New Lagrangian Heuristics for Ramp-constrained Unit Commitment Problems*

Antonio Frangioni §　　Claudio Gentile ♯

Fabrizio Lacalandra ¶

June 21, 2006

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 can be efficiently solved by Dynamic Programming (DP) techniques. Ramp constraints have historically eluded efficient exact DP approaches; however, this has recently changed [11]. We show that the newly proposed DP algorithm for ramp-constrained (1UC) problems, together with a new heuristic phase that explicitly takes into account ramp limits on the maximum and minimum available power at each hour, can reliably find good-quality solutions to even large-scale (UC) instances in short time.

Keywords: *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

*this paper has been partially supported by the Italian Government PRIN project 2005015592 "Decision Support Models for the Operators of the Italian Electricity Market and their Effects on the System Security", subproject "Optimization Problems in Unit Commitment and in the definition of the energy plan" Research Unit of University of Bologna.

§Università di Pisa, Dipartimento di Informatica, Largo B. Pontecorvo 3, 56127 Pisa – Italy, e-mail: `frangio@di.unipi.it`

♯Istituto di Analisi dei Sistemi ed Informatica "Antonio Ruberti" del C.N.R., Viale Manzoni 30, 00185 Rome – Italy, e-mail: `gentile@iasi.cnr.it`

¶MBI S.r.L., e-mail: `lacaland@di.unipi.it`

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 [15], which is of increasing importance in the ongoing liberalization of the electricity market in many countries [12]. 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 [12, 6].

Despite having attracted the interest of researchers for over 30 years, (UC) cannot be considered yet 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 optimization problem. Among the most efficient algorithmic approaches for (UC), Lagrangian Relaxation (LR) methods [1, 4, 5, 8, 14, 16] 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 either network-flow techniques—for hydro units—or Dynamic Programming (DP) techniques—for thermal units.

However, Lagrangian approaches critically depend on the ability of optimally solving the (1UC) problems efficiently, which in turn depends on the specific details of the operational constraints of the generating units that are represented in the mathematical model. 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. This assumption has been motivated by the fact that ramp constraints make the classical DP approaches for thermal (1UC) problems unusable, since they link together variables representing the power output at different timestamps, which are then no longer independent once that commitment (on/off) decision has been taken. Previous attempts to address this problem have used discretization of the power variables space [3, 2], piecewise-linear approximation of the cost function [7] or two-stage Lagrangian techniques [13] where ramp constraints are dualized; however, all these approaches increase the computational burden as the level of approximation is decreased, and do not obtain, in general, optimal solutions to the subproblem. Recently, a DP algorithm for thermal (1UC) with ramping constraints has been proposed [11] that can solve to optimality problems on a time horizon of n timestamps in $O(n^3)$ for "simple" convex cost functions, such as the quadratic ones typically used in operational settings. In this paper, we report on the use of that algorithm within LR approaches to (UC). We show that the sole use of ramp-constrained subproblems within existing Lagrangian approaches already allows to find good-quality solutions for most instances; coupled with a new heuristic phase that explicitly takes into account ramp limits on the maximum and minimum available power at each hour, the approach can reliably solve even "difficult" ramp-constrained (UC) instances in short time.

2 Lagrangian approaches to UC

Consider a set P of thermal units and a set H of hydro cascades, each comprising one or more basin units, over a discretized time horizon $\mathcal{T} = \{1, \ldots, n\}$. 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 $\sum_{i \in P} c^i(\mathbf{p}^i, \mathbf{u}^i) = \sum_{i \in P} \left(c^i(\mathbf{u}^i) + \sum_{t \in \mathcal{T}} c_t^i(p_t^i) \right)$; 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, e.g., see [15]) 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. For the sake of compactness of the presentation we do not explicitly report all them; the interested reader is referred to [5, 11] for further details. For system wide constraints, we only consider demand constraints (see [12, 6] for applications)

$$\sum_{i \in P} p_t^i + \sum_{h \in H} \sum_{j \in H(h)} q_t^j \geq \bar{d}_t \qquad \text{for each } t \in \mathcal{T},$$

where $\bar{d}_t$ is the forecasted load to be satisfied, $\mathbf{q}^j$ are the power production variables for the hydro plant j, and $H(h)$ is the set of hydro plants belonging to cascade h.

We will denote the feasible set of a given thermal unit $i \in P$ as $\mathcal{U}^i$, and the feasible set for a given hydro cascade $h \in H$ by $\mathcal{H}^h$. Then the *Lagrangian Relaxation* of (UC) is obtained by dualizing each of the global ("coupling") constraint via a Lagrangian multiplier λ_t, obtaining

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

where

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

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 finding the $\boldsymbol{\lambda}^*$ which gives the best (greatest) lower bound, i.e., in the optimal solution of the *Lagrangian Dual*:

$$\max\{\ \mathcal{L}(\boldsymbol{\lambda})\ :\ \boldsymbol{\lambda} \in \mathbb{R}^n\ \} \quad . \tag{2}$$

Since $\mathcal{L}(\cdot)$ is a non differentiable function, proper algorithms must be chosen for solving (2); *bundle* methods [9], particularly in their *disaggregated* variant [1], have been repeatedly reported to be quite efficient in solving (2), much more so [5] than alternative algorithms such as subgradient methods [2, 16].

However, solving (2) is not, in general, enough to solve (UC); even for $\boldsymbol{\lambda} = \boldsymbol{\lambda}^*$, the optimal solution to (1) is not guaranteed to—and will not in general—satisfy the relaxed constraints. Therefore, a number of *Lagrangian heuristics* [10] have been devised [1, 2, 5, 8, 16] that attempt to convert the (sequence of) unfeasible solution(s) provided by (2) in a (sequence of) feasible, hopefully "good", one(s). This kind of approach has proven to be capable of solving non ramp-constrained large-scale (UC) problems within a provable high degree of accuracy in relatively short time. However, extending the approach to ramp constraints requires, in principle, two 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^i_t variables.

Of those, (i) has only been made possible—without approximations or an excessive growth of the computational burden—by the development in [11]. In this paper we report about point (ii), showing that even the heuristics originally developed for (UC) without ramp constraints are efficient in the ramp-constrained case also, only provided that ramp-constrained (1UC)s are solved. Furthermore, a relatively simple modification of the heuristic allows to tackle even "difficult" ramp-constrained instances which are not solvable by the standard approach.

3 The combinatorial heuristic

Here we briefly recall the combinatorial heuristic, proposed in [5], which is run at each step of the iterative method used to solve (2). The heuristic uses as starting points the current value of the Lagrangian multipliers $\bar{\boldsymbol{\lambda}}$, the corresponding (integer) optimal solution $\bar{s} = [\bar{\mathbf{p}}, \bar{\mathbf{u}}, \bar{\mathbf{q}}]$ of (1) which violates the system-wide constraints, and the "convexified" solution $\tilde{s} = [\tilde{\mathbf{p}}, \tilde{\mathbf{u}}, \tilde{\mathbf{q}}]$ produced by the bundle approach [10], where $\tilde{\mathbf{u}}$ is not integral but system-wide constraints 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* 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) the actual power production $[\hat{\mathbf{p}}, \hat{\mathbf{q}}]$ of thermal and hydro units is determined by solving an *Economic Dispatch* problem, that is, the quadratic programming problem resulting from (UC) after having fixed values $\hat{\mathbf{u}}$ for the commitment variables.

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 (depending on the minimum up- and down-time constraints) of time instants centered on t, 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) initially sets $\hat{\mathbf{p}} = \bar{\mathbf{p}}$, and then checks for each timestamp t whether the residual demand

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

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

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

and $\bar{p}_{min}^i$, $\bar{p}_{max}^i$ are respectively the minimum and maximum power production of unit i (if committed). 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.

Clearly, 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 actually it is not because, due to ramping, the maximum (minimum) amount of power that can in fact be produced in t, given the chosen commitment, is smaller (larger) than $\bar{u}_t^+$ ($\bar{u}_t^-$). Our computational results show that this does not happen too often; this is due to the fact that both starting solutions $\bar{s}$ and $\tilde{s}$ are actually ramp-feasible, thus they "embed" enough information about ramp rates to provide the heuristic enough guidance to produce (good quality) feasible solutions.

However, the heuristic may be made more "conservative", hence more capable of finding solutions in instances where ramp constraints are very tight, by *explicitly* exploiting information about ramp constraints to determine $\bar{u}_t^-$ and $\bar{u}_t^+$. This is actually very easy at the first timestamp ($t = 1$), since the thermal power level of each unit prior the beginning of the time horizon ($\breve{p}_0^i$) is among the input data of the (UC) instance. Thus, for $t = 1$ an obvious improvement on the above formula is given by

$$\bar{u}_t^+ = \sum_{i \in P} \min(\ \bar{p}_{max}^i\ ,\ \breve{p}_{t-1}^i + \Delta_+^i\ )\hat{u}_t^i \tag{3}$$

where Δ_+^i is the maximum ramp-up rate for unit i; clearly, an analogous formula holds for $\bar{u}_t^-$, using the maximum ramp-down rate Δ_-^i. However, this does not immediately extend to case $t > 1$ because $\hat{p}_{t-1}^i$—which should be used as $\breve{p}_{t-1}^i$ in (3) to obtain a better $\bar{u}_t^+$—is not known, but rather a *final* result of the heuristic, only available *after* that commitment variables have been fixed (*using* $\bar{u}_t^-$ and $\bar{u}_t^+$).

A possible approach here is to *arbitrarily fix* some tentative values of the p_t^i variables, and use them as $\breve{p}_t^i$ in (3) to compute "more feasible" estimates of the maximum and minimum available power at each timestamp corresponding to currently committed units. Because the greedy heuristic (step (ii)) is "forward", i.e., fixes the u_t^i variables in order $t = 1, 2, \ldots, n$, we choose to fix the $\breve{p}_t^i$ values iteratively in the same order. Among several other possibilities, one can, e.g., chose $\breve{p}_t^i$ as the optimal solution of the following (convex) Quadratic Continuous Knapsack problem

$$\min \sum_{i \in \mathcal{Z}_t} c_t^i(p_t^i) \tag{4}$$

$$\max(\bar{p}_{min}^i, \hat{p}_{t-1}^i - \Delta_-^i) \leq p_t^i \leq \min(\bar{p}_{max}^i, \hat{p}_{t-1}^i + \Delta_+^i) \qquad i \in \mathcal{Z} \tag{5}$$

$$\sum_{i \in \mathcal{Z}_t} p_t^i = \tilde{d}_t \tag{6}$$

where $\mathcal{Z}_t$ is the set of the units currently committed at timestamp t (i.e., such that $\hat{u}_t^i = 1$). Actually, the formula should be slightly modified to account for units that have just been switched on at timestamp t (i.e., such that $\breve{p}_{t-1}^i = 0$); the modification is trivial and we omit it.

Problem (4)—(6) can be easily solved in $O(k \log k)$ ($k = |\mathcal{Z}_t| \leq n$) by a simple dual-based procedure, which is also capable to detect whether it is empty—and therefore t is (possibly) either overcommitted or undercommitted; this information is in fact used instead of $\bar{u}_t^-$ and $\bar{u}_t^+$ to decide whether the current commitment $\hat{u}_t^i$ need to be changed. Contrarily to the original procedure, this one may be "too conservative" by declaring a timestamp over- or undercommitted even if it could actually be possible to satisfy the demand; this is due to the arbitrary choice of the $\check{\mathbf{p}}$ among an actually much larger set of possibilities. Note that, actually, even the original heuristic may turn out to be too conservative for an analogous reason: the residual demand $\tilde{d}_t$ is computed using the (arbitrarily) fixed power production $\tilde{\mathbf{q}}$, so even if the procedure finds a "terminally over- or undercommitted" timestamp (one for which no units can be turned on or off to satisfy the demand), it is not necessarily true that the commitment $\hat{\mathbf{u}}$ is unfeasible. All this calls for a nontrivial combination of the heuristics, in order to exploit their respective strengths and weaknesses.

4 Computational Experiences

In this section we present some preliminary numerical results aimed at showing the efficiency and the effectiveness of the proposed approach. Our algorithm has been coded within a C++ commercial code, $\mathbb{P}$**ower**$\mathbb{S}$**ched**$\mathbb{O}$. We compared three versions of the code for solving ramp-constrained UC problems:

- a version using the "classical" DP disregarding ramps and using the "standard" combinatorial heuristic (basically, the code of [6]), only including ramp constraints in the ED (V1);
- a version using the ramp-constrained DP of [11] but still using the "standard" combinatorial heuristic (V2);
- a version using the ramp-constrained DP of [11] and using the new combinatorial heuristic (V3)

For our tests, we have randomly generated several sets of realistic pure thermal instances with a number of units ranging from 20 to 200. All the instances have *time-invariant start-up costs*; this is a "worst case" situation for the ramp-constrained DP, in that it 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.

Results are summarized in Table 1. Each row of the table reports averaged results of instances of the same size (number of generating units) "$p = |P|$", on a daily problem ($n = 24$). For each of the three variants, column "time" reports the required running time (in seconds), column "iter" reports the total number of iterations of the bundle method used to solve (2), and column "gap" reports the obtained gap (in percentage); the number in parenthesis next to the gap, if any, is the percentage 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.

It is clear from the table that using the ramp-constrained DP is of paramount importance for obtaining a successful overall Lagrangian approach; for small sizes V1 cannot even find one solution for many small instances, and even if it regularly does as the size increases, the obtained gaps are very large. From this preliminary results, the new heuristic (V3) appears to improve, albeit sometimes

p	V1 time	V1 iter	V1 gap	V2 time	V2 iter	V2 gap	V3 time	V3 iter	V3 gap
10	3.55	193	11.94(50)	5.23	190	1.77	6.32	190	1.50
20	5.85	212	15.40(33)	6.16	176	1.20	8.99	176	1.08
50	13.95	257	6.58(66)	13.73	206	0.48	21.69	206	0.42
75	18.52	291	13.26	24.66	217	0.80	37.64	217	0.58
100	21.64	256	12.60	55.67	193	0.48	29.94	193	0.37
150	48.30	378	8.94	52.90	280	4.13	77.17	280	4.13
200	63.68	372	10.41	134.04	304	0.12	96.83	304	0.09

Table 1: Results for pure thermal systems

only slightly, over the "standard" heuristic (V2); the only exception are instances with $p = 150$, which appear surprisingly "easy" for V1 and "difficult" for V2 and V3, which are equivalent. The exact reason of this behavior is as yet unknown, and further experiments are required to validate these results. Also, note that running times are somewhat different between V2 and V3 (but with no clear dominance), mainly because of the different number of EDs solved.

5 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 [11], 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. Our preliminary computational results show that the proposed heuristic is capable of efficiently solving with high provable accuracy very large, realistic, ramp-constrained instances in reasonable computational time on low-end hardware.

References

[1] 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.

[2] Bard, J.F. Short-term scheduling of thermal-electric generators using lagrangian relaxation. *Operations Research*, 36(5):765–766, 1988.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] Fan, W., Guan, X., and Zhai, Q. A new method for unit commitment with ramping constraints. *Electric Power Systems Research*, 62:215–224, 2002.

[8] 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.

[9] Frangioni, A. Generalized Bundle Methods. *SIAM Journal on Optimization*, 13(1):117–156, 2002.

[10] Frangioni, A. About Lagrangian Methods in Integer Optimization. *Annals of Operations Research*, 139:163–193, 2005.

[11] Frangioni, A. and Gentile, C. Solving Nonlinear Single-Unit Commitment Problems with Ramping Constraints. *Operations Research*, to appear, 2006.

[12] Hobbs, B.F., Rothkopf, M., ONeill, R.P., and Chao, H.P. *The next generation of unit commitment models*. Kluwer Academic Press, Boston, 2001.

[13] Lai, S.-Y. and Baldick, R. Unit Commitment with Ramp Multipliers. *IEEE Transactions on Power Systems*, 14(1):58–64, 1999.

[14] Madrigal, M. and Quintana, V.H. An Interior-Point/Cutting-Plane Method to Solve Unit Commitment Problems. In *Proceedings IEEE-PES Power Industry Computer Applications Conference*, pages 179–185, 1999.

[15] Wood, A.J. and Wollemberg, B.F. *Power generation operation and control*. John Wiley & Sons, 1996.

[16] 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.