Programs. *Opns. Res.* 36, 316-329.

BERTSEKAS, D.P., G.S. LAUER, N.R. SANDELL, AND T.A. POSBERGH. 1983. Optimal Short-Term Scheduling of Power Systems. *IEEE Transactions on Automatic Control*. AC-28, 1-11.

BIRGE, J.R., AND F. LOUVEAUX. 1997. *Introduction to Stochastic Programming*. Springer.

BOLLERSLEV, T. 1986. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*. 31, 307-327.

BOX, G.E.P., AND G.M. JENKINS. 1976. *Time Series Analysis: Forecasting and Control*. Holden-Day, San Francisco.

CARINO, D.R., AND W.T. ZIEMBA. 1998. Formulation of the Russel-Yasuda-Kasai Financial Planning Model. *Opns. Res.* 46, 433-449.

CARPENTER, T., I. LUSTIG, AND J. MULVEY. 1991. Formulating Stochastic Programs for Interior Point Methods. *Opns. Res.* 39, 757-770.

CARPENTIER, P., G. COHEN, J.C. CULIOLI, AND A. RENAUD. 1996. Stochastic Optimization of Unit Commitment: A New Decomposition Framework. *IEEE Transactions on Power Systems*. 11, 1067-1073.

CPLEX. <http://www.ilog.com/products/cplex>.

DUPACOVA, J., N. GROWE-KUSKA, AND W. ROEMISCH. 2000. Scenario Reduction in Stochastic Programming: An Approach Using Probability Metrics. *Mathematical Programming*. 95, 493-511.

ENGLE, R. 1982. Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation. *Econometrica*. 50, 987-1008.

ESCUADERO, L.F., J.L. DE LA FUENTE, C. GARCIA, AND F.J. PRIETO. 1996. Hydropower Generation Management Under Uncertainty via Scenario Analysis and Parallel Computation. *IEEE Transactions on Power Systems*. 11, 683-689.

ESCUADERO, L.F., AND M. PEREIRA. 2000. Power Play Goals. *OR/MS Today*. April issue.

EYDELAND, A, AND H. GEMAN. 1999. Fundamentals of Electricity Derivatives. *Energy Modeling and the Management of Uncertainty*. Risk Publications. 35-43.

EYDELAND, A., AND K. WOLYNIEC. 2003. *Energy and Power Risk Management: New Developments in Modeling, Pricing, and Hedging*. Wiley.

FLETEN, S.E., K. HOYLAND, AND S.W. WALLACE. 2002. The Performance of Stochastic Dynamic and Fixed-Mix Portfolio Models. *European Journal of Operational Res.* 140, 37-49.

FLETEN, S.E., S.W. WALLACE, AND W.T. ZIEMBA. 2002. Hedging Electricity Portfolios via Stochastic Programming. *Decision-Making Under Uncertainty: Energy and Power* (A. RYSZCZYNSKI AND C. GREENGARD (eds.)). Springer.

GENC, T., S. REYNOLDS, AND S. SEN. 2003. Dynamic Oligopolistic Games Under Uncertainty: A Stochastic Programming Approach. Available at <http://www.sie.arizona.edu/MORE/papers.html>.

GEORGANTZAS, N. C., AND W. ACAR. 1995. *Scenario-Driven Planning*. Quorum Books, Westport, CT.

GROWE-KUSKA, N., K.C. KIWIEL, M.P. NOWAK, W. ROEMISCH, AND I. WEGNER. 2002. Power Management in a Hydro Thermal System Under Uncertainty by Lagrangian Relaxation. *Decision-Making Under Uncertainty: Energy and Power*. (A. RYSZCZYNSKI, AND C. GREENGARD (eds.)) Springer.

HEITSC, H., AND W. ROEMISCH. 2003. Scenario Reduction in Stochastic Programming. *Computational Optimization and Applications*. 24, 187-206.

HOBBS, B.F, M.H. ROTHKOPF, R.P. O'NEILL, AND H-P. CHAO. 2001. *The Next Generation of Electric Power Unit Commitment Models*. Kluwer Academic Publishers, Dordrecht, The Netherlands.

HULL, J.C. 1997. *Options, Futures, and Other Derivatives*. Prentice-Hall, NJ.

- HOYLAND, K., AND S.W. WALLACE. 2001. Generating Scenario Trees for Multistage Decision Problems. *Management Science*. 27, 295-307.
- HOYLAND, K., M. KAUT, AND S.W. WALLACE. 2003. A Heuristic for Moment Matching in Scenario Generation. *Computational Optimization and Applications*. 24, 169-185.
- JACOBS, J., G. FREEMAN, J. GRYGIER, D. MORTON, G. SCHULTZ, K. STASCHUS, AND J. STEDINGER. 1995. SOCRATES-A System for Scheduling Hydroelectric Generation Under Uncertainty. *Annals of Opns. Res.* 59, 99-133.
- KENDALL, M.G., A. STUART, AND J.K. ORD. 1983. *The Advanced Theory of Statistics*. Charles Griffen, London.
- LULLI, G., AND S. SEN. 2004. A Branch-and-Price Algorithm for Multistage Stochastic Integer Programming with Application to Stochastic Batch-Sizing Problems. *Management Science*. 50, 786-796.
- NOWAK, M.P., AND W. ROEMISCH. 2000. Lagrangian Relaxation Applied to Power Scheduling in a Hydro Thermal System Under Uncertainty. *Annals of Operations Research*. 100, 251-272.
- TAKRITI, S., J.R. BIRGE, AND E. LONG. 1996. A Stochastic Model for the Unit Commitment Problem. *IEEE Trans. of Power Systems*. 11, 1497-1508.
- TAKRITI, S., B. KRASENBRINK, AND L. WU. 2000. Incorporating Fuel Constraints and Electricity Spot Prices into the Stochastic Unit Commitment Problem. *Opns. Res.* 48, 268-280.
- TAYLOR, S. 1986. *Modeling Financial Time Series*. John Wiley & Sons, New York.
- VEHVILAINEN, I. 2002. Basics of Electricity Derivative Pricing in Competitive Markets. *Applied Mathematical Finance*. 9, 45-60.
- WU, J., AND S. SEN. 2000. A Stochastic Programming Model for Currency Option Hedging. *Annals of Opns. Res.* 100, 227-250.