



ISTITUTO DI ANALISI DEI SISTEMI ED INFORMATICA
CONSIGLIO NAZIONALE DELLE RICERCHE

A. Frangioni, C. Gentile, F. Lacalandra

**SOLVING UNIT COMMITMENT PROBLEMS
WITH GENERAL RAMP CONSTRAINTS**

R. 637 Gennaio 2006

Antonio Frangioni – Dipartimento di Informatica, Largo B. Pontecorvo 3, 56127 Pisa (Italy).
Email: frangio@di.unipi.it.

Claudio Gentile – Istituto di Analisi dei Sistemi ed Informatica del CNR, Viale Manzoni 30 -
00185 Roma, Italy. Email: gentile@iasi.cnr.it. This author has been partly supported
by the EU “Marie Curie” Research Training Network no. 504438 ADONET.

Fabrizio Lacalandra – *OptiME*, E-mail: fabrizio.lacalandra@fastwebnet.it.

ISSN: 1128–3378

Collana dei Rapporti dell'Istituto di Analisi dei Sistemi ed Informatica, 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

Lagrangian Relaxation (LR) algorithms are among the most successful approaches for solving large-scale hydro-thermal Unit Commitment (UC) problems; this is largely due to the fact that the Single-Unit Commitment (1UC) problems resulting from the decomposition, incorporating many kinds of technical constraints such as minimum up- and down-time requirements and time-dependent startup costs, can be efficiently solved by Dynamic Programming (DP) techniques. Ramp constraints have historically eluded efficient exact DP approaches; however, this has recently changed [16]. We show that the newly proposed DP algorithm for ramp-constrained (1UC) problems allows to extend existing LR approaches to ramp-constrained (UC); this is not obvious since the heuristic procedures typically used to recover a primal feasible solution are not easily extended to take into account ramp limits. However, dealing with ramp constraints in the subproblems turns out to be sufficient to provide the LR heuristic enough guidance to produce good feasible solutions even with no other modification of the approach; this is due to the fact that (sophisticated) LR algorithms to (UC) duly exploit the primal information computed by the Lagrangian Dual, which in the proposed approach is ramp feasible.

Key words: Hydro-Thermal Unit Commitment, Ramp Limits, 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 [24], which is of increasing importance in the ongoing liberalization of the electricity market in many countries [17]. Indeed, while the (UC) problem 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) problem can usually be easily extended to forms of the problem arising in a market environment [17, 9].

Despite having attracted the interest of researchers for over 30 years, (UC) still cannot be considered a well-solved problem for all practical sizes and operating environments; this should not be surprising, since it is a large-scale, Mixed-Integer NonLinear Program (MINLP). Despite the ever-increasing availability of cheap computing power and the advances in off-the-shelf software for MINLP, solving (UC) by general-purpose software, even using the most advanced approaches available, is not feasible when the number of units [13] and/or the length of the time horizon [16] grows large; this is also confirmed by the results reported in § 6 of the present paper. Among the most efficient specialized algorithmic approaches for (UC), Lagrangian Relaxation (LR) methods [2, 6, 7, 8, 12, 21, 25] 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. These problems are typically easily solvable by network flow techniques—for hydro units—or Dynamic Programming (DP) techniques—for thermal units.

However, the success of LR methods critically depend on two factors:

- being able to optimally solving the (1UC) problems efficiently;
- being able to turn the unfeasible solution(s) produced at each step of the approach into “good”, feasible ones.

In turn, both these depends on the specific details of the operational constraints of the generating units that are represented in the mathematical model; each time that new constraints have to be included in the model, a (usually, nontrivial) modification is required to the solution method to (1UC) (e.g., [19, 20, 10]), and possibly of the *heuristic* used to produce feasible solutions. Thus, here appears the typical trade-off between the accuracy of the mathematical models employed and the efficiency of the corresponding solution process. For hydro units, for instance, it is customary to assume a linear relationship between the discharged water and the generated power, disregarding the nonlinearities corresponding to the variation of the water head and the

minimum discharge necessary to keep the turbines operational; these assumptions are reasonable when dealing with short term operation, and dramatically simplify the hydro (1UC) problems. For thermal units, it is usually assumed that the dynamic of the generating plant does not pose restrictions (other than on maximum and minimum power levels) on the amount of power generated at each timestamp of the time horizon; unfortunately, this is not realistic for large units or if the timestamps are to be taken small (e.g., 15 minutes), since then *ramp constraints* need to be considered. These limit the maximum increase or decrease of generated power from one timestamp to the next, reflecting the thermal and mechanical inertia that has to be overtaken in order for the unit to increase or decrease its output. Unfortunately, ramp constraints complicate the structure of thermal (1UC) problems, rendering the classical DP approaches unusable; this is due to the fact that the variables representing the power output at different timestamps are no longer independent once that commitment (on/off) decision has been taken. Four approaches have been proposed in the literature to overcome this difficulty:

- Discretizing the power variables space one may keep using a standard DP procedure [5, 4]; however, the computational burden increases as the granularity of the discretization is decreased, and the obtained solution is in general an approximated one.
- Approximating the cost function with a piecewise-linear (convex) function one may use the DP approach of [11], in turn based on the efficient algorithm the *Single-Unit Economic Dispatch* (1ED) problems with ramp constraints of [3, 23]; however, the computational burden increases with the accuracy of the approximation (the number of linear pieces in the cost function), and the obtained solution is in general an approximated one.
- Solving (1UC) problems only approximately, e.g., by using two-stage Lagrangian techniques as in [18], where ramp constraints are dualized and the corresponding Lagrangian multipliers updated (keeping those of the system balance constraints fixed) until a solution that is ramp-feasible is attained; however, the computational burden increases as multiple solutions of each (1UC) are required for each LR iteration, and the obtained solution is in general an approximated one.
- Using general-purpose, off-the-shelf Mixed-Integer Quadratic Programming (MIQP) solvers for solving (1UC) in alternative to DP approaches. This has the advantage of allowing to incorporate several other constraints in the formulation, e.g. related to the electrical market; since MIQP effective solvers have only recently become available, this approach has usually been coupled to approximation of the cost function by a piecewise-linear (convex) function, thus allowing use of (more common) Mixed-Integer Linear Programs (MILP) solver [1], although this comes at the cost of significantly increasing the number of continuous variables in the (1UC) problem, and therefore its solution cost. However, the solution time for general-purpose solvers (as it should be expected) grows very rapidly with the size of the instance, making this approach unusable for longer time horizons [16], even if a state-of-the-art MIQP solver is used.

Recently, a DP algorithm for thermal (1UC) with ramping constraints has been proposed [16] that can solve to optimality problems on a time horizon of n timestamps in $O(n^3)$ for convex cost functions whose unconstrained minimum can be found in $O(1)$, such as the quadratic ones typically used in operational settings. The approach allows to solve thermal (1UC) problems with a very general definition of ramping constraints and complex, time-dependent startup costs. In this paper, we report on the use of that algorithm within LR approaches to (UC).

We show that the increased complexity of the (1UC) subproblems—w.r.t. $O(n^2)$ of the case without ramp constraints—does not represent a serious computational bottleneck for this class of approaches. Also, we show that existing LR approaches that duly use the primal information provided by the Lagrangian Dual (e.g. [12, 8]) are fairly “resistant” to the introduction of ramp constraints, once these are properly taken into account in the Lagrangian subproblem. This is not obvious, since the heuristic procedures typically used to recover a primal feasible solution are not easily extended to take into account ramp constraints. Yet, solving ramp-constrained (1UC) subproblems turns out to be sufficient to provide the LR heuristic enough guidance to produce good feasible solutions even with no other substantial modification of the approach, due to the fact that the primal information computed by the Lagrangian Dual always remains ramp-feasible. To the best of our knowledge, this is the first LR approach that has been shown capable of solving—very large-scale—(UC) problems with ramp constraints with exact solution of the (1UC) subproblems and without any form of approximation of either the power production levels or the objective function. We show on a set of realistic thermal and hydro-thermal instances that the proposed approach consistently obtains much better solutions than a comparable approach solving non-ramp-constrained (1UC) problems. Also, we show that good quality solutions are obtained in a small fraction of the time required for the same task by state-of-the-art, general-purpose MIQP technology.

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, and let $\mathcal{T} = \{1, \dots, n\}$ be the set of timestamps defining the time horizon (the timestamp “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

$$\sum_{i \in P} c^i(\mathbf{p}^i, \mathbf{u}^i) = \sum_{i \in P} \left(c^i(\mathbf{u}^i) + \sum_{t \in \mathcal{T}} c_t^i(p_t^i) \right) ; \quad (1)$$

that is, while the power production cost at each hour is typically approximated via a (convex) quadratic separable form (neglecting for instance the so called valve points, see e.g. [24]) in the p_t^i variables, the cost function is nonseparable per hour due to time-dependent *start-up costs*, whose exact form has no impact on the proposed approach and is not reported here for the sake of notational simplicity. 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}_{min}^i and \bar{p}_{max}^i 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 timestamp where it is committed or decommitted, respectively. Then, the local constraints corresponding to each unit $i \in P$

are

$$\bar{p}_{min}^i u_t^i \leq p_t^i \leq \bar{p}_{max}^i u_t^i \quad t \in \mathcal{T} \quad (2)$$

$$p_t^i \leq p_{t-1}^i + u_{t-1}^i \Delta_+^i + (1 - u_{t-1}^i) \bar{l}^i \quad t \in \mathcal{T} \quad (3)$$

$$p_{t-1}^i \leq p_t^i + u_t^i \Delta_-^i + (1 - u_t^i) \bar{u}^i \quad t \in \mathcal{T} \quad (4)$$

$$u_t^i \geq u_r^i - u_{r-1}^i \quad t \in \mathcal{T}, r \in [t - \tau_+^i, t - 1] \quad (5)$$

$$u_t^i \geq 1 - u_{r-1}^i - u_r^i \quad t \in \mathcal{T}, r \in [t - \tau_-^i, t - 1] \quad (6)$$

$$u_t^i \in \{0, 1\} \quad t \in \mathcal{T} \quad (7)$$

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_0^i , the generated power p_0^i ; for the sake of minimum up- and down-time constraints (5), (6), it is also necessary to know how long it has been on or off at timestamp 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_t^j , v_t^j and w_t^j represent respectively discharged water, the volume of the reservoir and the spilled water at timestamp $t \in \mathcal{T}$. Constants \bar{v}_{min}^j and \bar{v}_{max}^j represents respectively the minimum and maximum volume for the reservoir, \bar{q}_{max}^j 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}_t^j represent the natural inflows at timestamp $t \in \mathcal{T}$. 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:

$$0 \leq q_t^j \leq \bar{q}_{max}^j \quad t \in \mathcal{T} \quad (8)$$

$$\bar{v}_{min}^j \leq v_t^j \leq \bar{v}_{max}^j \quad t \in \mathcal{T} \quad (9)$$

$$v_t^j - v_{t-1}^j = \bar{w}_t^j - w_t^j - q_t^j - \sum_{k \in \mathcal{S}(j)} \left(q_{t-t_{kj}}^k + w_{t-t_{kj}}^k \right) \quad t \in \mathcal{T} \quad (10)$$

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

- *Global constraints:* for each timestamp $t \in \mathcal{T}$, 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, again to avoid nonlinearities); then, the system-wide constraints—linking the different units among themselves—are:

$$\sum_{i \in P} p_t^i + \sum_{h \in H} \sum_{j \in H(h)} \alpha^j q_t^j = \bar{d}_t \quad t \in \mathcal{T} \quad (11)$$

We refer to (UC) as the problem of minimizing (1) subject to constraints (2)—(11). This is a large-scale, mixed-integer nonlinear program; however, it should be noted that when the u_t^i

variables are fixed, the remaining *Economic Dispatch* (ED) problem is a linearly constrained convex problem, hence easily solvable. 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 .

The UC model here proposed can be particularly suitable in the free markets where GenCos need to optimize their production schedule once that their own load profile has been (somehow) established by the market procedures.

We remark that several of the (widely accepted) simplifying assumptions in the above model can be relaxed without hindering the applicability of the proposed approach; in particular, more sophisticated models of hydro cascades, e.g. taking into account nonlinear effects of the water head on the power-to-discharged-water efficiency and/or nonzero technical minima for discharged water, could be used at the cost of more difficult hydro subproblems and (ED)s. Conversely, thermal units are modeled in a fairly sophisticated way, and the approach can be easily extended (cf. [16]) to handle issues such as:

- Data dependent on the timestamp t , i.e., different ramp limits Δ_{t+}^i and Δ_{t-}^i or start-up and shut-down power output \bar{l}_t^i and \bar{u}_t^i for each $t \in \mathcal{T}$, modeling, e.g., the reaction to varying operating parameters such as the external temperature.
- Data dependent on the history of the unit, i.e., different ramp limits, start-up and shut-down power output, maximum and minimum power levels, and even (coefficients of the) cost function according to *how long the unit has remained consecutively committed* before the current timestamp.
- Different discretization intervals for commitment and power variables, e.g., power production levels to be specified for each 15-minutes period of the time horizon whereas the commitment decisions can only be changed hourly.

3. The Lagrangian relaxation

The LR approach is based on dualizing each of the coupling constraint (11) via a Lagrangian multiplier λ_t , thereby forming the *Lagrangian Relaxation* of (UC)

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

where

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

Is is well known (e.g., [15]) that for each $\boldsymbol{\lambda} \in \mathbb{R}^n$, $\mathcal{L}(\boldsymbol{\lambda})$ gives a lower bound on the optimal value of (UC). Therefore, one is interested in the $\boldsymbol{\lambda}^*$ such that this lower bound is best (greatest), i.e., in the optimal solution of the *Lagrangian Dual* of (UC):

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

Since $\mathcal{L}(\cdot)$ is a non differentiable function, proper algorithms must be chosen for solving (13); *bundle* methods [14], particularly in their *disaggregated* variant [2], have been repeatedly reported to be quite efficient in solving (13), much more so [8] than alternative algorithms such as subgradient methods [4, 25].

However, solving (13) is not, in general, enough to solve (UC); even for $\lambda = \lambda^*$, the optimal solution to (12) is not guaranteed to—and will not in general—satisfy the relaxed constraints (11). Therefore, a number of *Lagrangian heuristics* have been devised that attempt to convert the (sequence of) unfeasible solution(s) provided by (13) in a (sequence of) feasible, hopefully “good”, one(s). Roughly speaking, there are two possible ways for doing this:

- either combinatorial heuristics are run at every step of the iterative solution of (13), thus for several different values of λ , that use the unfeasible optimal solution of the corresponding (12) to produce a feasible solution for (UC);
- or the Lagrangian problem is modified with further terms that try to enforce feasibility of the iterates.

Most approaches from the literature belong to the first group; it is outside the scope of the present paper to describe them in details, the interested reader being referred, e.g., to [2, 4, 8, 7, 12, 25]. These approaches suffer from the drawback that the combinatorial heuristic ran at each iteration is dependent on the exact details of the model to be solved, and may be difficult to adapt to different situations; this motivated the development of the second group of approaches, where a further quadratic term (either an Augmented Lagrangian one [6] or a proximal one [10]) is added to the problem to try to enforce feasibility of the iterates. Since the quadratic term may make the Lagrangian relaxation harder to solve, approximate solution via sequential linearization may have to be employed.

It is relevant for the present context that all the most recent and efficient approaches, both of the first [8, 12] and of the second [6, 10] type, make use of the *continuous* primal information generated by the solution of (13). In fact, it is well-known that the Lagrangian Dual of a problem with convex objective function and nonconvex feasible set is equivalent to the *convexified relaxation* of the problem where the original objective function is optimized over the intersection of the convex hull of the (integer) solutions of the LR and the relaxed constraints. Most importantly, this equivalence is *computational* in the sense that basically all algorithms for solving the Lagrangian Dual either naturally compute, or can be modified to do so, a solution to the convexified relaxation [15]. In particular, bundle methods applied to (13) compute at each iteration a continuous solution that satisfies all constraints (2)—(9), and that pretty rapidly becomes almost feasible w.r.t. the relaxed constraints (11). This continuous, almost feasible solution is very useful to drive the heuristic search for integer, entirely feasible solutions to (UC), as described in more details in §5.

The LR approach has proven to be capable of solving non ramp-constrained large-scale (UC) problems with a provable high degree of accuracy in relatively short time. However, extending the approach to ramp constraints requires, in principle, three steps:

- i) solving (possibly approximately) ramp-constrained (1UC) subproblems;
- ii) modifying the heuristics in order to take into account the ramping constraints while fixing the commitment u_t^i variables;
- iii) solving ramp-constrained Economic Dispatch problems at each iteration to recover the p_t^i and q_t^j variables once that the u_t^i variables have been (heuristically) fixed.

Of those, (i) has only been made possible—without approximations or an excessive growth of the computational burden—by the development in [16], while (iii) is trivial in that it only requires

to add some linear constraints to the classical (ED) formulation. As for the point (ii), the rationale behind the “augmented” approaches [6, 10] was exactly to provide a general substitute for specialized combinatorial heuristics in order to avoid the problem.

The aim of this paper is to show that point (ii) is—at least for ramp constraints, but possibly also in other cases—much less of a problem than often perceived, that is, the combinatorial heuristics originally developed for (UC) without ramp constraints are efficient in the ramp-constrained case also, only provided that ramp-constrained (1UC) (and, obviously, ramp-constrained (ED)s) are solved. This should not be too surprising considering that, when a ramp-constrained (12) is solved, the latest combinatorial heuristics construct the feasible solutions using, as starting points, two solutions—an integer unfeasible one and the continuous almost feasible one—that are both ramp-feasible. Hence, ramp-constrained (1UC) provide—through the ramp-feasibility of the primal solutions, as well as through better dual multipliers—the standard LR heuristic with enough guidance to produce good feasible solutions even with no other modification of the approach.

4. Solving the (1UC) problem

In this section we briefly sketch the algorithm of [16] for solving the ramp-constrained (1UC) problem. We assume that each term c_t^i in (1) is a “simple” convex function, in the sense that unconstrained minimization of c_t^i can be carried upon in $O(1)$. Then, the overall cost function c^i of (1) is the sum of c_t^i plus a start-up cost function that only depends (but possibly in sophisticated ways) on how long the unit has been turned off. Because we are only interested in one unit at a time, for notational simplicity in the remainder of the section we will drop the superscript i .

The approach of [16] is a DP procedure, that is, it requires the computation of a shortest path in a suitable acyclic directed *state space* graph. Following an idea originally proposed in [11], the state space graph comprises a node for all pairs (h, k) for $h, k \in \mathcal{T}$ and $k \geq h$, plus a source s and a sink d . Each state (h, k) represents the unit being turned on at timestamp h (i.e., it being uncommitted at timestamp $h - 1$), and being turned off again at timestamp k ; clearly, all states that correspond to operations violating the minimum up-time requirements are not constructed. The graph has arcs between any node (h, k) and (r, q) such that it is feasible to turn on the unit at time instant r given that it has been turned off at time instant k ; each of these arcs are labeled with the start-up cost of the unit at time instant r . There are also arcs from the source s to all nodes (h, k) compatible with the initial state of the unit; finally, there is a zero-cost arc from each node to the sink d . Clearly, every $s - d$ path on this graph represents a feasible solution to (1UC), arc costs representing the contribution of start-up costs to the objective function (fixed generating costs, if any, can also be easily included). It is now necessary to compute, for each node (h, k) , the optimal contribution of the variable generating costs, that depend on the p_t variables. This requires the solution of the *Economic Dispatch with ramp constraints* problem for the interval $[h, k]$, that is, the minimization of (the continuous part of) (1) subject to (2)—(7) if *all u_t variables for $t \notin [h, k]$ are fixed to zero and all those for $t \in [h, k]$ are fixed to one*; we will denote this problem as (ED_{hk}) . Since all the binary variables are fixed, this is an optimization problem with convex objective function and linear constraints, hence easily solvable. By assigning its optimal objective function value $z_{hk}^* = z(ED_{hk})$ to the corresponding node, the (node plus arc) cost of each $s - d$ path on the graph is that of the feasible solution it represents. Thus, (1UC) can be solved with a visit of a graph with $O(n^2)$ nodes and $O(n^4)$ arcs, once that all the data has been computed; exploiting some structural

properties of the state-space graph, the complexity of the visit can be reduced to $O(n^3)$ [16].

However, computing the node costs requires the solution of $O(n^2)$ convex problems, each with up to n variables. Hence, solving (ED_{hk}) efficiently—or, more to the point, solving *all* the $O(n^2)$ of them efficiently—is crucial. Indeed, the main contribution of [16] is the proposal of an efficient dynamic programming algorithm that can solve all $O(n^2)$ Economic Dispatch problems in $O(n^3)$ for cost functions whose unconstrained minimization is $O(1)$, i.e., given by a closed formula, such as the quadratic ones. The algorithm is based on the fact that, defining $z_{hk}(\bar{p})$ to be the optimal solution value of (ED_{hk}) if the further constraint $p_k = \bar{p}$ is imposed, one has a constructive way to compute $z_{h(k+1)}$ given z_{hk} . In particular, one can prove [16] the following fact:

Proposition 4.1. *The function z_{hk} is finite-valued only in $v + 1$ intervals $[m_0, m_1], [m_1, m_2] \dots [m_v, m_{v+1}]$, with $v \leq 2(k - h)$, in which*

$$z_{hk}(\bar{p}) = z^j(\bar{p}) \quad \text{if } \bar{p} \in [m_i, m_{i+1}] \text{ for } i = 0, \dots, v$$

where each function z^j is the sum of at most $k - h + 1$ functions c_t for $t \in \{h, h + 1, \dots, k\}$.

The proof of the above proposition is constructive, and provides an algorithm for explicitly computing $z_{h(k+1)}$ given the piecewise representation of z_{hk} (the case z_{hh} is trivial). The overall complexity of the procedure depends on the actual form of the functions c_t ; for quadratic functions, all the problems $(ED_{hh}), (ED_{h(h+1)}), \dots, (ED_{hk})$ can be solved in $O((k - h)^2)$, hence solving all the $O(n^2)$ (ED) problems is $O(n^3)$. Note that this complexity is only for computing the optimal objective function values; however, the optimal solutions of all the (ED_{hk}) problems corresponding to all nodes in the optimal path are also needed. Those solutions are easily found with a “backward pass”, using the available information constructed in the “forward pass”: the (both primal and dual) optimal solution of any (ED_{hk}) can be found in $O(k - h)$, so computing the optimal solution to (1UC), in terms of the power variables, after that the “upper” dynamic programming procedure is computed, is $O(n)$. Thus, with the approach of [16] one can solve (1UC), for the quadratic case, in $O(n^3)$ overall.

5. The combinatorial heuristic

Here we briefly recall the combinatorial heuristic, proposed in [8], which is ran at each step of the iterative method used to solve (13). The heuristic uses as starting points the current value of the Lagrangian multipliers $\bar{\lambda}$, the corresponding optimal solution $\bar{s} = [\bar{\mathbf{p}}, \bar{\mathbf{u}}, \bar{\mathbf{q}}]$ of (12)—in which $\bar{\mathbf{p}}$ is integral and constraints (2)—(10) are satisfied, but (11) are *not*—and the “convexified” solution $\tilde{s} = [\tilde{\mathbf{p}}, \tilde{\mathbf{u}}, \tilde{\mathbf{q}}]$ produced by the bundle approach, where $\tilde{\mathbf{u}}$ is not integral, constraints (2)—(10) are satisfied and constraints (11) are typically almost satisfied. Then, the following three steps are performed:

- i) the “convexified” hydro power production $\tilde{\mathbf{q}}$ is considered as fixed, and the total power demand is decreased accordingly;
- ii) a *greedy-type* heuristic is used to set a commitment status $\hat{\mathbf{u}}$ of the thermal units in order to try to guarantee that the remaining power demand can be satisfied;

- iii) finally, the actual power production $[\hat{\mathbf{p}}, \hat{\mathbf{q}}]$ of thermal and hydro units is determined by solving the *Economic Dispatch* problem (1)—(11) for the fixed value $\hat{\mathbf{u}}$ of the commitment variables; note that (ED) is large-scale convex quadratic program, hence theoretically “easy” but still relatively costly to solve in practice.

This heuristic is motivated by the fact that adjusting the commitment status of thermal units is relatively simple because the commitment decision at time t directly impacts only commitment decisions in a small set of time instants centered on t , depending on the minimum up- and down-time constraints, while changing the power output of some hydro units at a certain time instant potentially impacts the hydro power output of the units in all the time horizon. However, once the combinatorial decisions have been taken, the remaining continuous problem (which, however, is not guaranteed to have a feasible solution) is “easy”. In particular, the greedy heuristic at step (ii) checks for each timestamp t whether the residual demand

$$\tilde{d}_t = \bar{d}_t - \sum_{h \in H} \sum_{j \in H(h)} \alpha^j \hat{q}_t^j$$

can be satisfied by the active thermal units in the integral solution $\bar{\mathbf{u}}$ by simply checking that it belongs to the range $[\bar{u}_t^-, \bar{u}_t^+]$, where

$$\bar{u}_t^- = \sum_{i \in P} \bar{p}_{min}^i \bar{u}_t^i \quad \bar{u}_t^+ = \sum_{i \in P} \bar{p}_{max}^i \bar{u}_t^i .$$

If $\tilde{d}_t > \bar{u}_t^+$, then the timestamp t is said *undercommitted*, while if $\tilde{d}_t < \bar{u}_t^-$ it is said *overcommitted*; in either case, the solution $\bar{\mathbf{u}}$ has to be modified by turning some units on or off at t . For this purpose, a priority list of units is formed to decide which ones are more “promising” at any given time instant; the list is based on a combination of the Lagrangian cost of turning on the unit and on the “convexified” commitment status \tilde{u}_t^i of the unit, interpreted as a “probability” that the unit i should be on at timestamp t in the optimal solution.

We remark that this is only one of the several possible implementations of the algorithmic scheme; several others are presented and compared in [8], showing that the one just described is usually the best. Furthermore, the heuristic has been developed for the non-ramp-constrained case: the definition of \bar{u}_t^- and \bar{u}_t^+ does *not* take into account the ramping constraints, and therefore may trick the heuristic into concluding that a timestamp is “feasible” while in actuality it is not because, due to ramping, the maximum (minimum) amount of power that can actually be produced in t , given the chosen commitment, is smaller (larger) than \bar{u}_t^+ (\bar{u}_t^-).

It would be possible (although nontrivial) to try to adapt the heuristic to take into account somewhat ramp rates. Furthermore, one may go the route of [6, 10] and avoid the heuristic altogether; finally, the two approaches may be combined. However, the purpose of this paper is to show that the information about ramp rates “embedded” into \bar{s} and \bar{s} is enough to provide the heuristic, with no modification, enough guidance to produce (good quality) feasible solutions.

6. Computational Experiences

In this section we present some numerical results aimed at showing the efficiency and the effectiveness of the proposed approach. Our algorithm has been coded within a C++ commercial code, `PowerSched`. The code uses custom-written routines for all tasks except the solution of the ED—a large-scale, linearly constrained convex quadratic program—in step (iii) of the

heuristic, for which `Mosek 3.2` was used. Apart from comparing two versions of the approach, identical except for the fact that one uses the ramp-constrained DP of [16] and the other uses the “classical” DP disregarding ramps, we also compared the LR approach with using state-of-the-art, general-purpose, MIQP technology. For this, we passed formulation (1)—(11) to the MIQP solver `Cplex 9.1`. All the results have been obtained on an Opteron 246 (2Ghz) computer with 2Gb RAM, running Linux Fedora Core 3.

For our tests, we have randomly generated several sets of realistic pure thermal and hydro-thermal instances, with number of thermal units ranging from 20 to 200 and number of hydro units ranging from 10 to 100. 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. All the instances have *time-invariant start-up costs*; this is a “worst case” situation for the new approach, in that the new DP procedure requires $O(n^3)$ regardless to the fact that start-up costs are time-dependent or time-invariant, while the “classical” DP is $O(n^2)$ in the former case, but only $O(n)$ in the latter. Furthermore, introducing time-dependent startup costs in the MIQP formulation increases the number of continuous variables and constraints (although not the number of binary variables [22]) in the model, thereby making it more difficult to solve, while—apart from the possible effect on the DP procedure—time-dependent startup costs are basically handled for free in the LR approach. The (UC) instances are freely available at

<http://www.di.unipi.it/optimize/Data>

We first report results about pure thermal problems in Table 1. Each row of the table reports averaged results of 5 instances of the same size (number of generating units) “ $p = |P|$ ”, on a daily problem ($n = 24$). Columns “RCDP” report results for the LR approach using the ramp-constrained DP, while columns “UDP” report results for the LR approach using the standard unconstrained DP. In both cases, column “time” reports the required running time (in seconds), column “iter” reports the total number of iterations of the bundle method used to solve (13)—that is, the number of times that the heuristic is attempted—, column “sol” reports the total number of feasible solutions found and column “gap” reports the obtained gap

$$\frac{\text{best feasible solution} - \text{best lower bound}}{\text{best lower bound}}$$

(in percentage); the number in parenthesis next to the gap, if any, is the number of instances in that group for which no feasible solution at all has been found, so that the reported gap is the average among those for which at least a solution was found. Note that we have used the best overall lower bound for computing the gap, as opposed to the best lower bound produced by each single algorithm, so that any difference in gaps is only due to the quality of the obtained feasible solutions. For UDP, column “ Δlb ” reports the gap (in percentage) between the obtained lower bound and that of RCDP; the latter is always higher because the Lagrangian Relaxation is “more constrained”, and the Lagrangian Dual is solved to optimality.

Columns “`Cplex`” reports results obtained by the MIQP solver `Cplex 9.1`. The code was ran with a time limit of one hour, which is already unrealistic in most production environments; however, none of the instances were solved to optimality within that time. Thus, we report in column “first” the (average) time required by the MIQP solver for finding any integer solution, in column “best” the time required for finding its best solution, and in column “eqv” the time

required for finding a solution with equivalent or better quality than the best one found by RCDP (counting one hour if this never happens). Finally, columns “gap” and “ Δlb ” have the same meaning as for UDP.

p	RCDP				UDP					Cplex				
	time	iter	sol	gap	time	iter	sol	gap	Δlb	first best	eqv	gap	Δlb	
20	8	189	34	0.44	6	202	1	11.30(3)	2.49	24	2229	858	0.29	-0.14
50	17	195	33	0.26	16	247	1	5.25(3)	1.48	249	1491	2563	0.22	0.34
75	30	206	33	0.38	22	278	1	9.25	2.38	447	1514	2195	0.10	0.44
100	46	213	21	0.48	29	285	1	8.69	2.21	940	2327	1236	0.13	0.47
150	72	277	23	0.20	54	341	1	7.66	2.31	2348	2483	3280	0.24(1)	0.42
200	134	317	67	0.06	78	369	1	8.53	2.46	3600	3600	3600	* (5)	0.44

Table 1: Results for pure thermal systems

It is clear from the table that RCDP is much more effective than UDP; it obtains much smaller gaps—always at least one order of magnitude smaller, and up to three order of magnitude smaller in the largest instances—and it is always capable of finding at least a feasible solution, while UDP fails to solve several of the smallest instances. This is clearly explained by looking at column “sol”: while RCDP finds feasible solutions at least once in roughly ten dual iterations, UDP only finds on average *one* solution during all the course of the algorithm. This is not surprising, since the heuristic in UDP has absolutely no information about ramp constraints, that only come up in the solution of the ED problem, after that the critical (binary) decisions have been taken; on the contrary, RCDP benefits from much more accurate information coming from the ramp-constrained Lagrangian subproblems. This also explains why UDP is faster (though never more than a factor of two): it is not due to the extra cost of the DP procedure, that is negligible in the context of the overall approach, but rather to the fact that UDP solves a much smaller number of ED problems, which are rather costly. Clearly, the extra time required by RCDP is very well-spent in this case, especially considering the much higher robustness of the approach; the fact that RCDP also improves the lower bound, most often to the tune of more than 2%, is also a nice extra effect of using the more sophisticated DP procedure.

As for Cplex, given enough time the MIQP approach may deliver more accurate solutions than the LR approach on smaller-size instances. For very small instances ($p = 20$) this is true both for the lower bound and for the upper bound, while for instances with p up to 100 the lower bound is always dominated by the Lagrangian one, while the upper bound may be better. However, this comes at a hefty cost in terms of time, as Cplex takes anywhere between 25 and 150 times the running time of RDCP to deliver solutions of the same quality (even when it succeeds in doing this); we should also remark that RDCP may—and typically does—find its best solution way before the actual termination of the approach. Furthermore, the MIQP solver requires a long time just to find *any* feasible solution (possibly of bad quality), which is unacceptable in many operating environments. Finally, for $p > 100$ it may fail to find a feasible solution at all within an hour of running time; this *always* happens for the largest instances ($p = 200$), which are, incidentally, those that are better solved by the LR approach.

Results for hydrothermal instances are reported in Table 2; the meaning of the rows and columns is the same as for pure thermal instances, except that column “ h ” reports the total number of hydro units. Due to the previous results, we have avoided to report the results of Cplex, since the much larger size of the formulation (for the same number of thermal units)

make it even less efficient at solving these instances.

p	h	RCDP			UDP					
		time	iter	sol	gap	time	iter	sol	gap	Δlb
20	10	16	162	159	0.22	7	206	3	3.80	1.50
50	20	41	165	146	0.07	16	231	6	0.63	1.19
75	35	89	209	166	0.02	28	274	5	1.73	1.19
100	50	135	218	143	0.04	38	301	1	1.86	1.27
150	75	222	223	164	0.01	71	318	1	4.10(1)	1.20
200	100	353	244	192	0.05	90	305	2	4.38	1.25

Table 2: Results for hydrothermal systems

The obtained results confirm those of the pure thermal case. Hydrothermal instances are easier to solve for LR approaches, since the flexibility of hydro units allow the system to react better to changes of the demand, possibly without requiring (too many) thermal units to switch state. This is testified by the much smaller gaps, by the much larger fraction of dual iterations that actually produce a feasible solution, and by the fact that—barring an occasional (150, 75) instance—even UDP consistently finds feasible solutions. However, the gaps obtained by RCDP are in this case *most often* better than those of UDP by three orders of magnitude, and (with a comparable number of dual iterations) the former finds feasible solutions in the hundredths, while the latter only finds a few of them. Due to the much smaller number of ED solved, UDP is also roughly three times faster than RCDP; however, the latter never requires more than 6 minutes even for the largest instances, so the running time is compatible with even the most demanding operating environments, while the obtained solutions are of excellent quality.

7. Conclusions and directions for future work

In this paper, we have proposed a Lagrangian Relaxation (LR) approach for solving large-scale hydro-thermal Unit Commitment (UC) problems with ramp constraints on the thermal generating units. The keys of the effectiveness of the approach are the efficient algorithm for Single-Unit Commitment (1UC) problems with ramp constraints recently proposed in [16], that allows to exactly solve the Lagrangian subproblems without resorting to any form of approximation, and the sophisticated heuristics for producing an integer ramp-feasible and demand-feasible solution out of the two unfeasible ones (the integer demand-unfeasible and the continuous demand-almost-feasible) computed by the LR approach. Despite the fact that the heuristic has been developed for problems without ramping constraints, and does not take them into account, solving ramp-constrained subproblems turns out to be sufficient for solving with high provable accuracy very large realistic instances in reasonable computational time on low-end hardware.

Our approach compares favorably with using general-purpose Mixed-Integer Quadratic Program solvers. Although these may obtain better solutions and/or better bounds if let ran for enough time, the LR approaches are much more efficient in producing good quality solutions, even when restricting to time-invariant start-up costs, and are the only feasible alternative for larger-size instances. Furthermore, our LR approach could be improved in at least three ways:

- by improving the heuristic in order to take into account ramp constraints directly in the logic of the greedy approach, as opposed to only indirectly by the use of ramp-constrained starting solutions;

- by using an “augmented” approach a-la [6, 10] in alternative to, or in combination with, the combinatorial heuristic;
- by embedding the whole process in a Branch&Bound approach, like that of `Cplex`, in order to refine both the upper and the lower bound at the cost of a longer running time.

A different line of research involves increasing the accuracy of the employed mathematical models of hydro and thermal units to improve the real-life quality of the obtained solution, for instance by taking into account *reserve constraints*, i.e., reserving spare capacity on the active thermal units in order to be able to cope with the inherent uncertainty of the load and system failures. These constraints are necessary in a monopolistic context, and may be useful in a liberalized regime since extra available capacity is a valuable commodity that is normally traded in a separated market. Incorporating reserve constraints in the “classical” (UC) problem is usually not difficult; indeed, they are often considered in the literature [8]. However, extending our approach to reserve-constrained versions of (UC) is not straightforward, and it will be the subject of a future work.

References

- [1] Arroyo, J.M. and Conejo, A.J. Optimal response of a thermal unit to an electricity spot market. *IEEE Transactions on Power Systems*, 15(3):1098–1104, 2000.
- [2] Bacaud, L., Lemaréchal, C., Renaud, A., and Sagastizábal, C. Bundle Methods in Stochastic Optimal Power Management: A Disaggregated Approach Using Preconditioners. *Computational Optimization and Applications*, 20:227–244, 2001.
- [3] Bannister, C.H. and Kaye, R.J. A rapid method for optimization of linear systems with storage. *Operations Research*, 39(2):220–232, 1991.
- [4] Bard, J.F. Short-term scheduling of thermal-electric generators using lagrangian relaxation. *Operations Research*, 36(5):765–766, 1988.
- [5] Bechert, T.E. and Kwatny, H.G. On the optimal dynamic dispatch of real power. *IEEE Transactions on Power Apparatus and Systems*, PAS-91(1):889–898, 1972.
- [6] Belloni, A., Diniz, A., Maceira, M.E., and Sagastizábal, C. Bundle relaxation and primal recovery in unit commitment problems. The brazilian case. *Annals of Operations Research*, 120:21–44, 2003.
- [7] Borghetti, A., Frangioni, A., Lacalandra, F., Lodi, A., Martello, S., Nucci, C.A., and Trebbi, A. Lagrangian Relaxation and Tabu Search Approaches for the Unit Commitment Problem. In Saraiva, J.T. and Matos, M.A., editors, *Proceedings IEEE 2001 Power Tech Porto Conference*, 2001. Paper n. PSO5-397.
- [8] Borghetti, A., Frangioni, A., Lacalandra, F., and Nucci, C.A. Lagrangian Heuristics Based on Disaggregated Bundle Methods for Hydrothermal Unit Commitment. *IEEE Transactions on Power Systems*, 18:313–323, 2003.
- [9] Borghetti, A., Frangioni, A., Lacalandra, F., Nucci, C.A., and Pelacchi, P. Using of a cost-based Unit Commitment algorithm to assist bidding strategy decisions. In Borghetti, A., Nucci, C.A., and Paolone, M., editors, *Proceedings IEEE 2003 Power Tech Bologna Conference*, 2003. Paper n. 547.

- [10] Dubost, L., Gonzalez, R., and Lemaréchal, C. A primal-proximal heuristic applied to the French Unit-commitment problem. *Mathematical Programming*, 104(1):129–151, 2005.
- [11] Fan, W., Guan, X., and Zhai, Q. A new method for unit commitment with ramping constraints. *Electric Power Systems Research*, 62:215–224, 2002.
- [12] Feltenmark, S. and Kiwiel, K.C. Dual Applications of Proximal Bundle Methods, Including Lagrangian Relaxation of NonConvex Problems. *SIAM Journal on Optimization*, 10(3):697–721, 2000.
- [13] A. Frangioni and C. Gentile. Perspective cuts for a class of convex 0–1 mixed integer programs. *Mathematical Programming*, 106(2):225 – 236, 2006.
- [14] Frangioni, A. Generalized Bundle Methods. *SIAM Journal on Optimization*, 13(1):117–156, 2002.
- [15] Frangioni, A. About Lagrangian Methods in Integer Optimization. *Annals of Operations Research*, 139:163–193, 2005.
- [16] Frangioni, A. and Gentile, C. Solving Nonlinear Single-Unit Commitment Problems with Ramping Constraints. *Operations Research*, 54(4), to appear, 2006.
- [17] Hobbs, B.F., Rothkopf, M., O’Neill, R.P., and Chao, H.P. *The next generation of unit commitment models*. Kluwer Academic Press, Boston, 2001.
- [18] Lai, S.-Y. and Baldick, R. Unit Commitment with Ramp Multipliers. *IEEE Transactions on Power Systems*, 14(1):58–64, 1999.
- [19] Li, Z. and Shahidehpour, M. Generation Scheduling With Thermal Stress Constraints. *IEEE Transactions on Power Systems*, 18(4):1402–1409, 2003.
- [20] Lu, B. and Shahidehpour, M. Short-term scheduling of combined cycle units. *IEEE Transactions on Power Systems*, 19(3):1616–1625, 2004.
- [21] Madrigal, M. and Quintana, V.H. An Interior-Point/Cutting-Plane Method to Solve Unit Commitment Problems. *IEEE Transactions on Power Systems*, 15(3):1022–1027, 2000.
- [22] Nowak, M.P. and Römisich, W. Stochastic Lagrangian relaxation applied to power scheduling in a hydro-thermal system under uncertainty. *Annals of Operations Research*, 100:251–272, 2000.
- [23] Travers, D.L. and Kaye, R.J. Dynamic dispatch by constructive dynamic programming. *IEEE Transactions on Power Systems*, 13:72–78, 1998.
- [24] Wood, A.J. and Wollemborg, B.F. *Power generation operation and control*. John Wiley & Sons, 1996.
- [25] Zhuang, F. and Galiana, F.D. Towards a more rigorous and practical unit commitment by Lagrangian relaxation. *IEEE Transactions on Power Systems*, 3:763–773, 1988.