ISTITUTO DI ANALISI DEI SISTEMI ED INFORMATICA

CONSIGLIO NAZIONALE DELLE RICERCHE

A. Frangioni, C. Gentile, F. Lacalandra

TIGHTER APPROXIMATED MILP FORMULATIONS FOR UNIT COMMITMENT PROBLEMS

R. 660   Aprile 2007

Antonio Frangioni − Dipartimento di Informatica, Università di Pisa, Largo B. Pontecorvo 3 - 56127 Pisa, Italy. Email: frangio@di.unipi.it.

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

Fabrizio Lacalandra − 𝕆*pti*𝕄𝔼 , c/o MBI, via Scornigiana 7, 56121 Loc. Ospedaletto - Pisa, e-mail: fabrizio.lacalandra@fastwebnet.it.

Abstract

The short-term Unit Commitment (UC) problem in hydro-thermal power generation is a large-scale, Mixed-Integer NonLinear Program (MINLP), which is difficult to solve efficiently, especially for large-scale instances. It is possible to approximate the nonlinear objective function of the problem by means of piecewise-linear functions, so that UC can be approximated by a Mixed-Integer Linear Program (MILP); applying the available efficient general-purpose MILP solvers to the resulting formulations, good quality solutions can be obtained in a relatively short amount of time. We build on this approach, presenting a novel way to approximating the nonlinear objective function based on a recently developed class of *valid inequalities* for the problem, called "Perspective Cuts". At least for many realistic instances of a general basic formulation of UC, a MILP-based heuristic obtains comparable or slightly better solutions in less time when employing the new approach rather than the standard piecewise linearizations, while being not more difficult to implement and use. Furthermore, "dynamic" formulations, whereby the approximation is iteratively improved, provide even better results if the approximation is appropriately controlled.

Key words: Hydro-Thermal Unit Commitment, Mixed-Integer Linear Program Formulations, Valid Inequalities.

Nomenclature

The notation used throughout the paper is stated below. For unit consistency, note that hourly intervals are considered.

Constants:

n	number of time intervals (hours)
$\mathcal{T}$	set of all time periods
P	set of thermal units
H	set of hydro cascades, each comprising one or more basin units
$H(h)$	set of individual hydro units cascade $h \in H$ is composed of
$\mathcal{B}(j)$	set of the immediate predecessors of hydro unit j
t_{kj}	water time delay from plant $k \in \mathcal{B}(j)$ to the basin feeding hydro unit j
$\bar{v}_{min}^{j}$	minimum volume for the reservoir of hydro unit j [m^3]
$\bar{v}_{max}^{j}$	maximum volume for the reservoir of hydro unit j [m^3]
$\bar{q}_{max}^{j}$	technical maximum of discharged water of hydro unit j [m^3] (the technical minimum is assumed to be zero)
$\bar{w}_{t}^{j}$	natural inflows of the reservoir of hydro unit j at time period t [m^3]
τ_{+}^{i}	minimum up-time of thermal unit i [hours]
τ_{-}^{i}	minimum down-time of thermal unit i [hours]
Δ_{+}^{i}	maximum ramp-up rate of thermal unit i [MW/hours]
Δ_{-}^{i}	maximum ramp-down rate of thermal unit i [MW/hours]
$\bar{p}_{min}^{i}$	minimum power output of thermal unit i when operating in steady state [MW]
$\bar{p}_{max}^{i}$	maximum power output of thermal unit i when operating in steady state [MW]
$\bar{l}^{i}$	maximum power output of thermal unit i at the first hour of a commitment period, i.e., if the unit was off the previous hour [MW] (also referred as *startup ramp limit*)
$\bar{u}^{i}$	maximum power output of thermal unit i at the last hour of a commitment period, i.e., if the unit is going to be off the next hour [MW] (also referred as *shutdown ramp limit*)
a_{t}^{i}	quadratic term of power cost function of thermal unit i at period t [€ / MW^2]
b_{t}^{i}	linear term of power cost function of thermal unit i at period t [€ / MW]
c_{t}^{i}	constant term of power cost function of thermal unit i at period t [€]
$\bar{d}_{t}$	forecasted load to be satisfied at period t [MWh]
α^{j}	power-to-discharged-water efficiency of hydro unit j [MWh / m^3]

Variables:

u_{t}^{i}	status of thermal unit i at period t [0 / 1]
p_{t}^{i}	power output of thermal unit i at end of period t [MW]
q_{t}^{j}	discharged water of hydro unit j at time period t [m^3]
v_{t}^{j}	volume of the reservoir of hydro unit j at time period t [m^3]
w_{t}^{j}	spilled water of hydro unit j at time period t [m^3]
z_{t}^{i}	auxiliary function for expressing the objective function cost of thermal unit i at time period t [€]

Functions:

$s^{i}(\mathbf{u}^{i})$	start-up costs of unit i, possibly time-dependent

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 [16], which is of increasing importance in the ongoing liberalization of the electricity market in many countries [13]. 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, 13, 6].

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). In spite of 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 [10] and/or the length of the time horizon [11] grows large.

Recently, approximated Mixed-Integer Linear Program (MILP) formulations of UC have been proposed [7, 8, 14, 17] which exploit the efficient general-purpose available MILP solvers to compute good quality solutions in relatively small time, especially for low- to mid-size instances, although specialized approaches, e.g. based on Lagrangian Relaxation, are still competitive for very-large-scale instances and/or when very fast running times are required by the operational environment [12].

In [10], it has been shown that the efficiency and effectiveness of approaches using a MIQP solver can be consistently improved by adding to the MIQP formulation a properly chosen set of *valid inequalities* for the UC problem, called "Perspective Cuts", which "tighten" the formulation by cutting away parts of the feasible region of the continuous relaxation which do not belong to the convex hull of the integer feasible solutions. This amounts in practice to a piecewise linearization of the nonlinear part of the objective function of the problem, where the number of pieces need not to be chosen *a-priori*. While the resulting formulations are thus reminiscent of the previously-mentioned ones [7, 8, 14, 17], they differ in some relevant details, as discussed later on.

We show that, at least for one "classical" formulation of ramp-constrained hydro-thermal UC and on a set of realistic instances, MILP-based heuristics based on the new linearization obtain comparable or slightly better solutions in less time than analogous approaches using the classical linearization. This is particularly interesting in view of the fact that the new linearization is not more difficult to implement and use than the previously proposed ones, given the same underlying MILP solver. Furthermore, the new approach allows to more easily exploit the tools made available by all current MILP solvers in order to construct dynamic formulations

where the approximation is improved as needed during the solution process, leading to further improvements in the effectiveness and efficiency of the approach.

The structure of the paper is the following. In Section 2 we present the MINLP 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, 9]. In Section 3 we present the MILP approximation akin to those used, e.g., in [7, 17] and our alternative linearization based on "Perspective Cuts" [10]. Finally, in Section 4 we compare the linearizations within heuristic MILP-based approaches to UC, and we draw some conclusions.

2. The UC model

Given the constants and variables defined in the Nomenclature section, the objective function of UC, representing the total power production cost to be minimized, has the form

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

That is, the power production cost at each hour is customarily represented by a convex ($a_t^i > 0$) quadratic separable form in the power variables p_t^i, neglecting for instance the so called valve points [16]; fixed production costs are represented by the term $c_t^i u_t^i$. We do not dwell further upon the specific form of the (possibly time-dependent) *start-up costs* function $s^i(\mathbf{u}^i)$, only assuming that it can be properly represented within an MILP problem; the interested reader is referred to [15, 7] for details.

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*:

$$\bar{p}_{min}^i u_t^i \leq p_t^i \leq \bar{p}_{max}^i u_t^i \qquad\qquad t \in \mathcal{T} \tag{2}$$
$$p_t^i \leq p_{t-1}^i + u_{t-1}^i \Delta_+^i + (1 - u_{t-1}^i)\bar{l}^i \qquad\qquad t \in \mathcal{T} \tag{3}$$
$$p_{t-1}^i \leq p_t^i + u_t^i \Delta_-^i + (1 - u_t^i)\bar{u}^i \qquad\qquad t \in \mathcal{T} \tag{4}$$

$$u_t^i \geq u_r^i - u_{r-1}^i \qquad\qquad t \in \mathcal{T},\ r \in [t - \tau_+^i, t-1] \tag{5}$$
$$u_t^i \leq 1 - u_{r-1}^i + u_r^i \qquad\qquad t \in \mathcal{T},\ r \in [t - \tau_-^i, t-1] \tag{6}$$
$$u_t^i \in \{0, 1\} \qquad\qquad t \in \mathcal{T} \tag{7}$$

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. The time period "0" is used for indicating the initial conditions of the power system; 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 and its generated power p_0^i. 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*:

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

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

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

In order for the balance equations (10) to be well-defined, we assume knowledge of the volume of each reservoir at time period $t = 0$, 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 $k \in \mathcal{B}(j)$ such that $t < t_{kj}$).

- *Global constraints*: the system-wide constraints—linking the different units among themselves—are:

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

Note that the power-to-discharged-water efficiency is assumed constant, to avoid nonlinearities.

We refer to UC as the problem of minimizing (1) subject to constraints (2)—(11); this is a large-scale Mixed-Integer Nonlinear Program whose nonlinearities are all contained in the objective function. This formulation is a "basic" one, and it is less accurate than several previously proposed ones in some aspects, e.g. related to hydro units modeling [8, 14, 17] or to the so-called start-up and shut-down power trajectories [2]; however, the proposed technique can be easily adapted to many UC formulations. The only notable, although minor, characteristic of our formulation is the possibility to limiting the thermal units output to any prescribed value during the first and last hour of operations (cf. the constants $\bar{l}^i$ and $\bar{u}^i$ in (3) and (4), respectively), as first suggested in [11]. Our choice of simplifying assumptions appears to strike a balance between capturing the main aspects of practical UC problems and simplicity of the model, and is commonly accepted in the literature. For instance, spinning reserve constraints, either in the "standard" formulation (e.g. [5]) or in the more sophisticated one recently proposed in [7], could be easily included in the formulation, but they have not been used in the instances used in Section 4 since they are not likely to have any significant impact on the relative efficiency of the different approximate formulations tested in this paper. Several other 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 [8, 14, 17], could be used at the cost of more integer variables in the formulation. Analogously, valve points of thermal units [16] or cavitation points of hydro units can be easily modeled. Again, the introduction of these further elements should not impact on the relative efficiency of the different approximate formulations tested in this paper.

The UC model here considered, while having been historically motivated by the centralized decision environments prevalent in the past, is well-suited also 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, 13, 9, 6].

3. The piecewise-linear approximations

In order to make UC tractable by the efficient MILP solvers available, the nonlinear part of the objective function need be linearized. Since the nonlinear structure is identical for each time period and thermal unit, for notational simplicity in this section we consider both indices i and t fixed and we drop them. The issue is then how to best represent the quadratic objective function

$$f(p, u) = ap^2 + bp + cu \qquad (12)$$

by means of a piecewise-linear one. It is well-known that there are several different ways for doing this; one is represented in Figure 1, where $k + 1$ points $\bar{p}^0, \bar{p}^1, \ldots, \bar{p}^k$ in the interval $[\bar{p}_{min}, \bar{p}_{max}]$ are chosen (such that $\bar{p}^0 = \bar{p}_{min}$ and $\bar{p}^k = \bar{p}_{max}$), and a convex, piecewise-linear *upper* approximation of f, which coincides with the latter in the chosen points, is used to replace the original nonlinear objective function. This results in an MILP which differs from UC (for each i and t) only in the following details (let $f(p) = f(p, 1) = ap^2 + bp + c$):

- k new variables δ_l are introduced together with constraints

$$
\begin{aligned}
p &= \textstyle\sum_{l=1}^{k} \delta_l + \bar{p}_{min} u \\
0 &\leq \delta_l \leq \bar{p}^l - \bar{p}^{l-1} \qquad l = 1, \ldots, k
\end{aligned}
\qquad (13)
$$

- the cost coefficient of u in the objective function is changed to $f(\bar{p}_{min})$;

- each variable δ_l is given a linear cost F_l representing the linear function with value 0 when $\delta_l = 0$ and value $f(\bar{p}^l) - f(\bar{p}^{l-1})$ when $\delta_l = \bar{p}^l - \bar{p}^{l-1}$, i.e.,

$$F_l = \frac{f(\bar{p}^l) - f(\bar{p}^{l-1})}{\bar{p}^l - \bar{p}^{l-1}} = a(\bar{p}^l + \bar{p}^{l-1}) + b \ .$$

The MILP approximation of the quadratic function is therefore obtained by replacing (12) with

$$f(\bar{p}_{min})u + \sum_{l=1}^{k} F_l \delta_l$$

subject to the original constraints of the problem plus the extra constraints (13). We will refer to this approximated MILP formulation of UC as the Standard Piece-Wise Formulation (SPWF). There are different choices for the linearization; for instance, it is easy to construct a *lower* approximation which is tight to f in both function and derivative values in k points "in the middle" of the intervals. Most often, the articles where linearization is touted (e.g. [7, 8, 14, 17]) do not explicitly state how exactly the linearization is constructed, although sometimes this may be deduced; for instance, [7, Figure 1] most likely indicates the same upper approximation as in SPWF. For the purpose of the present paper, there is no substantial difference between an upper or a lower approximation, as discussed below. Indeed, both upper and lower approximations constructed in this way only work in the p-space; thus, when represented in the (p, u) space, as in Figure 1, one notices that the objective function of the new problem is linear along all
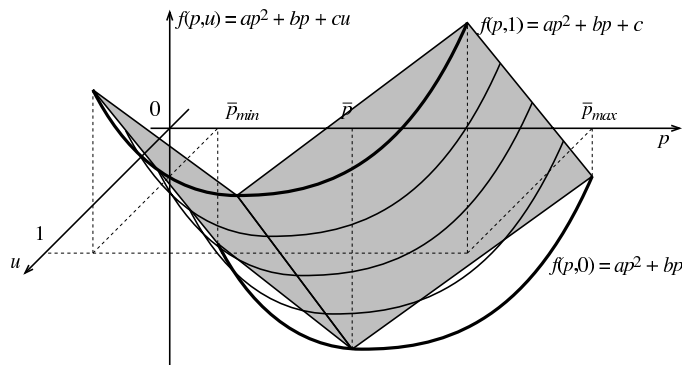
Figure 1: Piecewise-linear approximation of $f(p)$ in the (p, u) space

segments of extremes $(\bar{p}, 0)$ and $(\bar{p}, 1)$ for any feasible production level $\bar{p}$, always with the same slope c (in the Figure, $k = 2$, $\bar{p}^0 = \bar{p}_{min}$, $\bar{p}^1 = \bar{p}$, $\bar{p}^2 = \bar{p}_{max}$).

Although the previous linearization is quite natural, it arguably is not the best possible approximation of the objective function of UC; indeed, a different possibility is suggested in [10]. Arbitrarily choosing k points $\bar{p}^1, \ldots, \bar{p}^k$ in the interval $[\bar{p}_{min}, \bar{p}_{max}]$, a different way for producing an MILP which approximates UC is (for each i and t) as follows:

- each term of the form (12) is removed from the objective function and replaced with a corresponding new variable z; other terms in the objective function not containing p and u, e.g., those related to variable startup costs [15], are kept untouched;

- k constraints of the form

$$z \geq (2a\bar{p} + b)p + (c - a\bar{p}^2)u \tag{14}$$

  with $\bar{p} = \bar{p}^h$, $h = 1, \ldots, k$ respectively, are added to the formulation.

We will refer to the above as the Perspective-Cut (approximate) Formulation (PCF) of UC. This choice is justified by a sophisticated theoretical analysis which for the sake of clarity cannot be repeated here; the interested reader is referred to [10] for full details. Here we will just briefly illustrate the basic ideas underlying the construction, in order to clarify in what sense the above choice is, at least in theory, preferable to others, and what are its main differences w.r.t. the previous approach.

The function $f(p, u)$ in (12) is in principle only relevant at points (p, u) of its (disconnected) domain $\mathcal{D} = [0, 0] \cup [\bar{p}_{min}, \bar{p}_{max}] \times \{1\}$; however, standard Branch&Bound approaches typically solve the *continuous relaxation* of the provided formulation, where u is allowed to take values in $[0, 1]$ rather than $\{0, 1\}$, in order to derive lower bounds on the optimal value of the problem. It thus makes sense to study which formulation provides the best possible (workable) convex relaxation of UC.

While such a question does not admit any easy answer for the UC problem in its entirety, it can be answered if one restricts himself to the "basic blocks" of the problem; in fact, the *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 [10] to be

$$h(p, u) = ap^2/u + bp + cu \tag{15}$$

for all $u\bar{p}_{min} \leq p \leq u\bar{p}_{max}$, $u \in (0, 1]$ ($h(0, 0) = 0$, and $h(p, u) = +\infty$ anywhere else). This function is strongly related with a well-known object in convex analysis, the *perspective function*

$g(p, u) = u\, f(p/u)$ of $f(p)$. The epigraph of $g(p, u)$ defines a cone pointed in the origin and having as "lower shape" that of $f(p)$, as depicted in Figure 2; *epi h* is the section of the cone corresponding to $u \leq 1$.

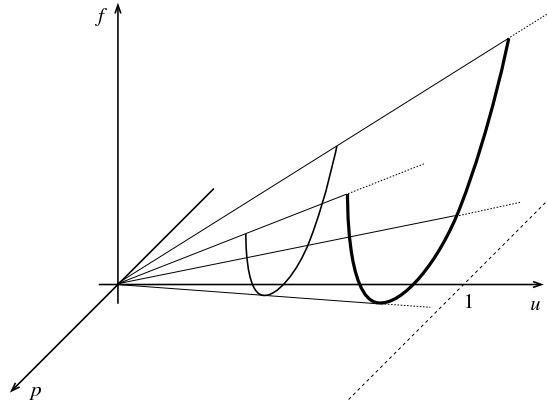

Figure 2: The perspective function of $f(p)$

Since $0 < u \leq 1$, it is immediate to verify that $h(p, u) \geq f(p, u)$ for all $(p, u) \in \mathcal{D}$, that is, h is a better objective function, for a continuous relaxation, than $f(p, u)$; indeed, elementary calculus shows that the maximum of $h(p, u) - f(p, u)$ over $\mathcal{D}$ is $a\bar{p}_{max}^2/4$, attained at $[\bar{p}_{max}/2, 1/2]$. However, using $h(p, u)$ as the objective function has a serious drawback: it is even a "more nonlinear" function than $f(p, u)$, which we already aim at at making "less nonlinear".

Yet, it is well-known that every convex function is the point-wise supremum of affine functions; for our case these can be easily characterized. Indeed, [10, Theorem 1] shows that the epigraph of h is composed of all and only triples (z, p, u) satisfying $u\bar{p}_{min} \leq p \leq u\bar{p}_{max}$, $0 \leq u \leq 1$ and the *infinite* system of linear inequalities (14), for *all* $\bar{p} \in [\bar{p}_{min}, \bar{p}_{max}]$. We refer to each inequality in (14) as a *perspective cut* (P/C); as illustrated in Figure 3, it defines the unique supporting hyper-plane to the function passing from $(0, 0)$ and $(\bar{p}, 1)$. Note that the epigraph of h is a cone, i.e., differently from the previous case (cf. Figure 1) the function is linear along all segments of extremes $(0, 0)$ and $(\bar{p}, 1)$ for any feasible production level $\bar{p}$ (with *varying* slope), as it is easy to verify algebraically. Thus, the PCF formulation corresponds to choosing

Figure 3: Piecewise-linear approximation of $h(p, u)$
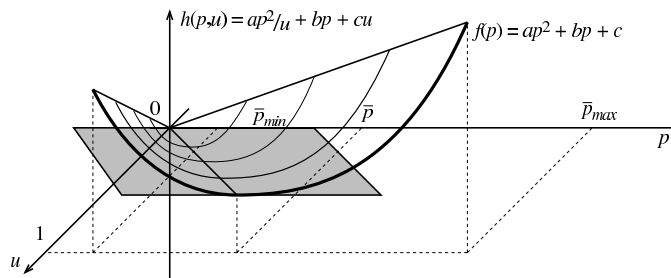
supporting hyper-planes tangent to the graph of $h(p, u)$ both in $(0, 0)$ and in the points $(\bar{p}^h, 1)$, and using as objective function the polyhedral function which is the point-wise maximum of the corresponding linear functions; this better describes the true behavior of the actual non-convex objective function, up to the extent possible to a convex approximation.

Given the standard MINLP formulation of UC, PCF is even slightly simpler to implement than SPWF. The differences between SPWF and PCF can be summarized as follows:

- Assuming k pieces are constructed for each i and t, SPWF has $k|P||\mathcal{T}|$ more continuous variables and $(2k+1)|P||\mathcal{T}|$ more constraints (counting box constraints) than UC, while PCF has only $|P||\mathcal{T}|$ more continuous variables and $k|P||\mathcal{T}|$ more constraints than UC; thus, PCF has significantly fewer variables and constraints than SPWF, especially as k grows, although the constraints (14) are slightly denser than box constraints.

- Since the objective function of PCF underestimates h, solving the continuous relaxation of PCF provides a valid lower bound to the optimal value of UC, and therefore the global lower bound provided by a Branch&Bound approach using PCF is valid for UC, unlike that provided by SPWF whose objective function is an upper estimate of f. This would make a difference, in theory, if the stopping criterion of the Branch&Bound would be computed by evaluating feasible solutions with the value of the "true" objective function (1); in this case, in fact, the solution found would be guaranteed to be optimal to the prescribed accuracy for PCF, but not for SPWF. This could be easily solved by using a lower approximation in SPWF, but anyway it is immaterial for the current approach, that in both cases is a heuristic one, as discussed below.

- PCF allows to easily work with a *dynamic k*, the number of constraints controlling how accurately the objective function is represented. In fact, one can choose a small set of initial constraints, solve the continuous relaxation of PCF and, if $u^* > 0$, check whether the solution $[v^*, p^*, u^*]$ satisfies the P/C (14) for $\bar{p} = p^*/u^*$; if not, the thus obtained cut can be added to the formulation, using the standard mechanisms that MILP solvers make available for implementing the so-called "Branch&Cut" approaches with user-defined cuts. 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. While this is in theory possible for SPWF too, adding constraints to a formulation during a Branch&Bound is now a standard feature to MILP solvers and is very effectively supported, which is not true for adding variables.

Apart from these differences, the two formulations share the largest part of their variables and constraints, and therefore once one of the two has been programmed, the other can be quite easily obtained with a few modifications, especially if using a high-level algebraic modeling system. Also, because the approach applies to a "very basic" portion of the UC problem, it can be easily applied to the numerous variants of the problem developed in the vast literature on the subject. Finally, the new formulation can be easily applied to the case where the cost function is piecewise-quadratic but non-convex, like for instance the case when valve points need to be taken into account; simply, the approach is applied separately to each segment where the function is convex.

4. Computational Experiences

In this section we present some numerical results aimed at testing the effectiveness of the P/C-based formulations within MILP-based heuristic approaches to UC. For this, we implemented three different approaches:

- SPWF: the MILP formulation, with $k = 4$ (as suggested in [7], while [17] uses $k = 3$) equidistant points, is constructed and passed to an MILP solver;

- PCF: same as before, but the P/C formulation is used;

- $PCFD_k$: initially, the P/C formulation with only *two* pieces, the ones corresponding with $\bar{p}_{min}$ and $\bar{p}_{max}$, is constructed; additional cuts, up to a maximum of k (a user-configurable parameter) per variable are then dynamically generated when needed as described in the previous paragraph.

The tests have been performed on an Opteron 246 (2 GHz) computer with 2 GigaBytes of RAM, running Linux Fedora Core 3, and using the highly regarded commercial solver `Cplex 9.1`. As all current commercial solvers, `Cplex` offers mechanisms (the `cut callback` functions) allowing easy implementation of the $PCFD_k$ approach.

A crucial parameter to be tuned for this kind of approaches is the prescribed relative accuracy obtained which the solver is allowed to stop: we tested all methods with two settings, a relatively "relaxed" one of 0.5% (the value used in [7, 12]), and the "tighter" 0.01% (the default value for `Cplex`, considered a very high accuracy). We should mention that for none of the approaches there is an *a-priori* guarantee that the obtained integer solution will in fact be accurate with that precision; this is because 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, for SPWF the lower bound is not *a-priori* valid, being the objective function of the MILP an upper approximation of (1); by the same token, however, the upper bound is a valid one. The converse obviously happens for PCF, since in that case the objective function of the MILP is a lower approximation of (1). All this is immaterial in practice, since the difference between the actual function value and its (both upper and lower) approximations was always very small, to the tune of 0.01%. However, to make the comparison absolutely fair the gaps reported in the following Tables 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), and comparing it with the best valid lower bound we know for each instance; since the same lower bound is used for both formulations (note that SPWF does not provide *any* valid lower bound for (1)), any difference in gaps is only due to the quality of the corresponding feasible solutions.

For our tests, we have used two sets of randomly generated realistic pure thermal and hydro-thermal instances, with a number of thermal units ranging from 10 to 200 and a number of hydro units ranging from 10 to 100, on a daily problem ($n = 24$). These have been generated with a modified version of the procedure described in [4], which 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. The procedure has only been modified to also randomly generate realistic ramping restrictions, 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*; introducing time-dependent startup costs in the MILP formulations is done in the same way for both, and results in the same increase of the number of constraints, thereby it should not materially impact on the comparison between SPWF and PCF. The UC instances are freely available at the OR-Library [3], and have already been used in [11, 12] for testing Lagrangian Relaxation approaches and MIQP- and MILP-based ones. The size of the different MILP formulations tested is reported in Table 1; 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 table, with $h = 0$, is therefore composed by "pure thermal" instances; each row reports averaged results of 5 instances of the same size. Column "bvar" reports the number of binary variables (equal for all formulations), while columns "cvar" report the number of continuous variables

for, respectively, SPWF and all the PC-based formulations. Then, columns "const" report the number of (non-box) constraints for, respectively, SPWF, PCF, and the PCFD$_k$ formulations. Finally, columns "P/Cs" report the number of P/Cs dynamically generated by PDFD$_4$ and PCFD$_\infty$, respectively, when the optimality tolerance is set to the "tight" value of 0.01% (the number is clearly lower with the "relaxed" tolerance of 0.5%).

Table 1: Dimensions of the different MILP formulations

p	h	bvar	cvar		const			P/Cs	
			SPWF	PCF*	SPWF	PCF	PCFD	$k=4$	$k=\infty$
10	0	240	1440	480	4326	4086	3126	683	805
20	0	480	2880	960	8607	8127	6207	1422	1821
50	0	1200	7200	2400	21412	20212	15412	3416	4247
75	0	1800	10800	3600	32709	30909	23709	5057	5876
100	0	2400	14400	4800	44829	42429	32829	6458	7461
150	0	3600	21600	7200	66651	63051	48651	9652	10642
200	0	4800	28800	9600	89442	84642	65442	12263	14011
20	10	480	3600	1680	8172	7692	5772	1190	1354
50	20	1200	8640	3840	20073	18873	14073	2802	3488
75	35	1800	12060	4800	30238	28438	21238	4443	4642
100	50	2400	18000	8400	40031	37631	28031	5941	6122
150	75	3600	27000	12600	61055	57455	43055	8161	8508
200	100	4800	36000	16800	81098	76298	57098	10488	10977

4.1. Comparing static formulations at lower accuracy

We first analyze the results obtained by comparing SPWF and PCF with stopping criterion at 0.5%. The results are displayed in Table 2; columns "SPWF" report results for the SPWF formulation, while columns "PCF" report results for the PCF formulation. In both cases, column "time" reports the required running time (in seconds), column "nd" reports the number of visited nodes in the enumeration tree, and column "LPs" reports the total number of LP solved; this is much larger than the number of nodes because Cplex 9.1 employs a sophisticated "Branch & Cut" approach where valid inequalities are automatically derived and added to the formulation to improve the lower bound. Furthermore, 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. Finally, column "rgap" reports the obtained gap (in percentage) between the lower bound obtained at the root node of the enumeration tree (solving the continuous relaxation of the MIP formulation), compared to the best valid upper bound we know for each instance; since the same upper bound is used for both formulations, the gaps can be compared.

The table shows that both SPWF and PCF obtain good quality solutions; most often, PCF attains solution of slightly better quality than SPWF. Furthermore, PCF most often terminates significantly faster. This is partly due to the fact that solving the continuous relaxation of PCF is slightly but noticeably faster than solving that of SPWF (this fact is not reported in the table due to space reasons), and to a larger extent due to the better root node gap (cf. column "rgap"). Although the difference may look minor, the reduction in root node gap is significant enough to diminish the total number of LPs solved, and often the number of Branch&Bound nodes, too, finally yielding a consistently reduced running time. This confirms the better quality of the

Table 2: Comparing SPWF and PCF for low accuracy

		SPWF					PCF				
p	h	gap	nd	LPs	time	rgap	gap	nd	LPs	time	rgap
10	0	0.31	0	23	0.95	1.61	0.28	0	16	0.76	1.50
20	0	0.34	0	23	3.72	1.34	0.36	8	25	3.56	1.25
50	0	0.21	0	25	21.93	1.38	0.21	0	12	12.09	1.26
75	0	0.20	10	59	56.31	1.43	0.18	14	51	45.88	1.30
100	0	0.17	16	76	94.09	1.39	0.15	0	19	43.55	1.27
150	0	0.12	16	115	218.69	1.32	0.11	2	42	146.80	1.20
200	0	0.09	6	87	267.78	1.37	0.08	0	52	234.97	1.25
20	10	0.21	140	258	93.53	0.82	0.20	0	70	3.71	0.69
50	20	0.06	0	60	17.98	0.70	0.10	0	65	18.93	0.63
75	35	0.11	170	300	96.86	0.57	0.07	70	224	64.52	0.52
100	50	0.06	180	266	130.86	0.58	0.07	35	155	81.41	0.53
150	75	0.06	300	554	467.62	0.58	0.05	90	316	293.50	0.52
200	100	0.05	205	321	427.71	0.56	0.03	35	168	314.00	0.51

lower bound produced by the PCF forulation w.r.t. that produced by the SPWF formulation, despite the fact that the latter is not even a guaranteed lower bound since the original objective function is upper approximated. The Table also shows that hydro-thermal instances typically have smaller gaps than pure thermal ones; this has always been the case in our experience (e.g. [12]). Intuitively, the reason is likely to be that hydro units give the model more flexibility to adapt to the discontinuities caused by the combinatorial nature of thermal units' operations.

4.2. Static vs. dynamic formulations at lower accuracy

Having proven that PCF is a worthy competitor for SPWF, we now proceed at testing the impact of dynamic vs. static generation of the P/C. For this, we compare PCF with two variants of $PCFD_k$, for $k = 4$ and $k = \infty$ respectively. $PCFD_4$ has the same maximum size as PCF, but cuts are generated only when needed, and therefore can "concentrate" on some "critical" variables, while leaving others (e.g., those that always attain zero value in the continuous relaxation) with a less accurate, but still sufficient, approximation of the objective function; furthermore, the points where the cuts are evaluated are chosen dynamically by the approach instead of *a-priori*. $PCFD_\infty$ allows for arbitrarily accurate approximations of the objective function, possibly paying a high price in terms of the size of the linear programs that need be solved at each node of the enumeration tree. The results are displayed in Table 3, where the meaning of the columns is the same as in the previous one.

The table shows interesting results. Both dynamic approaches are competitive with the static one. In particular, it appears that $PCFD_\infty$ is remarkably effective for small- to mid-scale instances, while $PCFD_4$ is more effective on the large-scale ones; for the largest hydro-thermal instances it provides slightly better solutions in half of the time required by PCF. This is probably due to the fact that for moderate size instances the more accurate approximation leads to finding a better solution quicker, but as the size of the instances grow large the increase in the computational cost of the solution of the linear programs corresponding to the many more P/C added overbalances the improvements in accuracy of the objective function. All in all, however, the results clearly show that an appropriate choice of the parameter k leads to

Table 3: Comparing PCF, PCFD$_4$ and PCFD$_\infty$ for low accuracy

		PCF				PCFD$_4$				PCFD$_\infty$			
p	h	gap	nd	LPs	time	gap	nd	LPs	time	gap	nd	LPs	time
10	0	0.28	0	16	0.76	0.30	0	27	0.86	0.28	0	24	0.80
20	0	0.36	8	25	3.56	0.36	0	22	2.51	0.33	0	29	3.00
50	0	0.21	0	12	12.09	0.19	0	33	14.17	0.18	0	27	13.08
75	0	0.18	14	51	45.88	0.19	2	39	36.62	0.22	0	21	22.58
100	0	0.15	0	19	43.55	0.17	0	22	34.31	0.20	0	23	36.51
150	0	0.11	2	42	146.80	0.11	4	76	104.68	0.12	10	87	169.68
200	0	0.08	0	52	234.97	0.10	0	63	183.01	0.14	12	99	235.60
20	10	0.20	0	70	3.71	0.30	5	75	4.18	0.15	0	45	2.51
50	20	0.10	0	65	18.93	0.10	10	112	19.06	0.13	0	56	10.93
75	35	0.07	70	224	64.52	0.05	115	357	70.55	0.03	95	368	64.80
100	50	0.07	35	155	81.41	0.05	15	110	47.62	0.04	40	181	60.78
150	75	0.05	90	316	293.50	0.05	115	374	194.10	0.05	115	428	216.33
200	100	0.03	35	168	314.00	0.02	0	85	155.36	0.03	135	466	342.69

substantially better results w.r.t. the static formulation.

4.3. Results with higher accuracy

Finally, we analyze the impact of the optimality threshold by presenting the results for all four approaches (SPWF and the three P/C-based ones) with the "tighter" stopping tolerance of 0.01%. Since attaining such a high accuracy may require a very long time, the search is stopped after 10000 seconds and the best solution obtained so far is returned. The results are displayed in Table 4; the meaning of the columns in this table is the same as in the previous ones.

Table 4: Comparing SPWF and PCF* for high accuracy

		SPWF		PCF		PCFD$_4$		PCFD$_\infty$	
p	h	gap	time	gap	time	gap	time	gap	time
10	0	0.01	22	0.01	15	0.01	12	0.01	16
20	0	0.01	3480	0.02	2969	0.02	3614	0.01	3481
50	0	0.09	10000	0.09	10000	0.08	10000	0.09	10000
75	0	0.09	10000	0.09	10000	0.08	10000	0.08	10000
100	0	0.07	10000	0.06	10000	0.06	10000	0.06	10000
150	0	0.07	10000	0.05	10000	0.05	10000	0.05	10000
200	0	0.07	10000	0.06	10000	0.05	10000	0.05	10000
20	10	0.01	288	0.01	383	0.01	238	0.01	317
50	20	0.01	9613	0.00	6855	0.00	7772	0.01	8326
75	35	0.01	10000	0.01	10000	0.01	10000	0.01	10000
100	50	0.01	10000	0.01	10000	0.01	10000	0.01	10000
150	75	0.01	10000	0.01	10000	0.01	10000	0.01	10000
200	100	0.01	10000	0.01	10000	0.01	10000	0.01	10000

As the table shows, allowing the search to continue decreases the final gap by a significant factor; it does not necessarily bring it down to 0.01%, even in the (few) instances that are solved up to the prescribed accuracy, due to the fact that the MILP formulations are only approximations of the "true" MIQP one. However, the improvement in accuracy comes at the expense of a dramatic increase of running times; all but the smallest instances are stopped

by the time limit, not a surprising result in view of the experiments reported in [10]. All the formulations attain similar results; however, for the small-scale instances that are solved up to the prescribed accuracy within the allotted time limit the P/C-based formulations are most often (slightly but noticeably) faster, while providing comparable or better solutions. For the other instances, within the same total running time the P/C-based formulations are able to attain slightly better final solutions on large-scale pure thermal instances, and are competitive on all other cases. The P/C-based formulations are also competitive for hydro-thermal instances, which however are solved with a very high degree of accuracy by both methods; the final gaps are only fractionally larger than 0.01%, and the algorithms cannot stop only because the lower bound computed by the MILP formulations is not as accurate as the one used for computing the Table, which is based on sophisticated Lagrangian techniques [12]. Among the P/C-based formulations, PCFD_∞ appears to be the more "robust", as it almost always reports—for a given running time—solutions of equivalent or (slightly) better quality than all the others.

In general, the results show that the PCF formulation provides, with the same implementation effort, a better description of the feasible region (objective function) of the "true" MIQP problem, which finally leads, *ceteris paribus*, to shorter running times and/or better feasible solutions. Allowing the number of P/Cs used, and the points where they are generated, to be dynamic further significantly improves the efficiency of the approach, especially if the allowed maximum number of cuts is properly managed.

5. Conclusions and directions for future work

In this paper, we have proposed a new way for constructing MILP approximated formulations for hydro-thermal Unit Commitment problems. While being not more difficult to implement than previously proposed formulations, the new approach significantly improves the performances of MILP-based heuristics to the problem, either in terms of required running time, or in terms of quality of the obtained solutions. With a limited additional implementation effort dynamic versions of the approach can be implemented which may lead to further significant improvements of the results. While the formulation is tested only on a "standard" form of the UC problem, the underlying concept can be applied to many other variants of the problem, where analogous results should be expected. All in all, these results show that appropriate formulations of UC problems can be used to find good-quality solutions in relatively short time by using off-the-shelf, general-purpose optimization software.

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] J. Arroyo and A. Conejo, "Modeling of Start-Up and Shut-Down Power Trajectories of Thermal Units," *IEEE Transactions on Power Systems*, vol. 19, no. 3, pp. 1562–1568, 2004.

[3] J. Beasley, "`http://people.brunel.ac.uk/~mastjjb/jeb/orlib/info.html`," 1990.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

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

[12] A. Frangioni, C. Gentile, and F. Lacalandra, "Solving Unit Commitment Problems with General Ramp Contraints," *International Journal of Electrical Power & Energy Systems*, vol. to appear, 2007.

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

[14] 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.

[15] 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.

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

[17] 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.