

**A LAGRANGIAN RELAXATION APPROACH
COUPLED WITH HEURISTIC TECHNIQUES FOR
SOLVING THE UNIT COMMITMENT PROBLEM**

SAMER TAKRITI

Department of Industrial and Operations engineering
University of Michigan, Ann Arbor, MI 48109, USA

JOHN R. BIRGE

Department of Industrial and Operations engineering
University of Michigan, Ann Arbor, MI 48109, USA

ERIK LONG

Department of Industrial and Operations engineering
University of Michigan, Ann Arbor, MI 48109, USA

Technical Report 97-18
December 12, 1997

A Lagrangian Relaxation Approach Coupled with Heuristic Techniques for Solving the Unit Commitment Problem

Samer Takriti, John R. Birge, and Erik Long

The University of Michigan

Department of Industrial and Operations Engineering

Ann Arbor, Michigan 48109-2117

Abstract— We develop a fast technique for solving the problem of scheduling the generating units of an electrical power system. This technique combines the traditional Lagrangian relaxation method and a heuristic method, such as simulated annealing or genetic algorithms, to obtain a fast solution for the unit commitment problem. Numerical results indicate a significant improvement in the quality of the solution and in the calculation time when this technique is used instead of pure Lagrangian relaxation.

Keywords— Unit commitment problem, Lagrangian relaxation, Dynamic programming, Heuristics, Simulated annealing, Genetic algorithms, Iterated hill-climbing.

I. INTRODUCTION

ONE of the most important problems in electrical power generation is the unit commitment problem. The primary concern of electrical power 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 the system reservoirs may not prove to be sufficient to respond to high demands. Therefore, costly thermal generating-units are often used to make up for the supply shortage.

The unit commitment problem refers to the task of finding an optimal schedule and 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 a scheduling horizon [6]. This problem becomes a multi-stage program with 0/1 variables. In this paper, we develop a fast method for solving the unit commitment problem. Our approach is a combination of Lagrangian relaxation, which is traditionally used to solve this problem, and a heuristic technique, such as simulated annealing. The heuristic builds on the solutions provided by the Lagrangian relaxation technique. The execution time is reduced since the Lagrangian dual does not have to be solved to optimality. The resulting solution is superior to that obtained from optimizing the dual.

To fix the notation throughout this paper, we assume that there are n generating units and that the duration of the

study horizon is T periods. A 24-hour horizon may be sufficient, but a longer horizon, one week or more, is often needed if pump or other storage units are considered. The state of each unit, i , at a time period, t , is represented by the 0-1 variable, u_t^i . A unit is on at time t if $u_t^i = 1$, and off if $u_t^i = 0$. The power output 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, i , 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 non-decreasing 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 [7]. 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 $\bar{f}_i(x_t^i, u_{t-1}^i, u_t^i)$. 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 must be off for at least l_i periods.

The mathematical formulation of the previous model is

$$\begin{aligned} \min_{x,u} \quad & \sum_{i=1}^n \sum_{t=1}^T \bar{f}_i(x_t^i, u_{t-1}^i, u_t^i) \\ \text{s.t.} \quad & \sum_{i=1}^n x_t^i \geq d_t, \end{aligned} \quad (1)$$

where $d_t \geq 0$ represents the total demand for electricity during period t . We assume that d_t is known in advance. In addition to the previous constraint, each unit must satisfy the minimum on-time, minimum off-time, and minimum and maximum operating level constraints. Usually, another constraint, $\sum_{i=1}^n G_i u_t^i \geq \bar{d}_t$, is added to the previous formulation. The value of \bar{d}_t represents the maximum demand that can occur at time t . The constraint represents the spinning reserve of the system. It states that the total planned operating capacity must be greater than the highest anticipated load [6], [8]. Without loss of generality, we do not include the spinning reserve constraint in our formulation. It can be treated in a similar fashion to the demand constraint.

The problem in (1) is a large-scale mixed-integer quadratic program. Many approaches have been proposed to solve this problem [2]. They can be classified into branch-and-bound methods, dynamic programming, priority ordering, and the Lagrangian relaxation method. The first two techniques are satisfactory from the theoretical point of view,

This work was partially supported by the National Science Foundation under Grant ECS-9216819 and the Electric Power Research Institute under Grant RP8030-13.

but they are practically intractable due to the large storage size required to implement them on a computer. The third approach is a greedy strategy and does not guarantee an optimal solution in general. The Lagrangian relaxation technique seems to be the most efficient, because it attempts to solve the problem indirectly by solving the dual problem.

Muckstadt and Koenig [6] appear to have first addressed the unit commitment problem with mathematical programming and to suggest a sound technique for solving it. They use a Lagrangian relaxation technique which decomposes the given problem into smaller subproblems. Each subproblem corresponds to 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 the minimum on-time and minimum off-time requirements. Dynamic programming is then used to solve the subproblems. A lower bound on the optimal cost of the primal problem is obtained.

Given this lower bound, Muckstadt and Koenig 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, 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 is guaranteed to find an optimal 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 [3] 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 to reach an acceptable answer. To accelerate 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.

Bertsekas et al.'s main contribution is the provision of 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 $\frac{2T(S^*+TC^*)}{\mathcal{L}^*}$ 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 \mathcal{L}^* is the optimal dual functional value.

In this paper, we use the Lagrangian dual to provide a set of schedules for the cycling units of the system. To find the best combination of the resulting schedules, an integer program must be solved. We use a heuristic to solve this integer program approximately. The Lagrangian relaxation approach is presented in Section II. Section III describes the approximate mathematical model and Section IV describes the suggested heuristic techniques for solving this model. We present numerical results in Section V.

II. A LAGRANGIAN RELAXATION APPROACH FOR SOLVING THE UNIT COMMITMENT PROBLEM

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 subproblems. 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

$$\max_{\lambda \geq 0} \mathcal{L}(\lambda),$$

where

$$\mathcal{L}(\lambda) = \min_{x,u} \sum_{i=1}^n \sum_{t=1}^T \bar{f}_i(x_t^i, u_{t-1}^i, u_t^i) - \frac{\lambda_t}{n} (x_t^i - d_t),$$

subject to unit minimum and maximum operating levels, and minimum on-time and off-time constraints. For a given λ , the value of $\mathcal{L}(\lambda)$ is computed by solving n mixed-integer quadratic programs,

$$F_i(\lambda) = \min_{x,u} \sum_{t=1}^T \bar{f}_i(x_t^i, u_{t-1}^i, u_t^i) - \frac{\lambda_t}{n} (x_t^i - d_t), \quad (2)$$

and adding the resulting values of $F_i(\lambda)$ over all generators. The optimization problem in (2) is called the i^{th} single-generator sub-problem since it only depends on the specifications of generator i . It can be solved efficiently using dynamic programming [8].

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 optimal objective function value. The difference between the optimal value of the original program and the optimal value of the Lagrangian dual problem is called the duality gap. It 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 that it is 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 a concave function. Note

2. Calculate the total generation, $\sum_{i=1}^n x^i$, corresponding to ω . There is no need to evaluate the objective function value.
3. If $|\sum_{i=1}^n x^i - d| < \epsilon$, terminate.
4. If $\sum_{i=1}^n x^i < d$, $\underline{\omega} \leftarrow \omega$. Otherwise, $\bar{\omega} \leftarrow \omega$.
5. Go to step 1.

In the following section, we give a brief description of three heuristic techniques to approximate (3): simulated annealing, genetic algorithm, and iterated hill-climbing.

IV. APPLYING HEURISTICS TO THE APPROXIMATE MATHEMATICAL FORMULATION

The main difficulty in applying heuristics to constrained optimization problems is in handling the constraints efficiently. Given a solution, $v^{i,k}$, the demand constraints are incorporated into the objective function by adding the term $\max\{0, \alpha(d_t - \sum_{i=1}^n x^i)\}$. Choosing very high penalties forces feasible solutions and ignores the value of the objective function. On the other hand, small penalties may result in infeasible solutions. In the power generation case, the penalty coefficient, α , represents the cost of producing/buying emergency electricity. For any utility company, the value of α varies within a small range which simplifies the task of choosing α .

The optimization problem of (3) takes on the form

$$\begin{aligned} \min_{\mu} \quad & \sum_{i=1}^n g(\mu_i) \\ \text{s.t.} \quad & 1 \leq \mu_i \leq K_i, \text{ and } \mu_i \text{ is integer,} \end{aligned}$$

where $g(\cdot)$ is evaluated using the technique described in Section III, and μ_i represents the number of the schedule to be used for generator i . Note that $v^{i,\mu_i} = 1$ and $v^{i,k \neq \mu_i} = 0$. To represent the decision variables of the problem, we use an integer vector of dimension n . Each element, i , in that vector is forced to be within the range 1 to K_i . The operators are designed to preserve this requirement.

A. Simulated Annealing

Starting from the best solution provided by the Lagrangian relaxation, a new solution is generated by randomly changing the value of a variable, μ_i , in the range 1 to K_i . If the new solution has a better objective function value, F^* , than that of the current solution, F , then the new solution is used. A new solution can also replace the current solution if

$$\text{rand}[0,1) < e^{(F^* - F)/T}.$$

Here, rand is a random number in the range $[0,1)$ and T is the temperature parameter of the solution technique. After a few iterations, we assume that a ‘‘thermal equilibrium’’ (following the annealing analogy) is reached and the temperature is lowered. The process is then repeated until T is small; that is, the system is frozen. The performance of the method depends on the parameter T and the rate at which it is changed. A more detailed description of this technique can be found in [1].

B. Genetic Algorithms

The algorithm starts with a population that contains randomly generated ‘‘chromosomes’’ (a description of a solution using the genetics analogy) as well as the chromosomes provided by the Lagrangian iterations. New generations are then produced using changes called mutation and single-point crossover. In the case of mutation, a gene, i , of a randomly selected chromosome, μ , is altered arbitrarily in the range 1 to K_i . In the single-point crossover, we form two new offsprings from two randomly selected parents by swapping corresponding segments of the parents’ chromosomes. The crossover point is selected randomly. After each iteration, the best member in the population survives to the next generation. The interested reader may refer to [5] for more details.

C. Iterated Hill-Climbing

This is a greedy strategy. Starting from a feasible solution, a certain number of solutions are produced by altering the value of a randomly selected variable, μ_i , in the appropriate range. We choose the best among all of these solutions and then repeat the process for a number of iterations. The resulting solution is very sensitive to the starting solution.

V. NUMERICAL RESULTS

We implemented the Lagrangian relaxation approach and the three heuristic techniques using the C language in a UNIX environment. All the results presented in this section were produced on a SUN SPARCstation 20 under SunOS 2, release 4.1.3. The test data was supplied by the Michigan Electric Power Coordination Center. We used a 168-hour planning horizon. The execution time for evaluating $g(\cdot)$ for a given set of schedules, μ , depends on the threshold, ϵ , when the quadratic program of Section III is solved. We chose ϵ to be less than 0.5×10^{-3} of the lowest demand per period. In this case, the CPU time for evaluating $g(\cdot)$ was 0.056 seconds.

Since our goal is to produce a quick solution for the unit commitment problem, we keep the number of functional evaluations, $g(\mu)$, very small in all heuristics. For the simulated annealing, we start with a temperature parameter of 20 and reduce it by a factor of 0.85 after reaching the thermal equilibrium. The temperature is reduced 20 times. In the case of genetic algorithms, the size of our population is 20 and the number of generations is 30. The hill-climber attempts 20 different random directions and repeats the process 30 times. In all cases, the number of functional evaluations is close to 600. This approach may seem to violate the spirit of these techniques of repeating the process many times. However, using a small number of functional evaluations is justified if we take a closer look at the progress of the objective function value when a heuristic is used.

Table II displays the progress of the best objective function value, $\min \sum_{i=1}^n g(\mu_i)$, for four different problems when simulated annealing is used. In the first column, we have the number of functional evaluations performed so far.

TABLE II
THE PROGRESS OF THE SIMULATED ANNEALING HEURISTIC FOR
FOUR TEST PROBLEMS

	1	2	3	4
1	21584668	24707732	21687471	18830124
30	21340819	24410338	21088045	18730916
60	21163972	24265624	21088045	18376573
90	21019058	24258771	21088045	18376573
120	21019058	24256922	21076796	18376573
150	20969437	24252298	21076796	18154308
180	20957253	24249676	21020033	18154308

From Table II, the objective function values corresponding to different solutions appear relatively close to each other. Hence, searching in the neighborhood of the current solution may yield an acceptable sub-optimal solution.

In Table III, we solve the four test problems using simulated annealing, genetic algorithms, and iterated hill-climbing. The heuristic techniques begin after 200 Lagrangian iterations or a duality gap of less than 1% is reached. The number of possible solutions for each problem is given in the row "State Space". The "Lagrangian Relaxation" row provides the best primal solution corresponding to the optimum of the Lagrangian dual. Note that the solution provided by simulated annealing is superior to that of the dual in most cases. As Table III indicates, the choice of the initial Lagrange multipliers is significant. It affects the size of the state space, i.e., the number of schedules produced for each generator. It also improves the quality of the solution obtained. Here, we approximate the initial Lagrange multipliers using the technique described in [8].

VI. CONCLUSIONS

We develop a fast technique for solving the unit commitment problem. It starts by solving the Lagrangian dual up to a certain point. Then, it builds on the solutions obtained from the Lagrangian iterations to approximate the original problem. The approximation is a mixed-integer program that can be solved accurately and efficiently using heuristic techniques. We use simulated annealing, genetic algorithms, and iterated hill-climbing as heuristics to solve the previous program. The execution time as well as the solution obtained improve when our approach is used.


REFERENCES

- [1] D. H. Ackley. An empirical study of bit vector function optimization. In L. Davis, editor, *Genetic Algorithms and Simulated Annealing*, pages 170–204, Los Altos, California, 1987. Morgan Kaufmann Publishers.
- [2] 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.
- [3] 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.
- [4] P. O. Lindberg and S. Feltenmark. Solving detailed structured duals of unit commitment problems. In *15th International Symposium on Mathematical Programming*, University of Michigan, Ann Arbor, August 1994.


Table III
SOLUTIONS OF THE APPROXIMATE MATHEMATICAL FORMULATION FOR FOUR TEST PROBLEMS

Problem	Gap \leq 1% and initial λ is 0				Gap \leq 1% and approximation of initial λ			
	1	2	3	4	1	2	3	4
Simulated Annealing	20947987	24238373	20851877	18154308	20863734	24113105	20828779	18140447
Genetic Algorithm	21094499	24456744	20916272	18547278	20863734	24113105	20828779	18140447
Hill-Climbing	21188071	24247237	20972102	18747809	20863734	24113105	20828779	18140447
Lagrangian Relaxation	21016929	24246000	20930767	18204403	20863734	24121146	20829199	18140447
State Space	512×10^6	84×10^6	159×10^6	115×10^6	2	18	8	1

- [5] Z. Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs*. Springer-Verlag, 1994.
- [6] 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.
- [7] 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.
- [8] S. Takriti, J. R. Birge, and E. Long. Intelligent unified control of unit commitment and generation allocation. Technical Report 94-26, Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, 1994.




Samer Takriti is a Visiting Research Scientist and an Adjunct Assistant Professor in the Department of Industrial and Operations Engineering at the University of Michigan. He holds a Ph.D. degree and a Master's degree in Industrial and Operations Engineering, and a Master's degree in Civil and Environmental Engineering from the University of Michigan. He is interested in large-scale systems and applied operations research.



finance, transportation, and public policy.

John R. Birge is Professor and Chair of Industrial and Operations Engineering at the University of Michigan, where he has been since 1980. He received Master's and Ph.D. degrees from Stanford University in Operations Research. His undergraduate degree is in Mathematics from Princeton University. His research concentrates on optimization of stochastic systems. Besides power systems, his application interest includes energy and environmental modeling, production scheduling, fi-



Erik Long is a Master's Candidate in the Department of Industrial and Operations Engineering at the University of Michigan. He holds a B.S. degree in Industrial and Operations Engineering and a B.A. degree in Economics from the University of Michigan. His research interests are in electric power systems and financial modeling.