ISTITUTO DI ANALISI DEI SISTEMI ED INFORMATICA
"Antonio Ruberti"
CONSIGLIO NAZIONALE DELLE RICERCHE

A. Frangioni, C. Gentile, F. Lacalandra

HYBRID LAGRANGIAN-MILP APPROACHES FOR UNIT COMMITMENT PROBLEMS

R. 668 Settembre 2007

Antonio Frangioni – Dipartimento di Informatica, Polo Universitario della Spezia, Via dei Colli 90, 19121 La Spezia – Italy. Email: frangio@di.unipi.it.

Claudio Gentile – Istituto di Analisi dei Sistemi ed Informatica, CNR, Viale Manzoni 30 - 00185 Roma, Italy. Email: gentile@iasi.cnr.it..

Fabrizio Lacalandra – $\mathbb{O}$*pti*$\mathbb{ME}$, c/o MBI, via Scornigiana 7, 56121 Loc. Ospedaletto - Pisa. E-mail: fabrizio.lacalandra@fastwebnet.it.

ISSN: 1128–3378

Collana dei Rapporti dell'Istituto di Analisi dei Sistemi ed Informatica "Antonio Ruberti", CNR

viale Manzoni 30, 00185 ROMA, Italy

tel. ++39-06-77161
fax ++39-06-7716461
email: iasi@iasi.rm.cnr.it
URL: http://www.iasi.rm.cnr.it

Abstract

The short-term Unit Commitment (UC) problem in hydro-thermal power generation is a fundamental problem in short-term electrical generation scheduling. Historically, Lagrangian techniques have been used to tackle this large-scale, difficult Mixed-Integer NonLinear Program (MINLP); this requires being able to efficiently solve the Lagrangian subproblems, which has only recently become possible (efficiently enough) for units subject to significant ramp constraints. In the last years, alternative approaches have been devised where the nonlinearities in the problem are approximated by means of piecewise-linear functions, so that UC can be approximated by a Mixed-Integer Linear Program (MILP); in particular, using a recently developed class of *valid inequalities* for the problem, called "Perspective Cuts", significant improvements have been obtained in the efficiency and effectiveness of the solution algorithms. These two different approaches have complementary strengths; Lagrangian ones provide very good lower bounds quickly, but require sophisticated heuristics—which may need to be changed every time that the mathematical model changes—for producing actual feasible solutions. MILP approaches have been shown to be able to provide very good feasible solutions quickly, but their lower bound is significantly worse. We present a hybrid approach which combines the two methods, trying to exploit each one's strengths; we show, by means of extensive computational experiments on realistic instances, that the hybrid approach may exhibit significantly better efficiency than either of the two basic ones, depending on the degree of accuracy requested to the feasible solutions.

Key words: Hydro-Thermal Unit Commitment, Mixed-Integer NonLinear Program Formulations, Lagrangian Relaxation.

1. Introduction

The short-term Unit Commitment (UC) problem in hydro-thermal power generation systems requires to optimally operate a set of hydro—possibly cascade connected—and thermal generating units, over a given time horizon (typically one day or one week), in order to satisfy a forecasted energy demand at minimum total cost. The generating units are subject to some technical restrictions, depending on their type and characteristics; for hydro units typical constraints concern the discharge rate, spillage limits, reservoir storage and effect on downstream units. As for the thermal units, they must usually satisfy minimum up- and down-time constraints and upper and lower bounds over the produced power when the unit is operational, besides having complex power production and start-up costs. Closely representing the actual operating behavior of generating units within mathematical optimization models is crucial for being able to effectively coordinate the production of the generating system taking into account each unit's characteristics [30], which is of increasing importance in the ongoing liberalization of the electricity market in many countries [24]. Indeed, while UC, in the form treated in this paper originated from the era of monopolistic producers, it has numerous applications even in the liberalized regime; furthermore, algorithmic approaches developed for the "classical" UC can usually be easily extended to forms of the problem arising in a market environment [1, 24, 7].

Several variants of UC have been investigated during the last 30 years, and several specialized algorithmic approaches have been proposed for solving, possibly approximately, this problem; in fact, being UC a large-scale, Mixed-Integer NonLinear Program (MINLP), it is rather difficult to solve instances of realistic size within the time limits required by operational environments. Among the most efficient specialized algorithmic approaches for (UC), Lagrangian Relaxation (LR) methods [2, 4, 5, 6, 12, 27, 32] surely play a major role. These approaches exploit the *spatial structure* of the problem, that is, the fact that removing the constraints that tie the different units together one obtains a set of disjoint Single-Unit Commitment (1UC) problems, requiring to optimally operate one single (hydro or thermal) unit over the time horizon. Thus, the applicability of LR methods critically depend on being able to optimally solving the 1UC problems efficiently; in turn, this depend on the specific details of the operational constraints of the generating units that are represented in the mathematical model. When new constraints have to be included in the model, they can be relaxed in the same way as "linking" ones, still yielding a correct lower bound. If this results in too "loose" a relaxation, several approaches can be tried to include the constraints in the Lagrangian subproblem without changing too much its structure, such as by embedding it into a Simulated Annealing approach [25] or by deriving proper Benders cuts [26]. Of course, when all else fails it may be necessary to develop entirely new solution methods [17]. Also, LR approaches typically require heuristics for producing solutions which actually respect the relaxed constraints, and these heuristics may also need to be updated each time the underlying mathematical model changes, although this is not always crucial [22] and "augmented" approaches [4, 11] can be used to devise algorithms which do not require a combinatorial heuristic at all.

While LR approaches are still the method of choice for very-large-scale instances and/or when very fast running times are required by the operational environment [22], their inherent "rigidity" justifies the interest towards methods that are more resilient to (large or small) changes of the mathematical model of the generating units. In particular, in recent years several authors have reported that Mixed-Integer Linear Program (MILP) *approximated* formulations of UC, where the nonlinearities of the problem are approximated by piecewise-linear function, allow to harness the very efficient available general-purpose, off-the-shelf MILP solvers to produce very good solutions in relatively short time [8, 9, 28]. It ought to be remarked that piecewise-linearization (apparently a very crude and trivial technique) can be performed with different degrees of sophistication; for instance, in [21] it is shown that the use of a set of *valid inequalities* for the UC problem, proposed in [16] and based on a nonlinear convex approximation of the nonconvex objective function of the problem, do improve the performances of MILP-based approaches with respect to using more ordinary linearizations.

These two different approaches—LR and MILP-based ones—have complementary strengths. Being able to solve the combinatorial 1UC problem exactly, due to the sophisticated specialized algorithms available [17], and owing to the powerful bound-computing properties of Lagrangian methods [14], LR approaches provide very good lower bounds; coupled with appropriate approaches for minimizing the

(nondifferentiable) Lagrangian function [13], i.e., solving the corresponding Lagrangian dual, these good bounds can be obtained efficiently. On the other hand, by their nature LR approaches do not provide feasible solutions; even when modified for doing so, either by heuristics [22] or by other means [4, 11], they may not be able to provide solutions whose quality matches the very good accuracy of the lower bound. Conversely, even when improved by valid inequalities, either specialized for the problem [16] or general-purpose ones, the lower bound computed by MILP-based approaches is significantly worse than that of LR approaches, not only at the root node of the enumeration tree, but also after that a very significant number of nodes have been explored. On the other hand, MILP-based approaches have been shown to be able to provide very good feasible solutions quickly; often, the very first solution computed by the rounding heuristics embedded in the efficient, general-purpose, off-the-shelf MILP solvers turns out to be of very good quality when this gap is computed against the LR lower bound. Unfortunately, due to the weakness of the MILP lower bound, this is typically not recognized by the method until much later in the optimization process.

In this paper we present the results of a hybrid approach which combines the two methods, trying to exploit each one's strengths; this basically boils down to using the LR bound inside the optimization process of the MILP-based one. We show that the hybrid approach may exhibit significantly better efficiency than either of the two basic ones, depending on the degree of accuracy requested to the feasible solutions. While the hybrid approach is almost never significantly slower than the pure MILP-based one, due to the very fast computation of the LR lower bound, there are cases—typically, when the required accuracy to the final solution is either very coarse or very tight—when it does not deliver significant benefits. Whether or not this happens also depend on the size of the instance to be solved and on its type (whether it is pure thermal or hydro-thermal). By means of extensive computational experiments on realistic instances, we devise guidelines giving indications about the appropriateness of the new techniques in different situations.

The structure of the paper is the following. In Section 2 we present the formulation of the specific form of UC problem we consider; while we focus, for our results, on a quite "classical" formulation, the idea could easily be applied to a number of other UC problems, e.g. taking into account market constraints [1, 10]. In Section 3 we recall the basic ideas of the two algorithmic approaches. Then, in Section 4 we describe the hybrid approach, present the computational results aimed at evaluating its effectiveness in different conditions, and we draw some conclusions.

2. The UC model

Consider a set P of thermal units and a set H of hydro cascades, each comprising one or more basin units. We denote by $\mathcal{T} = \{1, \ldots, n\}$ the set of time periods defining the time horizon; the time period "0" will be used for indicating the initial conditions of the power system. Introducing status and power production variables of the thermal units, u_t^i and p_t^i, respectively, with $i \in P$, $t \in \mathcal{T}$, the objective function of UC, representing the total power production cost to be minimized, has the form

$$\sum_{i \in P} c^i(\mathbf{p}^i, \mathbf{u}^i) = \sum_{i \in P} \left(s^i(\mathbf{u}^i) + \sum_{t \in \mathcal{T}} \left(a_t^i (p_t^i)^2 + b_t^i p_t^i + c_t^i u_t^i \right) \right) \tag{1}$$

where $a_t^i > 0$. That is, the power production cost at each hour is customarily represented by a convex quadratic separable form in the power p_t^i variables, neglecting for instance the so called valve points [30]. Fixed production costs are represented by the term $c_t^i u_t^i$, while *start-up costs*, possibly time-dependent, are denoted by the nonseparable function $s^i(\mathbf{u}^i)$. We do not dwell further upon the specific form of the $s^i(\mathbf{u}^i)$, only assuming that it can be properly represented within a MILP problem [29, 8] and handled by the procedure used for solving the 1UC problem [17]; most forms of start-up costs arising in practical problems satisfy both assumptions.

The constraints of UC can be partitioned into three sets: local constraints for thermal units, local constraints for hydro units, and global (system wide) constraints.

- *Local constraints for thermal units*: for each thermal unit $i \in P$, let τ_+^i and τ_-^i be respectively the minimum up- and down-time requirements, Δ_+^i and Δ_-^i be respectively the maximum ramp-up

and ramp-down rates, $\bar{p}^i_{min}$ and $\bar{p}^i_{max}$ be respectively the maximum and minimum power output of unit when operating in steady state, and $\bar{l}^i$ and $\bar{u}^i$ be the maximum power that can be produced by the unit in the time period where it is committed or decommitted, respectively (this is usually less than $\bar{u}^i$, and may also be less than Δ^i_+ and/or Δ^i_-). Then, the local constraints corresponding to each unit $i \in P$ are

$$\bar{p}^i_{min} u^i_t \leq p^i_t \leq \bar{p}^i_{max} u^i_t \qquad t \in \mathcal{T} \tag{2}$$

$$p^i_t \leq p^i_{t-1} + u^i_{t-1}\Delta^i_+ + (1 - u^i_{t-1})\bar{l}^i \qquad t \in \mathcal{T} \tag{3}$$

$$p^i_{t-1} \leq p^i_t + u^i_t \Delta^i_- + (1 - u^i_t)\bar{u}^i \qquad t \in \mathcal{T} \tag{4}$$

$$u^i_t \geq u^i_r - u^i_{r-1} \qquad t \in \mathcal{T},\ r \in [t - \tau^i_+, t-1] \tag{5}$$

$$u^i_t \leq 1 - u^i_{r-1} + u^i_r \qquad t \in \mathcal{T},\ r \in [t - \tau^i_-, t-1] \tag{6}$$

$$u^i_t \in \{0, 1\} \qquad t \in \mathcal{T} \tag{7}$$

Constant τ^i_+ indicates how many further periods after a startup period unit i must remain online, in order to avoid excessive mechanical stress due to too frequent startup/shutdown procedures that would in the long term deteriorate the unit's conditions; analogously, τ^i_- indicates how many further periods after a shutdown period unit i must remain offline. Note that we assume knowledge of the *complete state* of each unit prior to the beginning of the current operation, that is, its commitment u^i_0 and its generated power p^i_0; for the sake of minimum up- and down-time constraints (5), (6), as well as for the computation of time-dependent startup costs (if any), it is also necessary to know for how long each unit has been on or off prior to time period 0.

- *Local constraints for hydro cascade units*: each cascade $h \in H$ is composed by a set $H(h)$ (possibly containing only one element) of individual hydro units; for each $j \in H(h)$, variables q^j_t, v^j_t and w^j_t represent respectively discharged water, the volume of the reservoir and the spilled water at time period $t \in \mathcal{T}$. Constants $\bar{v}^j_{min}$ and $\bar{v}^j_{max}$ represents respectively the minimum and maximum volume for the reservoir, $\bar{q}^j_{max}$ represents the technical maximum of discharged water (the technical minimum is assumed to be zero in order to avoid nonlinearities in the model), while $\bar{w}^j_t$ represents the natural inflows at time period $t \in \mathcal{T}$. Finally, let $\mathcal{S}(j)$ be the (possibly empty) set of the immediate predecessors of unit j—those whose discharge and spillage reaches j without passing through other reservoirs—and t_{kj} be the water time delay from plant $k \in \mathcal{S}(j)$ to the basin feeding plant j. Then, the local constraints corresponding to each unit $j \in H(h)$ are:

$$0 \leq q^j_t \leq \bar{q}^j_{max} \qquad t \in \mathcal{T} \tag{8}$$

$$\bar{v}^j_{min} \leq v^j_t \leq \bar{v}^j_{max} \qquad t \in \mathcal{T} \tag{9}$$

$$v^j_t - v^j_{t-1} = \bar{w}^j_t - w^j_t - q^j_t + \sum_{k \in \mathcal{S}(j)} \left(q^k_{t-t_{kj}} + w^k_{t-t_{kj}} \right) \qquad t \in \mathcal{T} \tag{10}$$

Note that, in order for the balance equations (10) to be well-defined, we assume knowledge of the volume of each reservoir at $t = 0$, as well as water discharged and spilled at all time periods prior to $t = 1$ for which the water is still arriving to one of the downstream basins (i.e., those $k \in \mathcal{S}(j)$ such that $t < t_{kj}$).

- *Global constraints*: for each time period $t \in \mathcal{T}$, let $\bar{d}_t$ be the forecasted load to be satisfied, and for each hydro unit j let α^j be the power-to-discharged-water efficiency (assumed constant to avoid nonlinearities); then, the system-wide constraints—linking the different units among themselves—are:

$$\sum_{i \in P} p^i_t + \sum_{h \in H} \sum_{j \in H(h)} \alpha^j q^j_t = \bar{d}_t \qquad t \in \mathcal{T} \tag{11}$$

While not present in the instances used in Section 4, *spinning reserve* constraints, which ensure that a fraction of the maximum power of active units is kept available to face unforeseen events in the electrical grid, can be easily added to the formulation either in the "standard" formulation (e.g. [6]) or in the more sophisticated one recently proposed in [8].

We refer to UC as the problem of minimizing (1) subject to constraints (2)—(11); this is a large-scale, Mixed-Integer Nonlinear Program, which is difficult to solve for the size required by practical applications. Although US has been historically motivated by the centralized decision environments prevalent in the past, it is still well-suited for being employed in today's free market regime, both at the stage where GenCos need to optimize their production schedule once that their own load profile has been established by the market procedures, and within approaches for computing optimal bidding strategies [1, 24, 10, 7]. For future reference, we will denote the set defined by constraints (2)—(7) for a given thermal unit $i \in P$ as $\mathcal{U}^i$, and the set defined by constraints (8)—(10) for a given hydro cascade $h \in H$ by $\mathcal{H}^h$.

It should be remarked that the above model contains several simplifying assumptions whereby several phenomenon occurring in reality are either disregarded or (most often, linearly) approximated. For instance, the power-to-discharged-water efficiency represented by the linear function of coefficient α^j in (11) is in reality a nonlinear function due to the effect of the water level in the basin on the water potential energy released in the turbines. Furthermore, hydro units typically have nonzero technical minima for discharged water (the smallest quantity of discharged water necessary to operate the turbines), analogous to the technical minimum $\bar{p}^i_{min}$ of thermal units (cf. 2)), as well as cavitation points, which make their feasible operating set nonconvex. As for thermal units, their power production cost increases dramatically in a small neighborhood of the so-called valve points [30]. Although the above simplifying assumptions are widely accepted in the literature, more sophisticated models including them are possible. They would typically require either more integer variables, or more nonlinear (and nonconvex) terms in the formulation (1)—(11), likely rendering it significantly more difficult to solve by MINLP (or MILP-based) approaches. As for LR ones, the impact of each modeling change on the 1UC subproblems to be solved would need to be assessed, and specialized solution approaches might need to be developed. While both avenues are possible, slight modifications of the mathematical models are more likely to be allowed (without too much impact on effectiveness) by a MILP-based approach, rather than by a LR one; this largely explains the practical interest on this kind approach. On the other hand, if the modeling improvements substantiate into additional linear constraints (possibly linking to additional variables), then relaxing all the added constraints in a Lagrangian fashion would leave the same 1UC problems as in the present case. Thus, the approach presented in this paper is likely to be applicable, and significant, also for more involved UC models.

3. The two algorithmic approaches

3.1. The Lagrangian Relaxation approach

The LR approach is based on dualizing the *coupling constraints* (11) via a vector of Lagrangian multiplier $\boldsymbol{\lambda} = [\lambda_t]_{t\in\mathcal{T}}$, thereby forming the *Lagrangian Relaxation* of (UC)

$$\mathcal{L}(\boldsymbol{\lambda}) = \sum_{i\in P} \phi^1_i(\boldsymbol{\lambda}) + \sum_{h\in H} \phi^2_h(\boldsymbol{\lambda}) + \sum_{t\in\mathcal{T}} \lambda_t \bar{d}_t \tag{12}$$

where

$$\begin{aligned} \phi^1_i(\boldsymbol{\lambda}) &= \min \{\, c^i(\mathbf{p}^i, \mathbf{u}^i) - \boldsymbol{\lambda}\mathbf{p}^i \;:\; (\mathbf{p}^i, \mathbf{u}^i) \in \mathcal{U}^i \,\} \\ \phi^2_h(\boldsymbol{\lambda}) &= \min \left\{ -\boldsymbol{\lambda} \textstyle\sum_{j\in H(h)} \alpha^j \mathbf{q}^j \;:\; [\mathbf{q}^j]_{j\in H(h)} \in \mathcal{H}^h \right\} \end{aligned}$$

Is is well known (e.g., [14]) that for each $\boldsymbol{\lambda} \in \mathbb{R}^n$, $\mathcal{L}(\boldsymbol{\lambda})$ is a lower bound on the optimal value of UC. Therefore, one is interested in the $\boldsymbol{\lambda}^*$ such that this lower bound is the best (maximum), i.e., in the optimal solution of the *Lagrangian Dual* of UC:

$$\max \{\, \mathcal{L}(\boldsymbol{\lambda}) \;:\; \boldsymbol{\lambda} \in \mathbb{R}^n \,\} \quad . \tag{13}$$

Since $\mathcal{L}(\cdot)$ is a convex but non differentiable function, proper algorithms must be chosen for solving (13); *bundle* methods [13], particularly in their *disaggregated* variant [2], have been repeatedly reported to be quite efficient in solving (13), much more so [6] than alternative algorithms such as subgradient methods [3, 32]. Since this approach requires the repeated solutions of 1UC problems, one for each unit and for each iteration of the bundle method, it can only be computationally effective if these problems can be solved very efficiently. For hydro units, $\mathcal{H}^h$ turns out to be a single-commodity flow problem, and therefore the corresponding 1UC can be solved by means of the very efficient available network flow techniques (e.g. [15, 18, 23] and the references therein). For thermal units, a Dynamic Programming (DP) procedure has recently been proposed in [17] which can solve 1UC problems with ramping constraints in $O(n^3)$ overall for the formulation of $\mathcal{U}^i$ used in this paper; note the DP procedure is quite general, and can work with any analytic convex function as well as under less restrictive assumptions on the behavior of the units than those in force here.

However, solving (13) is not, in general, enough to solve (UC); even for $\boldsymbol{\lambda} = \boldsymbol{\lambda}^*$, the optimal solution to (12) is not guaranteed to—and will not in general—satisfy the relaxed constraints (11). Two possible ways have been proposed for dealing with this problem:

- either combinatorial heuristics are run at every step of the iterative solution of (13), thus for several different values of $\boldsymbol{\lambda}$, that use the infeasible optimal solution of the corresponding (12) to produce a feasible solution for (UC) [2, 3, 6, 5, 12, 32, 22];
- or the Lagrangian problem is modified with further terms that try to enforce feasibility of the obtained solutions [4, 11].

Both approaches have been reported to obtain good results, especially for large-scale instances and when the maximum allotted running time is small. However, another class of approaches has recently been shown to hold promises for not-so-large scale instances and if running times are less of an issue, especially if the quality of the obtained solution is of paramount importance.

3.2. The MILP approximation approach

On the outset, MILP-based approaches are quite simple; the nonlinear part of the objective function is (piecewise) linearized, in order to make UC tractable by the efficient MILP solvers available. How the linearization is performed, however, is not a secondary issue. Since the nonlinear structure is identical for each time period and thermal unit, for notational simplicity in this paragraph we consider both indices i and t fixed and we drop them. The issue is then how to best represent the quadratic objective function

$$f(p,u) = ap^2 + bp + cu \tag{14}$$

by means of a piecewise-linear one. There are some well-known ways in which this can be done, but a more effective way has been recently proposed in [21] based on ideas developed in [16]. Because the function $f(p,u)$ in (14) is in principle only relevant at points (p,u) of its (disconnected) domain $\mathcal{D} = [0,0] \cup [\bar{p}_{min}, \bar{p}_{max}] \times \{1\}$, it makes sense to consider its *convex envelope* of $f(p,u)$ over $\mathcal{D}$, that is, the convex function with the smallest (in set–inclusion sense) epigraph containing that of f, can be shown [16] to be

$$h(p,u) = \begin{cases} 0 & \text{if } p = 0 \text{ and } u = 0, \\ ap^2/u + bp + cu & \text{if } u\bar{p}_{min} \le p \le u\bar{p}_{max}, u \in (0,1] \\ +\infty & \text{otherwise.} \end{cases} \tag{15}$$

This function is strongly related with a well-known object in convex analysis, the *perspective function* $g(p,u) = u\, f(p/u)$ of $f(p)$, whose epigraph defines a cone pointed in the origin and having as "lower shape" that of $f(p)$. Since $0 < u \le 1$, it is immediate to verify that $h(p,u) \ge f(p,u)$ for all $(p,u) \in \mathcal{D}$; therefore, h is a better objective function for a *continuous relaxation* of UC (the model obtained by replacing the integrality constraints (7) with just $u \in [0,1]$). Because $h(p,u)$ is a convex function, it is possible to solve such continuous relaxation efficiently; however, it is a "more nonlinear" function than $f(p,u)$, which we already aim at making "less nonlinear" in order to be able to employ MILP

techniques. It should be remarked at this point that a reformulation of (15) in terms of *second order cone constraints* is possible which would allow to exploit the growing availability of solvers for problems incorporating such constraints in order to tackle the problem [19]; however, for the class of problems of interest here the following alternative, first proposed in [16] and further developed in [21], turns out to be more appropriate. By defining a further variable v intended to denote the function value, (15) can be replaced by the following *infinite* set of linear inequalities:

$$v \geq (2a\bar{p} + b)p + (c - a\bar{p}^2)u \quad \bar{p} \in [\bar{p}_{min}, \bar{p}_{max}] \tag{16}$$

We refer to each inequality in (16) as a *perspective cut* (P/C); it simply defines the unique supporting hyperplane to the function passing from $(0, 0)$ and $(\bar{p}, 1)$. While it is clearly impossible to employ an exact representation of $h(p, u)$ in these terms, an *approximation* of the function can be easily defined as follows:

- each term of the form (14) is removed from the objective function and replaced with the corresponding new variable v; other terms in the objective function not containing p and u, e.g., those related to variable startup costs [29], are kept untouched;
- k points $\bar{p}^h$, $h = 1, \ldots, k$ are arbitrarily selected in the interval $[\bar{p}_{min}, \bar{p}_{max}]$, and the corresponding k constraints of (16) with $\bar{p} = \bar{p}^h$ are added to the formulation.

Starting from the *exact* MINLP formulation (1)—(11), this produces an *approximate* MILP formulation with $|P||\mathcal{T}|$ more continuous variables and $k|P||\mathcal{T}|$ more constraints. A relevant feature of this approach is that it is very easy to make it *dynamic*, using the standard mechanisms that MILP solvers make available for implementing the so-called "Branch&Cut" approaches. In fact, one can choose a small set of initial constraints, solve the continuous relaxation and, if $u^* > 0$, check whether the solution $[v^*, p^*, u^*]$ satisfies the P/C (16) for $\bar{p} = p^*/u^*$; if not, the thus obtained cut can be added to the MILP model, which is then re-solved. Thus, any required degree of approximation of the original objective function to UC can be obtained without starting with a formulation with a very large k. Furthermore, the values $\bar{p}$ which generate the constraints are automatically selected during the approach, which is arguably preferable to choosing them a-priori without any knowledge of the structure of the continuous solutions which will be generated.

4. The hybrid approach

4.1. Motivation

We now briefly recall some results obtained during previous computational experiments [22, 21], in order to motivate the introduction of the hybrid approach. We compare results obtained from the following algorithms:

- LR is the Lagrangian approach of §3.1 *without any heuristic*, i.e., which only computes a valid lower bound on the optimal value of the problem;
- LRH is the same Lagrangian approach where heuristics are ran to obtain feasible solutions (of course, this also attains the same valid lower bound as LR);
- PCFD$_k$ is the approach whereby the P/C formulation is initially constructed with only *two* P/Cs per variable, the ones corresponding with $\bar{p}_{min}$ and $\bar{p}_{max}$, and a Branch&Cut approach is initiated, using the highly regarded commercial solver `Cplex 9.1`, where additional P/Cs, up to a maximum of k (a user-configurable parameter) per variable, are dynamically generated as needed as described in the previous paragraph. As in [21], we test two values for k, namely $k = 4$ and $k = \infty$ (the latter corresponding to no preset limit to the number of P/Cs).

Note that *both* approaches are heuristic ones, since the MILP solver stops when its *perceived* gap is less than the given threshold, but that gap does not accurately measure the true one. In fact, PCFD$_k$

computes the upper bound using a lower approximation of the true cost function, and therefore the bound is not *a-priori* valid. In our results, however, all the reported gaps have been computed by re-evaluating the objective function value of the integer solution provided by the solver using the "true" quadratic objective function (1), in order to be able to compare the results of both approaches.

For our tests, we have used two sets of randomly generated realistic pure thermal and hydro-thermal instances, with number of thermal units ranging from 10 to 200 and number of hydro units ranging from 10 to 100, on a daily problem ($n = 24$). The generator produces a generating set with "small", "medium" and "large" thermal units in realistic proportions; the characteristics of each unit are then randomly generated within a set of realistic parameters, depending on the type of the unit. Ramping restrictions are also randomly generated within realistic measures, resulting in large units to require between two and three hours to ramp from the technical minimum to the technical maximum. For simplicity, all the instances have *time-invariant start-up costs*. The UC instances are freely available at

`http://www.di.unipi.it/optimize/Data`

and have already been used in [22, 21]. The tests have been performed on an Opteron 246 (2 GHz) computer with 2 GigaBytes of RAM, running Linux Fedora Core 3.

The results are displayed in Table 1 and Table 2. In the tables, column "p" reports the total number of thermal generating units, while column "h" reports the total number of hydro units. The first half of the tables, with $h = 0$, is therefore composed by "pure thermal" instances; each row reports averaged results of 5 instances of the same size. Columns "LR" and "LRH" report results for the two Lagrangian algorithms, while columns "PCFD_4" and "PCFD_∞" report results for the MILP-based ones. In all cases, column "time" reports the required running time (in seconds), column "gap" reports the obtained gap (in percentage) between the (true) objective function value of the integer feasible solution reported by the formulation and the *best* valid lower bound we know for each instance (LR does not produce any feasible solution, so this information is not reported), and column "dgap" reports the obtained gap (in percentage) between the (best) lower bound by each algorithm and the *best* valid upper bound we know for each instance. The same best upper bound and lower bounds are used for computing the gaps for both formulations, so that the gaps can be compared. Form the MILP-based approaches, column "nds" reports the total number of expored nodes in the Branch&Bound tree.

The results of the two tables differ only for one critical parameter, i.e., the stopping tolerance for `Cplex`; in Table 1 this is set to the relatively "relaxed" value of 0.5%, while in Table 2 this is left to the default value for `Cplex`, i.e. the much stricter 0.01%. Of course, this does not impact on the results of the Lagrangian approaches, which are reported twice only for reading convenience. In both cases we set a time limit of 10000 seconds to `Cplex`; however, while in the lower accuracy case this limit was never even approached, in the higher accuracy case the solver had most often to stop for having depleted the allotted time without reaching the desired accuracy threshold.

The results in Table 1 show that the approaches have complementary strengths. The Lagrangian approach is very efficient, especially if it is only required to compute the lower bound ("LR" columns); producing feasible solutions only come at a further significant expense, in relative terms, and the obtained gaps are often worse than those of the MILP approach, except for the largest instances, especially hydro-thermal ones. On the other hand, while the MILP approach often provides better-quality solutions with a very limited (if at all) recourse to enumeration (cf. columns "nds"), this comes at the expense of a significantly longer running time, especially as the size of the instances grows; furthermore, the quality of the global lower bound is significantly worse than that computed by the Lagrangian approach.

The above results are fully confirmed by these reported in Table 2 for the higher accuracy setting of 0.01%. While this setting produces solutions of much better quality, it does so at the expense of a dramatic increase of running times. It is particularly striking that, with the exception of the smallest instances, the lower bound computed by the MILP approaches improves much less than the upper bound; in most cases, it is much worse—up to almost an order of magnitude—than that computed by the Lagrangian approach in a tiny fraction of the time, despite enumeration of several tens of thousands of Branch&Bound nodes.

It should be remarked at this point that the gaps reported in the above tables are "abstract", in the sense that they are obtained with respect to the best known solutions; therefore, the "perceived" gaps of both approaches are typically higher. In particular, the MILP approach, while computing quite accurate

		LR		LRH		$PCFD_\infty$				$PCFD_4$			
p	h	dgap	time	time	gap	gap	dgap	nds	time	gap	dgap	nds	time
10	0	0.677	0.24	0.75	0.99	0.28	0.291	0	0.80	0.30	0.297	0	0.86
20	0	0.472	0.46	1.83	0.56	0.33	0.298	0	3.00	0.36	0.283	0	2.51
50	0	0.086	1.25	4.84	0.28	0.18	0.278	0	13.08	0.19	0.279	0	14.17
75	0	0.076	2.22	9.41	0.34	0.22	0.267	0	22.58	0.19	0.261	2	36.62
100	0	0.059	3.68	14.74	0.33	0.20	0.257	0	36.51	0.17	0.262	0	34.31
150	0	0.048	6.14	21.20	0.17	0.12	0.175	10	169.68	0.11	0.184	4	104.68
200	0	0.049	8.52	34.80	0.09	0.14	0.179	12	235.60	0.10	0.178	0	183.01
20	10	0.072	0.59	1.76	0.39	0.15	0.166	0	2.51	0.30	0.162	5	4.18
50	20	0.004	1.46	6.36	0.06	0.13	0.084	0	10.93	0.10	0.083	10	19.06
75	35	0.011	3.20	15.01	0.04	0.03	0.058	95	64.80	0.05	0.030	115	70.55
100	50	0.010	5.44	24.74	0.04	0.04	0.068	40	60.78	0.05	0.068	15	47.62
150	75	0.008	8.85	37.41	0.02	0.05	0.059	115	216.33	0.05	0.059	115	194.10
200	100	0.005	12.95	50.91	0.01	0.03	0.050	135	342.69	0.02	0.050	0	155.36

Table 1: Comparing LR, LRH, $PCFD_4$ and $PCFD_\infty$ at low accuracy

		LR		LRH		$PCFD_\infty$				$PCFD_4$			
p	h	dgap	time	time	gap	gap	dgap	nds	time	gap	dgap	nds	time
10	0	0.677	0.24	0.75	0.99	0.02	0.020	523	17	0.02	0.021	434	12
20	0	0.472	0.46	1.83	0.56	0.02	0.015	28795	3547	0.02	0.017	40837	3915
50	0	0.086	1.25	4.84	0.28	0.09	0.136	25266	10000	0.09	0.138	29980	10000
75	0	0.076	2.22	9.41	0.34	0.09	0.164	25067	10000	0.08	0.162	27185	10000
100	0	0.059	3.68	14.74	0.33	0.07	0.166	19126	10000	0.06	0.168	21066	10000
150	0	0.048	6.14	21.20	0.17	0.05	0.156	14570	10000	0.05	0.159	16093	10000
200	0	0.049	8.52	34.80	0.09	0.05	0.166	8912	10000	0.06	0.166	10123	10000
20	10	0.072	0.59	1.76	0.39	0.02	0.018	10666	268	0.02	0.018	11026	249
50	20	0.004	1.46	6.36	0.06	0.00	0.024	104880	7285	0.01	0.023	118062	6496
75	35	0.011	3.20	15.01	0.04	0.01	0.025	96879	10000	0.01	0.028	85757	10000
100	50	0.010	5.44	24.74	0.04	0.01	0.034	96148	10000	0.01	0.037	99305	10000
150	75	0.008	8.85	37.41	0.02	0.01	0.043	67230	10000	0.01	0.043	68664	10000
200	100	0.005	12.95	50.91	0.01	0.01	0.040	45682	10000	0.01	0.039	47011	10000

Table 2: Comparing LR, LRH, $PCFD_4$ and $PCFD_\infty$ at high accuracy

primal solutions, is typically unable to properly recognize their accuracy due to its weak lower bound (an occurrence not at all infrequent for difficult combinatorial optimization problems).

4.2. The hybrid approach

The previous results suggest trying to combine the strengths of the two approaches, that is, the very efficient computation of an extremely accurate lower bound of the Lagrangian approach with the effective heuristic of the MILP one. For this, a hybrid approach can be simply devised which works as follows:

- first, the Lagrangian lower bound is computed (without any heuristic for producing feasible solutions);
- then, the MILP approach is started, but the MILP solver is provided with knowledge of the lower bound computed in the previous section, which it then can use to terminate the search as soon as a solution with the prescribed accuracy is reached.

Given the sophisticated tools made available to the (expert) user by the currently available off-the-shelf MILP solvers, in particular, the `Cplex callback` functions have been used for this task.

An interesting feature of this approach is that it combines the flexibility of MILP-based algorithms with the efficiency that is only possible by using specialized algorithms like the Lagrangian ones. In

fact, Lagrangian approaches, while historically among the best options for solving UC problems of large size, have been reported to be difficult to adapt to the specific needs of actual operational environments. Despite several attempts to make up for this drawback (e.g. [11]), the interest in approaches based on more flexible tools like MILP solvers has been steadily increasing [8, 9, 31], especially since the combined effect of algorithmic advances and hardware improvements have brought large-scale UC problems on the verge of being routinely solvable by these tools. Being the UC model employed in this paper a quite general one, most adaptations required to tailor it to different operational environments would lead to further constraints (and possibly variables) w.r.t. those listed in (2)—(11); by relaxing these constraints, the approach described in §3.1 can still be used to solve the Lagrangian problem. Thus, at the cost of an increased number of Lagrangian multipliers, the lower bound computation can be extended to many different versions of UC problems, as shown e.g. in [7]. This could require substantial modifications in the heuristic for a Lagrangian-only approach [20], but it would be relatively simpler to implement if the upper bound is instead provided by a MILP-based approach; thus, the flexibility is retained while the effective available specialized approaches can still be used.

We now present results aimed at assessing the effectiveness of the hybrid approach over the standard MILP-based one. The results are presented in Tables 3 to 6. In each table, the columns "LB" refer to the hybrid approach, whereas the columns "NoLB" refer to the standard MILP-based one; the meaning of "gap" and "time" is the same as the previous tables. The difference between each table only lies in the accuracy required to the MILP solver before stopping, which is 0.5% (the same value as in Table 1) for Table 3, 0.1% for Table 4, 0.05% for Table 5, and, finally, 0.01% (the same value as in Table 2) for Table 6. In all cases, same the upper limit of 10000 seconds of running time is kept.

		$PCFD_4$				$PCFD_\infty$			
		NoLB		LB		NoLB		LB	
p	h	gap	time	gap	time	gap	time	gap	time
10	0	0.28	0.80	0.34	0.97	0.30	0.86	0.37	1.06
20	0	0.33	3.00	0.32	3.60	0.36	2.51	0.36	3.16
50	0	0.18	13.08	0.19	27.46	0.19	14.17	0.20	16.39
75	0	0.22	22.58	0.25	28.82	0.19	36.62	0.22	28.05
100	0	0.20	36.51	0.15	41.44	0.17	34.31	0.16	60.16
150	0	0.12	169.68	0.10	148.88	0.11	104.68	0.11	136.18
200	0	0.14	235.60	0.08	323.36	0.10	183.01	0.08	258.57
20	10	0.15	2.51	0.17	4.21	0.30	4.18	0.24	6.34
50	20	0.13	10.93	0.10	26.96	0.10	19.06	0.10	12.51
75	35	0.03	64.80	0.06	59.47	0.05	70.55	0.10	75.23
100	50	0.04	60.78	0.04	44.95	0.05	47.62	0.05	66.61
150	75	0.05	216.33	0.02	244.05	0.05	194.10	0.04	228.32
200	100	0.03	342.69	0.03	253.59	0.02	155.36	0.02	217.56

Table 3: Comparing the hybrid and the standard approach, accuracy 0.5%

For the low accuracy of 0.5% (Table 3), the hybrid approach does not deliver any significant benefit; indeed, while at times (e.g. $p = 200, h = 100$) a positive effect can be seen, in general the improved lower bound is not useful. This is clearly explained by the fact that for this setting the MILP approach most often terminates at the root node (cf. Table 1), or after very little branching; thus, improving the bound cannot have any significant effect. Especially for the small-sized instances, the small *absolute* increase in running time due to the Lagrangian bound computation most often results in a sizable *relative* worsening of the efficiency of the approach.

The situation is radically different for all mid- to large-scale pure thermal instances when the accuracy is set to 0.1%, as shown in Table 4. While the standard approach cannot solve any instance with $p > 50$ (and very few of those with $p = 50$) with the prescribed accuracy within the allotted time limit, the hybrid approach solves very efficiently all instances with $p \geq 100$. The reason, as evident from the results reported in Tables1 and 2, is that the Lagrangian lower bound accuracy actually increases as n does.

		PCFD$_4$				PCFD$_\infty$			
		NoLB		LB		NoLB		LB	
p	h	gap	time	gap	time	gap	time	gap	time
10	0	0.10	12.45	0.10	12.70	0.10	9.77	0.10	9.97
20	0	0.10	1295.28	0.10	2201.87	0.10	1169.94	0.10	1157.22
50	0	0.09	8279.78	0.11	4084.79	0.10	10000.00	0.11	4014.01
75	0	0.07	10000.00	0.09	3974.94	0.07	10000.00	0.09	2286.03
100	0	0.07	10000.00	0.09	289.01	0.06	10000.00	0.09	94.56
150	0	0.05	10000.00	0.06	193.38	0.05	10000.00	0.08	207.86
200	0	0.05	10000.00	0.07	337.33	0.06	10000.00	0.07	315.88
20	10	0.07	31.38	0.09	14.31	0.07	41.08	0.08	30.01
50	20	0.02	41.86	0.05	27.22	0.02	47.62	0.04	12.92
75	35	0.03	64.45	0.06	57.95	0.04	81.77	0.06	71.03
100	50	0.03	40.61	0.04	41.42	0.04	60.20	0.05	62.85
150	75	0.02	232.99	0.02	235.04	0.04	191.52	0.04	203.18
200	100	0.03	240.38	0.03	231.35	0.02	198.25	0.02	206.66

Table 4: Comparing the hybrid and the standard approach, accuracy 0.1%

However, hydro-thermal instances typically attain an accuracy of 0.1% (or higher) very quickly even when the threshold is set to 0.5%, which means that the Lagrangian lower bound cannot be of much help in this case; this is in fact confirmed by the results in the table.

		PCFD$_4$				PCFD$_\infty$			
		NoLB		LB		NoLB		LB	
p	h	gap	time	gap	time	gap	time	gap	time
10	0	0.06	15.42	0.06	15.72	0.06	11.63	0.06	11.85
20	0	0.06	2473.11	0.06	2440.86	0.06	2470.49	0.06	2499.97
50	0	0.09	10000.00	0.09	8113.35	0.09	10000.00	0.10	8489.08
75	0	0.09	10000.00	0.09	10002.22	0.08	8256.79	0.08	8259.00
100	0	0.07	10000.00	0.07	8018.89	0.06	10000.00	0.06	6538.84
150	0	0.05	10000.00	0.06	5151.71	0.05	10000.00	0.06	6151.20
200	0	0.05	10000.00	0.05	6255.99	0.06	10000.00	0.06	6271.77
20	10	0.06	73.26	0.06	73.00	0.06	71.19	0.06	68.40
50	20	0.01	623.95	0.02	34.44	0.01	269.34	0.03	44.53
75	35	0.02	177.50	0.03	59.37	0.02	124.85	0.03	100.47
100	50	0.02	438.39	0.04	39.45	0.02	665.37	0.05	60.00
150	75	0.02	1669.30	0.02	224.67	0.01	1144.10	0.04	201.31
200	100	0.02	1082.41	0.03	238.81	0.01	451.98	0.02	202.94

Table 5: Comparing the hybrid and the standard approach, accuracy 0.05%

Further increasing the accuracy to 0.05% confirms the trend previously seen; the higher accuracy clearly requires a longer running time to be achieved, but still the hybrid approach is significantly more efficient than the standard one, as Table 5 shows. The effect is now visible also on hydro-thermal ones; while the increased accuracy results in increased running time for both approaches, the standard one now requires up to an order of magnitude more (cf. $p = 100, h = 50$) than the hybrid one.

The highest setting of 0.01% (the same as Table 2) confirms the basic trend. The two approaches are now equivalent on the pure thermal instances; the better lower bound provided by the Lagrangian is no longer enough to allow the hybrid approach to terminate before the time limit (for $p > 20$), hence the only effect of the Lagrangian computation is to make the hybrid approach ever so slightly slower than the standard one. However, on hydro-thermal instances the hybrid approach is again significantly more effective than the standard one; while the latter fails to solve all large-scale instances, the former does, resulting in a speedup of up to a factor of 50 (cf. $p = 50, h = 20$).

		$PCFD_4$				$PCFD_\infty$			
		NoLB		LB		NoLB		LB	
p	h	gap	time	gap	time	gap	time	gap	time
10	0	0.02	16.66	0.02	16.84	0.02	12.49	0.02	12.80
20	0	0.02	3547.24	0.02	3699.26	0.02	3914.50	0.02	3946.51
50	0	0.09	10000.00	0.09	10001.25	0.09	10000.00	0.09	10001.25
75	0	0.09	10000.00	0.09	10002.22	0.08	10000.00	0.08	10002.22
100	0	0.07	10000.00	0.07	10003.68	0.06	10000.00	0.06	10003.68
150	0	0.05	10000.00	0.05	10006.14	0.05	10000.00	0.05	10006.14
200	0	0.05	10000.00	0.05	8248.37	0.06	10000.00	0.06	10008.52
20	10	0.02	268.49	0.02	263.40	0.02	248.75	0.02	255.95
50	20	0.00	7285.00	0.01	841.26	0.01	6495.96	0.01	121.86
75	35	0.01	10000.00	0.01	5033.34	0.01	10000.00	0.01	5045.42
100	50	0.01	10000.00	0.01	1198.73	0.01	10000.00	0.01	5789.69
150	75	0.01	10000.00	0.01	3376.87	0.01	10000.00	0.01	1145.61
200	100	0.01	10000.00	0.01	1182.27	0.01	10000.00	0.01	463.46

Table 6: Comparing the hybrid and the standard approach, accuracy 0.01%

5. Conclusions

In this paper, we have proposed an hybrid MILP-Lagrangian approach to hydro-thermal Unit Commitment problems; this require the nontrivial combination of several algorithmic techniques including convex nondifferentiable optimization, nonlinear optimization, and polyhedral techniques for mixed-integer optimization. By means of computational experiments, we have shown that for an appropriate combination of type (pure thermal or hydro thermal) and size of the instance, and for specific choices of the required accuracy threshold, the hybrid approach is substantially more effective than the standard MILP-based one. Thus, the newly proposed approach allows for more finely tuning of the trade-off between accuracy of the obtained solution and required running time with respect to the pure Lagrangian one; hence, we believe that the proposed algorithm can be valuable for those operating environments where the accuracy of the obtained solution is a crucial parameter. An interesting feature of the approach is that it combines the flexibility of MILP-based algorithms with the efficiency that is only possible by using specialized algorithms, thereby being flexible enough to suit the needs of most operating environments while retaining, at least in part, the traditional superior performances of Lagrangian approaches.

References

[1] J. Arroyo and A. Conejo, "Optimal Response of a Thermal Unit to an Electricity Spot Market," *IEEE Transactions on Power Systems*, vol. 15, no. 3, pp. 1098–1104, 2000.

[2] L. Bacaud, C. Lemaréchal, A. Renaud, and C. Sagastizábal, "Bundle Methods in Stochastic Optimal Power Management: A Disaggregated Approach Using Preconditioners," *Computational Optimization and Applications*, vol. 20, pp. 227–244, 2001.

[3] J. Bard, "Short-term Scheduling of Thermal-electric Generators using Lagrangian Relaxation," *Operations Research*, vol. 36, no. 5, pp. 765–766, 1988.

[4] A. Belloni, A. Diniz, M. Maceira, and C. Sagastizábal, "Bundle Relaxation and Primal Recovery in Unit Commitment Problems. The Brazilian Case," *Annals of Operations Research*, vol. 120, pp. 21–44, 2003.

[5] A. Borghetti, A. Frangioni, F. Lacalandra, A. Lodi, S. Martello, C. Nucci, and A. Trebbi, "Lagrangian Relaxation and Tabu Search Approaches for the Unit Commitment Problem," in *Proceedings IEEE 2001 Powerteck Porto Conference* (Saraiva, J.T. and Matos, M.A., eds.), vol. Paper n. PSO5-397, 2001.

[6] A. Borghetti, A. Frangioni, F. Lacalandra, and C. Nucci, "Lagrangian Heuristics Based on Disaggregated Bundle Methods for Hydrothermal Unit Commitment," *IEEE Transactions on Power Systems*, vol. 18, pp. 313–323, 2003.

[7] A. Borghetti, A. Frangioni, F. Lacalandra, C. Nucci, and P. Pelacchi, "Using of a Cost-based Unit Commitment Algorithm to Assist Bidding Strategy Decisions," in *Proceedings IEEE 2003 Powerteck Bologna Conference* (Borghetti, A., Nucci, C.A., and Paolone, M., eds.), vol. Paper n. 547, 2003.

[8] M. Carrión and J. Arroyo, "A Computationally Efficient Mixed-integer Linear Formulation for the Thermal Unit Commitment Problem," *IEEE Transactions on Power Systems*, vol. 21, no. 3, pp. 1371–1378, 2006.

[9] G. Chang, M. Aganagic, J. Waight, J. Medina, T. Burton, S. Reeves, and M. Christoforidis, "Experiences with Mixed Integer Linear Programming Based Approaches on Short-term Hydro Scheduling," *IEEE Transactions on Power Systems*, vol. 16, no. 4, pp. 743–749, 2001.

[10] S. de la Torre, J. Arroyo, A. Conejo, and J. Contreras, "Price Maker Self-Scheduling in a Pool-Based Electricity Market: A Mixed-Integer LP Approach," *IEEE Transactions on Power Systems*, vol. 17, no. 4, pp. 1037–1042, 2002.

[11] L. Dubost, R. Gonzalez, and C. Lemaréchal, "A Primal-proximal Heuristic Applied to the French Unit Commitment Problem," *Mathematical Programming*, vol. 104, no. 1, pp. 129–151, 2005.

[12] S. Feltenmark and K. Kiwiel, "Dual Applications of Proximal Bundle Methods, Including Lagrangian Relaxation of NonConvex Problems," *SIAM Journal on Optimization*, vol. 10, no. 3, pp. 697–721, 2000.

[13] A. Frangioni, "Generalized Bundle Methods," *SIAM Journal on Optimization*, vol. 13, no. 1, pp. 117–156, 2002.

[14] A. Frangioni, "About Lagrangian Methods in Integer Optimization," *Annals of Operations Research*, vol. 139, pp. 163–193, 2005.

[15] A. Frangioni and C. Gentile, "New Preconditioners for KKT Systems of Network Flow Problems," *SIAM Journal on Optimization*, vol. 14, no. 3, pp. 894 – 913, 2004.

[16] A. Frangioni and C. Gentile, "Perspective Cuts for a Class of Convex 0–1 Mixed Integer Programs," *Mathematical Programming*, vol. 106, no. 2, pp. 225 – 236, 2006.

[17] A. Frangioni and C. Gentile, "Solving Nonlinear Single-Unit Commitment Problems with Ramping Constraints," *Operations Research*, vol. 54, no. 4, pp. 767 – 775, 2006.

[18] A. Frangioni and C. Gentile, "Prim-based Support-Graph preconditioners for Min-Cost Flow Problems," *Computational Optimization and Applications*, vol. 36, no. 2-3, pp. 271 – 287, 2007.

[19] A. Frangioni and C. Gentile, "A Computational Comparison of Reformulations of the Perspective Relaxation: SOCP vs. Cutting Planes," 2008. Manuscript in preparation.

[20] A. Frangioni, C. Gentile, and F. Lacalandra, "New Lagrangian Heuristics for Ramp-Constrained Unit Commitment Problems," in *Proceedings 19th Mini-EURO Conference in Operational Research Models and Methods in the Energy Sector – ORMMES 2006*, Coimbra, 2006.

[21] A. Frangioni, C. Gentile, and F. Lacalandra, "Tighter Approximated MILP Formulations for Unit Commitment Problems," *submitted to IEEE Transactions on Power Systems*, 2007.

[22] A. Frangioni, C. Gentile, and F. Lacalandra, "Solving Unit Commitment Problems with General Ramp Contraints," *International Journal of Electrical Power and Energy Systems*, vol. 30, pp. 316 – 326, 2008.

[23] A. Frangioni and A. Manca, "A Computational Study of Cost Reoptimization for Min Cost Flow Problems," *INFORMS Journal on Computing*, vol. 18, no. 1, pp. 61–70, 2006.

[24] B. Hobbs, M. Rothkopf, R. O'Neill, and H. Chao, *The Next Generation of Unit Commitment Models*. Boston: Kluwer Academic Press, 2001.

[25] Z. Li and M. Shahidehpour, "Generation Scheduling With Thermal Stress Constraints," *IEEE Transactions on Power Systems*, vol. 18, no. 4, pp. 1402–1409, 2003.

[26] B. Lu and M. Shahidehpour, "Short-term scheduling of combined cycle units," *IEEE Transactions on Power Systems*, vol. 19, no. 3, pp. 1616–1625, 2004.

[27] M. Madrigal and V. Quintana, "An Interior-Point/Cutting-Plane Method to Solve Unit Commitment Problems," *IEEE Transactions on Power Systems*, vol. 15, no. 3, pp. 1022–1027, 2000.

[28] O. Nilsson and D. Sjelvgren, "Mixed-integer Programming Applied to Short-term Planning of a Hydro-thermal System," *IEEE Transactions on Power Systems*, vol. 11, no. 1, pp. 281–286, 1996.

[29] M. Nowak and W. Römisch, "Stochastic Lagrangian Relaxation Applied to Power Scheduling in a Hydro-thermal System Under Uncertainty," *Annals of Operations Research*, vol. 100, pp. 251–272, 2000.

[30] A. Wood and B. Wollemberg, *Power Generation Operation and Control*. John Wiley & Sons, 1996.

[31] Z. Yu, F. Sparrow, B. Bowen, and F. Smardo, "On convexity issues of short-term hydrothermal scheduling," *Electrical Power and Energy Systems*, vol. 20, pp. 451–457, 2000.

[32] F. Zhuang and F. Galiana, "Towards a More Rigorous and Practical Unit Commitment by Lagrangian Relaxation," *IEEE Transactions on Power Systems*, vol. 3, pp. 763–773, 1988.