Intelligent Unified Control of Unit Commitment and Generation Allocation

Samer Takriti John R. Birge Erik Long

Department of Industrial and Operations Engineering
University of Michigan
Ann Arbor, Michigan 48109-2117

September 1994

Technical Report 94-26

Intelligent Unified Control of Unit Commitment and Generation Allocation

Samer Takriti John R. Birge Erik Long

October 23, 1994

Contents

1 Unit Commitment Problem

One of the most important problems in electrical power generation is the unit commitment problem. The primary concern of system operators is having enough capacity to meet demands during their peak load periods. The limited amount of hydro-electric energy stored in the dams and system reservoirs may not prove to be sufficient to respond to high demands. Therefore, costly thermal generation units must be used in order to make up for the supply shortage. The unit commitment problem refers to the task of finding an optimal schedule, and a production level, for each generating unit over a given period of time. The unit commitment decision indicates which generating units are to be in use at each point in time over the scheduling horizon [8]. Clearly, this problem is a multi-stage program with some 0/1 variables.

Throughout this paper, we assume that there are n generating units and that the duration of the study horizon is T. Normally, a 24-hour horizon is sufficient, but a longer horizon, one week, is needed if pump-storage units are to be considered. The state of each unit, i, at a time period, t, is represented by the variable u_t^i. A unit is on at time t if $u_t^i = 1$, and off if $u_t^i = 0$. The level at which unit i operates during a period, t, is $x_t^i \geq 0$. The minimum and maximum operating levels for each unit are g_i and G_i, respectively. The cost function, f_i', of operating a unit, i, at a level, x_t^i, is assumed to be a convex quadratic function of x_t^i. This function measures the total fuel and maintenance cost associated with each output level, x_t^i, in the feasible operating range. A start up cost, S_i, is incurred whenever the state of a unit changes from zero to one. The cost function, f_i', of each unit, i, is modified so that it takes into consideration the start up cost S_i. We will refer to this modified function by $f_i(x_t^i, u_{t-1}^i, u_t^i)$. Note that the function f_i is not continuous, hence it is non-differentiable and non-convex. When a unit is switched on, there is a minimum on-time requirement, that is, it has to be on for at least L_i periods. A similar constraint applies for the case when a unit is switched off; it has to be off for at least l_i periods.

The mathematical formulation of the model is

$$\begin{array}{lrcl} \min_{x,u} & \sum_{i=1}^{n} \sum_{t=1}^{T} f_i(x_t^i, u_{t-1}^i, u_t^i) & & \\ \text{subject to} & \sum_{i=1}^{n} x_t^i & \geq & d_t, \; t = 1, \ldots, T, \end{array} \tag{1}$$

where $d_t \geq 0$ represents the demand for electricity during period t. In addition to the previous constraints, each unit must satisfy the minimum on-time, minimum off-time, and minimum and maximum operating level constraints. The problem in (1) is a large-scale mixed-integer quadratic program. Many approaches have been proposed to solve this problem; they can be classified into branch-and-bound methods, dynamic programming, priority ordering, and Lagrangian relaxation method [1]. The first two techniques are satisfactory from the theoretical point of view, but they are practically intractable due to the large size of storage required to implement them on a computer. The third approach is a greedy strategy and does not guarantee an optimum solution in general. The Lagrangian relaxation technique seems to be the most efficient; it attempts to solve the problem indirectly by solving the dual problem.

Muckstadt and Koenig [8] appear to have been the first people to address the unit commitment problem and to suggest a sound technique for solving it. They use a Lagrangian relaxation, which decomposes the given problem into smaller sub-problems. Each sub-problem corresponds to the problem of minimizing the cost of operating a generator in the electrical system over the study horizon. Dealing with individual generators simplifies the task of representing the constraints that depend on the state of the generator from period to period, such as minimum on-time and minimum off-time requirements. Dynamic programming is then used to solve the sub-problems, and a lower bound on the optimal cost of the primal problem is obtained. The authors use branch-and-bound to enumerate all possible states of the system efficiently. At each node of the branch-and-bound tree, the states of some generators at certain time periods are fixed, and the Lagrangian dual of the problem at that node provides a lower bound that helps in pruning the search tree. However, in order to maintain high efficiency, they do not allow more than a small number of Lagrange multiplier updates at each node. The update is performed using a subgradient method that approximates the steepest-ascent direction since the dual gradient is not unique at some points. The dual solution at some nodes may provide a feasible solution which further reduces the size of the search tree. This strategy, that is, Lagrangian relaxation coupled with branch-and-bound, is common in integer programming and guaranteed to find an optimum solution.

The previous technique may fail in practice, as shown by the authors, due to the large number of nodes that need to be studied. Bertsekas, Lauer, Sandell, and Posbergh [5] use a similar Lagrangian relaxation technique but, rather than using branch-and-bound, they update the Lagrange multipliers and resolve the problem. The process is repeated until the duality gap is small enough. In order to accelerate the calculations, the cost function of each generator's sub-problem is approximated by a differentiable function. The approximate dual problem is solved using a quadratically convergent constrained version of Newton's method, which makes use of the gradient and the Hessian matrix of the approximate dual function. The main contribution of this paper is providing an upper bound on the size of the duality gap when the number of generators is greater than the study horizon. Their bound is given by $\frac{2T(S^*+TC^*)}{\pounds^*}$ where T is the length of the planning horizon, S^* is the maximum start-up cost over all generators, C^* is the maximum cost of operating a generating unit for one period, and $\pounds^*$ is the optimum dual functional value. The previous result is consistent with a well known property of Lagrangian relaxation: "the duality gap is typically quite small (in relative terms) if the number of separable terms is large, and in fact becomes smaller as the number of separable parts increases" [4].

A very similar approach is suggested by Merlin and Sandrin [7]. They use the same Lagrangian relaxation to separate the problem. Pump-storage hydro units are treated as any other generator. They then use the subgradient algorithm to update the Lagrange multipliers as proposed by Poljak [10]. Their numerical results are very impressive given the computing power at that time; they solved a system with 172 units over 48 hour horizon to within 0.42% in two minutes. These results may be the consequence of using linear cost functions for the generators while other references use piecewise-linear cost functions.

Zhuang and Galiana [14] provide a heuristic that can be used at termination of the dual maximization procedure if the resulting primal solution is infeasible. The main idea is to increase, with a moderate step size, the penalties on the violated constraints iteratively until a feasible solution is found.

In this paper, we use a Lagrangian relaxation technique to decompose the problem into smaller sub-problems. Each sub-problem is solved using dynamic programming. We use the special structure of the single-generator sub-problem to reduce the size of the state space of the dynamic program, and to obtain a more efficient formulation than that used classically. We use a subgradient technique to solve the dual maximization problem [10, 3].

2 Lagrangian Relaxation Approach

In order to make the program in (1) separable, a Lagrange multiplier, $\lambda_t \geq 0$, is associated with each of the constraints $\sum_{i=1}^{n} x_t^i \geq d_t$. This choice of relaxation decomposes the problem into n single-generator sub-problems. Constraints that depend on the change in the state of a generator from period to period, such as, minimum up-time and minimum down-time, become easy to implement. The Lagrangian dual problem has the form

$$\max_{\lambda \geq 0} \pounds(\lambda)$$

Figure 1: The Lagrangian Relaxation Approach described in Section 2

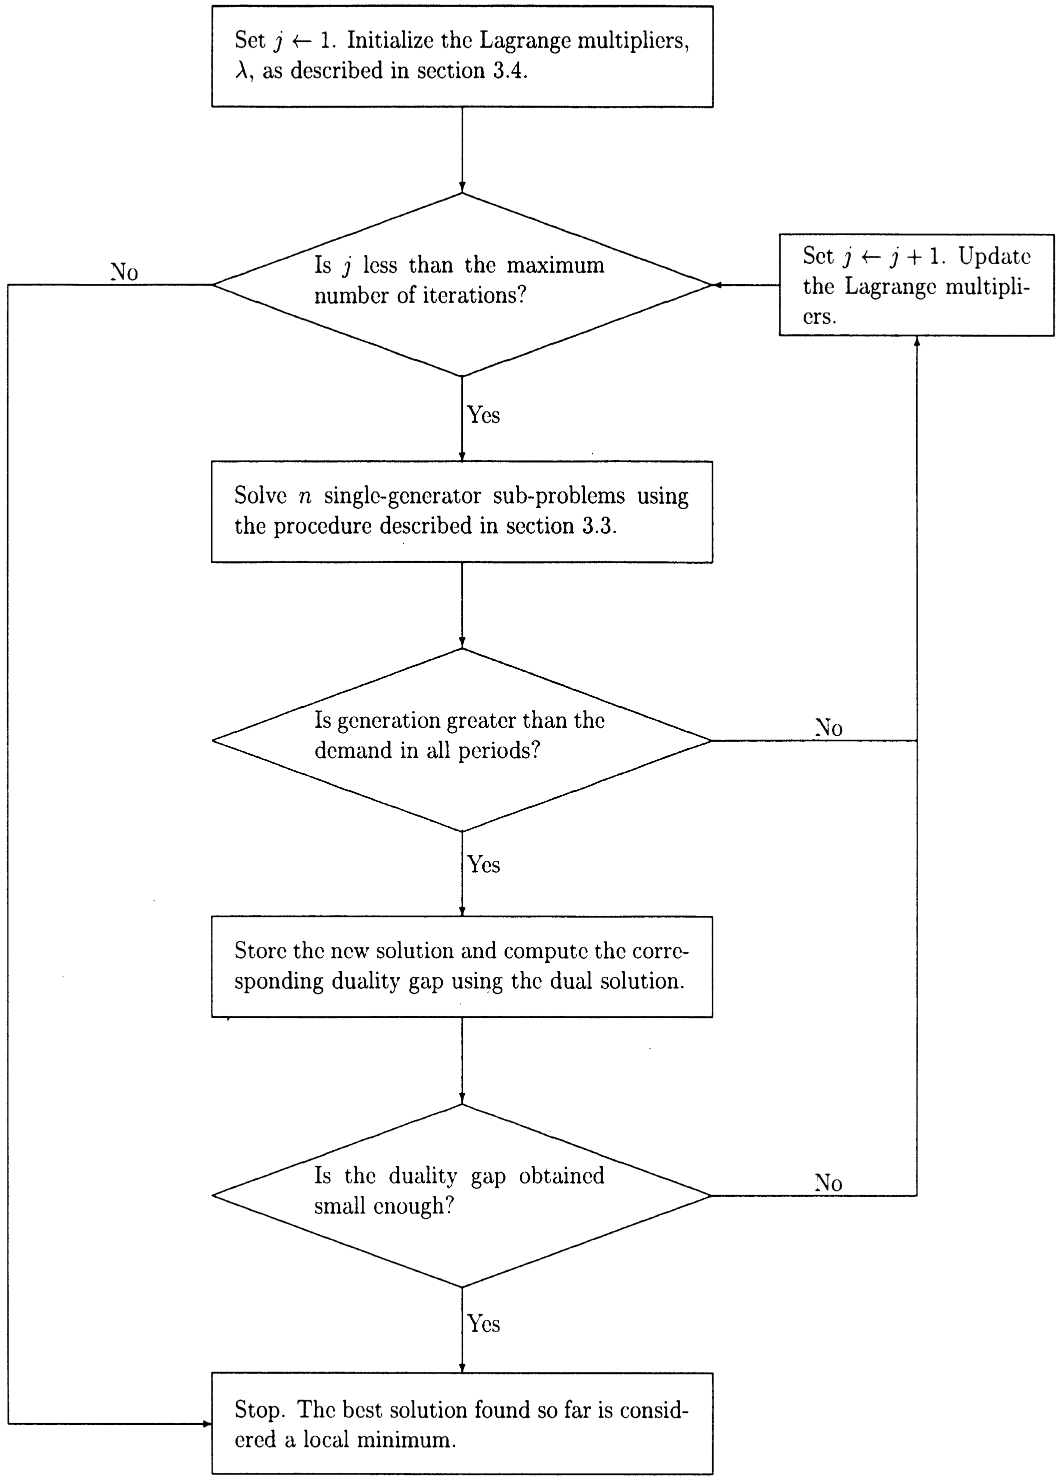

where

$$\pounds(\lambda) = \min_{x,\ u} \sum_{i=1}^{n} \sum_{t=1}^{T} f_i(x_t^i, u_{t-1}^i, u_t^i) - \sum_{t=1}^{T} \lambda_t (\sum_{i=1}^{n} x_t^i - d_t)$$

subject to unit minimum and maximum operating levels, and minimum on-time and off-time constraints. The Lagrangian dual function can be rewritten as

$$\pounds(\lambda) = \min_{x,u} \sum_{i=1}^{n} \sum_{t=1}^{T} f_i(x_t^i, u_{t-1}^i, u_t^i) - \frac{\lambda_t}{n}(x_t^i - d_t).$$

For a given λ, the value of $\pounds(\lambda)$ is computed by solving n mixed-integer quadratic programs

$$F_i(\lambda) = \min_{x,\ u} \sum_{t=1}^{T} \left(f_i(x_t^i, u_{t-1}^i, u_t^i) - \frac{\lambda_t}{n}(x_t^i - d_t) \right),\ i = 1,\ \ldots,\ n, \tag{2}$$

and adding the resulting values of $F_i(\lambda)$ over all generators. Note that each of these programs is independent of the others, which motivates a parallel implementation of this step. The optimization problem in (2) is called the i^{th} single-generator sub-problem since it only depends on the specifications of generator i.

If the primal solution corresponding to a given λ is feasible, then the Lagrange function value is a lower bound on the primal objective function value (weak duality). It follows that the maximum dual objective function value is a lower bound on the primal optimum objective function value. The difference between the optimum value of the original program and the optimum value of the Lagrangian dual problem is called the duality gap, and is expected to be strictly positive since the feasible region of the relaxed problem is not convex. One can also show that the Lagrange function is concave, hence continuous, in the parameter λ and bounded, which implies that a global optimum can be reached by using an appropriate convex programming method. The two previous remarks are the main attraction in this technique since one can replace a hard primal problem by that of maximizing concave function. Note that the Lagrange function is not differentiable at all points λ, which complicates the process of maximizing the dual function.

In order to solve the dual problem, a starting point, λ, is chosen according to some criterion, then the value of the Lagrange function, $\pounds(\lambda)$, is computed by solving n minimization problems. If the resulting primal solution is feasible, then $\pounds(\lambda)$ is a lower bound and $\sum_i F_i(\lambda)$ is an upper bound on the optimum value of (1). If the difference between the upper and lower bounds is relatively small, the procedure terminates. Otherwise, the process is repeated by choosing a better Lagrange multiplier. A more detailed description of the steps of this process is provided in the following sections.

3 Solving Single-Generator Sub-Problem

We used dynamic programming to solve the mixed-integer quadratic program in (2). At each time period or stage, a generator can be in one of two states: on or off. The formulation needs to take into consideration the minimum on-time and minimum off-time constraints; that is, a variable must represent the number of periods spent so far in the current state. Many researchers suggested using $L + l$ nodes at each stage of the dynamic programming graph [8, 7, 14]; no change in the generator state is permitted at nodes where a generator has been on/off for less than L/l periods. The following observations reduce the number of nodes in the dynamic programming graph substantially:

- Since one can only make decisions at those nodes in which a generator has been on/off for at least L/l periods, one needs only to consider these nodes in the dynamic programming formulation.
- Furthermore, the decision set of a generator that has been on/off for longer than L/l is the same as that of a generator that has been on/off for exactly L/l periods.

So, it suffices to consider two nodes at each stage: on for at least L periods and off for at least l periods. In other words, the number of nodes in the dynamic programming graph is $2T$ which seems to be best in this case.

Figure 2: Traditional Dynamic Programming Formulation for Single-Generator Sub-Problem

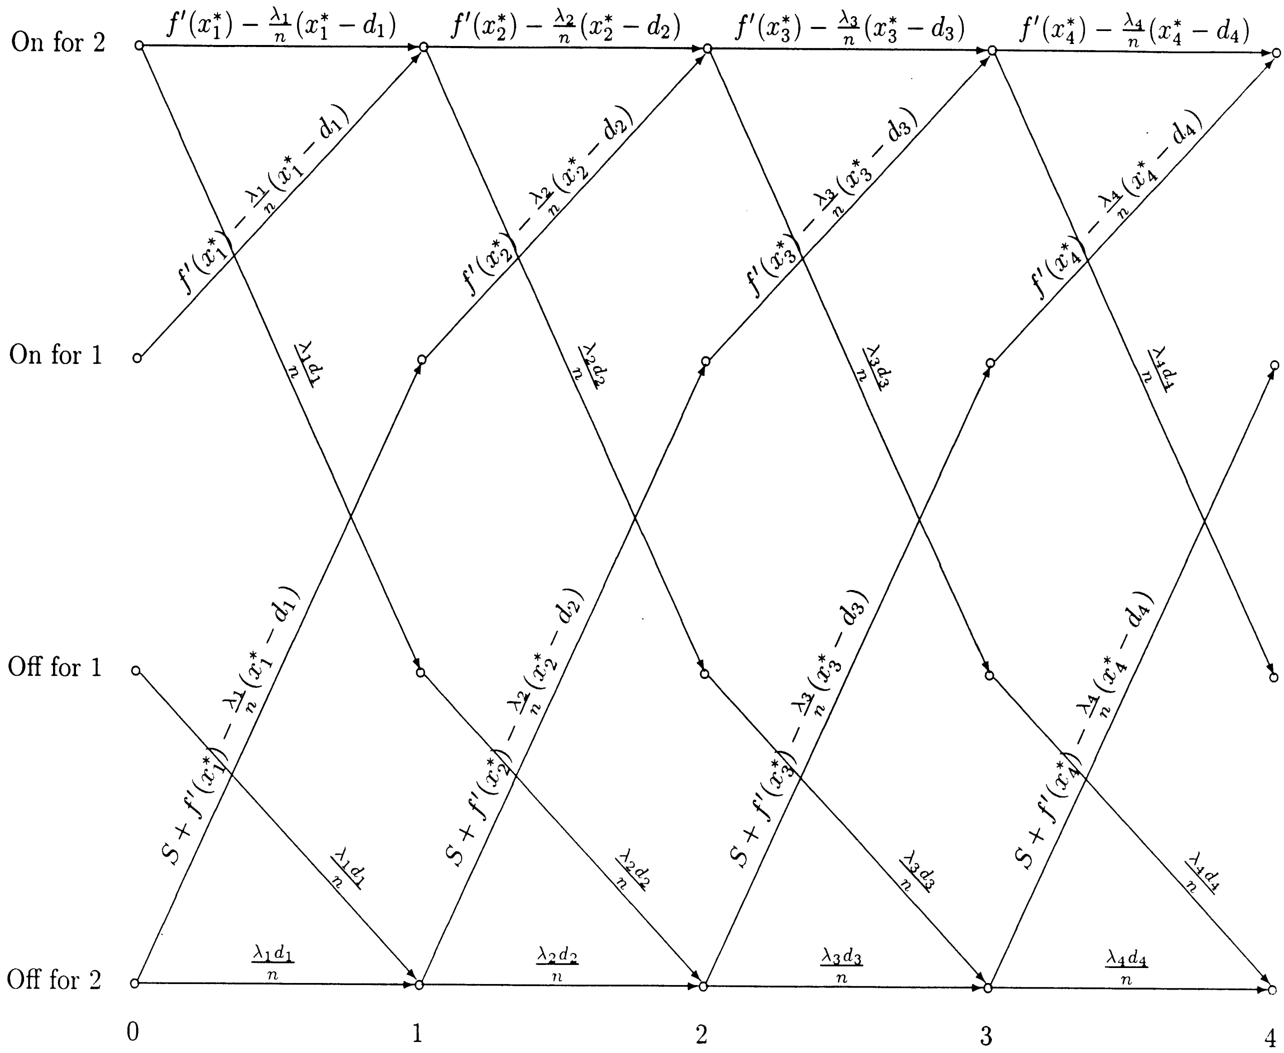

Each node is represented by two indices. The first index indicates the state of the generator: on/off, and the second one indicates the time or the stage variable. For instance, the node (off, 3) indicates that at time period 3, the current generator has been off for at least l periods. Hence, this generator can be switched on, that is, move to node $(\text{on}, 3+L)$, at a cost of

$$c(\text{off},3) + \sum_{t=4}^{3+L} \min_{x_t^i} \left(f_i'(x_t^i) - \frac{\lambda_t}{n}(x_t^i - d_t) \right) + S_i.$$

On the other hand, we can keep the generator off, that is, move to the node (off, 4) at a cost of $c(\text{off},3) + \lambda_t d_t / n$. The cost at any node, c, represents the minimal cost needed to reach this node in the graph given the initial state at time 0.

Given a Lagrange multiplier, λ, the following algorithm describes the steps required to find an optimum solution for the single generator problem[1].

- Initialization. Let x_t^*, $t = 1, \ldots, T$, be an optimum solution for the convex optimization problem:
$$\min_{x_t} f'(x_t) - \frac{\lambda_t}{n}(x_t - d_t), \; t = 1, \ldots, T,$$
such that x_t^* is within the operating range of that generator. Let $c(\text{off}, t) = \infty$ for all $t = 0, \ldots, T+l-1$, and $c(\text{on}, t) = \infty$ for all $t = 0, \ldots, T+L-1$. If the initial state is off then let $c(\text{off}, 0) = 0$; otherwise, let $c(\text{on}, 0) = 0$. Set $t \leftarrow 0$.

- General Step.

  1. $c(\text{off}, t+1) = \min\left\{c(\text{off}, t+1),\; c(\text{off}, t) + \frac{\lambda_{t+1}}{n} d_{t+1}\right\}$.
  $c(\text{on}, t+L) = \min\left\{c(\text{on}, t+L),\; c(\text{off}, t) + \sum_{\tau=t+1}^{\min\{t+L,T\}} \left(f'(x_\tau^*) - \frac{\lambda_\tau}{n}(x_\tau^* - d_\tau)\right) + S\right\}$.
  $c(\text{on}, t+1) = \min\left\{c(\text{on}, t+1),\; c(\text{on}, t) + f'(x_{t+1}^*) - \frac{\lambda_{t+1}}{n}(x_{t+1}^* - d_{t+1})\right\}$.
  $c(\text{off}, t+l) = \min\left\{c(\text{off}, t+l),\; c(\text{on}, t) + \sum_{\tau=t+1}^{\min\{t+l,T\}} \frac{\lambda_\tau}{n} d_\tau\right\}$.
  2. Set $t \leftarrow t+1$.
  3. If $t < T$ then go to step 1.
  4. Check the nodes (off, t), $t = T, \ldots, T+l-1$, and (on, t), $t = T, \ldots, T+L-1$, to find the optimum value of the objective function corresponding to the given λ.

Note that the number of functional evaluations needed to compute the summation in step 1 is only one. A predecessor pointer is associated with each node. Whenever the cost, c, of a node changes, the pointer is updated so that it points to that node from which the transition has the lowest cost. At termination, one can use the predecessor pointers recursively to find an optimal strategy.

[1]In order to simplify the notations, the index i will be skipped throughout the steps of the algorithm.

Figure 3: Suggested Dynamic Programming Formulation for Single-Generator Sub-Problem

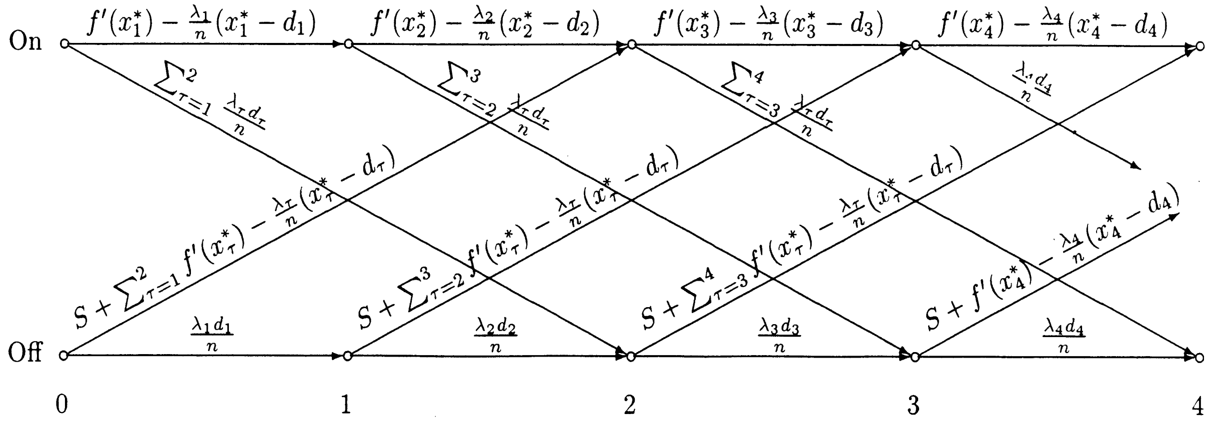

4 Initial Lagrange Multipliers

Clearly, the choice of the initial values of the Lagrange multipliers is very crucial since it affects the number of iterations, hence, the execution time of the algorithm. The following two sections describe techniques that we found to provide very good approximations for λ efficiently.

4.1 Quadratic Approximation

Since the cost function of each generator, $f_i'(x)$, is a quadratic convex function, one can obtain a starting λ by solving the quadratic program:

$$\begin{array}{rrcl} \min_x & \sum_{i=1}^{n}\sum_{t=1}^{T} f_i'(x_t^i) + \frac{S_i}{G_i}x_t^i & & \\ \text{subject to} & \sum_{i=1}^{n} x_t^i & \geq & d_t,\ t=1,\ \ldots,\ T \\ & 0 \leq x_t^i & \leq & G_i. \end{array} \tag{3}$$

In other words, the minimum on and minimum off time requirements, and the lower bounds on generation capacities are relaxed. The term $\frac{S_i}{G_i}x_t^i$ takes into consideration the start up cost of each unit.

The program in (3) is a convex quadratic program that can be solved efficiently using any quadratic programming algorithm, such as Lemke's. Note that the previous quadratic program is separable, and can be written as the sum of T quadratic programs, each of which has exactly one constraint

$$\begin{array}{lrcl} \min_x & \sum_{i=1}^n f_i'(x_t^i) + \frac{S_i}{G_i} x_t^i & & \\ \text{subject to} & \sum_{i=1}^n x_t^i & \geq & d_t \\ & 0 \leq x_t^i & \leq & G_i. \end{array} \tag{4}$$

The dual variables associated with the demand constraints, $\sum_{i=1}^n x_t^i \geq d_t$, are used as initial Lagrange multipliers of (1).

We can make the program in (4) separable by associating a Lagrange multiplier, $\mu_t \geq 0$, with the demand constraint. Now, we have n different quadratic convex programs; each of which has one bounded decision variable, x_t^i. The solution is found easily by finding the point, $\overline{x}_t^i(\mu_t)$, at which the derivative of the objective function vanishes, and then choosing the closest feasible point, $x_t^{i*}(\mu_t)$, to $\overline{x}_t^i(\mu_t)$ as an optimum solution. The value of the dual multiplier, μ_t is then updated, and the process is repeated until $\sum_{i=1}^n x_t^{i*}(\mu_t)$ is equal to d_t. The final values of μ_t are used as an approximation for the Lagrange multipliers for the mixed-integer quadratic program in (2). In order to improve the efficiency of the procedure, one can build a table that contains the demand values corresponding to different marginal costs, μ_t, at the beginning of the solution process.

4.2 Linear Approximation

Since the coefficients of the quadratic terms in the cost functions are relatively small, one can approximate $f_i'(x)$ using a linear function, $c_i x$, such that

$$c_i = \frac{f_i'(G_i) + S_i}{G_i}.$$

Here again, the term $\frac{S_i}{G_i}$ is added to take the start up cost into consideration. Since each generator has a constant marginal cost in this case, that is, the cost per Megawatt-Hour does not depend on the load on the generator, we can sort the generators in the ascending order of their marginal costs, c_i. If we ignore the minimum generation capacity requirements, and the minimum on and off time constraints, we can find easily which units to put on line to meet any given demand at a minimal cost. The problem simply is to minimize a piecewise-linear convex function subject to the constraint of meeting the demand, and is solved by choosing the generators in the order of c_i.

The previous procedure can be further simplified and accelerated by building a table that contains the slopes, c_i, in ascending order in one column, and the corresponding cumulative maximum generation capacities in the other column. For any given demand at time t, one can find the range in which the demand falls, and use the corresponding c_i as an approximation for the Lagrange multiplier at that time period. The resulting function is convex and denoted by $\Lambda(.)$ for use in subsequent sections. This approximation was implemented and, surprisingly, the dual objective function value of the initial solution was always within two percent of the optimum dual objective function value.

5 Updating Lagrange Multipliers

If the primal solution corresponding to the current dual solution, λ, is infeasible, or if the primal solution is feasible but the duality gap is not small enough, then a new dual vector must be used so that a better primal feasible solution is obtained. Since the dual objective function is concave, but not everywhere differentiable, it may not have a gradient everywhere and a subgradient technique is necessary to maximize this function. A vector ξ is called a subgradient of $\pounds(.)$ at $\overline{\lambda}$ if

$$\pounds(\lambda) \leq \pounds(\overline{\lambda}) + (\lambda - \overline{\lambda})^T \xi.$$

If the subgradient is unique at a point λ, then it is the gradient at that point. The set of all subgradients at λ is called the subdifferential, $\partial £(\lambda)$, and it is a closed convex set. A necessary and sufficient condition for optimality in subgradient optimization is $0 \in \partial\lambda$.

We used the subgradient optimization algorithm; it generates a sequence of points using the rule

$$\lambda^{l+1} = \lambda^l + \alpha_l \xi^l$$

where ξ^l is any subgradient of $£(.)$ at λ^l [13]. The step size, α_l, has to be chosen carefully in order to achieve a good performance by the algorithm. Poljak [10] has shown that the necessary and sufficient conditions to guarantee the convergence of the sequence λ^l to an optimum solution is

$$\alpha_l ||\xi^l|| \to 0 \text{ and } \sum_l \alpha_l ||\xi^l|| \to \infty, \tag{5}$$

where $||.||$ denotes the Euclidean norm. In fact, a geometric convergence rate can be achieved if we choose

$$\alpha_l = \frac{£^* - £(\lambda^l)}{||\xi^l||^2},$$

where $£^*$ is the optimal value of the Lagrange function. A problem here is that $£^*$ is not known in advance. We adopted a simple updating rule that satisfies (5) but has no known theoretical convergence rate

$$\alpha_l = \frac{1}{a + bl},$$

where a and b are constants; they are chosen according to the given data. We also imposed an upper bound on the step size in order to maintain dual feasibility.

Recall that our problem is to maximize

$$£(\lambda^l) = \min_{x,\, u} \sum_{i=1}^{n} \sum_{t=1}^{T} f_i(x_t^i, u_{t-1}^i, u_t^i) - \sum_{t=1}^{T} \lambda_t^l (\sum_{i=1}^{n} x_t^i - d_t)$$

subject to $\lambda^l \geq 0$. A subgradient of the function $\pounds(\lambda^l)$ with respect to λ^l is given by

$$\nabla_t \pounds(\lambda^l) = d_t - \sum_{i=1}^{n} x_t^i, \; t = 1, \ldots, T.$$

If $\nabla_t \pounds(\lambda^l)$ is not equal to zero, then it is an ascent direction and the value of $\pounds(\lambda^l)$ will increase by moving to $\lambda_t^l + \alpha_l(d_t - \sum_i x_t^i)$, $\alpha_l \geq 0$. Since we want to maintain dual feasibility, that is, $\lambda_l \geq 0$, the maximum step length is

$$\alpha_{\max} = \frac{\lambda_t^l}{\sum_{i=1}^{n} x_t^i - d_t} \tag{6}$$

when $\sum_{i=1}^{n} x_t^i \geq d_t$. In the case where the primal solution is infeasible, $\sum_{i=1}^{n} x_t^i \leq d_t$, there is no upper bound on the value of α_l, but we imposed a maximum step length in the implementation. The previous points can be summarized as follows:

- If $\lambda_t^l = 0$ and $\sum_i x_t^i \geq d_t$, then set $\lambda_t^{l+1} \leftarrow \lambda_t^l$.
- If $\lambda_t^l > 0$ and $\sum_i x_t^i \geq d_t$, then set $\lambda_t^{l+1} \leftarrow \lambda_t^l + \alpha_l(d_t - \sum_i x_t^i)$, where α_l does not exceed $\frac{\lambda_t^l}{\sum_i x_t^i - d_t}$.
- If $\sum_i x_t^i \leq d_t$, then set $\lambda_t^{l+1} \leftarrow \lambda_t^l + \alpha_l(d_t - \sum_i x_t^i)$.

The previous procedure is very efficient and simple to implement.

6 Pump-Storage Hydro Unit

The problem of optimizing pump-storage units has been studied extensively. Bannister and Kaye [2] discuss the case of a single reservoir and a single production facility that represents all available generating units. The operating cost of the production facility is assumed to be piecewise linear and convex. The authors present a linear programming formulation, and then exploit the special structure of the constraint matrix in order to develop an efficient dynamic programming formulation, which in turn is used to solve the model over the study horizon. However, this paper assumes that the cost function of operating the hydro unit is the same for all time periods, and hence, independent of the load on the system. In reality, the cost associated with pumping water into the reservoir of the hydro unit, and the savings resulting from releasing water from the storage depend on the system's load and its current state. The previous difficulty is a result of the non-linearity of the generator's cost functions and the start up and shut down costs.

Merlin and Sandrin [7], and Aoki, Itoh, Satoh, Nara, and Kanezashi [1] provide a more general solution by assuming that the pump-storage unit is a special generating unit which can produce and consume electricity. The cost function of the hydro unit is assumed to be a linear function of the amount of water released from, or pumped into, the reservoir. The cost coefficient at each time period, t, is the marginal cost, λ_t, of producing electricity at that time period. The water utilization problem is then solved using dynamic programming. In each iteration, the values of λ change, which consequently changes the cost associated with the hydro unit, hence its policy. Our numerical experience with this technique indicates that the number of iterations needed to solve the problem is relatively large. The reason may be the strong relationship between all time periods which is introduced into the model by treating the hydro unit as any other generator. Another problem with using a linear cost function for the hydro unit is that, given a small marginal cost, λ_t, at period t, the dynamic programming technique tends to pump a lot of water and may fill the reservoir in that period. On the other hand, if λ_t is relatively large, the dynamic programming solution may use most of the water at that period.

7 Efficient Pump-Storage Unit Formulation

In order to reduce the number of iterations, one can ignore the pump-storage hydro unit, solve the thermal unit commitment problem to obtain the marginal cost, λ_t, at each period,

t, and then solve the water allocation problem using dynamic programming. Now, the load on the thermal units is updated, and the unit commitment problem is solved again to obtain a new set of λ. The process is repeated until the values of λ converge. Note that at each iteration, a large mixed-integer quadratic program needs to be solved.

The previous procedure is not very efficient since the initial values of λ are far from being optimal, and the number of iterations needed is large. One can obtain a better initial set of Lagrange multipliers by using the approximating procedure described in section 4.2. Notice that instead of solving a unit commitment problem in each iteration, we need to perform a table look-up in order to come up with a reasonable λ. Then, the demand is changed and the new values of λ are used and so on. We formulated the previous problem as a dynamic program by discretizing the water level in the reservoir.

Assume that the water level at the beginning of the horizon is h_0 and that the terminal water level at the end of the horizon is h_T. The maximum water level in the reservoir is $\overline{h}$. The cost of pumping water into the reservoir, $p(h_{t-1}, h_t, d_t)$, is a function of the water levels, h_{t-1} and h_t, and the load on the system, d_t, during time period t. Let us denote the amount of electricity needed to change the water level in the reservoir from h_{t-1} at the beginning of period t, to h_t at the end of that period by $H(h_{t-1}, h_t)$. The function $H(h_{t-1}, h_t)$ is assumed to be positive for $h_{t-1} < h_t$; that is, pumping, and negative for $h_{t-1} > h_t$ which represents generating electricity. Clearly, the ratio $\frac{-H(h+\Delta,h)}{H(h,h+\Delta)}$, $\Delta > 0$, is strictly less than one; it represents the efficiency of the hydro unit.

Given a certain demand, d_t, our approximation of the cost of generating one unit of electricity, λ_t, is given in section 4.2 by the convex function $\Lambda(d_t)$. Note that $\Lambda(d_t)$ is a function of the load and does not depend on the time period. The cost of changing the water level from h_{t-1} to h_t is

$$p(h_{t-1}, h_t, d_t) = \int_{h=h_{t-1}}^{h_t} \Lambda(d_t - H(h_{t-1}, h))dh. \quad (7)$$

Since we decided to discretize the variable h, a summation over all intervals between h_{t-1} and h_t should replace the previous integral.

The recursive equation of the dynamic program is

$$P_t(h_t) = \min_{0 \leq h_{t-1} \leq \bar{h}} \{P_{t-1}(h_{t-1}) + p(h_{t-1}, h_t, d_t)\}, \; t = 1, \; \ldots, \; T. \quad (8)$$

The boundary conditions are $P_0(h_0) = 0$ and $P_0(h) = \infty$ for $h \neq h_0$. The optimal objective value is $P_T(h_T)$ and an optimal strategy can be found easily by backward recursion.

One can look at the previous dynamic program as an attempt to smoothen the demand for electricity in different periods using the hydro unit. During low demand periods, water is pumped into the reservoir at a relatively low cost, which increases the load on the thermal units. When the marginal cost is large; that is, the demand is high, water is released to absorb part of the demand which reduces the load on the thermal units.

After solving the dynamic program in (8), we update the load on the thermal units

$$\tilde{d}_t = d_t + H(h_{t-1}, h_t), \; t = 1, \; \ldots, \; T,$$

and solve the resulting thermal unit commitment problem. Using the marginal values obtained from the Lagrangian relaxation of the unit commitment problem, we again minimize the cost of operating the hydro unit, which is assumed to be a linear function of $\tilde{d}_t$ with λ_t as its cost coefficient. As mentioned earlier, having a constant cost, λ_t, at each time period for any change in the water level may cause the hydro unit to over pump or release. This extreme behavior is dampened by restricting the maximum amount of water that can be

Figure 4: The Procedure Used in Finding an Optimal Solution for the Hydro Unit

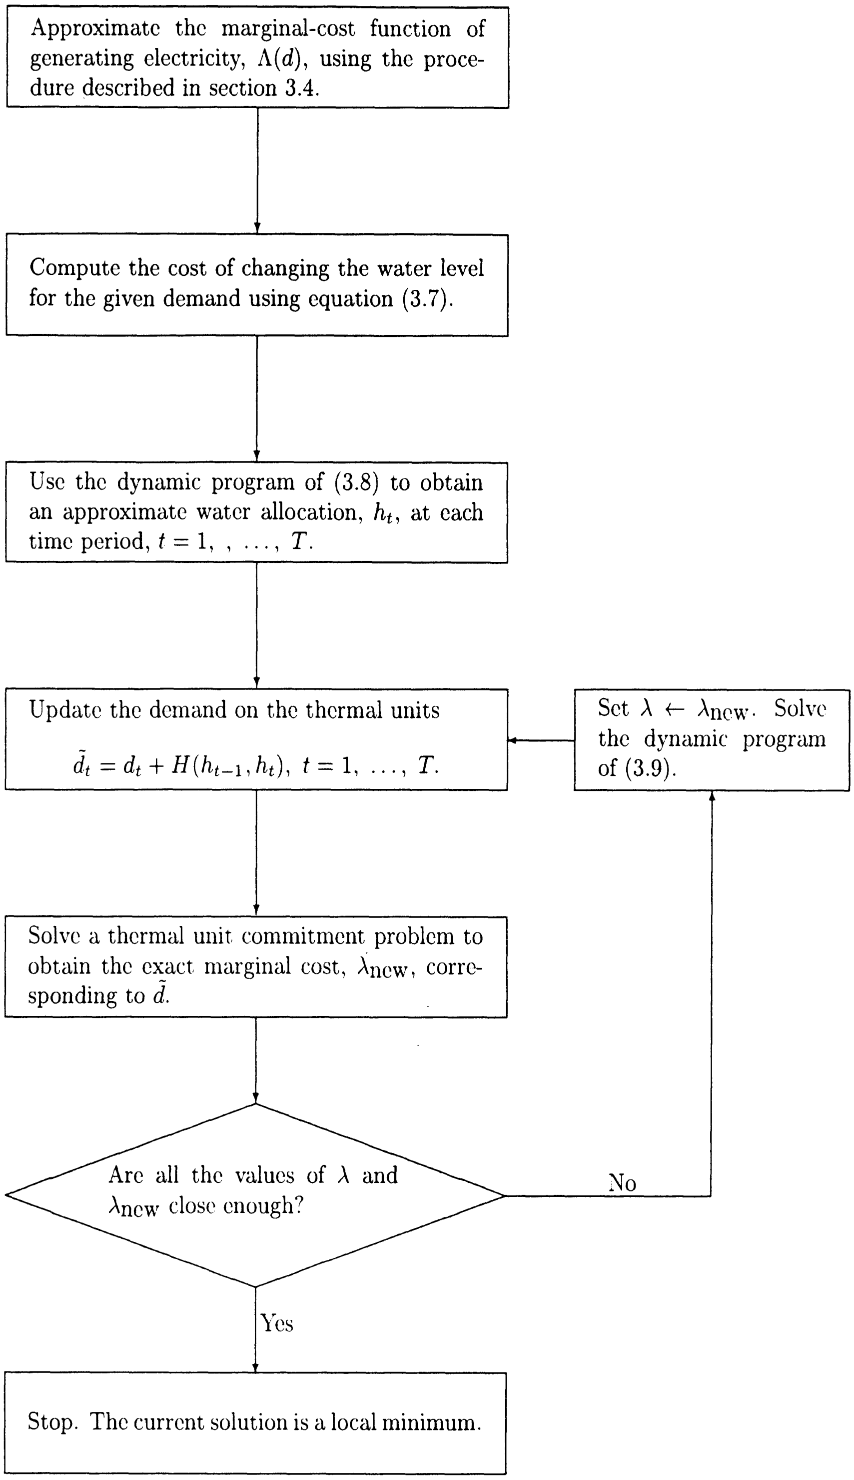

pumped or released in one period to $\overline{H}$. The recursive equation of the resulting dynamic program takes the form

$$P_t(h_t) = \min_{h_{t-1}}\{P_{t-1}(h_{t-1}) + \lambda_t(\tilde{d}_t - H(h_{t-1}, h))\}, \; t = 1, \; \ldots, \; T, \tag{9}$$

where h_{t-1} takes all possible water levels in the range $[\max\{0, h_t - \overline{H}\}, \min\{\overline{h}, h_t + \overline{H}\}]$. The dynamic program in (9) leads to a new hydro policy. The process is repeated until the values of λ converge.

8 Demand Uncertainty

Throughout the previous sections, $d_t \geq 0$ represented the total demand for electricity during time period t. Clearly, d_t is a random variable that fluctuates according to the weather, the day of the week, the time of day, and many other factors.

The main difficulty is that the demand for electricity is continuously varying. In practice, the demand during each period, which is usually one hour, is estimated using the weather forecast and the data that has been already collected from similar periods in previous years. This estimation is done for each period in the problem horizon which is generally one week; the approximated demands are then used to find an optimum strategy, which determines the status of each generator during each time period.

Since the actual demand at any point in time may be different than the approximated demand, experienced operators observe the changes in the demand, and make the necessary real-time decisions. For instance, if the demand for electricity increases above a certain level, the load on the generator with the lowest cost per electricity unit, that is, the smallest slope or marginal cost, is increased. In this case, the operator is trying to meet the demand at a low cost using a greedy strategy. But after using the flat parts of the cost functions for all the running generators, it may actually be more beneficial to start up another generator rather than using the generator with the smallest slope. In some cases, all running generators may reach their upper limits, and still another generator may need to be started.

A widely used approach to handle demand uncertainty is to develop forecasts of both the average and the peak demand for each time period. The unit commitment problem formulation is then altered so that it takes into consideration the worst case scenario; that is, the peak demand at each period, $\bar{d}_t$.

$$\begin{array}{lrcl} \min_{x,u} & \sum_{i=1}^{n} \sum_{t=1}^{T} f_i(x_t^i, u_{t-1}^i, u_t^i) & & \\ \text{subject to} & \sum_{i=1}^{n} x_t^i & \geq & d_t, \ t = 1, \ \ldots, \ T \\ & \sum_{i=1}^{n} G_i u_t^i & \geq & \bar{d}_t, \ t = 1, \ \ldots, \ T. \end{array}$$

Planners require excess reserve capacity not only to guarantee the existence of enough generating capacity, but also to protect the system against the inability to satisfy the demand when generating equipment failures occur. Excess system capacity is often called spinning reserve [8].

The following section describes a unit commitment formulation that takes into consideration the non-deterministic nature of the demand. An efficient technique for solving this problem is also presented.

9 Stochastic Unit Commitment Formulation

Rather than solving the unit commitment problem for the expected demand vector, one can consider a set of possible scenarios and solve the unit commitment problem for the demand of each of these scenarios. Each scenario is assigned a weight, P_s, that reflects the possibility of its occurrence in the future. In other words, the uncertainty about future demand is modeled by a number of deterministic sub-problems; this approach is called scenario analysis [11].

Here, one hopes to discover similarities and trends by studying different instances of the sub-problems, and eventually, come up with a "well-hedged" solution to the original problem.

The policy we are looking for has to satisfy the constraint that if two different scenarios s and s' are indistinguishable at time period t on the basis of information available about them at time t, then the decision made for scenario s must be the same as that of scenario s'. The previous constraint is modeled by partitioning the scenario set at each time period into disjoint subsets that are called scenario bundles. Clearly, a bundle at time t is refined in subsequent time periods into smaller disjoint bundles [11].

Figure 5: Scenario bundles

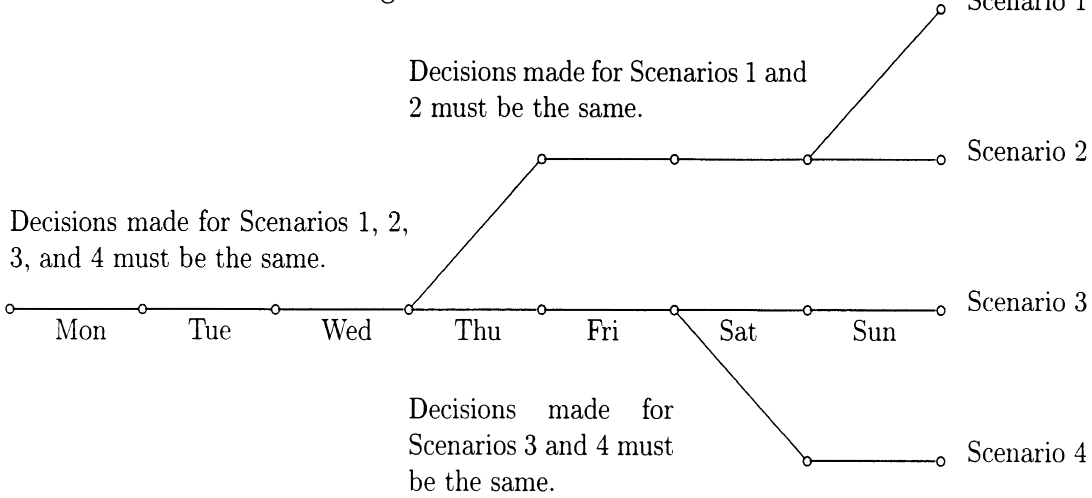

To clarify the previous concept, assume that the unit commitment problem was solved for S demand vectors, d^s, $s = 1, \ldots, S$, which resulted in a three-dimensional array representing the status of each unit at each time period under each scenario, $u_t^{i,s}$. The system administrator needs to make a decision concerning the status of each unit, i, during the first time period based on the solutions obtained. In general, the values of $u_1^{i,s}$, $s = 1, \ldots, S$, are not equal, and an optimal decision cannot be made without reformulating the problem so that the decisions $u_1^{i,s}$ are the same for all $s = 1, \ldots, S$ and $i = 1, \ldots, n$.

One does not only have to consider the first time period, but must also take into account all subsequent bundles that could affect the decisions made throughout the study horizon. Two scenarios are members of the same bundle, Ω_k, at time period t if both of them have the same demand values for all time periods $1, \ldots, t$. Note that each scenario is a member of exactly one bundle at each time period, which motivates the notation $B(s,t)$. If two scenarios are members of the same bundle at time t, then their bundles in time periods $1, \ldots, t-1$ are also the same. In other words,

$$B(s_1,t) = B(s_2,t) \Rightarrow B(s_1,\tau) = B(s_2,\tau),\ \tau = 1, \ldots, t-1.$$

Mathematically, a bundle at time period t is represented as a constraint on the decision variables, $u_t^{i,s}$, of its scenarios. Adding the bundle constraints results in a large-scale mixed-integer quadratic program that combines S unit commitment problems together. The objective function is to minimize the weighted sum of the objective functions of the smaller problems, that is, to minimize the expected cost over all possible scenarios. Here is the mathematical formulation:

$$\begin{array}{ll} \min_{x,u} & \sum_{s=1}^{S} P_s \sum_{i=1}^{n} \sum_{t=1}^{T} f_i(x_t^{i,s}, u_{t-1}^{i,s}, u_t^{i,s}) \\ \text{subject to} & \sum_{i=1}^{n} x_t^{i,s} \geq d_t^s,\ t = 1, \ldots, T,\ s = 1, \ldots, S \end{array} \tag{10}$$

and

$$\forall s_1, s_2 : B(s_1,t) = B(s_2,t) = \Omega_k \Rightarrow u_t^{i,s_1} = u_t^{i,s_2} = c_k^i,\ t = 1, \ldots, T,\ i = 1, \ldots, n.$$

The previous program is solved in a Lagrangian relaxation-like technique: a multiplier, $\mu_t^{i,s}$, is associated with the bundle constraint on each variable $u_t^{i,s}$, and the corresponding penalty term, $\mu_t^{i,s}(u_t^{i,s} - c_k^i)$, is added to the objective function. The target value c_k^i is the same for all scenarios that share the same bundle, and is assumed to be the weighted average

Figure 6: The Progressive Hedging Algorithm Described in Section 9

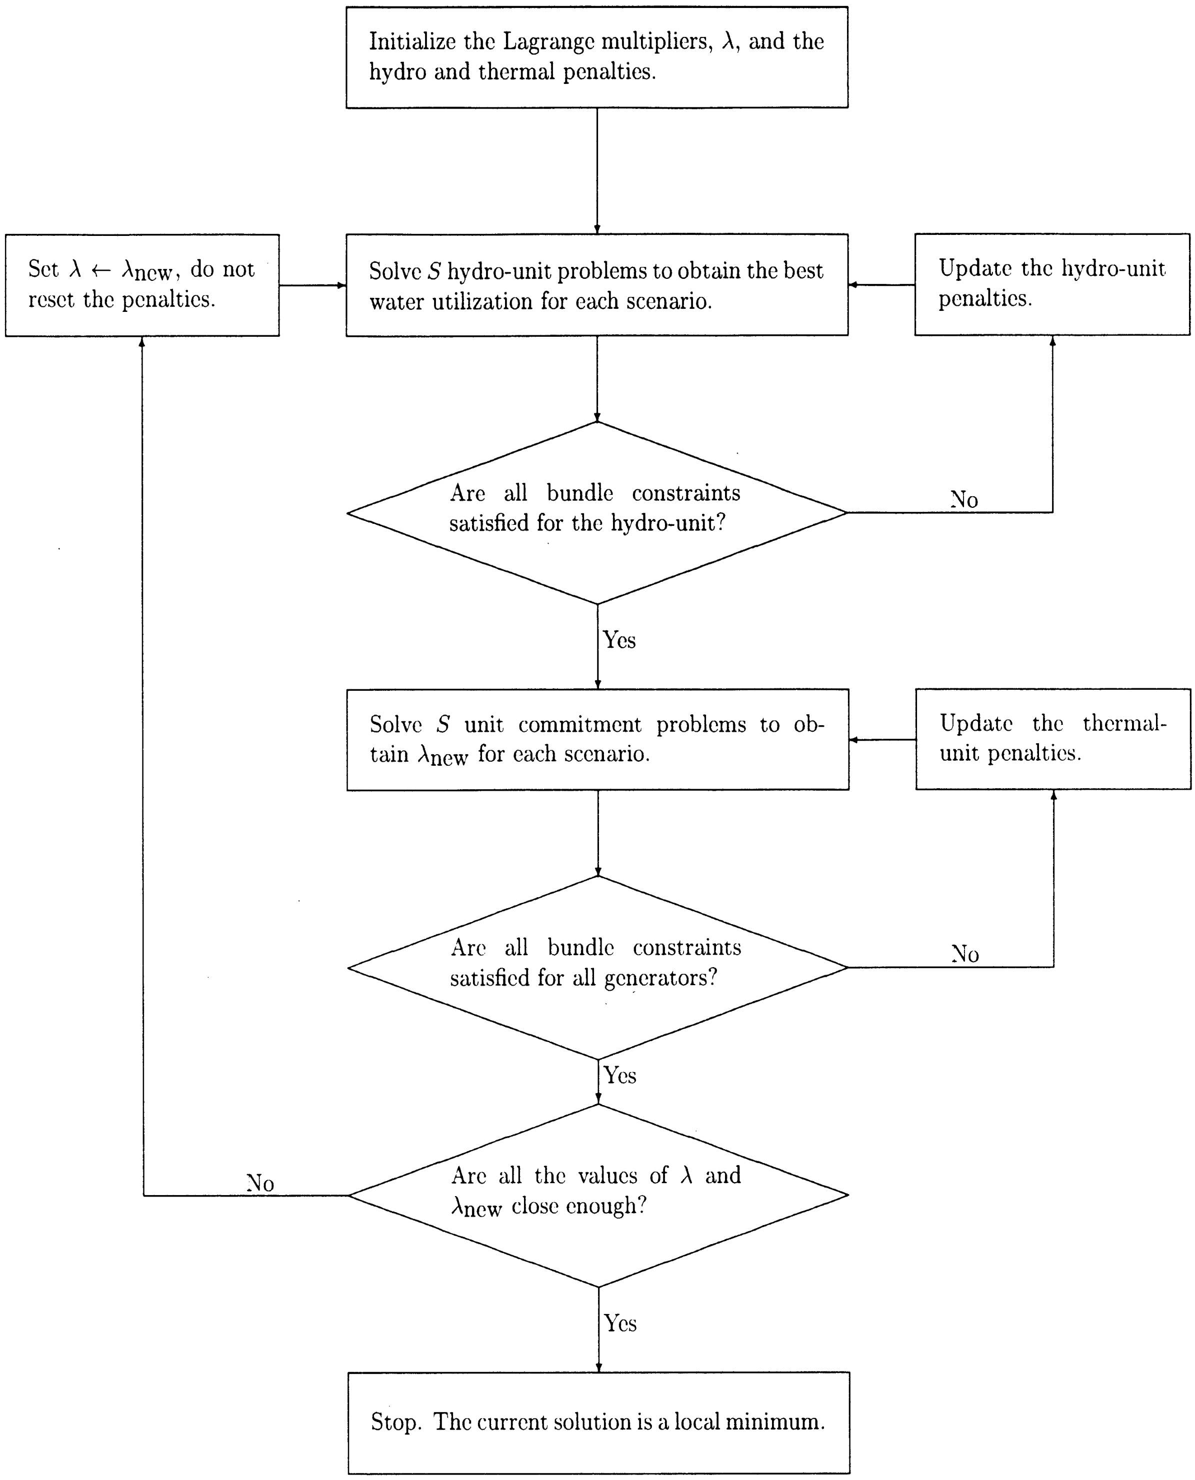

of the decisions made

$$c_k^i = \frac{\sum_{s:B(s,t)=\Omega_k} P_s u_t^{i,s}}{\sum_{s:B(s,t)=\Omega_k} P_s}.$$

The process is very similar to the one proposed for solving the deterministic unit commitment problem using Lagrangian relaxation. Starting from a set of penalties, S unit commitment problems are solved. If the resulting solutions satisfy the bundle constraints, then we stop with an admissible implementable policy or a feasible strategy. Otherwise, the penalties are updated and the process is repeated. We call the previous process progressive hedging [11]. Since the program in (10) is not convex, the feasible policy obtained at termination is a local minimum for the problem. In order to get better policies, different starting penalties and different updating strategies should be used.

10 Duality Gap Estimation

Our goal is to minimize the expected cost of generating electricity using n generators for a horizon of T periods over S possible scenarios. Mathematically, the problem can be written as:

$$\begin{array}{ll} \min_{x,u} & \sum_{s=1}^{S} P_s \sum_{i=1}^{n} \sum_{t=1}^{T} f_i(x_t^{i,s}, u_{t-1}^{i,s}, u_t^{i,s}) \\ \text{subject to} & \sum_{i=1}^{n} x_t^{i,s} \geq d_t^s, \end{array} \tag{11}$$

and subject to minimum on-time, minimum off-time, minimum and maximum operating levels, and the appropriate bundle constraints. In order to represent the bundle constraints, we define $\tau(s)$ to be the first period in which a scenario, $s \geq 2$, does not share a bundle with another scenario $s' < s$. We also define $\sigma(s) < s$ to be a scenario that shares the same bundle with s at all time periods prior to and including $\tau(s) - 1$. The bundle constraints are then written as

$$u_t^{i,s} - u_t^{i,\sigma(s)} = 0, \; t = 1, \ldots, \tau(s) - 1, \; i = 1, \ldots, n, \; s = 2, \ldots, S.$$

The program in (11) is made separable into S deterministic unit commitment problems by using the appropriate Lagrangian multipliers. Let us relax the bundle constraints by associating a multiplier, $\mu_t^{i,s}$, with each equality constraint $u_t^{i,s} - u_t^{i,\sigma(s)} = 0$. Here is the resulting dual objective function:

$$\max_{\mu} \min_{x,u} \sum_{s=1}^{S} P_s \sum_{i=1}^{n} \sum_{t=1}^{T} f_i(x_t^{i,s}, u_{t-1}^{i,s}, u_t^{i,s}) - \sum_{s=2}^{S} \sum_{i=1}^{n} \sum_{t=1}^{\tau(s)-1} \mu_t^{i,s} \left(u_t^{i,s} - u_t^{i,\sigma(s)}\right) \tag{12}$$

In this section, we show that under certain conditions, the duality gap is relatively small; hence, the use of this relaxation is justified.

Proposition 10.1 *Under the assumption that there exists a feasible solution for the given unit commitment problem of each scenario, the duality gap between the minimum of (11) and the maximum of (12) is bounded above by KC. Here, K is the number of instances at which the scenario tree branches; i.e., $K = \| \bigcup_{s=2}^{S} \{\tau(s)\} \|$, and C is a constant determined by the generating units of the system.*

Proof Let us solve the unit commitment problem for all scenarios $s = 1, \ldots, S$ without taking the bundle constraints into consideration. Our goal is to find, using these optimal policies, a feasible policy that satisfies all bundle constraints. Note that if we force each generator, i, to be on for at least L_i periods prior to each branching point in the scenario tree, the resulting optimal solutions are feasible and satisfy the bundle constraints. In other words, scenarios that are in the same bundle have identical schedules for the generating units. This follows from the fact that scenarios contained in a given bundle at a time period t have the same demand for all periods prior to t, and that the states of all generators at each of the K branching points are forced to be identical.

After relaxing the bundle constraints and solving the unit commitment problem corresponding to each scenario optimally, the state of a generator, i, under a given scenario and prior to a specific branching point, t, can be in one of three states. Here, we describe a procedure that uses the optimal policies to produce a feasible schedule that satisfies all bundle constraints.

1. **on for at least L_i periods.**

   No action is necessary to make the problem feasible.

2. **on for $\alpha < L_i$ periods.**

   Here, the generator i is off for the $\beta \geq l_i$ periods prior to $t - \alpha$. We study two cases:

   - $\boldsymbol{\alpha + \beta - L_i \geq l_i}$

     Switch on generator i for the L_i periods prior to t. Note that i is off for $\alpha + \beta - L_i$ periods in the new schedule; i.e., the minimum-off time requirement is satisfied. The additional cost incurred is $S_i + L_i f_i(g_i, 1, 1)$.

   - $\boldsymbol{\alpha + \beta - L_i < l_i}$

     Switch on generator i for all $\alpha + \beta$ periods prior to t. Note that the resulting schedule satisfies the minimum-on time requirement since there are at least L_i periods prior to $t - \alpha - \beta$ in which the generator is on. The additional cost incurred is at most $(L_i + l_i) f_i(g_i, 1, 1)$.

3. **off for β periods.**

   We discuss two cases:

   - $\boldsymbol{\beta \geq L_i + l_i}$

     Switch on generator i for the L_i periods prior to t. The resulting schedule is feasible since it satisfies both the minimum on-time and minimum off-time constraints. The additional cost incurred is $S_i + L_i f_i(g_i, 1, 1)$.

   - $\boldsymbol{\beta < L_i + l_i}$

     Switch on generator i for all β periods prior to t. The resulting schedule is feasible since this generator is on for at least L_i periods prior to $t - \beta$. The additional cost incurred is at most $(L_i + l_i) f_i(g_i, 1, 1)$.

From the previous discussion, it costs at most $C = L_i f_i(g_i, 1, 1) + \max\{S_i, l_i f_i(g_i, 1, 1)\}$ to force a generator, i, to be on for L_i periods prior to a given time period t. Since we have at most K branching points for each scenario, an upper bound on the increase in the objective function cost that results from enforcing the bundle constraints is:

$$\sum_{s=1}^{S} P_s KC = KC.$$

Clearly, KC is an upper bound on the duality gap of the relaxation used; i.e., it is an upper bound on the difference between the minimum of (11) and the maximum of (12). ■

Note that the number of branching points, K, is bounded above by $S - 1$. The ratio of the duality gap to the number of scenarios, S, is called the relative gap. It is bounded above by $\frac{KC}{S}$. In practice, K does not increase linearly with the number of scenarios since the branches are created at certain periods in the planning horizon. If this is the case, the relative duality gap approaches zero as the number of scenarios increases.

11 Scenario Generation

In order to obtain a policy that is of practical use, we need to provide the progressive hedging algorithm with scenarios that truly reflect all possible future demands. Furthermore, the probabilities assigned to these scenarios must be calculated carefully. Clearly, this is not an easy task, and more research must be done in this area in order to develop a better understanding of the demand randomness and the related factors. One thing which must be considered in regards to scenario generation, is that the more scenarios which are created, the better the hedging policy of the algorithm. On the other hand, the execution time of the algorithm grows rapidly as the number of scenarios included increases and as their demands are more diverse.

Among the factors that randomly affect the load on the system are generator failures and unexpected or unusual demands. One should also consider the available electricity that can be purchased from other suppliers and the sales of electricity to other regions with emergency needs. Here is a closer look at some of these factors and the suggested methods of treating them in our model.

11.1 Generator Failures

A given generator may be unavailable for use at any time due to unforeseen circumstances, such as mechanical problems, unscheduled maintenance, etc. The current strategy of Detroit Edison and Consumers Power is to maintain adequate operating reserves for these types of situations. At any given time, some of this operating reserves is available from the pump-storage plant.

We have modeled this problem by creating a scenario which has demand increases equal to the unavailable generator's capacity. These demand increases occur, during the periods in which the generator is expected to be down. The scenario is then assigned a weight equal to the probability that this generator will be unavailable. Different scenarios can be created for different generator failures.

Another way to model this problem is to approximate the generating capacity loss over a certain period of time. The advantage of this technique is its independence from the individual generators. For instance, the Michigan Electric Power Coordination Center personnel estimate that the unexpected generation loss is approximately 400 Megawatt-Hours every three days. The previous statement is translated into an increase in the load of 400

Megawatt-Hours per day with a probability of one third, and no increase in the load with a probability of two thirds. We started with the expected scenario, and then created a scenario tree that covers all possible increases using the previous rule.

11.2 Inaccurate Forecast

Experienced schedulers forecast the load for each hour of the week on Sunday and, if necessary, update this forecast on each shift. In order to estimate the load on the system, a scheduler uses the data of the same week from previous years, the data of days with similar weather conditions, and his/her personal experience. The resulting load forecast is then used as the input for the unit commitment program which determines an optimum strategy.

Here again, we created several scenarios with high demand increases or decreases. The unusual changes in the demand are assigned probabilities according to the likelihood of their occurrence. This approach makes it easy for a scheduler to include all the relevant data of a certain week in the model after assigning a suitable weight to each of them.

11.3 Other Factors

Electricity can be bought from or sold to other electric companies. The price per Megawatt-Hour varies depending on the time of the day, the contracts governing the trade between these companies, the amount of electricity needed, the urgency of the electricity need, etc. All the resulting changes in the load on the system can be modeled as possible scenarios. Each of these scenarios has an additional generator with an appropriate cost function corresponding to the price of the available electricity. The minimum and maximum generating capacities of such a generator are set equal to the total amount of electricity that can be purchased in each period. Each generator of this type is a must-run unit in its corresponding scenario, and is assumed to be unavailable in all other scenarios. The only difficulty in applying this approach is that the probability assigned to each scenario needs to be estimated accurately. This estimation can be done easily if there is a large database that will permit a good approximation of these probabilities.

12 Computational Experience

The progressive hedging technique was implemented using the C language on an IBM RS/6000. The code is written so that it makes full use of the RISC architecture. A unit commitment problem with 2400 binary variables and 2400 continuous variables is solved in less than one second to within 0.1% of the optimum solution.

We used the same data that is used by Michigan Electric Power Coordination Center. The system has a pump-storage hydro plant with six units, and more than hundred thermal units. Water is pumped into the reservoir during low-demand periods so that it can be used during peak-demand periods. For this study, the power needed to pump the water into the storage facility is assumed to be linear to the change in the water level. In other words, to change the water level in the storage by one foot, one needs a constant amount of electricity that is independent of the level of water in the pump-storage plant. This constant was assumed to be 340 Megawatt-Hours per foot in our calculations. The efficiency of the storage facility, that is, the ratio of the power generated by one foot of water to the power needed to pump one foot of water into the reservoir is 70%. So, it makes sense to pump water for future use as long as the ratio of the marginal cost during the pumping period to the marginal cost during the generation period is less than 0.7.

The thermal units are divided into four different categories. The first is the must-run units which are assumed to be on-line everyday throughout the year. These units have relatively lower operating costs compared to the others. The second category is the unavailable units.

The unavailability is due to maintenance, mechanical failures, or other factors. Cyclers form another category. These units must be scheduled for use in advance, since they require a long time to get them on-line. The last category is peakers. Peakers are committed when no other economic resources are available, or when a system condition requires their quick start capability. For this study, the minimum on and off times for peakers are equal to one hour, so that they can be used when needed then turned off immediately after that.

12.1 Generation Shortage Example

We used 22 different scenarios in this example. The expected demand for each period over a one-week horizon was provided by Detroit Edison and Consumers Power; it is that demand for the time periods between Monday November 9, 1992 and Sunday November 15, 1992. We assumed that the demand over Saturday and Sunday for all scenarios is the same as that of the forecasted demand, scenario one. Saturdays and Sundays have low demand in general, hence, they are not the bottleneck in the unit commitment schedule. The other scenarios were generated from scenario one by adding 400 Megawatt-Hours with probability one third at the beginning of Tuesday, Wednesday, Thursday, and Friday. This increase in the demand takes into account the possible failure of some thermal units. Table 1 summarizes the demand increases of different scenarios relative to the expected demand.

The scenario tree and different scenario bundles can be seen in Figure 7 and Table 1. For instance, scenario one and two have the same demand throughout Monday, Tuesday, Wednesday, and Thursday, and therefore they are members of the same bundles up to the beginning of Friday. That means that our unit commitment schedules for these two scenarios must be the same for all periods prior to Friday. In other words, the water level and the thermal units' status, on or off, are the same. Scenarios three, four, and five have the same demand up to Friday morning, and therefore, our unit commitment schedules for these

Table 1: Demand Increases for Different Scenarios

Scen.	Prob.	Mon.	Tue.	Wed.	Thu.	Fri.	Sat.	Sun.
1	19.05%	0	0	0	0	0	0	0
2	9.52%	0	0	0	0	400	0	0
3	9.52%	0	0	0	400	400	0	0
4	4.76%	0	0	0	400	800	0	0
5	0.95%	0	0	0	+20%	+20%	0	0
6	9.52%	0	0	400	400	400	0	0
7	4.76%	0	0	400	400	800	0	0
8	4.76%	0	0	400	800	800	0	0
9	1.90%	0	0	400	800	1200	0	0
10	0.95%	0	0	400	800	+15%	0	0
11	0.95%	0	0	+20%	+20%	+20%	0	0
12	9.52%	0	400	400	400	400	0	0
13	4.76%	0	400	400	400	800	0	0
14	0.95%	0	400	400	400	+20%	0	0
15	4.76%	0	400	400	800	800	0	0
16	1.90%	0	400	400	800	1200	0	0
17	0.95%	0	400	400	+15%	+15%	0	0
18	4.76%	0	400	800	800	800	0	0
19	1.90%	0	400	800	800	1200	0	0
20	1.90%	0	400	800	1200	1200	0	0
21	0.95%	0	400	800	1200	1600	0	0
22	0.95%	0	400	+20%	+20%	+20%	0	0

scenarios must be the same for all periods prior to Friday. Note that all scenarios have the same demand on Monday, hence they are members of the same bundles for all periods prior to Tuesday morning. Figures 8 to 12 in section 12.1.1 show the demand comparisons between each scenario and the expected demand.

We used the progressive hedging technique to solve the previous problem. Note that for any bundle, penalties need to be applied only to the cyclers if their status are not the same. The must-run units and the unavailable units have the same status, and the decisions on the peakers' status can be made in real time since their response time is negligible. This remark reduces the calculations considerably since progressive hedging has to deal with the small number of cyclers only.

Table 2 shows the costs of applying the optimum strategy for a given scenario to each of the other scenarios. Obviously, the lowest cost for each scenario occurs when its own optimum strategy is used. The expected value of using each optimum strategy is computed and given in the column $E(x)$. The costs corresponding to applying the progressive hedging policy to all scenarios are provided in the last row of Table 2. The difference between the expected value of each optimum policy and that of the progressive hedging represents the savings incurred by using progressive hedging.

Note that scenario one is the forecasted demand and its optimum strategy is the strategy used by Michigan Electric Power Coordination Center to build their schedule. We call this strategy the deterministic strategy since it is computed using the expected demand. Applying this strategy to the other scenarios represents what Detroit Edison and Consumers Power may pay if the forecast is inaccurate, that is, if another scenario occurs. Clearly, the expected cost corresponding to the progressive hedging strategy, \$19,978,000, is lower than that of the deterministic strategy, \$20,124,000, yielding a savings of \$146,000 for this week. As a matter

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	$E(x)$	Save	%
Pr	19.05%	9.52%	9.52%	4.76%	0.95%	9.52%	4.76%	4.76%	1.90%	0.95%	0.95%	9.52%	4.76%	0.95%	4.76%	1.90%	0.95%	4.76%	1.90%	1.90%	0.95%	0.95%			
1	19247	19416	19583	19823	22812	19748	19989	20236	20584	22245	24573	19915	20156	23120	20403	20742	23727	20647	20986	21344	21689	29673	20124	145.12	0.73%
2	19284	19416	19613	19788	22268	19778	19953	20200	20409	21457	24032	19945	20120	22263	20367	20577	22938	20611	20821	21190	21426	28819	20075	96.87	0.48%
3	19270	19438	19583	19787	22242	19763	19951	20134	20438	21903	24003	19930	20117	22840	20301	20605	23153	20545	20848	21124	21482	29078	20072	93.63	0.47%
4	19300	19462	19626	19786	21939	19791	19963	20150	20348	21341	23700	19959	20131	22172	20317	20515	22586	20561	20759	21066	21312	28408	20056	77.01	0.39%
5	19429	19589	19752	19922	21324	19918	20088	20256	20450	21156	23085	20086	20256	21638	20424	20618	21928	20668	20862	21056	21264	26444	20123	144.57	0.72%
6	19280	19448	19608	19797	22270	19747	19955	20144	20442	21933	23746	19938	20122	22871	20312	20603	23186	20493	20803	21085	21439	28815	20071	92.83	0.46%
7	19311	19472	19636	19808	21960	19802	19950	20157	20357	21347	23441	19968	20140	22182	20324	20524	22604	20513	20708	21022	21266	28139	20055	76.10	0.38%
8	19343	19504	19668	19840	21599	19833	20006	20134	20370	21362	23087	20001	20173	22214	20343	20538	22290	20527	20723	20925	21162	27417	20062	83.33	0.42%
9	19344	19505	19669	19841	21594	19834	20007	20176	20348	21354	23076	20002	20174	22197	20344	20539	22294	20527	20723	20925	21159	27394	20062	83.53	0.42%
10	19395	19555	19719	19889	21471	19884	20053	20224	20419	21119	22947	20053	20223	21588	20392	20585	22056	20575	20769	20966	21171	26771	20090	111.04	0.56%
11	19519	19678	19841	20011	21413	20006	20176	20344	20537	21242	22347	20174	20344	21727	20512	20705	22010	20680	20873	21067	21276	24838	20180	201.68	1.01%
12	19291	19459	19618	19809	22293	19781	19967	20157	20455	21952	23774	19914	20131	22887	20318	20608	23209	20506	20806	21091	21453	28846	20082	103.86	0.52%
13	19319	19480	19644	19817	22003	19810	19981	20166	20375	21419	23484	19975	20116	22288	20334	20535	22679	20516	20725	21023	21288	28246	20066	88.00	0.44%
14	19396	19555	19719	19889	21840	19883	20052	20236	20429	21134	23316	20051	20219	21578	20404	20595	22370	20586	20779	21075	21279	27449	20111	132.44	0.66%
15	19349	19510	19674	19846	21658	19839	20011	20182	20385	21436	23147	20006	20178	22319	20301	20552	22387	20532	20736	20939	21198	27571	20073	94.92	0.48%
16	19354	19515	19678	19850	21621	19843	20017	20186	20381	21365	23110	20011	20182	22206	20353	20514	22317	20536	20731	20934	21173	27459	20072	93.94	0.47%
17	19454	19613	19777	19946	21348	19943	20113	20279	20472	21175	22829	20109	20278	21644	20447	20639	21927	20630	20823	21018	21225	26162	20135	156.58	0.78%
18	19379	19540	19703	19875	21688	19869	20040	20213	20414	21465	22829	20035	20207	22348	20378	20582	22416	20491	20752	20951	21207	26930	20092	113.26	0.57%
19	19383	19544	19707	19879	21651	19873	20046	20215	20411	21392	22791	20040	20212	22235	20382	20576	22340	20553	20708	20946	21182	26818	20091	112.05	0.56%
20	19387	19548	19711	19883	21635	19877	20050	20219	20413	21401	22757	20044	20216	22239	20386	20580	22324	20557	20751	20923	21181	26718	20093	114.17	0.57%
21	19420	19580	19744	19914	21497	19909	20079	20249	20442	21160	22637	20077	20246	21664	20417	20609	22093	20587	20780	20976	21156	26193	20105	126.20	0.63%
22	19555	19714	19877	20046	21446	20042	20212	20378	20571	21272	22380	20209	20378	21735	20546	20740	22039	20714	20906	21099	21309	24813	20214	235.95	1.18%
Policy	19247	19446	19598	19798	21324	19769	19972	20176	20371	21119	22347	19944	20150	21578	20348	20547	21942	20549	20747	20945	21193	24813	19978	0.00	0.00%

Table 2: The Cost in $1000 of Applying the Policy of Each Scenario to All Other Scenarios

of fact, in the long run, the cost of applying the progressive hedging strategy must be less than the expected cost of applying any other strategy to all scenarios.

12.1.1 Demand Comparison

Figure 7: Scenario Tree

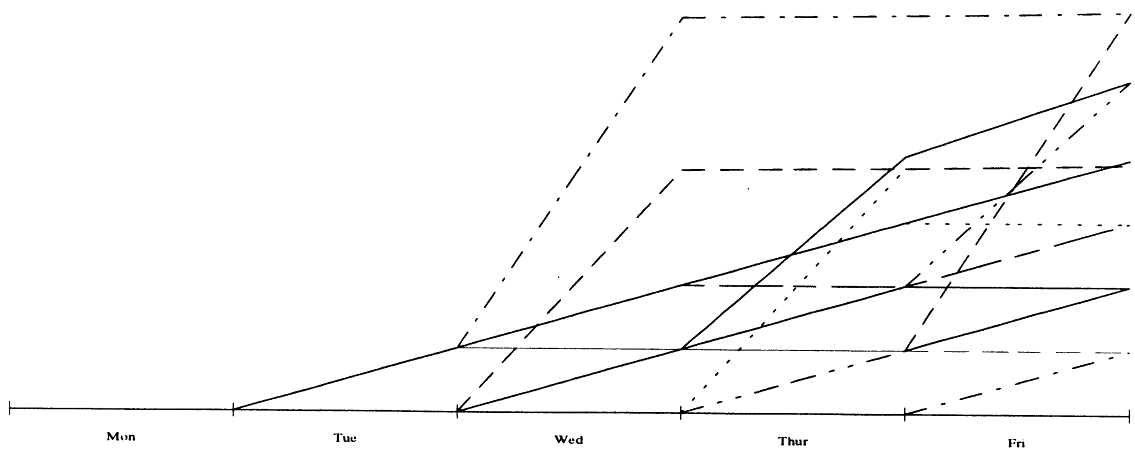

Figure 8: Electricity Demand Comparison for Scenarios 1 to 5

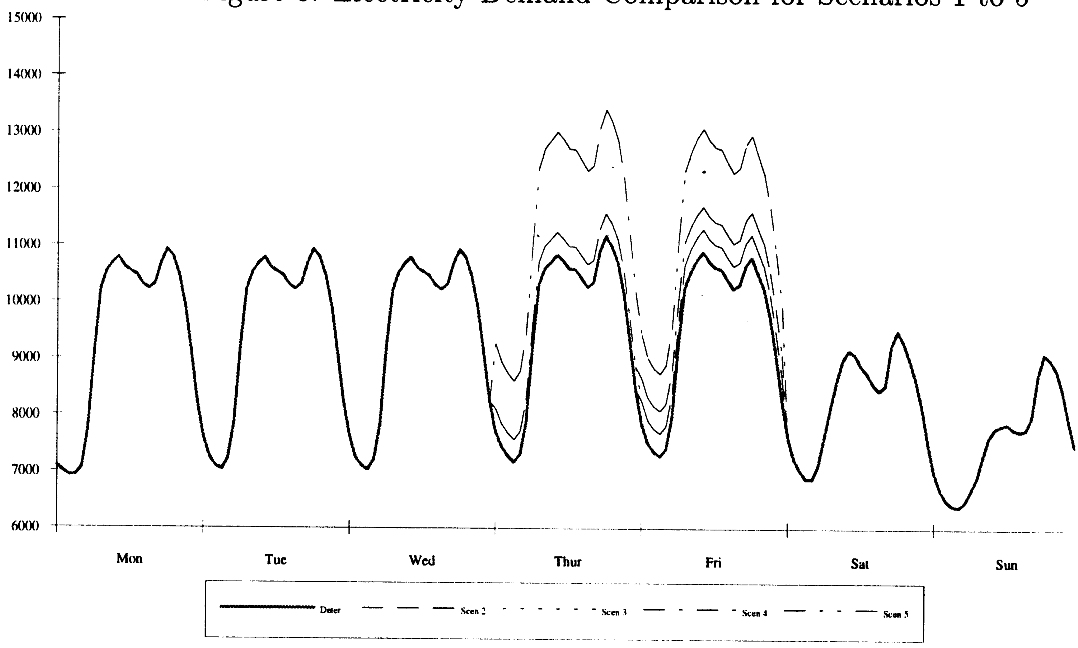

Figure 9: Electricity Demand Comparison for Scenarios 6 to 9

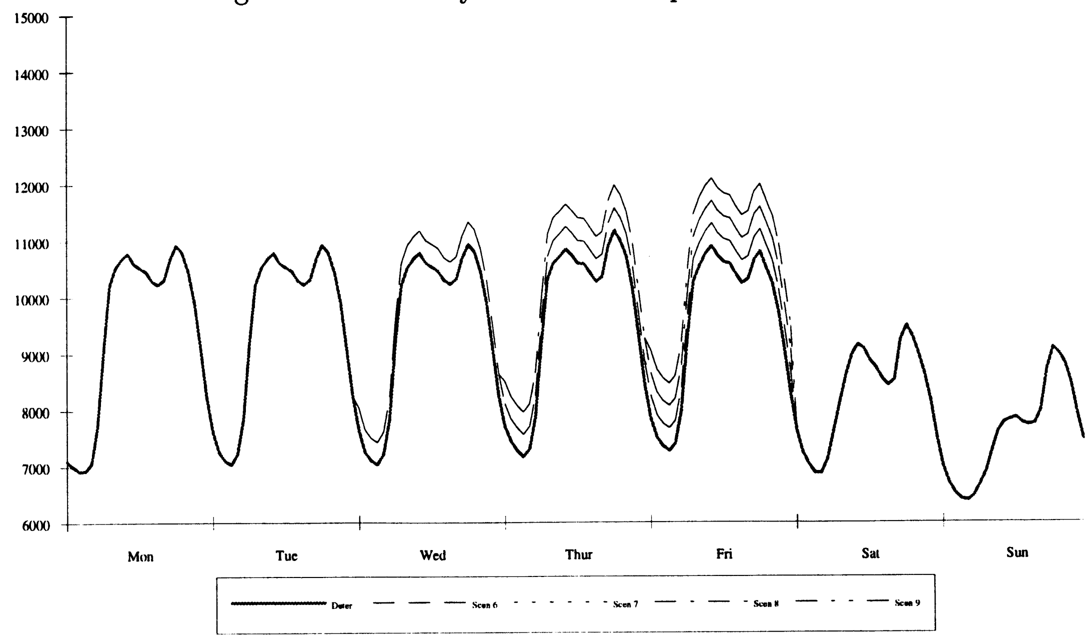

Figure 10: Electricity Demand Comparison for Scenarios 10 to 13

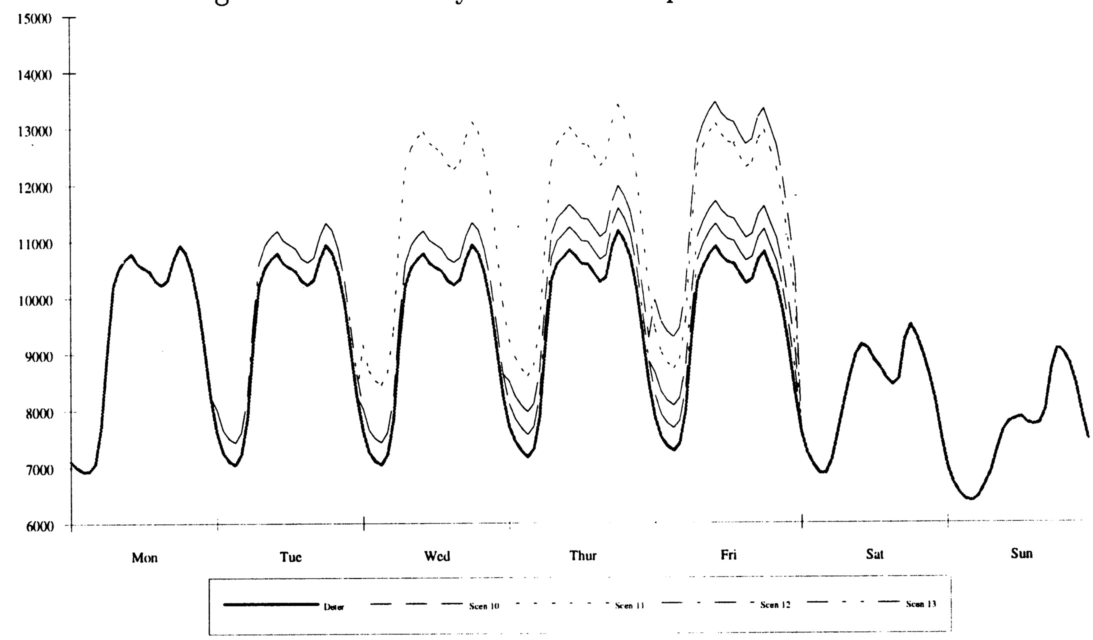

Figure 11: Electricity Demand Comparison for Scenarios 14 to 18

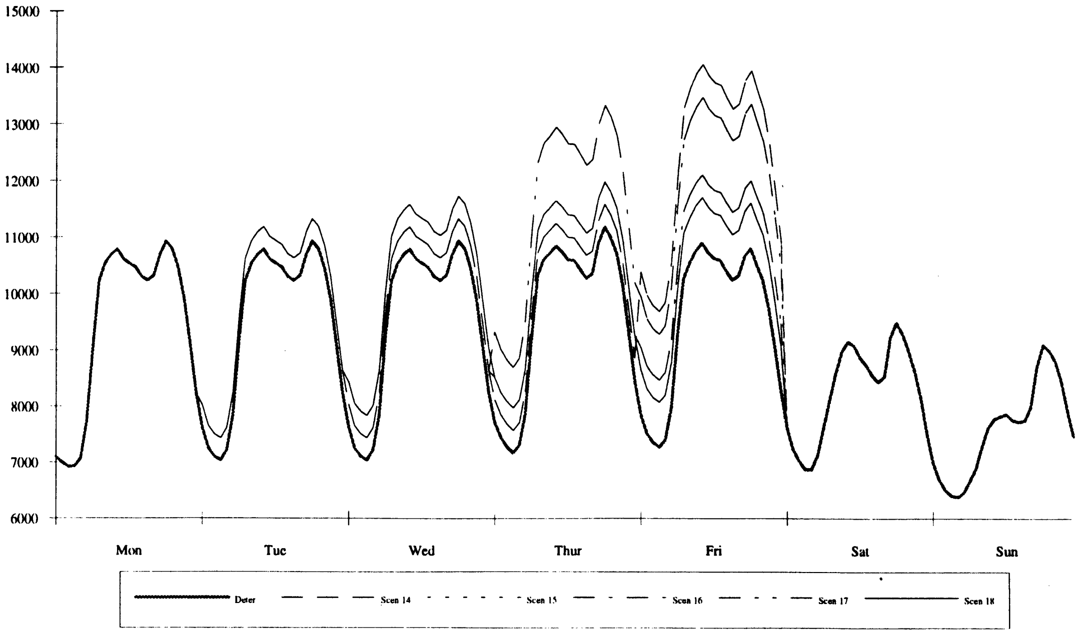

Figure 12: Electricity Demand Comparison for Scenarios 19 to 22

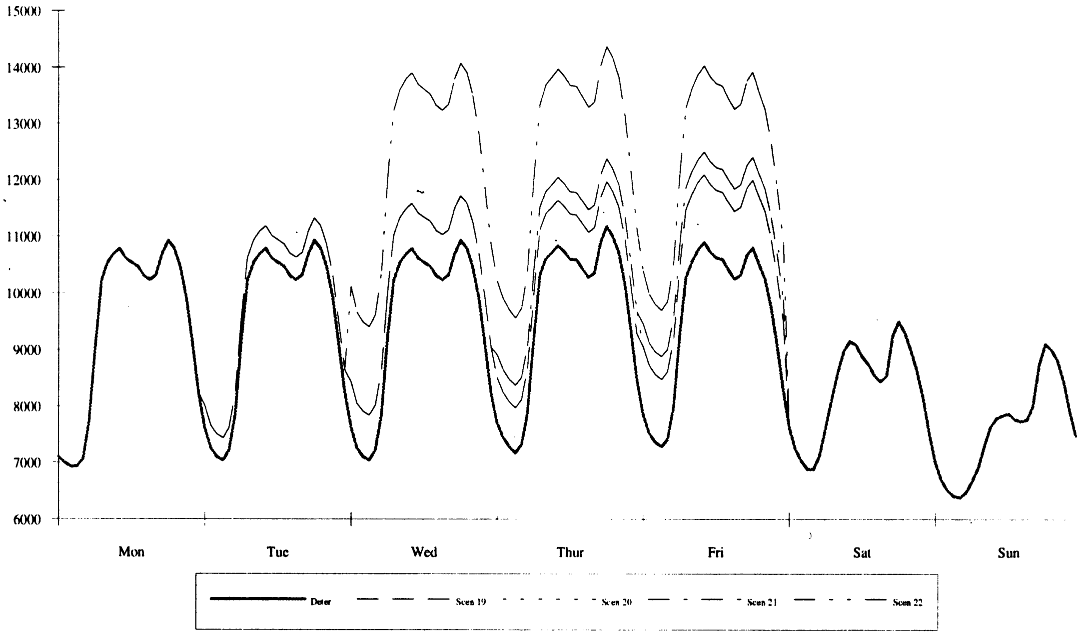

12.1.2 Pump-Storage Plant Water-Level Comparison

As we are trying to make the same decisions for scenarios that share the same bundles, water levels in the storage facility must be the same for such scenarios. Note how scenarios one and two maintain the same water level throughout the week up until Friday morning. It is interesting to compare the hedging strategy of scenario one to the deterministic strategy. The deterministic strategy uses more water in the periods between Monday and Thursday; while the hedging strategy is being more conservative by taking into consideration the possibility of demand increase on Friday.

Since it is very hard, for computational reasons, to make the difference in the water levels in a bundle equal to zero, we assumed that two policies are the same if the weighted difference in their water levels is less than 0.3 foot, that is, 100 Megawatt-Hours. This is why scenario five does not follow exactly the policy of other scenarios sharing the same bundle.

Figure 13: Water level Comparison for Scenarios 1 to 5

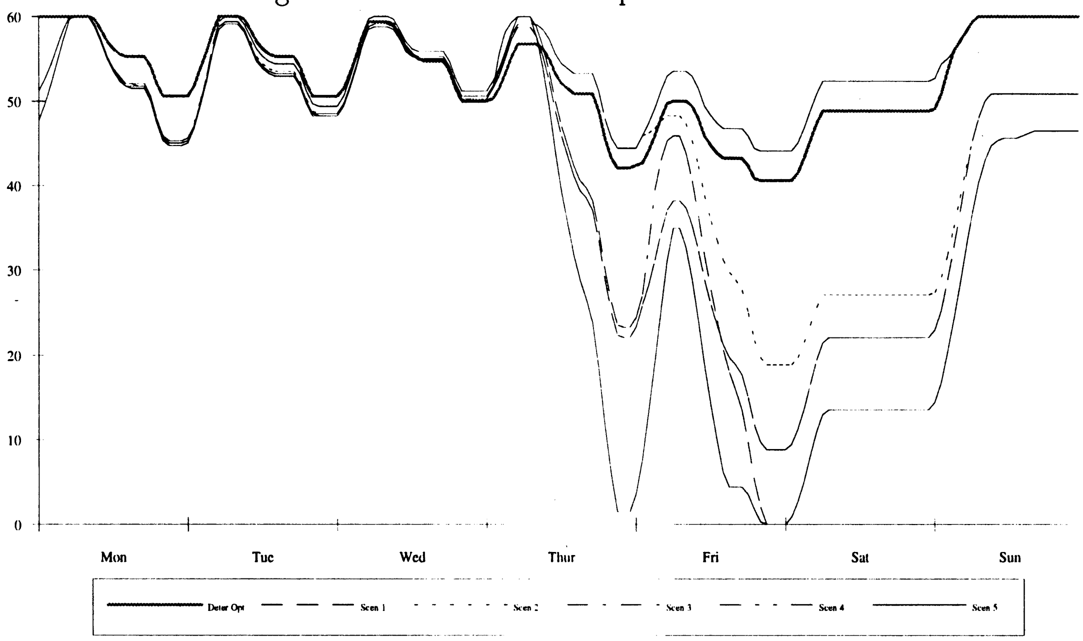

Figure 14: Water level Comparison for Scenarios 6 to 9

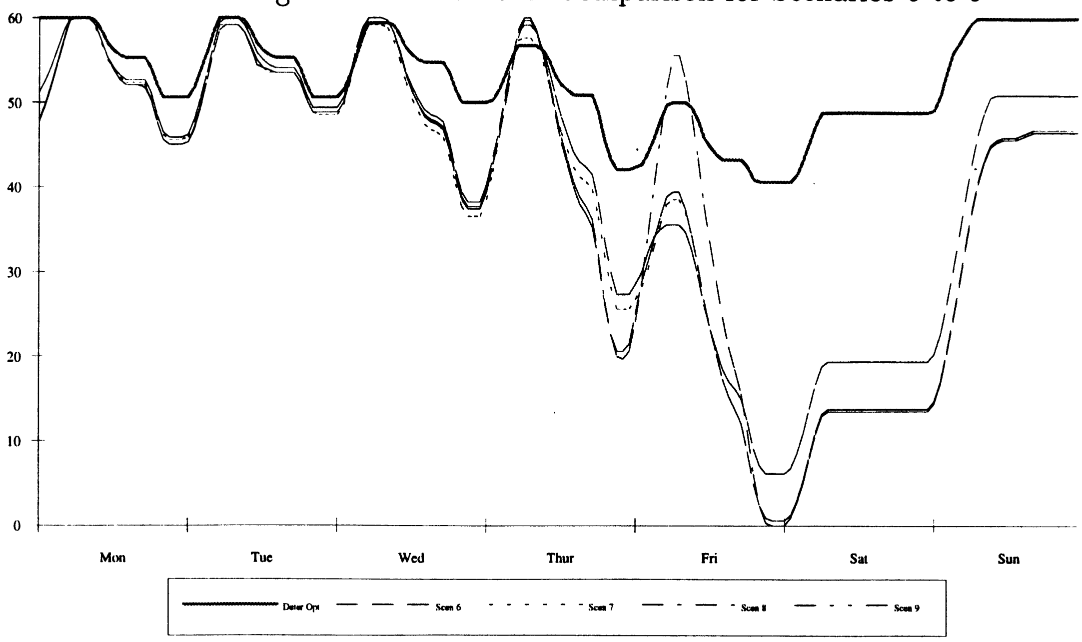

Figure 15: Water level Comparison for Scenarios 10 to 13

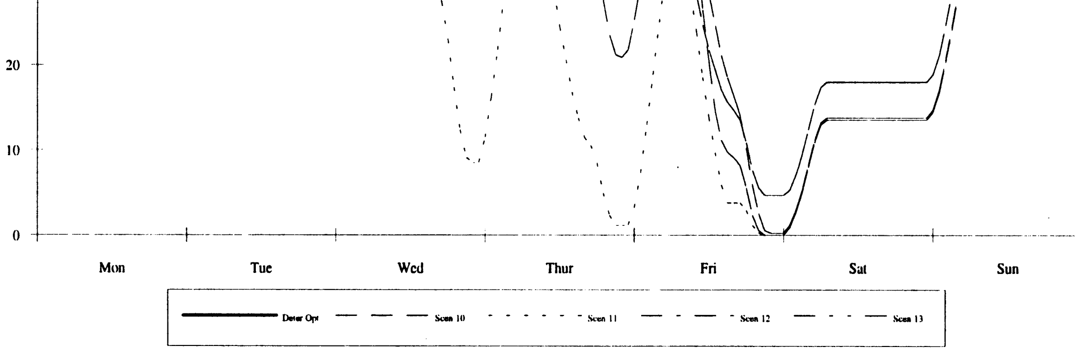

Figure 16: Water level Comparison for Scenarios 14 to 18

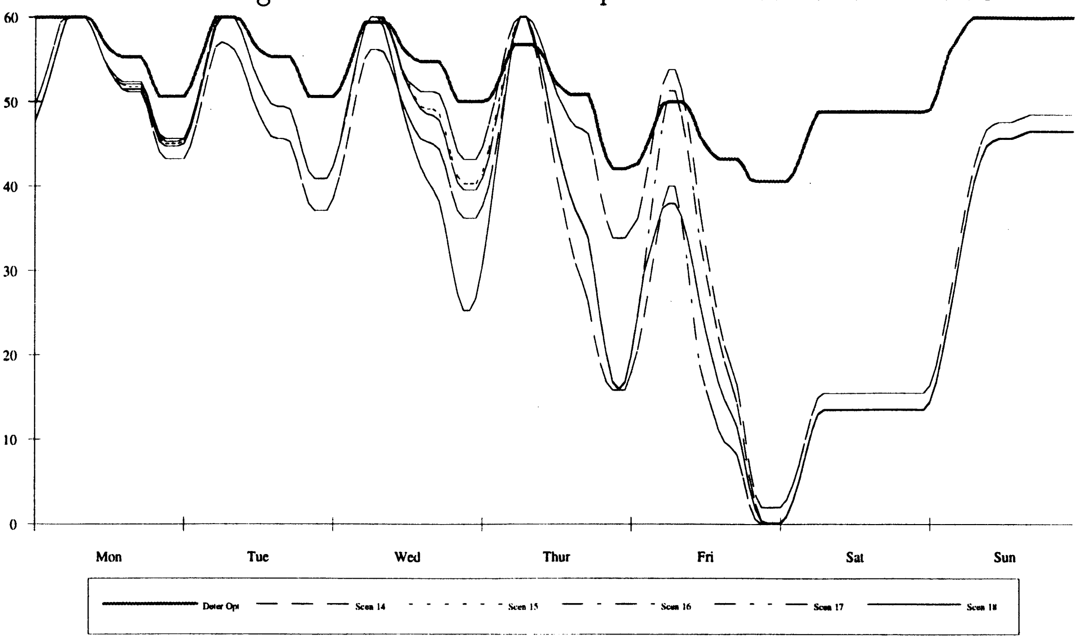

Figure 17: Water level Comparison for Scenarios 19 to 22

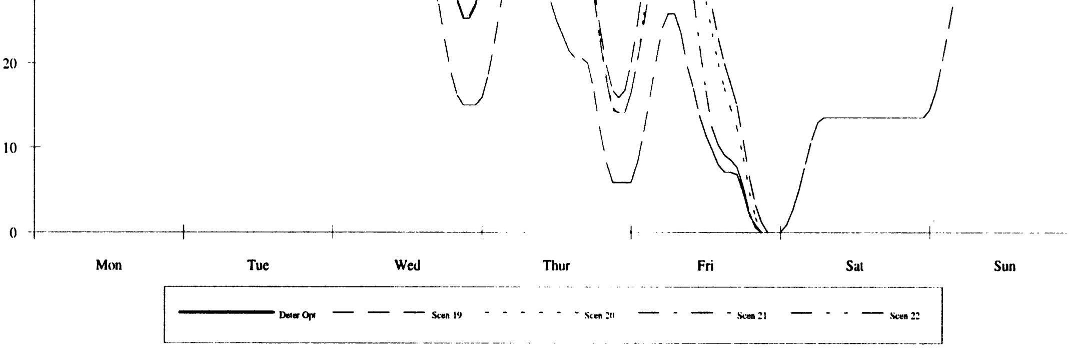

12.1.3 Pump-Storage Plant Generation/Consumption Comparison

As mentioned earlier, the relationship between the change in the water level and the amount of electricity generated is linear. A similar relationship holds in the case of pumping. We assumed that 340 Megawatt-Hours is needed to raise the water level in the storage facility by one foot. The efficiency of the hydro-unit is 70% which implies that 238 Megawatt-Hours are generated whenever the water level decreases by one foot. The following figures are obtained from the water levels in the storage. The positive values indicate generating electricity and the negative values indicate pumping of water into the facility. Note that the ratio of the area under the positive part of any curve, that is, total generation during the week, to the area under the negative part of that curve is equal to 0.7. The previous relation results from the condition that the water levels at the beginning and at the end of the week are the same.

Figure 18: Generation/Consumption Comparison for Scenarios 1 to 5

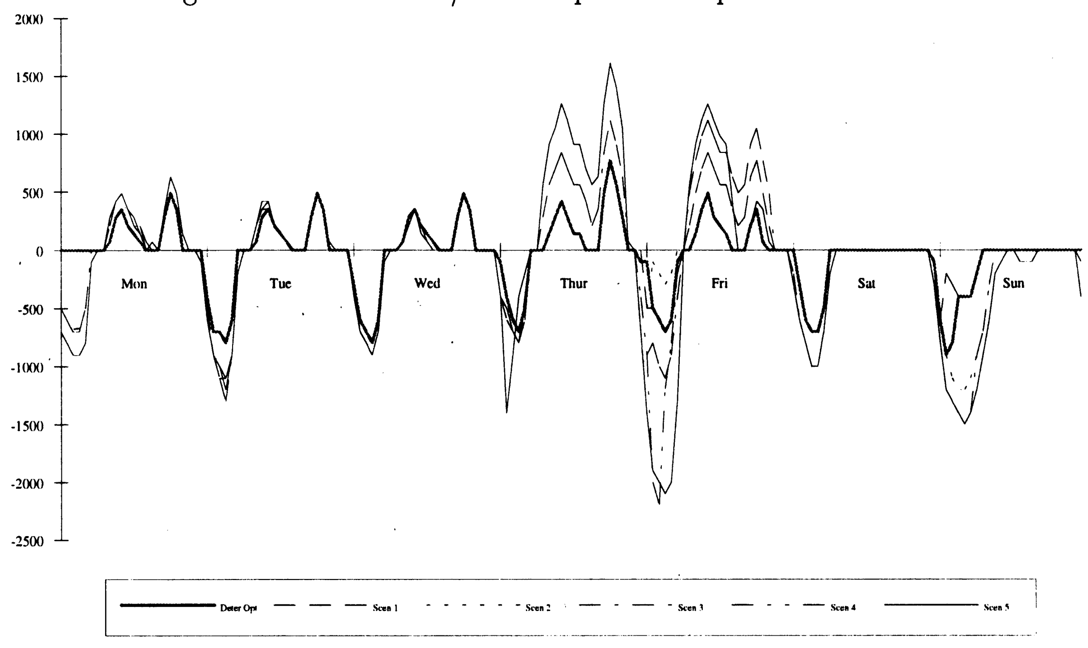

Figure 19: Generation/Consumption Comparison for Scenarios 6 to 9

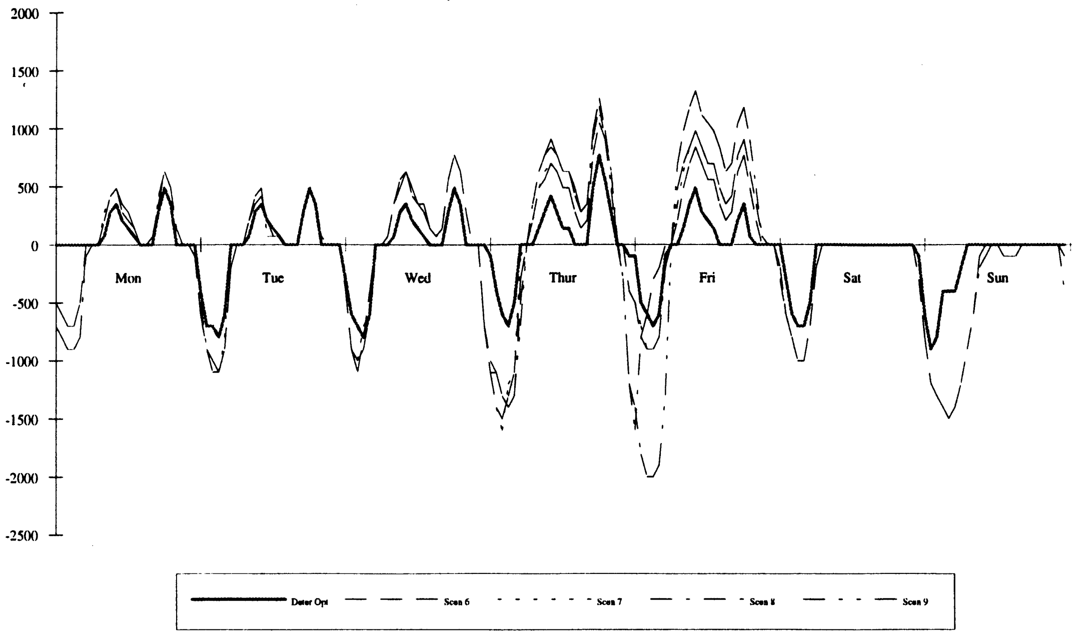

Figure 20: Generation/Consumption Comparison for Scenarios 10 to 13

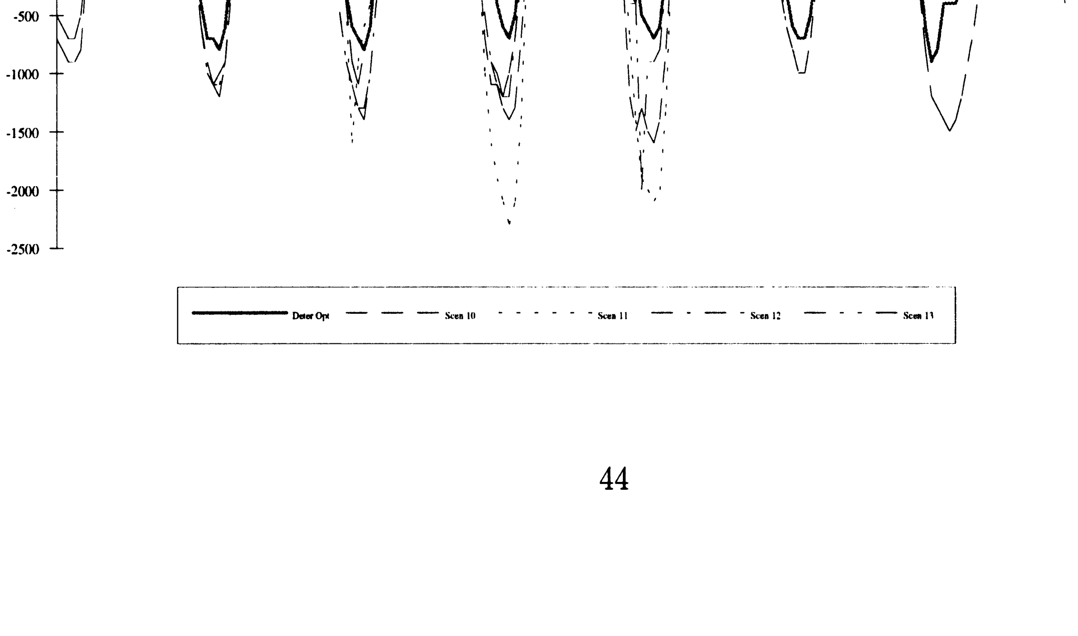

Figure 21: Generation/Consumption Comparison for Scenarios 14 to 18

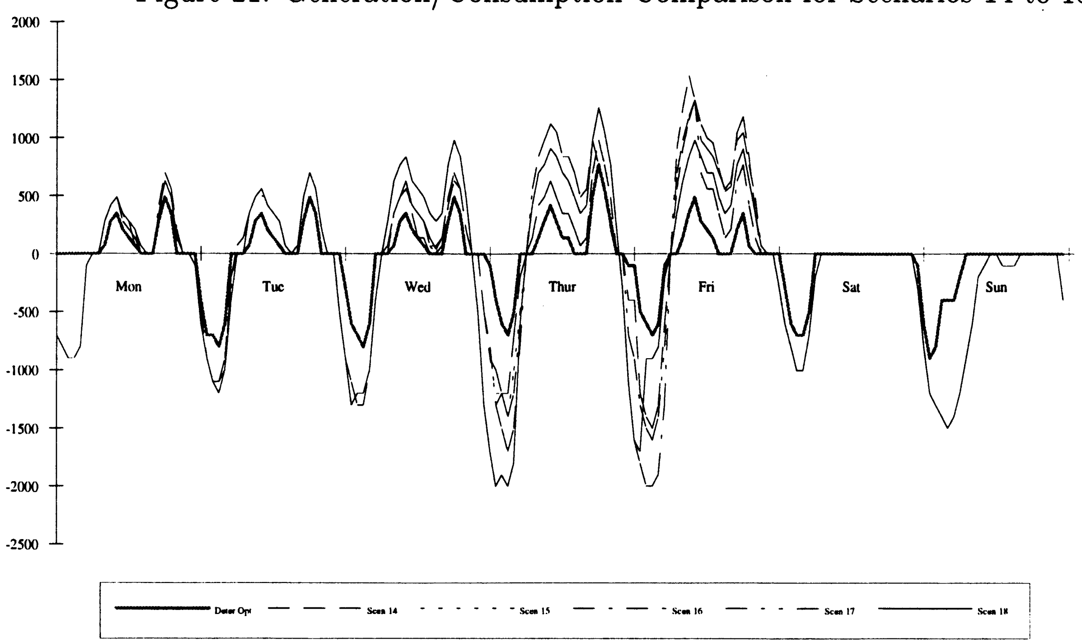

Figure 22: Generation/Consumption Comparison for Scenarios 19 to 22

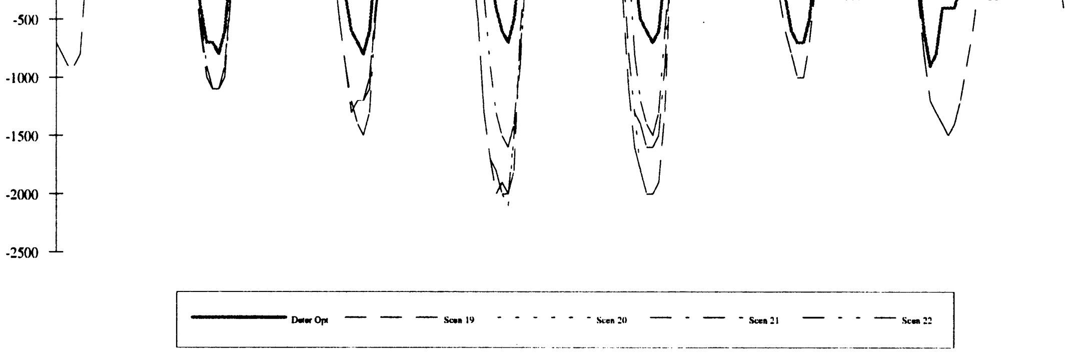

12.1.4 Smoothened Demand Comparison

The following figures show how the hydro-unit smoothens the demand of each scenario. Electricity is generated during peak periods since the marginal cost is very high at these times. On the other hand, the load is increased on the system during the night and the weekend in order to fill up the water storage. The original deterministic demand is displayed on each graph to serve as a measure of comparison.

In most cases, the hydro-unit starts generating electricity when the demand is 10,000 Megawatt-Hours. For these scenarios with high demands, the generation level increases depending on the future demand. We can observe a similar behavior for pumping: the hydro-unit starts pumping when the demand level is lower than 9,000 Megawatt-Hours. The previous demand levels are consistent with the marginal costs, λ, obtained from solving the unit commitment problem: $\lambda \approx \$17$ when the demand is 9,000 Megawatt-Hours and $\lambda \approx \$23$ when the demand is 10,500 Megawatt-Hours, yielding a ratio of 0.74.

Figure 23: Smoothened Demand Comparison for Scenarios 1 to 5

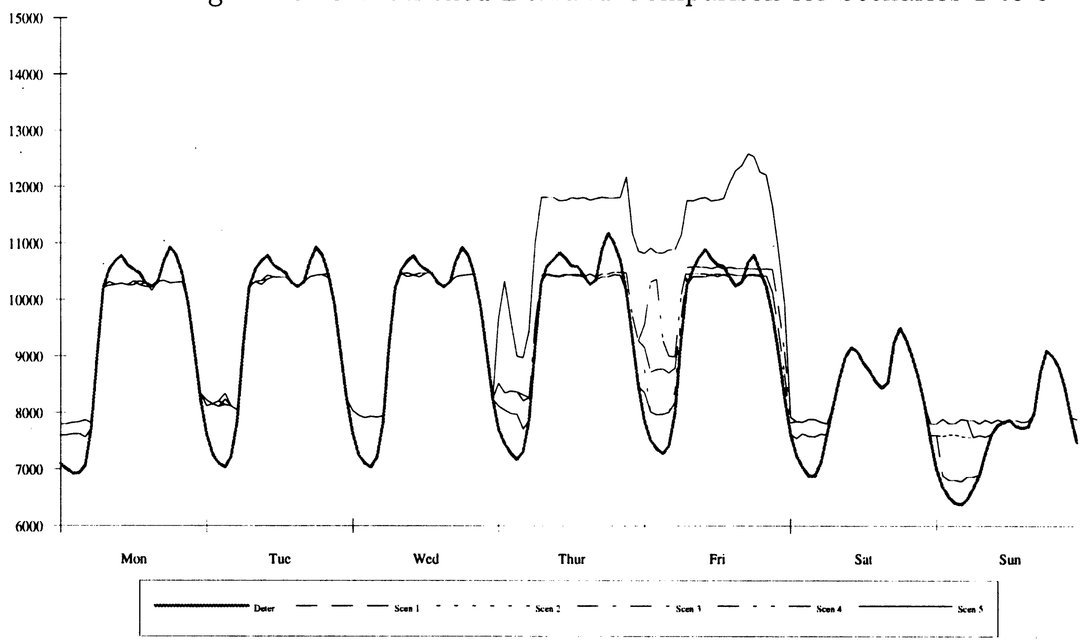

Figure 24: Smoothened Demand Comparison for Scenarios 6 to 9

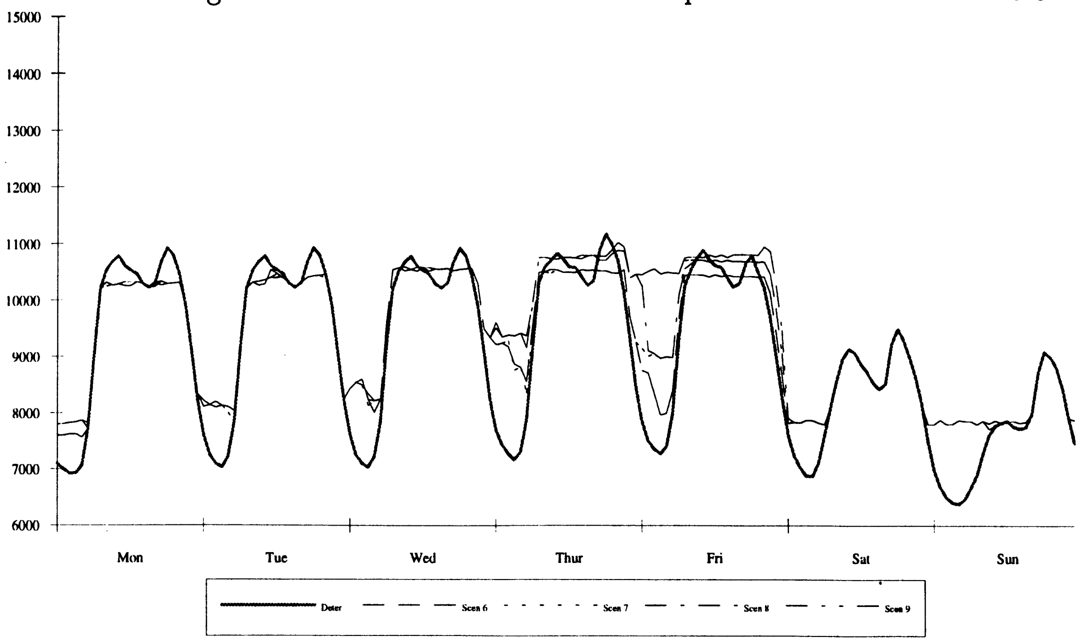

Figure 25: Smoothened Demand Comparison for Scenarios 10 to 13

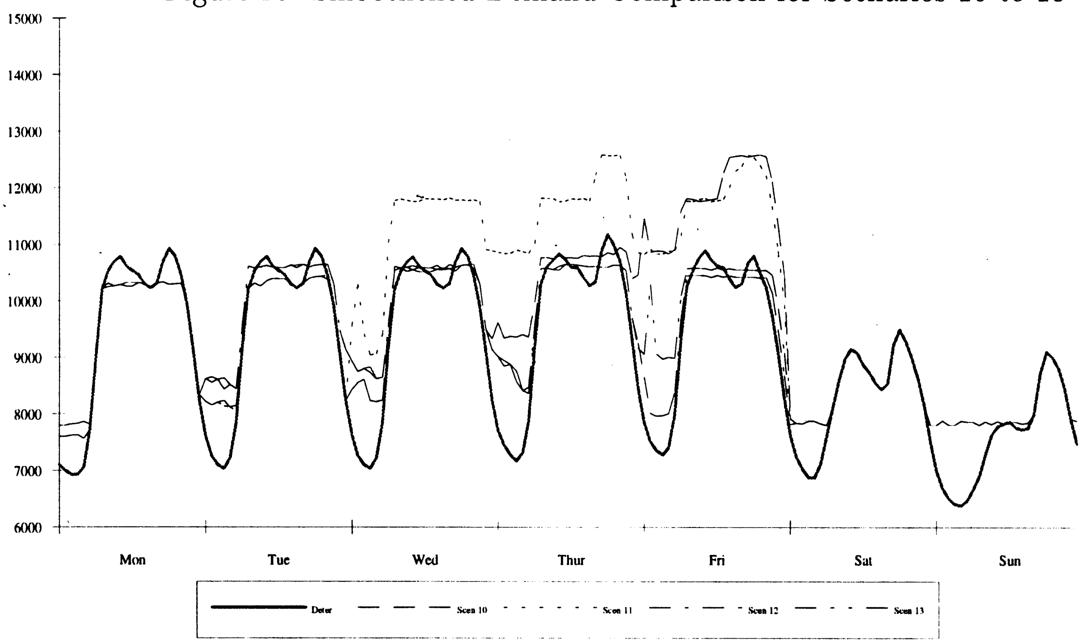

Figure 26: Smoothened Demand Comparison for Scenarios 14 to 18

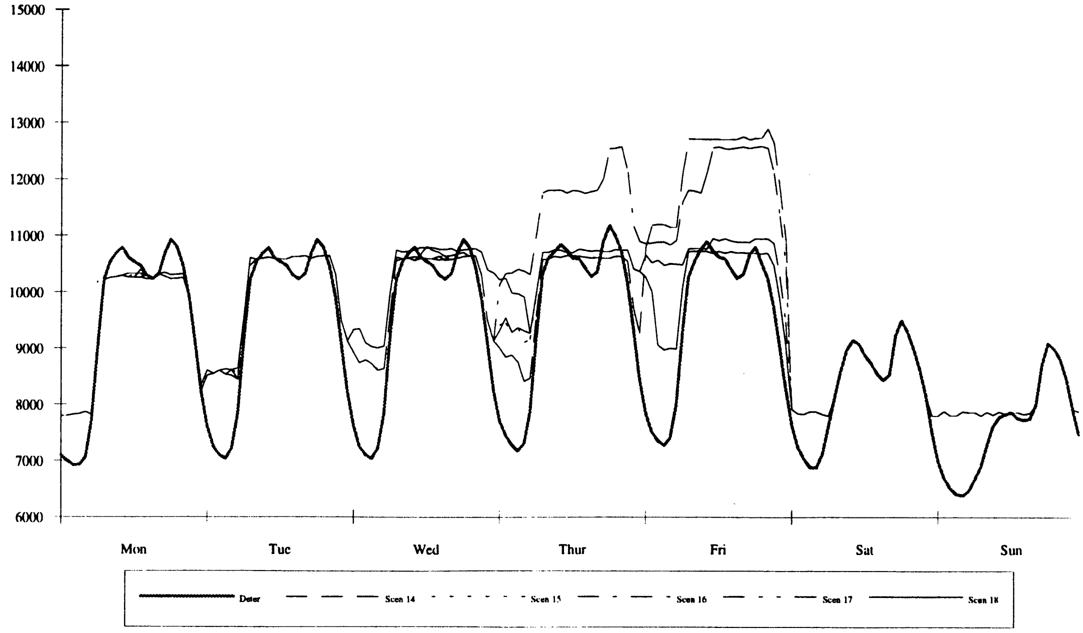

Figure 27: Smoothened Demand Comparison for Scenarios 19 to 22

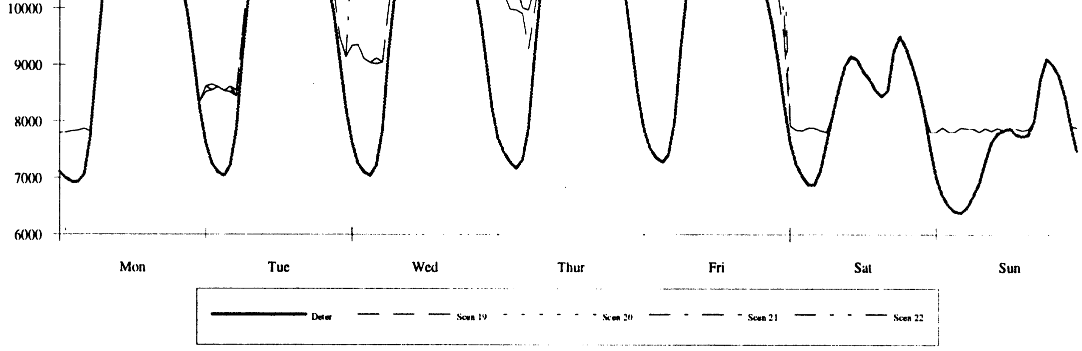

12.2 Inaccurate Forecast Example

We created sixteen different scenarios using the demand data provided by Detroit Edison and Consumers Power for the months of August in years 1990, 1991, 1992, and 1993. The four weeks of August in each of the four years are assumed to be probabilistically independent with identical demand distribution, hence, each week is considered a possible scenario with a weight of $\frac{1}{16}$. In order to create scenario bundles, we assumed that the demand on Monday, Tuesday, and Wednesday is known in advance, and equal to the average demand over these days in the given sixteen scenarios. On Thursday morning, there are four branches, each one of them is taken as the average of the demand on Thursdays in August in each of the four years. Then, each of Thursday's scenarios splits on Friday morning; the resulting sixteen branches have the same demand on Friday, Saturday, and Sunday as that of the sixteen observed scenarios. Figure 28 shows the scenario tree, and Figures 29, 30, 31, and 32 provide the electricity demand used in each one of the sixteen scenarios.

We solved the previous problem for each of the individual scenarios to obtain its optimal policy, which is then applied to other scenarios to compute the expected operating cost. We also solved the problem using the expected demand; the results are shown in the "Deterministic" row of Table 3. Progressive hedging is then used to minimize the expected cost over all scenarios; the results are shown in the last row of the same table. The hedging policy is expected to perform better than any other policy; it saves on the average $120,000 over the deterministic one.

The following sections provide all the results related to this example, namely, different water levels in the pump-storage plant, electricity consumed or generated by the hydro unit, and the smoothened demand curves for different scenarios. The deterministic scenario is drawn to compare its demand to the other scenarios. The optimal water level and the corresponding electricity generation/consumption of the deterministic scenario are also shown.

12.2.1 Demand Comparison

Figure 28: Scenario Tree

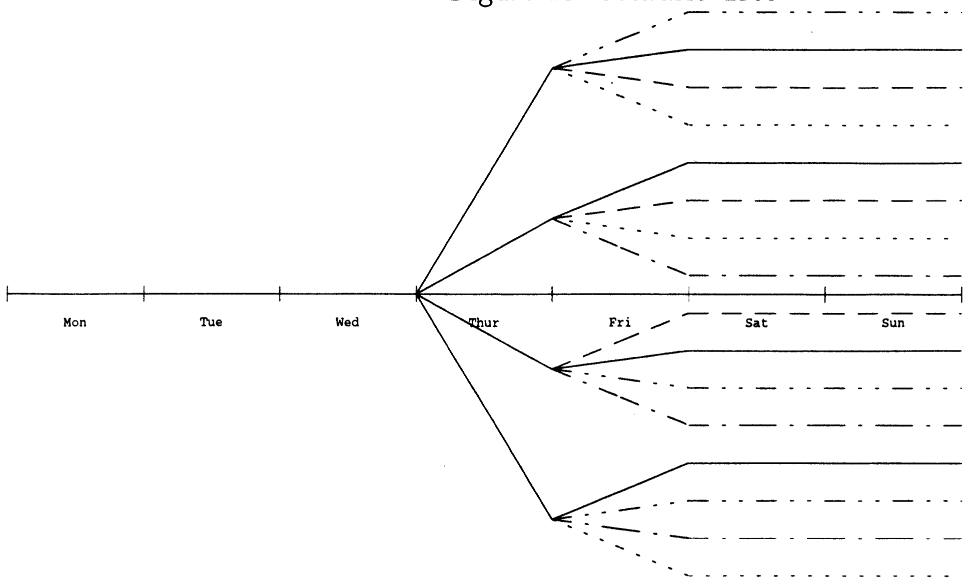

Figure 29: Electricity Demand Comparison for Scenarios 1 to 4

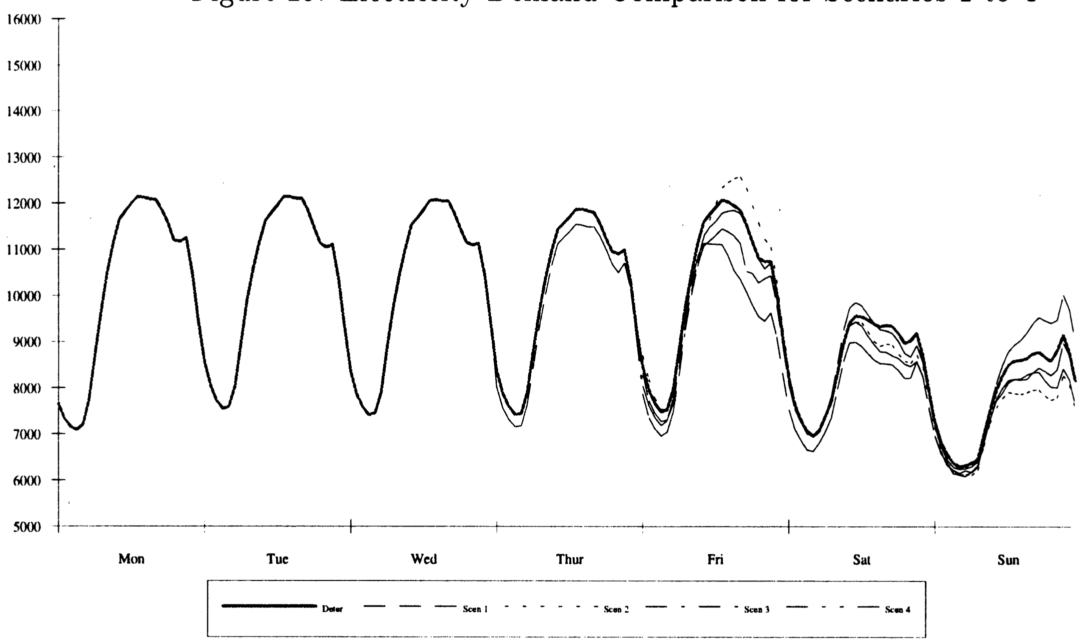

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	$E(x)$	Save	%
Pr	6.25%	6.25%	6.25%	6.25%	6.25%	6.25%	6.25%	6.25%	6.25%	6.25%	6.25%	6.25%	6.25%	6.25%	6.25%	6.25%			
1	20304	20529	20492	20003	20335	19893	19639	24138	21293	21162	20344	21695	21330	20861	19946	25659	21102	199.05	0.95%
2	20332	20496	20517	20027	20359	19919	19664	23672	21244	21064	20369	21551	21354	20885	19973	25134	21035	132.40	0.63%
3	20304	20566	20489	19999	20334	19890	19633	24234	21325	21194	20340	21759	21329	20858	19942	25757	21122	219.49	1.05%
4	20318	20580	20503	19999	20348	19903	19647	24225	21323	21192	20345	21740	21327	20856	19939	25724	21124	221.50	1.06%
5	20311	20534	20496	20008	20334	19896	19644	24142	21286	21155	20336	21687	21311	20842	19927	25639	21097	194.54	0.93%
6	20321	20583	20506	20017	20351	19890	19649	24251	21324	21193	20339	21759	21307	20837	19920	25736	21125	222.38	1.06%
7	20310	20571	20494	20005	20339	19894	19632	24240	21316	21185	20331	21751	21302	20832	19916	25729	21116	213.29	1.02%
8	20461	20621	20643	20156	20485	20047	19796	23150	21321	21150	20482	21487	21446	20981	20075	24107	21026	123.07	0.59%
9	20379	20540	20565	20073	20406	19965	19710	23555	21240	21071	20392	21480	21330	20864	19962	24824	21022	119.60	0.57%
10	20404	20565	20590	20098	20432	19990	19736	23569	21266	21063	20418	21508	21356	20889	19989	24814	21045	142.50	0.68%
11	20340	20599	20525	20036	20370	19925	19668	24270	21337	21207	20330	21772	21295	20826	19909	25725	21135	232.25	1.11%
12	20413	20572	20599	20107	20440	19999	19745	23443	21271	21103	20427	21433	21364	20907	19988	24329	21009	106.24	0.51%
13	20344	20597	20529	20034	20373	19930	19676	24228	21326	21196	20356	21732	21296	20828	19913	25683	21127	224.93	1.08%
14	20342	20595	20527	20038	20371	19928	19673	24228	21323	21194	20354	21730	21295	20826	19911	25681	21126	223.57	1.07%
15	20340	20600	20526	20036	20370	19926	19668	24270	21338	21207	20353	21773	21296	20827	19908	25726	21135	232.72	1.11%
16	20472	20646	20666	20177	20523	20068	19816	23347	21336	21163	20496	21498	21432	20976	20056	23970	21042	139.43	0.67%
Deter	20381	20563	20557	20082	20419	19987	19746	23090	21241	21187	20404	21426	21330	20885	20087	24977	21023	120.09	0.57%
Policy	20306	20497	20489	20000	20341	19894	19642	23162	21240	21075	20352	21433	21296	20826	19910	23977	20903	0.00	0.00%

Table 3: The Cost in $1000 of Applying the Policy of Each Scenario to All Other Scenarios

Figure 30: Electricity Demand Comparison for Scenarios 5 to 8

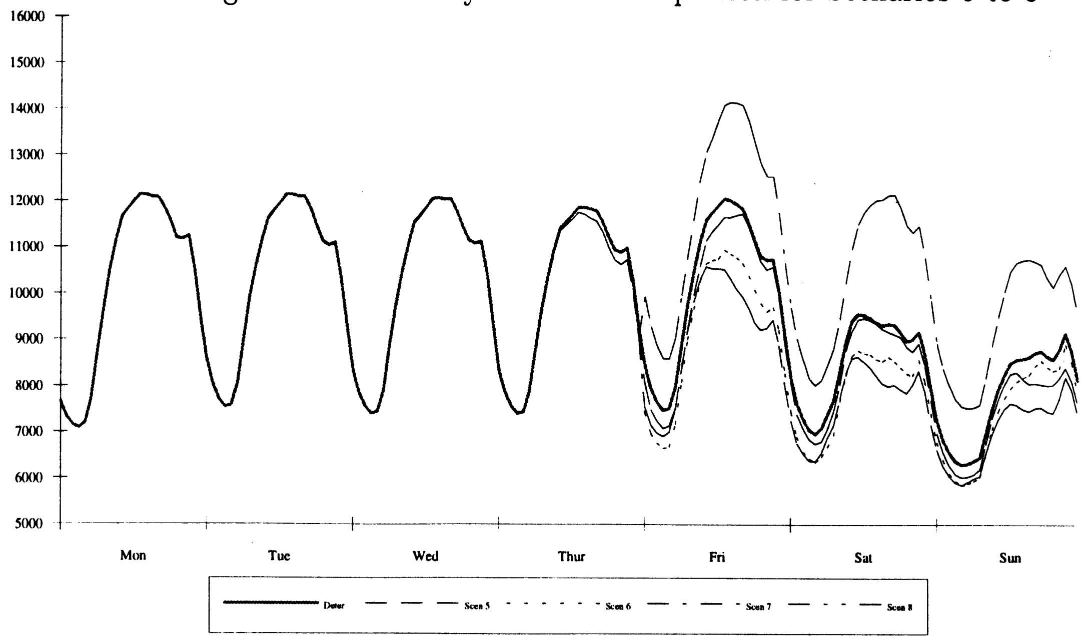

Figure 31: Electricity Demand Comparison for Scenarios 9 to 12

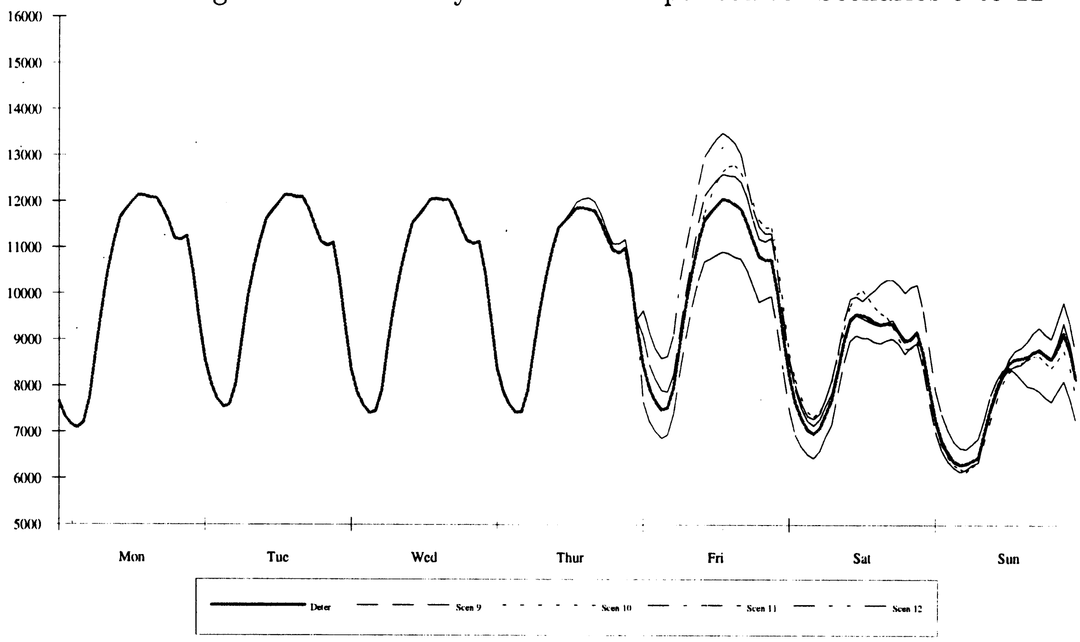

Figure 32: Electricity Demand Comparison for Scenarios 13 to 16

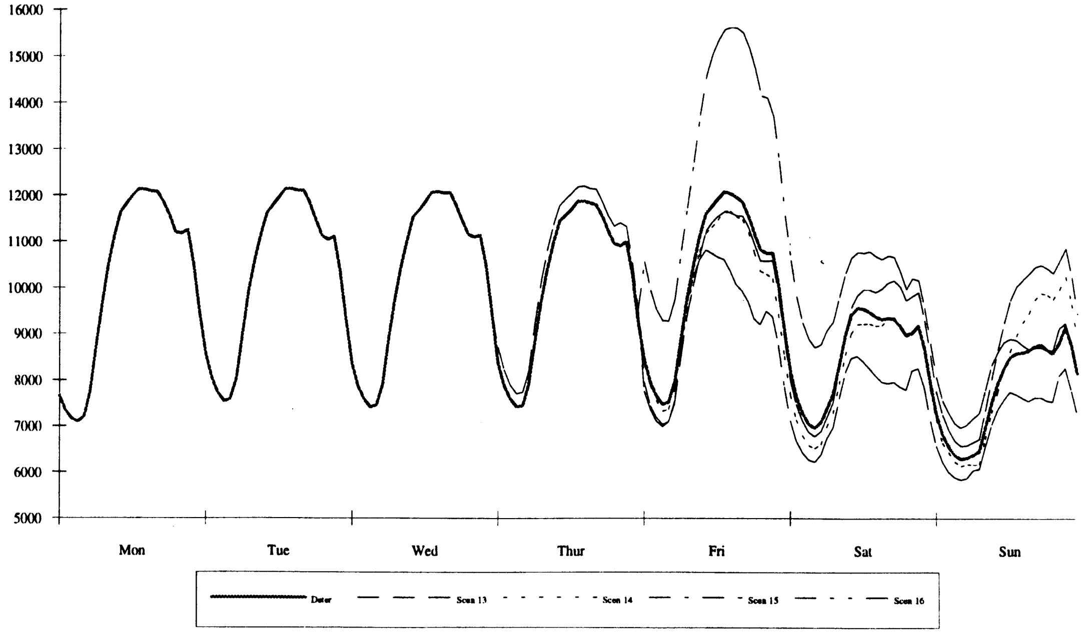

12.2.2 Pump-Storage Plant Water-Level Comparison

Figure 33: Water level Comparison for Scenarios 1 to 4

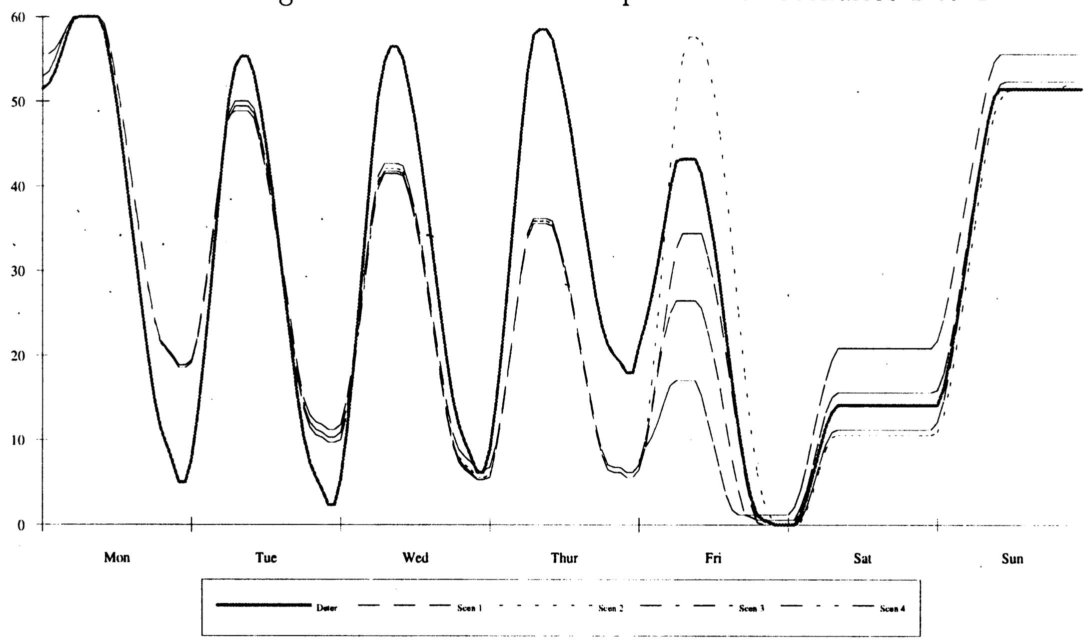

Figure 34: Water level Comparison for Scenarios 5 to 8

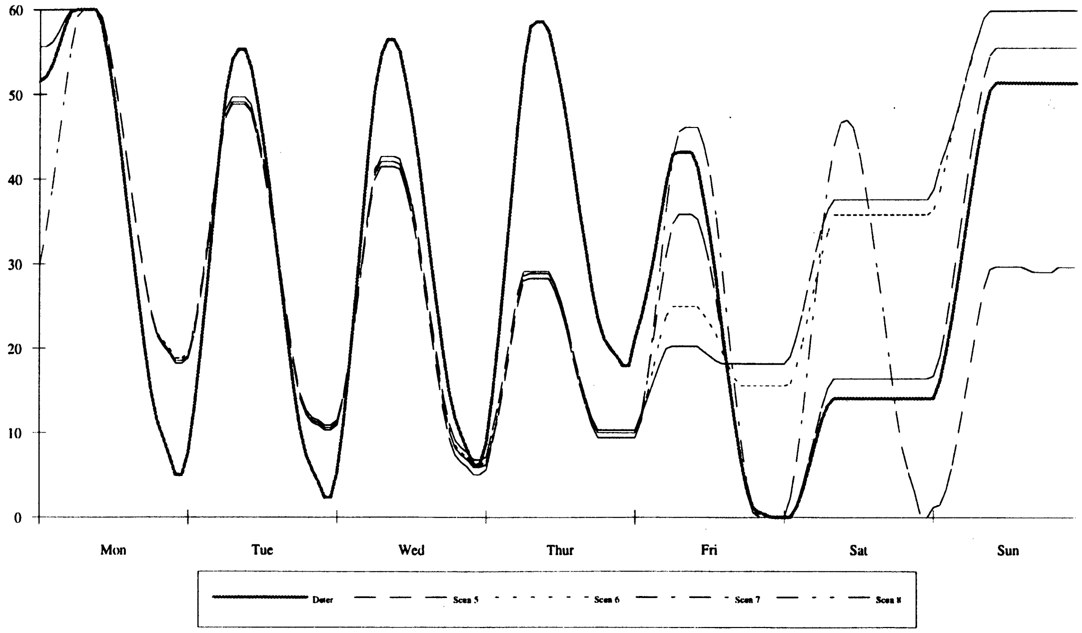

Figure 35: Water level Comparison for Scenarios 9 to 12

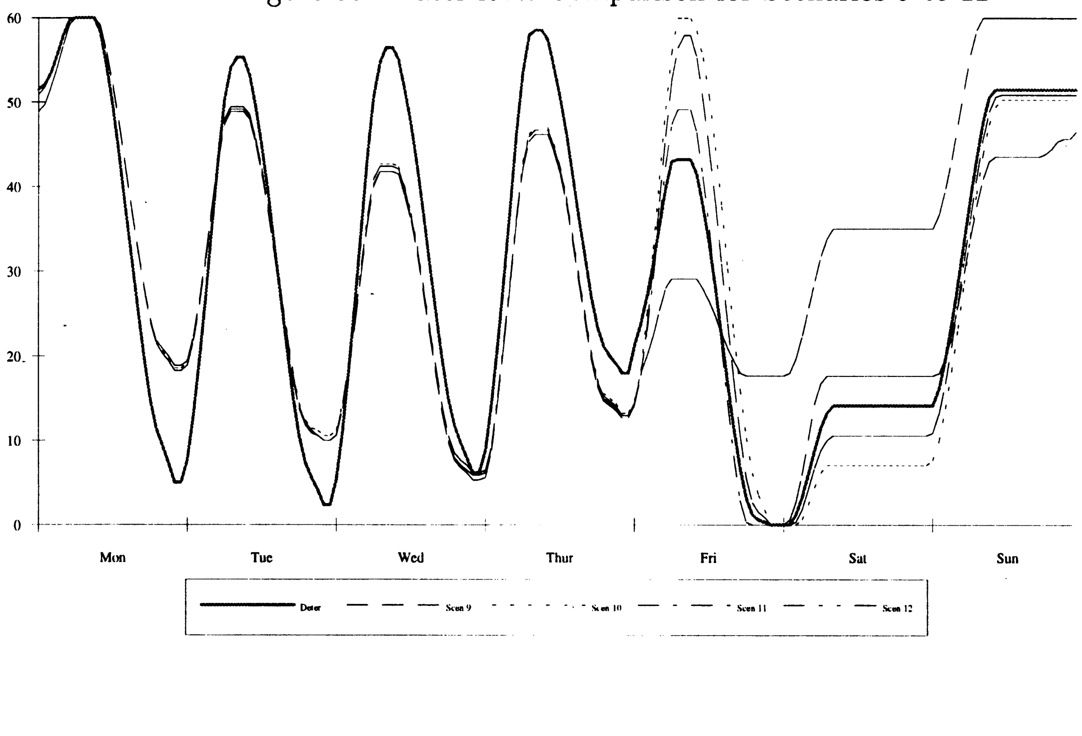

Figure 36: Water level Comparison for Scenarios 13 to 16

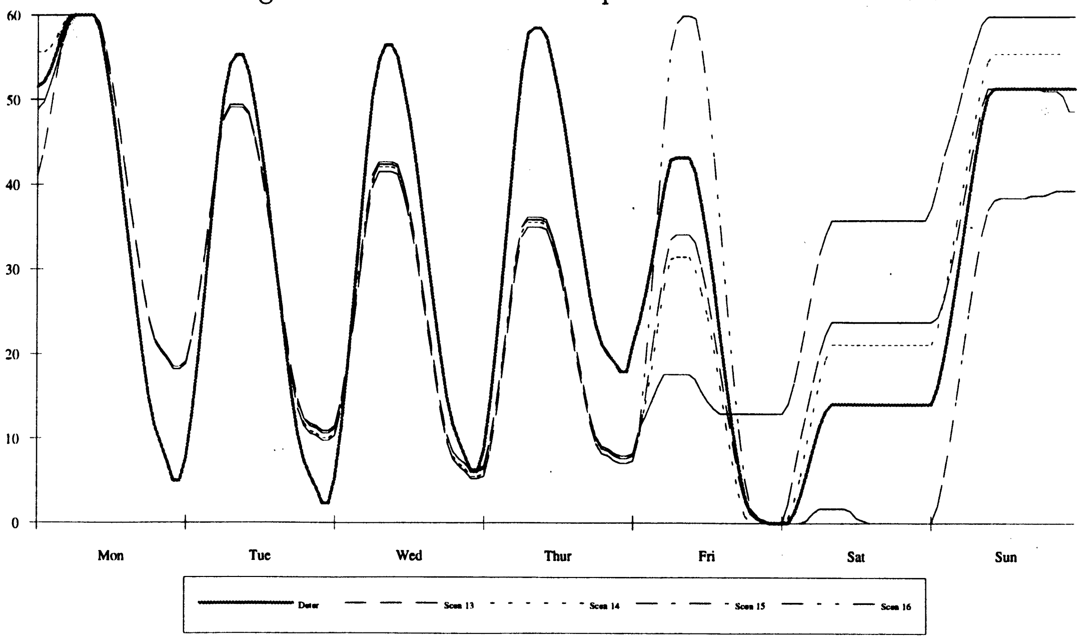

12.2.3 Pump-Storage Plant Generation/Consumption Comparison

Figure 37: Generation/Consumption Comparison for Scenarios 1 to 4

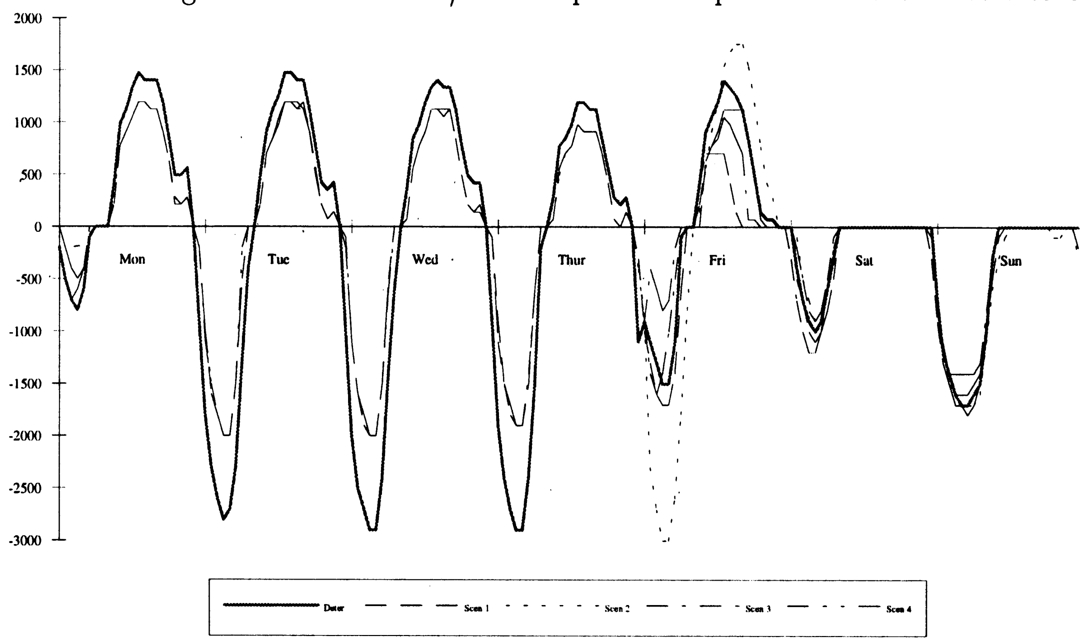

Figure 38: Generation/Consumption Comparison for Scenarios 5 to 8

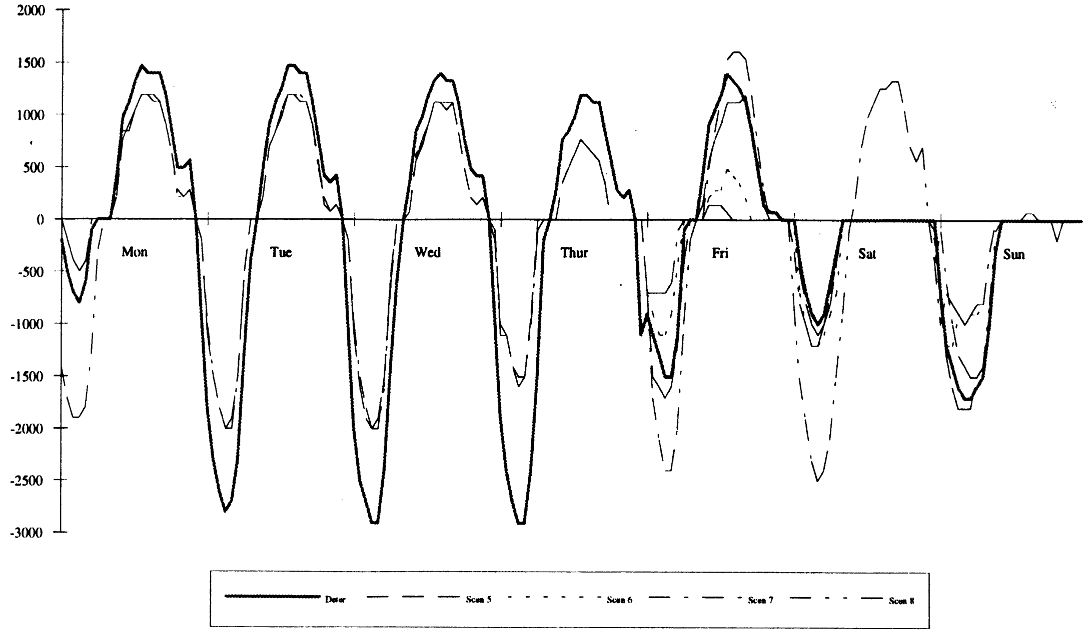

Figure 39: Generation/Consumption Comparison for Scenarios 9 to 12

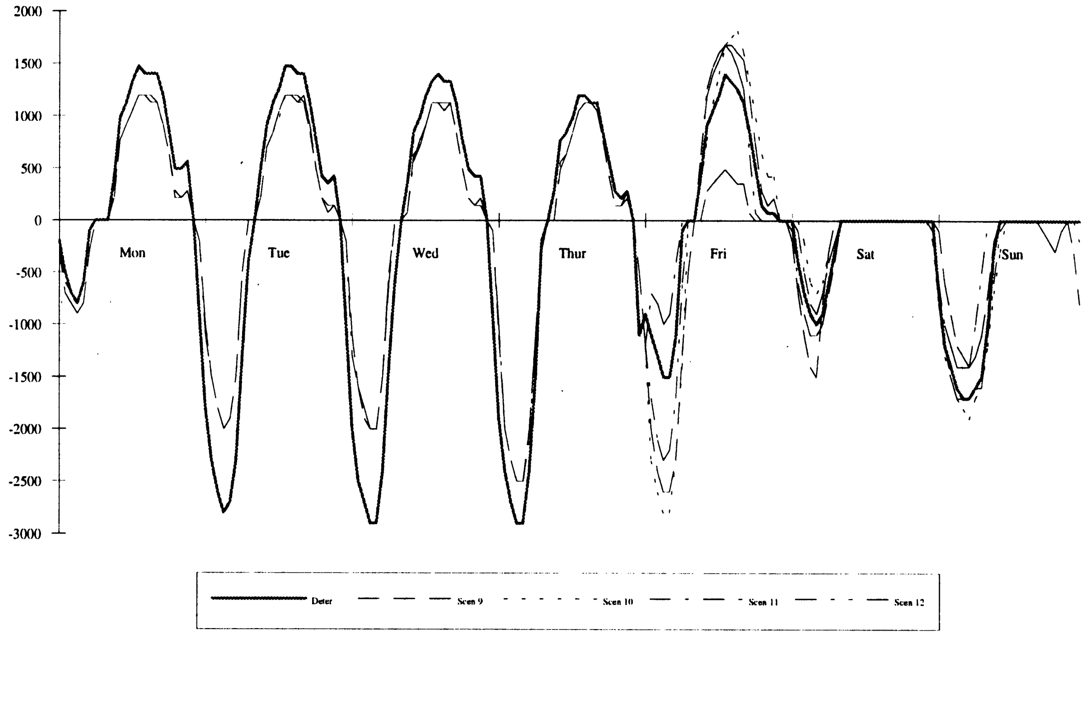

Figure 40: Generation/Consumption Comparison for Scenarios 13 to 16

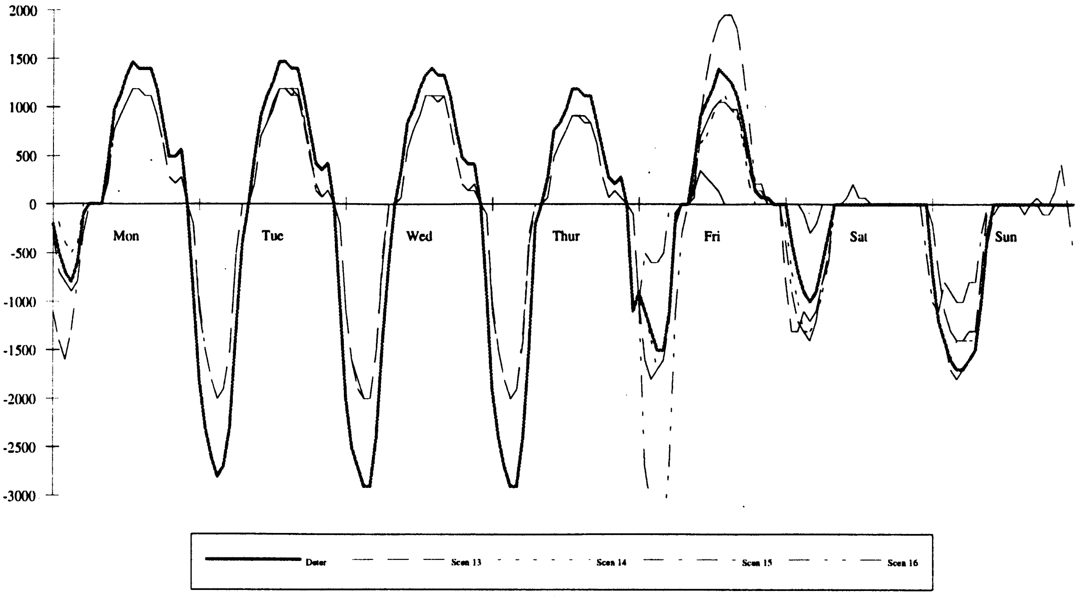

12.2.4 Smoothened Demand Comparison

Figure 41: Smoothened Demand Comparison for Scenarios 1 to 4

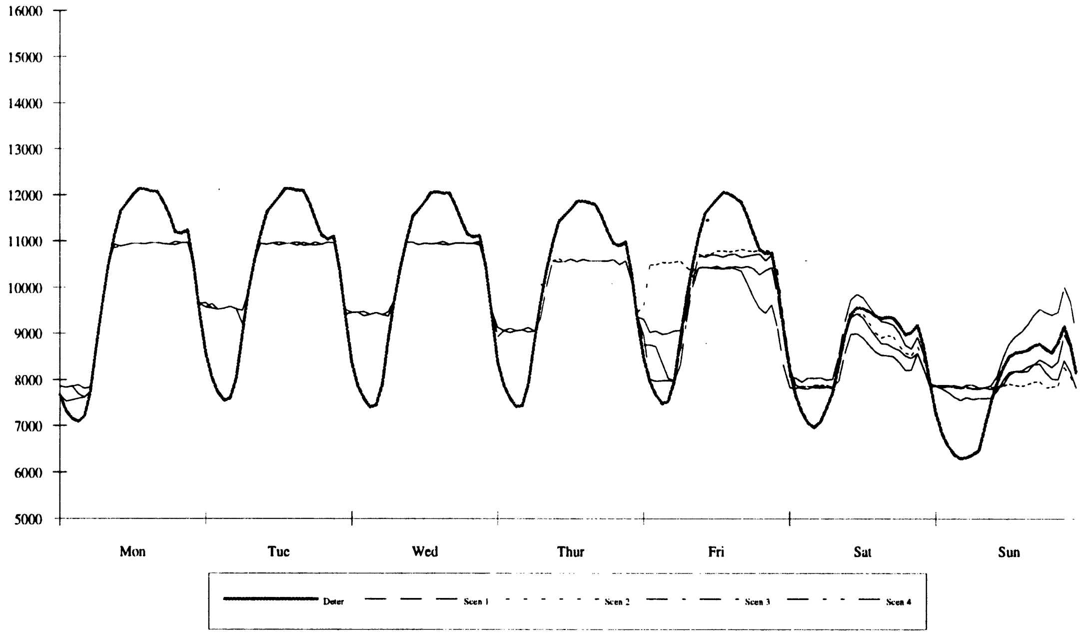

Figure 42: Smoothened Demand Comparison for Scenarios 5 to 8

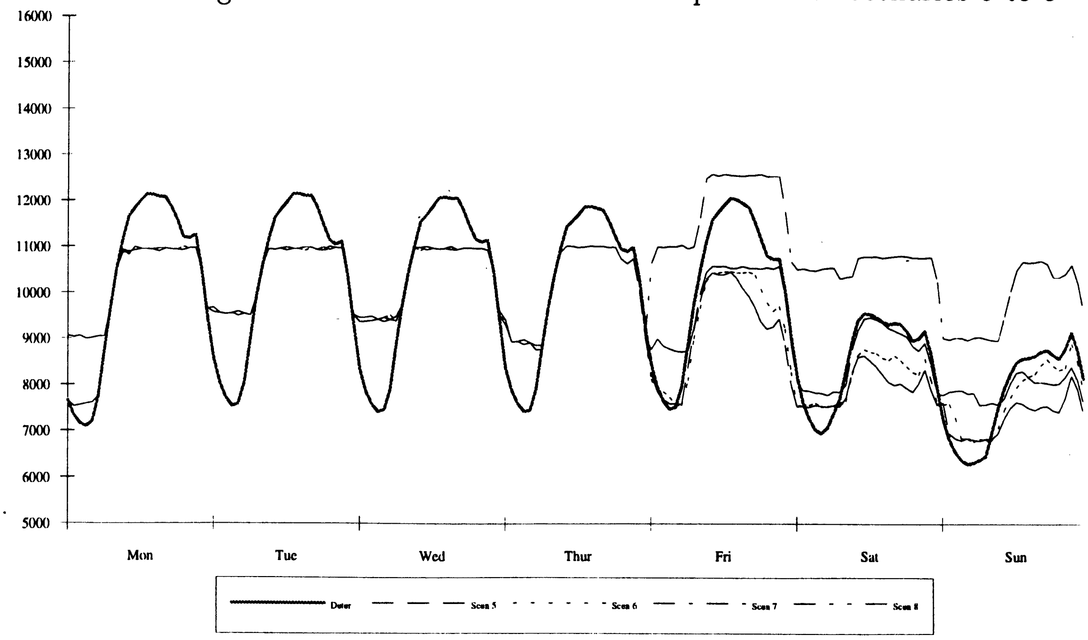

Figure 43: Smoothened Demand Comparison for Scenarios 9 to 12

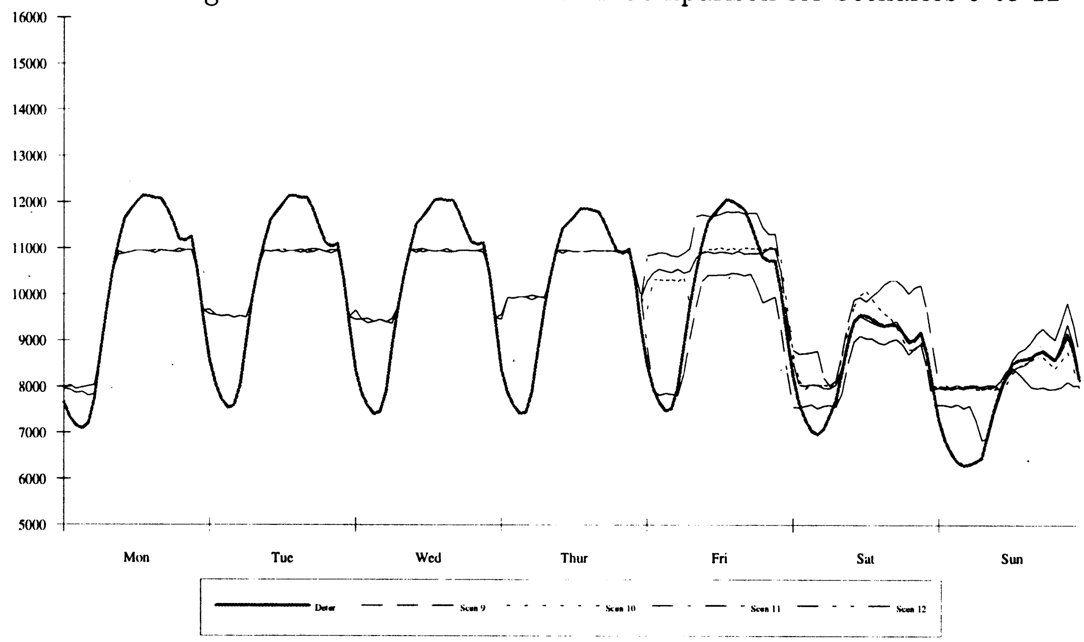

Figure 44: Smoothened Demand Comparison for Scenarios 13 to 16

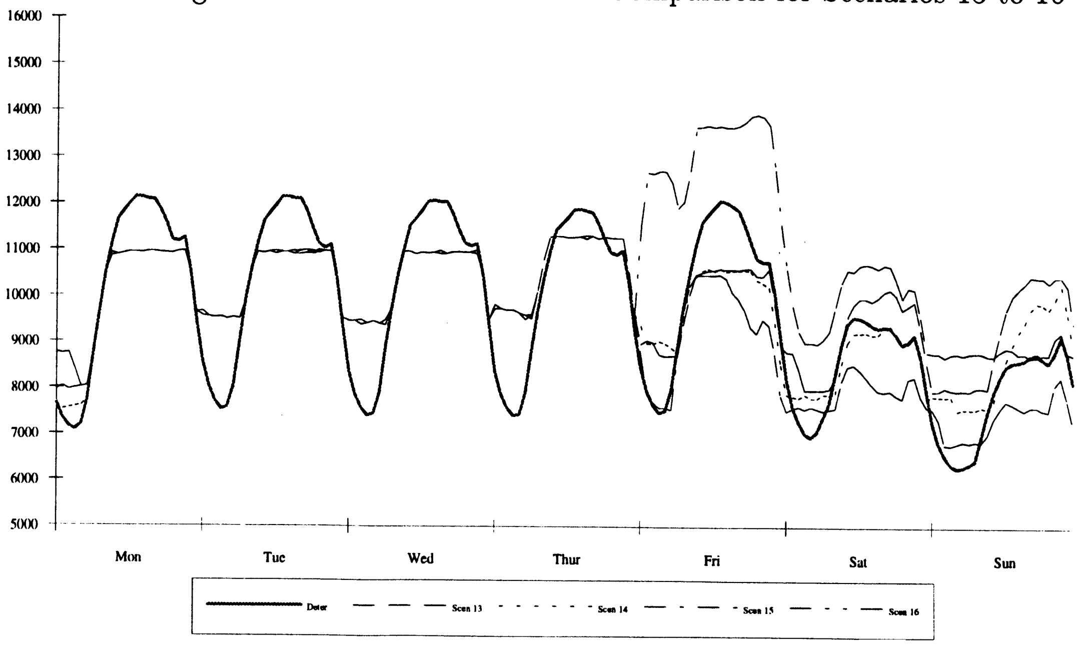

13 Conclusions

We improved the existing method for solving the unit commitment problem by selecting a good starting set of marginal costs, and improving the single-generator sub-problem dynamic programming formulation. In order to handle demand uncertainty, we suggested solving the unit commitment problem for different scenarios in order to obtain an optimum policy for each of them. Since the policies may differ, a penalty term is applied to each policy which violates the average policy. The problem is then resolved and new penalties are obtained. This process is repeated until a unique optimum policy is reached. The preliminary results indicate that using the progressive hedging technique can reduce the operating cost of the system by 1-2%.

References

[1] K. Aoki, M. Itoh, T. Satoh, K. Nara, and M. Kanezashi. Optimal long-term unit commitment in large scale systems including fuel constrained thermal and pumped-storage hydro. *IEEE Transactions on Power Systems*, 4(3):1065–1073, August 1989.

[2] C. H. Bannister and R. J. Kaye. A rapid method for optimization of linear systems with storage. *Operations Research*, 39(2):220–232, March-April 1991.

[3] M. S. Bazaraa and H. D. Sherali. On the choice of step size in subgradient optimization. *European Journal of Operational Research*, 7:380–388, 1981.

[4] D. P. Bertsekas. *Constrained Optimization and Lagrange Multiplier Methods*, chapter 5.6, page 368. Academic Press, 1982.

[5] D. P. Bertsekas, G. S. Lauer, N. R. Sandell, and T. A. Posbergh. Optimal short-term scheduling of large-scale power systems. *IEEE Transactions on Automatic Control*, 28(1):1–11, January 1983.

[6] East Central Area Reliability Group. Reliability criterion, 1981. Document #1.

[7] A. Merlin and P. Sandrin. A new method for unit commitment at electricité de france. *IEEE Transactions on Power Apparatus and Systems*, PAS-102(5):1218–1225, May 1983.

[8] J. A. Muckstadt and S. A. Koenig. An application of lagrangian relaxation to scheduling in power-generation systems. *Operations Research*, 25(3):387–403, May-June 1977.

[9] M. V. F. Pereira and L. M. V. G. Pinto. Stochastic optimization of a multireservoir hydroelectric system: A decomposition approach. *Water Resources Research*, 21(6):779–792, June 1985.

[10] B. T. Poljak. A general method for solving extremal problems. *Soviet Mathematics Doklady*, 8:593–597, 1967.

[11] R. T. Rockafellar and R. J.-B. Wets. Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of Operations Research*, 16(1):119–147, 1991.

[12] S. Ruzić and N. Rajaković. A new approach for solving extended unit commitment problem. *IEEE Transactions on Power Systems*, 6(1):269–277, February 1991.

[13] J. F. Shapiro. *Mathematical Programming: Structures and Algorithms.* John Wiley & Sons, Inc., 1979.

[14] F. Zhuang and F. D. Galiana. Towards a more rigorous and practical unit commitment by lagrangian relaxation. *IEEE Transactions on Power Systems*, 3(2):763–773, May 1988.