A Mixed Integer Programming Solution for Market Clearing and Reliability Analysis

Dan Streiffert, Russ Philbrick, and Andrew Ott

***Abstract*— This paper presents a description of a Mixed Integer Programming (MIP) solution for solving the PJM Unit Commitment problem. Included is a description of the Day-ahead market clearing problem and the Reliability Analysis problem. This is followed by an overview of the MIP development process and some selected comparisons with our previously existing Lagrangian Relaxation (LR) algorithm. The paper describes many of the inherent problems associated with MIP solutions and illustrates how these issues were dealt with to provide a fast, accurate, and robust MIP solution.**

***Index Terms*—Electricity market, Power Generation Dispatch, Mixed Integer Programming, Unit Commitment**

I. Introduction

AREVA's T&D division has developed a Mixed Integer Programming (MIP) solution to the Unit Commitment problem for use in Day-ahead market clearing and Reliability Analysis studies at large Regional Transmission Organizations (RTOs). This implementation has been tested and refined for over a year and was put into production use at PJM in August, 2004 [1].

The benefits of the MIP formulation compared to the Lagrangian Relaxation (LR) include: 1) Global optimality. 2) a more accurate measure of optimality, 3) improved modeling of security constraints, and 4) enhanced modeling capabilities and adaptability.

Another major benefit in using a MIP formulation is that the developer's focus is on problem definition (i.e. codification of the requirements) rather than algorithmic development. Addition of new constraints and variables does not require continual enhancement of complex scheduling methods (e.g. addition of new LR multipliers). Even relatively simple ideas such as unit ramping constraints can result in extensive research and algorithmic development in case of the LR [2], while the corresponding MIP formulation is straightforward.

There are many issues associated with the MIP that, if not dealt with carefully, can lead to a poor implementation. These include: increased memory requirements, large variations in run times, and complex constraint formulations. This paper describes how we have dealt with these issues to provide a fast, accurate, and robust MIP solution.

The following sections describe the motivation for the MIP development, a brief description of the RTO Unit Commitment problem, an overview of the MIP method and some selected comparisons with our Lagrangian Relaxation (LR) algorithm. These are followed by a description of some newly developed MIP-Based models (e.g. combined cycle modeling).

It is not the intent of this paper to provide detailed descriptions of LR [3] and MIP algorithms as these are well documented elsewhere.

II. Motivation for MIP Development

As the operator of the world's largest wholesale market for electricity, PJM must ensure that market-priced electricity flows reliably, securely and cost-effectively from more than 1100 Generating resources to serve a peak load in excess of 100,000 MW. In doing so, PJM must balance the market's needs with thousands of reliability-based constraints and conditions before it can schedule and commit units to generate power the next day. The PJM market design is based on the Two Settlement concept [4]. The Two-Settlement System provides a Day-ahead forward market and a real-time balancing market for use by PJM market participants to schedule energy purchases, energy sales and bilateral contracts. Unit commitment software is used to perform optimal resource scheduling in both the Day-ahead market and in the subsequent Reliability Analysis.

As the market was projected to more than double its original size, PJM identified the need to develop a more robust approach for solving the unit commitment problem. The LR algorithm was adequate for the original market size, but as the market size increased, PJM desired an approach that had more flexibility in modeling transmission constraints. In addition, PJM has seen an increasing need to model Combined-cycle plant operation more accurately. While these enhancements present a challenge to the LR formulation, the use of a MIP formulation provides much more flexibility. For these reasons, PJM began discussion with its software vendors, in late 2002, concerning the need to develop a production grade MIP-based approach for large-scale unit commitment problems.

Dan Streiffert and Russ Philbrick are with AREVA T&D Corporation, 11120 NE 33rd Place, Bellevue, WA 98004 USA (e-mail: dan.streiffert@areva-td.com and russ.philbrick@areva-td.com)

Andrew Ott is with PJM Interconnection LLC, Norristown, PA 19403-2497 USA.

The partnership with AREVA T&D in the development of the production grade MIP-based unit commitment has been very successful from PJM's perspective. The resulting MIP-based approach has given PJM more flexibility in deploying enhancements in the Day-ahead market analysis and in the daily reliability analysis that is performed to ensure that adequate resources are committed to reliably meet demand. In practice, unit commitment in the PJM Day-ahead energy market often solves faster than the LR-based approach.

Since the MIP-based commitment provides the ability to more accurately incorporate the transmission system in the unit commitment problem, it provides the opportunity to develop enhanced system operator tools to reduce the cost of maintaining grid reliability. In the future, PJM expects to employ the MIP-based unit commitment approach in the real-time energy market to manage the optimal deployment of Combustion Turbines.

III. RTO Unit Commitment Problem

The Day-ahead market clearing problem includes next-day generation offers, demand bids, virtual bids and offers, and bilateral transactions schedules. The objective of the problem is to minimize costs subject to system constraints. The Day-ahead market is a financial market that provides participants an operating plan with known compensation: If their generation (or load) is the same in the real-time market, their revenue (or cost) is the same. Compensation for any real-time deviations is based on real-time prices, providing participants with opportunities to improve profit (or reduce cost) if they have flexibility to adjust their schedules.

The Reliability Analysis problem includes the same data as provided to the Day-ahead market, with modification as permitted during the re-bidding period after the Day-ahead market closes (for PJM, re-bidding is permitted only for generators not committed in the Day-ahead market). The goal of the Reliability Analysis problem is to purchase sufficient and appropriate additional capacity to ensure reliable operations, regardless of the capacity cleared in the Day-ahead financial market. When PJM calls on additional capacity, it provides a guarantee (known as the "make-whole payment") that the affected generators will receive additional payments if real-time prices are insufficient to cover their bid-in costs. Consequently, a secondary goal of the Reliability Analysis problem is to minimize the cost of potential make-whole payments by minimizing the cost of capacity (versus the cost of energy) committed by PJM.

In both problems, unit commitment accepts data that define bids (e.g., generator constraints, generator costs, and costs for other resources) and the physical system (e.g., load forecast, reserve requirements, security constraints). In real-time, the limited responsiveness of units and additional physical data (e.g., state estimator solution, net-interchange forecast) further constrains the unit commitment problem.

In all problems, resources are selected based on minimum cost based on bid prices and on physical deliverability by the transmission system.

IV. Problem Formulation

The following describes the Day Ahead Unit Commitment Problem:

Minimize Objective Cost[1] =
- Unit Energy Cost + Startup Cost
- \+ Transaction Cost + Virtual Bid Cost
- \+ Demand Bid Cost + Wheeling Cost

Subject to:
- Area Constraints:
  - Demand + Net Interchange
  - Spinning and Operating Reserves
- Zonal Constraints:
  - Spinning and Operating Reserves
- Security Constraints
- Unit Constraints:
  - Minimum and Maximum Generation limits
  - Reserve limits
  - Minimum Up/Down times
  - Hours up/down at start of study
  - Must run schedules
  - Pre-scheduled generation schedules
  - Ramp Rates
  - Hot, Intermediate, & Cold startup costs
  - Maximum starts per day and per week
  - Maximum Energy per day and per study length

A. *Security Constraints*

Two significant requirements for this project were: 1) To include the impact of all significant security constraints within the commitment process; and 2) To accurately model the impact of congestion on the Unit Commitment (UC).

In the production implementation, security constraints are modeled within the Unit Commitment program using a set of linearized sensitivities (distribution factors) to represent the change in line flow with respect to injection and Phase Angle Regulator. There are three sources for these data: 1) Manual entries; 2) the system Topology Processor; and 3) Contingency Analysis (also called Simultaneous Feasibility Test or SFT).

Manual entries and the Topology Processor are used to define an initial set of constraints based on operator experience and history. Additional constraints are identified by iterating between the UC and SFT applications. Each SFT is a single period (hour) analysis over multiple contingencies. This means the SFT analysis can be decomposed both by hour and contingency as shown in Figure 1. The implementation is flexible in that any number of CPUs can be allocated to the

[1] Revenue from transaction sales, virtual bids and demand bids are added as negative costs so that by minimizing the objective cost the profit is maximized. For Day Ahead studies, this results in a large negative objective cost.

SFT analysis. SFT jobs are put in a system queue which allocates them to specific processors as they become available.

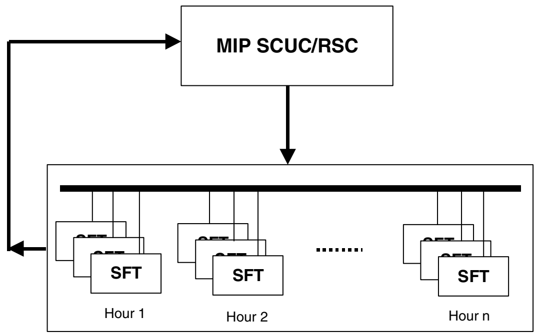

Figure 1. Iteration with Parallel SFT Applications

Each SFT solution may result in the identification of new active constraints, which are then introduced into the subsequent UC execution. This process generally converges within 2-3 iterations.

V. The Unit Commitment Problem

The Unit Commitment problem is a large-scale non-linear mixed integer programming problem. Integer variables are required for modeling: 1) Generator hourly On/Off-line status, 2) generator Startups/Shutdowns, 3) conditional startup costs (hot, intermediate & cold). Due to the large number of integer variables in this problem, it has long been viewed as an intractable optimization problem. Most existing solution methods make use of simplifying assumptions to reduce the dimensionality of the problem and the number of combinations that need to be evaluated. Examples include priority-based methods, decomposition schemes (LR) and stochastic (genetic) methods. While many of these schemes have worked well in the past, there is an increasing need to solve larger (RTO-size) problems with more complex (e.g. security) constraints, to a greater degree of accuracy.

Over the last several years, the number of units being scheduled by RTOs has increased dramatically. PJM started with about 500 units a few years ago, and is now clearing over 1100 each day. MISO cases will be larger still.

As the RTO markets have increased in size, there has been an increasing requirement for more accurate and complete models of security constraints. This requires iterating between the unit commitment software and network security analysis applications, resulting in the need to solve the UC problem multiple times to obtain good results. Security constraints are particularly difficult to model with LR methods due to the large number of multipliers that must be added to the problem [5]. As a result, finding optimal security constrained solutions with the LR is a very demanding problem for large RTO-sized problems. On the other hand, these constraints pose little problem for the MIP and may actually speed up the convergence process by reducing the size of the feasible region.

Market clearing applications have a greater need for transparency and true optimality than traditional UC implementations can provide. Each commitment result is subject to review by the participants for correctness. Solutions that are even 1% away from optimal will not stand up to this type of analysis. Ideally the problem should be solved to optimality. In practice, solutions within 0.1% of optimality appear to be acceptable for Day-ahead and 0.5% for Reliability Analysis. While this level of accuracy has been obtained in the past using the LR (See Section IX.), maintaining this level has been programmatically and computationally difficult with the increasing number of security constraints.

VI. Recent Advances In Mixed Integer Programming

The classical MIP implementation utilizes a Branch and Bound scheme. This method attempts to perform an implicit enumeration of all combinations of integer variables to locate the optimal solution. In theory, the MIP is the only method that can make this claim. It can, in fact, solve non-convex problems with multiple local minima.

Since the MIP methods utilize multiple Linear Programming (LP) executions, they have benefited from recent advances in both computer hardware and software [6].

Bixby [7] describes several significant advances in MIP algorithms, including Heuristics, Node Presolve, and Cutting Planes, that were introduced into CPLEX™ 6.5 [8]. These methods intelligently introduce numerous redundant constraints (cutting planes [9]) into the problem with the goal of finding solutions that are integer feasible at the LP corner points. This idea attempts to transform the combinatorial problem into a series of constrained LPs. When successfully implemented, this can lead to dramatic reductions in solution times for MIP problems.

Johnson [10] provides examples applying the MIP to a small 17 unit problem using CPLEX versions 3.0 through 6.5. Table I shows the results of those studies for the same one day problem[2].

Table I.
One Day Unit Commitment using MIP.

CPLEX Version	Seconds	Nodes
3.0	1687	15637
4.0	2285	19789
5.0	1700	18488
6.0	1253	22258
6.5	98	281

The following sections will illustrate the use of more recent versions of CPLEX on RTO-size UC problems.

[2] This problem was solved on a 400 MHz Pentium™ II PC.

A. Solving a RTO Day-ahead Problem

This section presents results from using the CPLEX 7.1 and CPLEX 9.0 MIP solvers on a large-scale RTO Day Ahead Unit Commitment problem. This problem has 593 units and a 48 hour time horizon.

Figure 2 illustrates the results from a Day-ahead[3] study solved[4] using CPLEX 7.1. Included are the objective cost of the current integer solution ("Best Integer"), the current lower bound ("Best LP"), and the CPU time. The MIP Gap = ABS[(Best Integer-Best LP)/Best LP] * 100 (%) provides a good measure of optimality that can be used as a convergence criteria. If the MIP Gap goes to zero, then the solution represents the true global minima.

The MIP solution process begins by solving the "relaxed" problem in which all integer variables are temporarily allowed to have continuous values. This takes 134 seconds and establishes an initial lower bound on the objective cost. The solver then proceeds to the cutting plane logic. This process introduces several hundred cutting planes, but fails to find an integer feasible solution. The solver then enters the Branch and Bound phase where it produces the results shown in Figure 2. The solver does *not* find a solution within the convergence criteria of 0.1%, and after reaching a specified time limit of 1800 seconds, the solver terminates with a best integer solution of -$260,251,896[5].

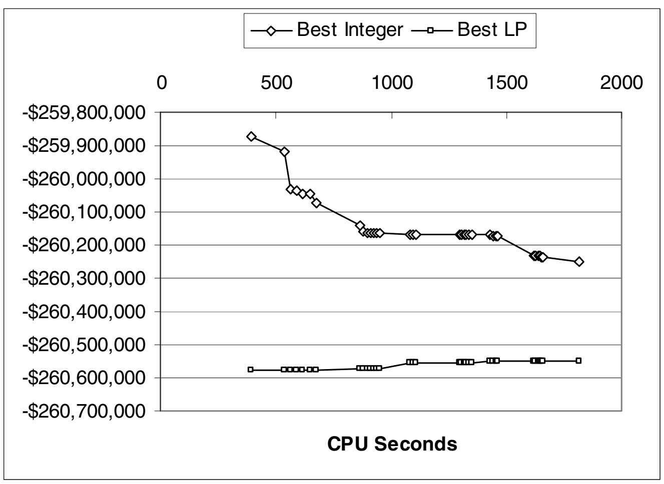

Figure 2. Day-ahead Solutions vs. CPU Time with CPLEX 7.1

The long solution times and slow convergence characteristics of earlier MIP solvers is why many experienced people question the ability of MIP algorithms to solve these "NP Hard"[6] problems. Actually, these results are quite remarkable when compared with earlier CPLEX versions, which were unable to solve this problem at all.

Figure 3 shows the results from solving the same problem using CPLEX 9.0. Notice that in this case the first integer feasible solution with CPLEX 9.0 is better than all but the last solution using CPLEX 7.1. Also, CPLEX 9.0 never enters the Branch and Bound phase, and converges to the specified convergence tolerance of 0.1 % in 386 seconds with an objective cost of -$260,351,581.

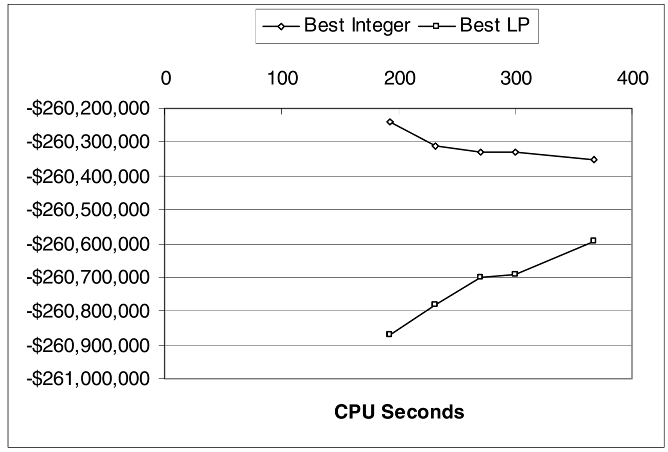

Figure 3. Day-ahead Solutions vs. CPU Time with CPLEX 9.0

The conclusion reached from this type of analysis was that one would want to avoid entering the Branch and Bound phase if at all possible. The development efforts were focused on supplementing the cutting plane methodology where possible. This was accomplished using alternative constraint formulations, making them as tight as possible, and through additional redundant constraints that would (hopefully) aid in driving the corner points towards integer feasible solutions. This effort has been so successful that virtually all of the current problems solve to within reasonable convergence tolerances before entering the Branch and Bound phase.

B. Use of Parallel MIP Algorithms

It should be noted that Branch and Bound algorithms can take advantages of parallel processing [11]. Since our implementation generally converges before utilizing the Branch & Bound step, we do not see any benefit to using parallel MIP solvers and have not pursued the use of this feature.

VII. A Hybrid Lagrangian Relaxation Method

AREVA's LR algorithm is a hybrid approach that makes use of the fundamental LR concepts, but avoids many of the convergence problems associated with these methods. The approach utilizes a procedure described by Zhuang [12] that directly computes the exact commitment cost for each candidate (uncommitted) unit, in each hour of deficit reserve. These costs are determined through multiple executions of a Single Unit Dynamic Programming (SUDP) module. For each hour of deficit reserve, each candidate unit is

[3] This problem contains 615,792 constraints, 637,943 variables, and 3,222,740 non-zeros.

[4] Note that all RTO case studies in this paper were run on an Intel® Xeon™ 3.06 GHz. Processor with 3.5 GB of RAM

[5] The objective cost includes revenues from sales, which are modeled as negative costs (i.e. a profit). This is equivalent to maximizing the profit.

[6] NP-hard – non-deterministic Polynomial-time hard.

provisionally forced on (must run) and a new SUDP is performed to obtain a new profit based on committing the unit. The commitment price of each candidate unit can then be computed as the difference between the new and old profit divided by the difference in available reserves. Zhuang utilized this commitment price as a way to systematically adjust the commitment multipliers in increments that are precisely based on the cost to commit the next unit.

Unfortunately, this does not deal explicitly with the problem of identical units. We have improved on this by using the unit commitment price to sequentially (in order of increasing price) commit candidate units in each deficit hour until a reserve-feasible solution is obtained. This is similar to the "Sequential Bidding" method developed by Prof. Fred Lee [13]. The key difference is that Lee's method utilizes a single priority order for all hours in the study whereas ours can theoretically utilize different orders in each hour.

The process is computationally intensive in that numerous SUDP solutions must be performed to compute these commitment prices. This time is partially offset by eliminating the need to iterate (search) for correct values of these multipliers.

Once a reserve-feasible commitment schedule is found, a Linear Programming (LP) model is used to solve the Economic Dispatch (ED) problem. The ED is solved for all time periods simultaneously, allowing for inclusion of both security and temporal constraints directly into the formulation. While the LP formulation is computationally intensive, it eliminates many algorithmic difficulties and provides shadow prices for all active constraints. These shadow prices are utilized to drive subsequent iterations in a direction that will tend to alleviate system constraints and minimize the objective cost.

One significant disadvantage of this hybrid approach is that a dual cost or duality gap cannot be computed.

VIII. The Transition from LR to MIP

Since we were already solving the ED using an LP formulation, the extension of this to a MIP formulation consisted of formulating only the unit constraints enforced in the SUDP as MIP constraints. Most of the remaining system constraints required little or no modification. This resulted in a single application in which one can run and compare the two different algorithms and be reasonably confident that identical problems are being solved.

Our original LR was implemented using the modeling system AIMMS [14]. This significantly expedited the conversion process to a MIP formulation, and greatly improved our communication process with a number of MIP consultants that were employed in this project.

A. Alternative Constraint Formulations

The majority of the work involved in this project was testing alternative constraint formulations. The most difficult constraints to deal with in this context are conditional constraints. For example, in the PJM model, each unit may have a hot, intermediate and cold startup cost, depending on how many hours the unit has been shut down. The classical way of dealing with conditional constraints is to introduce additional integer variables [15], one for each modeled condition. Consequently, we have three additional integer variables for each unit in each hour. The addition of these variables can have dramatic impacts on the MIP solution times, and in many cases can result in problems that will not solve in reasonable times.

In some cases, one can deal with conditional constraints by simply eliminating them. For example, the LR formulation did not enforce unit ramping constraints except as an upper bound constraint when units are initially started up. To include this same constraint in the MIP formulation required additional integer variables to determine which hours this constraint should be enforced. This problem was finally eliminated by simply enforcing the ramping constraint in all hours in the MIP formulation. This resulted in an improved ramping model that actually solved in less time than one without ramp rates. In fact, we observed several times that inclusion of additional constraints can improve MIP performance.

B. Improved MIP Solvers

We chose to utilize the commercially available MIP CPLEX solver linked with the AIMMS system. This technology allowed us to upgrade to newer versions of this solver without any significant additional time investment. With this approach we were able to transition easily from CPLEX version 7.0, to 8.1, and 9.0. Each of these releases resulted in major performance improvements.

IX. LR vs. MIP Comparisons

In this section we present some comparisons between our Lagrangian Relaxation (LR) algorithm and the MIP.

A. LR vs. MIP for a Day-ahead Market Case

This section presents results from a Day-ahead market case consisting of 885 units over a 48 hour study horizon. The LR was run for 20 iterations (until no further improvement could be obtained), and the MIP[7] was run with a zero convergence tolerance and a time limit of 600 seconds.

Figure 4 shows the integer solutions for the LR and MIP runs. The LR found its best solution (-$209,477,011) in 15 iterations. Additional iterations[8] resulted in no improvement in this cost.

The MIP terminated with a solution of -$209,432,693, which had a MIP Gap of 0.02%.

The difference between the best LR and the best MIP solution is $5,682 or 0.003 %! This is quite remarkable, and is a strong validation that both formulations are producing

[7] This Day-ahead MIP problem consists of 611,272 constraints, 647,540 variables, and 3,304,707 non-zeros.

[8] The LR sequence of solutions is not monotonic decreasing as. When an increase in objective cost occurs, it indicates we have over-stepped a minima. The step size is then reduced for subsequent iterations.

excellent results for this problem with similar CPU time.

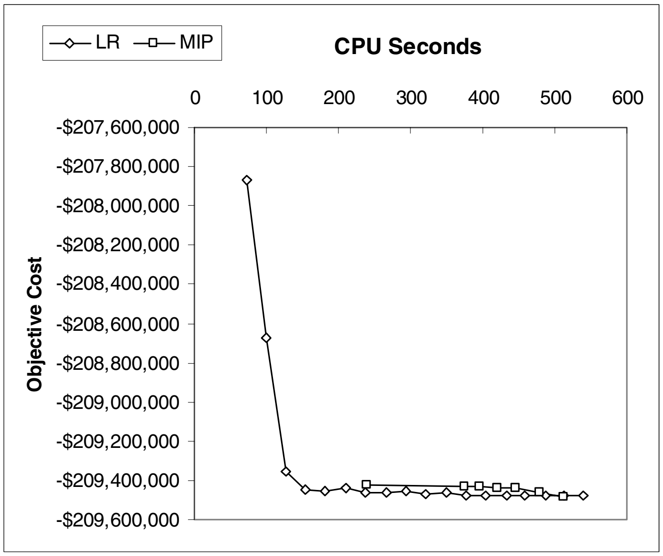

Figure 4. LR and MIP Solutions for a Day-ahead Case

The MIP takes longer than the LR to find an initial feasible solution, but the initial MIP solution tends to be closer to optimal than the initial LR solutions. Results viewed through a case comparison function show that the two solutions are similar in terms of total number of units committed, constraint prices, etc.

B. LR vs. MIP for a Reliability Analysis Case

This section illustrates how the LR can get caught in local minimums. Results are presented from a Reliability Analysis case consisting of 593 units over a seven day time horizon. The LR was run for 20 iterations and the MIP[9] was run to a convergence tolerance of zero and a time limit of 1800 seconds.

Figure 5 shows the integer solutions for the LR and MIP runs. The first 3 iterations of the LR were infeasible due to violations of the security constraints. The LR found its best feasible solution in 10 iterations of $118,979,243. Additional iterations resulted in no improvement in this cost.

The MIP case took 827 seconds to solve the initial relaxed solution. Shortly thereafter an initial feasible solution was found with a MIP gap of 0.58%. The third feasible solution was $99,909,400 with a MIP gap of 0.44%. The MIP terminated after the time limit of 1800 seconds with no additional integer feasible solutions. The lower bound on the LP solution was raised so that the final MIP Gap was 0.34%.

Typically, the LR and MIP solutions are much closer for Reliability Analysis runs than for this case. One could easily argue that a more sophisticated LR could avoid this particular problem. Guan [16] provides an analysis of these problems and presents a penalty-based method that may produce better schedules.

[9] This Reliability Analysis MIP problem consists of 1,401,056 constraints, 1,337,911 variables and 7,076,702 non-zeros.

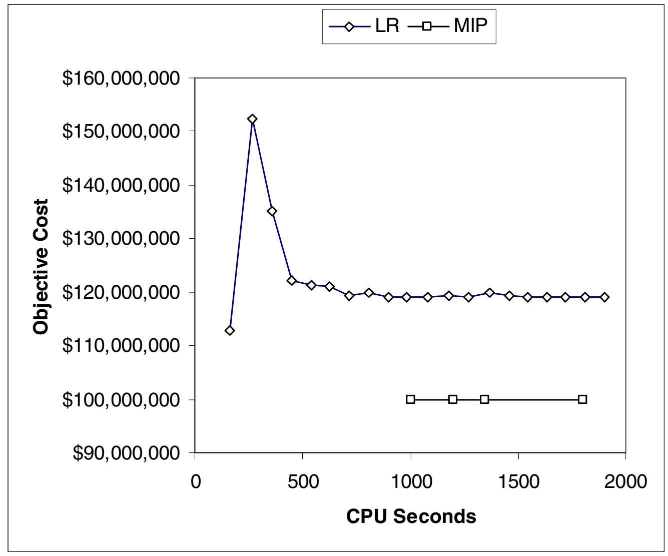

Figure 5. LR and MIP Solutions for Reliability Analysis Case

C. Solution Characteristics of LR vs. MIP

For large problems, the MIP solver often takes much longer to find an initial feasible solution than the LR. This is primarily related to the solution time for the initial relaxed LP problem, which is much larger than any of the LPs solved in the LR. For smaller problems, the MIP can be much faster than the LR.

LR memory requirements and run times grow linearly with respect to problem size, whereas the MIP growth is more exponential. This is illustrated in Figure 6.

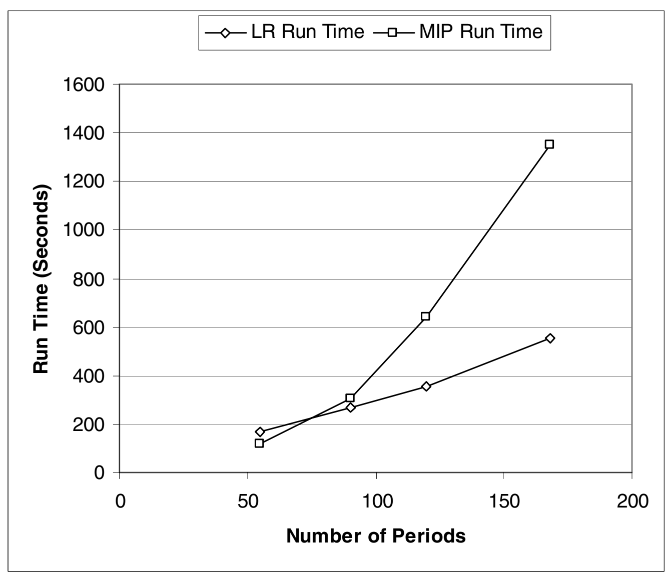

Figure 6. Run Time vs. Number of Periods in Study

The MIP deals directly with security constraints, and is more likely to find good feasible results when the system is highly constrained. Also, the MIP Gap provides a direct measure of optimality. The solver can guarantee that the current best solution is within the resulting MIP Gap of the global optimum.

D. Day -ahead vs Reliability Analysis

The Day-ahead market problem tends to be much easier to solve than the Reliability Analysis. One reason for this is the large number of dispatchable demands and virtual bids available in the Day-ahead market problem. This results in a very large set of integer feasible solutions. In the Reliability Analysis problem, these dispatchable demands are replaced with a fixed load forecast and the set of integer feasible solutions is much smaller. When solving the RA problems, the LR must determine many more multipliers than is required with the Day-ahead runs to find a feasible solution.

Another significant consideration for solving the Reliability Analysis problem is that the study horizon is typically 3-7 days. Memory requirements for the MIP tend to grow exponentially with the number of study periods. This is perhaps the single most significant barrier to solving the larger MIP problems. Seven day studies for PJM are very close to exceeding the limits of 32 bit systems. Moving to 64 bit hardware can alleviate the memory problem, but not necessarily the run-time problem.

X. Variable Time Step Modeling

The main barrier to solving large MIP problems relates to the exponential growth in solution times and memory requirements. One relatively straight-forward way of dealing with this problem is to reduce the number of modeled time steps. For example, if we are running a 7 day study, we could model the first and second days with hourly intervals, the 3rd and 4th days with two hour intervals, and the remaining days at 4 hour intervals (Daily Period Length = 1,1,2,2,4,4,4). This achieves a reduction in the study hours from 168 (at hourly intervals) to 24 + 24 + 12 + 12 + 6 + 6 + 6 = 90 with the above configuration. To obtain the correct energy usage, the period length is used when computing energy usage (e.g. Mwh = Mw * Period Length).

The impact of several combinations of Daily Period Lengths on run times and objective cost on a seven day study using a convergence tolerance of 0.5 % is illustrated in Table II. Each solution, except the last, was within the convergence tolerance relative to the 168 period solution.

With this approach, significant reductions in run times (and memory requirements) can be achieved while preserving the optimality of the solution in terms of objective cost. This indicates one can achieve very good solutions that recognize major constraints (including temporal constraints) with the reduced problem sizes. Through this type of analysis one can determine how much of a reduction can be achieved while still preserving optimality. Other impacts of using variable time steps can be viewed through a case comparison function.

Table II
MIP Solutions with Variable Step Sizes

Daily Period Length	Periods	Run Time (Sec.)	Objective Cost	% Change
1,1,1,1,1,1,1	168	1348	$99,909,400	
1,1,1,2,2,2,2	120	641	$99,952,757	0.04%
1,1,2,2,4,4,4	90	307	$99,938,465	0.03%
1,2,4,6,8,8,8	55	120	$99,115,120	-0.80%

This scheme focuses the computational effort on the periods of greatest interest (and accuracy), while still taking into account the effects of long minimum up and down time constraints on future days and any multi-hour energy constraints in the problem. When the primary concern is with the commitments made in the first day or two of a study horizon, then very good results can be achieved with a dramatic reduction in run times by modeling future days at a lower level of resolution.

XI. Enhanced Modeling with the MIP

Use of the MIP formulation to solve the Unit Commitment problem opens up many opportunities to deal directly with a number of constraints and models that tend to be very difficult to implement with the LR formulation. These include modeling of combined cycle plants, hydro unit commitment, forbidden zones, multi-area and zonal constraints, ancillary service markets, and many more.

A. Combined Cycle Plants

Scheduling of combined cycle units violates the LR assumption that the problem may be decomposed by unit. This has resulted in numerous complex algorithms for dealing with these units. Shahidehpour [17] presents an excellent scheme that utilizes a Dynamic Programming algorithm to evaluate pre-defined state-transitions for the combined cycle plant.

The MIP method eliminates the need for these schemes, and it is possible to develop relatively complex models by adding a few additional constraints to the model. The MIP can implicitly evaluate the commitments of all permissible combinations of units at the plant. Transitions between commitments can include startup costs and all other unit constraints. AREVA T&D is currently prototyping a MIP-based combined cycle model that is expected to go into production in the near future.

B. Hydro Scheduling and Unit Commitment

Hydro unit commitment violates the LR decomposition assumption when scheduling cascading river systems, where the availability of water at a particular plant and time period is dependent on the upstream and downstream unit commitment decisions. This is extremely difficult to deal with using LR approaches, but relatively straight forward using the MIP. AREVA T&D has successfully developed MIP models for Pumped Storage plants and cascading river systems.

At this time, all hydro is pre-scheduled for the PJM system, but the ability to integrate hydro scheduling into the Day Ahead market can improve scheduling, to the benefit of both PJM and the resource owners.

C. Ancillary Service Markets

As markets become more successful, there will be opportunities to expand these to include additional market products, such as ancillary services. While the LR formulation can theoretically be expanded for these additional products, there is considerable development risk and uncertainty to both the RTO and the associated vendor. The MIP formulation considerably reduces this risk since it virtually eliminates any algorithmic changes, and simply requires the addition of new constraints and variables (i.e. codification of requirements). These changes can be rapidly prototyped and validated using the MIP formulation.

XII. Conclusions

While LR and other methods can produce very accurate results, they rely on continual algorithmic development efforts that often involve simplifying assumptions and heuristic procedures. Recent advances in MIP algorithms, particularly in the area of cutting plane methods, have made this a viable alternative for large RTO-size Unit Commitment problems. The success of this effort opens many new opportunities for more sophisticated modeling (combined cycle units, hydro unit commitment, ancillary services, etc.) with a significant reduction in project risk.

We believe that these benefits, combined with expected future advances in MIP solvers, will result in an industry-wide move towards this methodology. This can potentially lead to significant improvements in system reliability, more transparent solutions, improved modeling capabilities and reduced cost to the consumers.

XIII. Acknowledgment

The success of this project resulted from valuable help we received from Johannes Bisschop, Marcel Hunting and Bianca Makkink (Paragon Decision Technology B.V., The Netherlands), and Laurence Wolsey (CORE, Belgium).

XIV. References

[1] "PJM'S NEW GENERATION-SCHEDULING SOFTWARE TO SAVE CUSTOMERS ESTIMATED $56 MILLION *Unit Commitment System Uses First of its Kind Technique",* PJM News Release, June 24, 2004

[2] Wei Fan, Xiaohong Guan, Qiaozhu Zhai, "A New Method for Unit Commitment with Ramping Constraints", *Electric Power Systems Research*, Vol. 63, No. 3, 2002, pp. 215-224.

[3] S. Virmani, E. Adrian, E. Lo, "Unit Commitment – Lagrangian Relaxation", in *Application of Optimization Methods for Economy/Security Functions in Power System Operations*, IEEE Tutorial Course 90EH0328-5-PWR, pp 37-46.

[4] Andrew L. Ott, "Experience with PJM Market Operation, System Design, and Implementation", *IEEE Transactions on Power Systems*, May, 2003, Digital Object Identifier 10.1109/TPWRS.2003.810698

[5] J. Shaw, "A Direct Method for Security-Constrained Unit Commitment", *IEEE/PES 1994 Summer Meeting*, 94 SM 591-8 PWRS

[6] Robert Bixby, "Solving Real-World Linear Programs: A Decade and More of Progress", *Operations Research*, Vol 50, Issue 1, Jan. – Feb. 2002

[7] Robert E. Bixby, Mary Fenolon, Zonghao Gu, Ed Rothberg, Roland Wunderling, "MIP: Theory and Practice – Closing the Gap", Publication of ILOG CPLEX Division, Available at http://www.ilog.com

[8] CPLEX, A product of ILOG, Inc.

[9] Laurence A. Wolsey, *Integer Programming*, , John Wiley & Sons, 1998

[10] R.B. Johnson, S.S. Oren, and A.J. Svoboda, "Equity and efficiency of unit commitment in competitive electric markets", Utilities Policy, 6(1): 9-20, 1997

[11] ILOG Parallel CPLEX, A product of ILOG, INC.

[12] F. Zhuang, F. Galiana, "Towards a More Rigorous and Practical Unit Commitment by Lagrangian Relaxation", *IEEE Transactions on Power Systems*, Vol. 3, No. 2, May 1988.

[13] F. Lee, "Thermal Unit Commitment by Sequential Methods", in *Application of Optimization Methods for Economy/Security Functions in Power System Operations*, IEEE Tutorial Course 90EH0328-5-PWR, pp 47-54.

[14] AIMMS Modeling Language, Paragon Decision Technology B. V., http://www.aimms.com

[15] Hillier and Lieberman, *Introduction to Operations Research*, Third Edition, Holden-Day, Inc.1980, p 732.

[16] X. H. Guan, Q.Z. Zhai, F. Lai, "New Lagrangian Relaxation Algorithm for Resource Scheduling with Homogeneous Subproblems", *Journal of Optimization Theory and Applications*: Vol. 113, No. 1. pp. 65-82, April 2002.

[17] Mohammad Shahidehpour, "Short-Term Scheduling of Combined Cycle Units, *IEEE Transactions on Power Systems*, August, 2004, ISSN 0885-8950.

XV. Biographies

Dan Streiffert joined AREVA T&D in December 1987. He received his M.S. in Applied Mathematics and B.S. in Geophysical Engineering from Michigan Technological University. He is currently principal developer of AREVA's Unit Commitment application.

Russ Philbrick joined AREVA T&D in July 1999. He received his M.S. and PhD from Stanford following service as a Nuclear Submarine Officer. He currently is Product Manager and a developer of AREVA's Unit Commitment and Market Monitoring applications. Prior to joining AREVA T&D, Dr. Philbrick was on the facility at Washington State University.

Andrew L. Ott joined PJM in October 1996. He received his B.S. in Electrical Engineering from Penn State University and his M.S. in Applied Statistics from Villanova University. He is currently Vice President of Market Services for PJM and is responsible for the oversight of PJM Market Operations and Market Development. He was responsible for implementation of the current PJM Locational Marginal Pricing system, the PJM Financial Transmission Rights Auction and the PJM Day-ahead Energy Market. Prior to joining PJM, Mr. Ott was employed by GPU for 13 years in transmission planning and operations.