



A Primal-Proximal Heuristic Applied to the Unit-Commitment Problem

Louis Dubost, Robert Gonzalez, Claude Lemaréchal

► **To cite this version:**

| Louis Dubost, Robert Gonzalez, Claude Lemaréchal. A Primal-Proximal Heuristic Applied to the
| Unit-Commitment Problem. [Research Report] RR-4978, INRIA. 2003. inria-00071600

HAL Id: inria-00071600

<https://hal.inria.fr/inria-00071600>

Submitted on 23 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

*A Primal-Proximal Heuristic Applied to the
Unit-Commitment Problem*

Louis Dubost — Robert Gonzalez — Claude Lemaréchal

N° 4978

Octobre 2003

THÈME 4



*Rapport
de recherche*

A Primal-Proximal Heuristic Applied to the Unit-Commitment Problem

Louis Dubost ^{*}, Robert Gonzalez [†], Claude Lemaréchal [‡]

Thème 4 — Simulation et optimisation
de systèmes complexes
Projet Bipop

Rapport de recherche n° 4978 — Octobre 2003 — 18 pages

Abstract: This paper is devoted to the numerical resolution of unit-commitment problems. More precisely we present the French model optimizing the daily production of electricity. Its resolution is done in two phases: first a Lagrangian relaxation solves the dual to find a lower bound; it also gives a primal relaxed solution. The latter is used in the second phase for a heuristic resolution based on a primal proximal algorithm. This second step comes as an alternative to an earlier approach based on augmented Lagrangian (i.e. a *dual* proximal algorithm). We illustrate the method with some real-life numerical results. A companion paper is devoted to a theoretical study of the heuristic in the second phase.

Key-words: Unit-commitment problem, proximal algorithm, Lagrangian relaxation, primal-dual heuristics

^{*} EDF R&D, 1 avenue du Général de Gaulle, 92141 Clamart; louis.dubost@edf.fr

[†] EDF, 9 route de la porte de Buc, 78005 Versailles; robert.gonzalez@rte-france.com

[‡] Inria Rhône-Alpes; claude.lemarechal@inrialpes.fr

Une heuristique primale-duale appliquée à la production électrique

Résumé : Cet article concerne l'optimisation de la production électrique en France. Nous présentons le modèle utilisé par EdF. Sa résolution se fait en deux phases: tout d'abord une relaxation lagrangienne donne une borne inférieure du coût optimal, ainsi qu'une solution du problème primal relaxé. Cette dernière est injectée dans la seconde phase pour une optimisation heuristique utilisant un algorithme proximal primal. Une approche antérieure pour cette seconde phase utilisait le lagrangien augmenté, c'est à dire un algorithme proximal *dual*. Nous illustrons la méthode par quelques résultats numériques tirés du réel. L'heuristique primale-duale a fait l'objet d'une étude théorique, publiée par ailleurs.

Mots-clés : Unit-commitment, optimisation de la production, algorithme proximal, relaxation lagrangienne, heuristique primale-duale

1 Introduction

This paper reports on a practical application: to optimize the daily production of the French electricity mix. We symbolize this *unit-commitment* problem by

$$\min c(x), \quad x \in X \subset \mathbb{R}^n, \quad g(x) = 0 \in \mathbb{R}^m. \quad (1)$$

Section 2 briefly outlines the unit-commitment model in operation at EdF (Électricité de France, the French Electricity Board).

In (1), $X = \prod_i X^i$, $c(x) = \sum_i c^i(x^i)$, where i denotes a production unit (nuclear, thermal or hydro-valley), with decision variable x^i varying in X^i ; also, $g(x) = \sum_i g^i(x^i)$ collects *linking constraints*: the problem is decomposable. It is therefore conveniently solved via Lagrangian relaxation (see [5, 3, 27] among others, and also the review [34]): to (1), is associated the Lagrangian problem

$$\begin{aligned} \theta(\lambda) &:= \min_{x \in X} L(x, \lambda), \quad \text{where} \\ L(x, \lambda) &:= c(x) + \lambda^\top g(x). \end{aligned} \quad (2)$$

Computing $\theta(\lambda)$ amounts to solving independently the local problems

$$\min_{x^i \in X^i} c^i(x^i) + \lambda^\top g^i(x^i), \quad (3)$$

where i ranges over the set of production unit.

A dual algorithm constructs iteratively λ_k to maximize θ (the dual problem), (2) being repeatedly solved for $\lambda = \lambda_k$. Section 3 recalls the approach and gives some illustrative results obtained with EdF's unit-commitment model, in which a bundle method maximizes θ [24], with the help of the quadratic solver of [19, 20].

The above technique provides a lower bound for the optimal cost. However, (1) itself is still unsolved: no optimal schedule need be produced while maximizing θ , not even an approximately feasible one. In its current implementation, EdF's model uses a heuristic based on augmented Lagrangian [8] to recover approximately feasible/optimal schedules. The present paper proposes another heuristic, which can roughly be outlined as follows.

After convergence, the dual algorithm produces a primal point \hat{x} , which solves a certain *convex relaxation* of (1) (think of column generation, see [10, 28], [22, Sections. 2.6, 4.2, 5.2], [38]). We call *pseudo-schedule* this \hat{x} ; unless there is no duality gap, the pseudo-schedule is a non-trivial convex combination of schedules, and therefore does not lie in the nonconvex set X . Nevertheless, \hat{x} tends to satisfy the linking constraints, a property which is enjoyed by none of the minimizers in (2) – here precisely lies the trouble with Lagrangian relaxation!

Starting from this idea, our heuristic perturbs θ to

$$\begin{aligned} \theta_r(\lambda) &:= \min_{x \in X} L_r(x, \lambda), \quad \text{where} \\ L_r(x, \lambda) &:= L(x, \lambda) + r \|x - \hat{x}\|^2. \end{aligned}$$

A crucial observation is that θ_r is readily computable because

$$L_r(x, \lambda) = \sum_i [c_i(x^i) + \lambda^\top g^i(x^i) + r \|x^i - \hat{x}^i\|^2].$$

It can therefore be maximized with just the same bundle algorithm that maximized θ , using essentially the same sub-algorithms that solved (3). By contrast, the augmented Lagrangian approach [15, 30, 32, 33, 4] requires to minimize $L(x, \lambda) + r \|g(x)\|^2$ and the decomposability property of (1), (2) is destroyed. Section 4 presents the heuristics and gives illustrative results suggesting the superiority of the present approach.

Actually, the above heuristic gives fairly good empirical results on the unit-commitment problem. A question of interest is therefore: can it be assessed by some theoretical study? This is the subject of [9], a sequel to the present work.

2 The French model for electricity production

The French production mix is composed of nearly 150 thermal units and 50 hydro-valleys. Depending on the period of the day and on the season of the year, the load demand oscillates between some 40 000 and 70 000 MW. A reasonably reliable demand forecast, as well as the list of available power units with their updated technical constraints and economic characteristics, are available each day by 4 pm. Appropriate production schedules must then be computed, to be sent to the local units by 5 pm.

The generation mix is composed of nuclear, thermal (fuel or coal) and hydro-plants, which represent respectively 61%, 16% and 23% in terms of total capacity. Of course these figures are very different from the actual distribution of the generation, which highly depends on the economic characteristics of each unit. Gas turbines are also present; they are not included in the model: due to their high operating costs and flexibility, they are generally viewed as a potential reserve, to be used in some emergency cases.

Calling T the time horizon, each schedule is represented by a vector $x^i \in \mathbb{R}^T$; each day is discretized in 48 half-hours. However, optimizing over one day only (which is the required exercise) is a myopic strategy. It may produce costly situations such as shutting down by the end of the day some thermal unit which will have to be switched on early the day after: a startup sequence is a costly transition. This is why the model takes $T = 96$, i.e. the optimization is performed over two days, even though only the first half of the schedules are actually sent to the local units.

Altogether, the unit-commitment problem is to compute feasible schedules x^i , minimizing the total operating cost over the period, while matching the predicted demand and satisfying some *reserve* constraints, present to face unexpected events (increase in demand, failure of a production unit). We now proceed to describe the technological and economic characteristics of EdF's model.

2.1 Nuclear and thermal units

Nuclear and thermal units have similar behaviours; their decision variable x^i used in the notation (1) is indeed the power delivered p^i .

The operating domain related to each thermal unit is defined by a particular set of dynamic constraints, a non-exhaustive list of which is given below:

- (i) The delay between two output variations must be longer than some value depending on the unit.
- (ii) Ramping constraints are associated to each thermal unit: jumping from one output level to another cannot be done instantaneously.
- (iii) After switching off a unit, a minimum down time must precede any following startup sequence.
- (iv) Particular startup curves must be used, depending on which cold start they apply to (e.g. depending on how long the unit has been switched off).
- (v) When a generator is on, it must remain online for a minimum amount of time.
- (vi) When a generator is on, its output level must remain between given maximum and minimum values (both positive).
- (vii) A generator may or may not be allowed to participate to the frequency regulation at time t (this is a device to ensure the automatically generated reserves, see §2.3 below).

If $p^i = (p_1^i, \dots, p_T^i)$ denotes the production vector of unit i , satisfaction of these various constraints can be symbolized by $p^i \in P^i$. Besides, a cost $c^i(p^i)$ is associated to the production vector (with respect to the notation (1), (3), remember that the decision variable is here $x^i = p^i$). When the linking constraints of §2.3 are dualized, unit i will aim at minimizing $c^i(p^i) + \lambda^\top g^i(p^i)$, subject to $p^i \in P^i$: a fairly complicated problem, with lots of Boolean constraints, and nonlinear cost.

A simplifying assumption is made at this point, replacing the continuous segment $[p^i, \bar{p}^i]$ introduced in (vi) by a set of discrete values: for instance, the output level of a thermal unit can take on 3 values only (\underline{p}^i , \bar{p}^i , and a nominal value in between). With this simplification, the problem becomes purely discrete, and can be given a formulation suitable to standard *dynamic programming*. For this, a number of *state variables* are used, to characterize the state of unit i ; for example

- (i) the time q_t^i elapsed since the last output variation,
- (iii) when the unit is off, the time r_t^i elapsed since the switch off (this is also useful to cope with (iv)),
- (v) when the unit is on, the time s_t^i elapsed since the switch on,
- (...) and so on.

These state variables evolve along time according to appropriate *state equations*, in which p^i acts as a *control variable*. The states are themselves subject to various constraints, such as $q_t^i \geq \underline{q}^i$, $r_t^i \geq \underline{r}^i$, $s_t^i \geq \underline{s}^i$, and so on. In a word, the problem can be solved efficiently by dynamic programming.

Remark 2.1 *As already said, our list of constraints above is not exhaustive, the main merit of dynamic programming being its flexibility. However there are also daily constraints, expressing for example that the number of output variations and startup sequences in a day is limited. A direct treatment of these constraints increases the number of states drastically, the cost in CPU becomes prohibitive. The current strategy is therefore to penalize or dualize them.*

For the record, "semi-local" constraints also exist. The units are actually clustered in production plants (a plant may contain up to 4 units) and constraints link units within the same plant. For example, there is a minimum delay between two startup sequences of two different units within the same plant. No rigorous method has been found to deal with this

kind of constraints; dualizing them would increase the size of the problem with little hope to compute schedules satisfying them, therefore they are ignored in the model. \square

2.2 Hydro-valleys

The French hydro-power system is composed of 50 independent hydro-valleys. A hydro-valley i is a set of interconnected reservoirs – we index them by r – and associated power plants – indexed by j ; the plants in turn have a certain number of turbines. The decision variable x^i in the notation (1) is now the set of flows w^j passing through the various plants, each of which resulting in a power $p^j = p^j(w^j)$. A flow w^j can take on finitely many values; these values correspond to the optimal rate of efficiency of the turbines in the given plant. Besides, some plants are equipped to pump water from their downstream reservoir up into their upstream reservoir (thus allowing to re-use water later in peak hours).

Again a simplifying assumption is made here: each w^j is allowed to vary continuously in a given interval¹ $[0, \bar{w}^j]$. Besides, the rate of each turbine is considered as fixed. With this simplification, the power delivered is a piecewise linear function, which turns out to be *concave*, as depicted on Fig. 1 (each break point corresponds to switching a turbine, and the most efficient turbines are of course switched on first). As a result, the power delivered can be written as the minimum of a number of affine functions; say

$$p^j(w^j) = \min \{ \rho_g^j w^j + y_g^j : g \text{ runs over the turbines} \}. \quad (4)$$

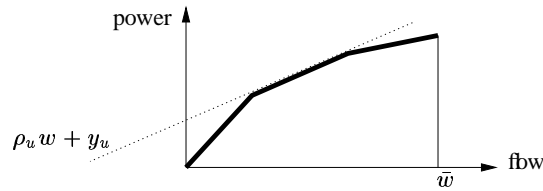


Figure 1: Hydro-production as a function of flow

Pumping is modelled as follows: every plant with the pumping capability is duplicated in two plants, say j and j' ; the (nonnegative) flow $w^{j'}$ of the “dummy” plant j' is counted upstream, and $p^{j'}(w^{j'}) \leq 0$ (pumping consumes energy!).

The volume V^r in a given reservoir r is then a *state variable*, which obeys an equation of the type

$$V_t^r = V_{t-1}^r - \sum_{j \in N_{\downarrow}(r)} w_{t+\tau(r,j)}^j + \sum_{j \in N_{\uparrow}(r)} w_{t-\tau(j,r)}^j + O_t^r, \quad t = 1, \dots, T.$$

Here O represents the exchange of water with the outside (rain, snow, spillage); $N_{\downarrow}(r)$ [resp. $N_{\uparrow}(r)$] is the set of neighboring plants down [resp. up] reservoir r ; $\tau(r, j)$ [resp. $\tau(j, r)$] is the travel time of water from reservoir r to plant j [resp. from plant j to reservoir r] – with the appropriate conventions for the dummy pumps.

Now the cost of the schedules in valley i is just the global loss in water:

$$c^i(w^i) = \sum_r \gamma^r (V_0^r - V_T^r), \quad \text{for each valley } i;$$

this is a linear function of the decision variable w . The marginal costs γ^r result from mid- and long-term models, assessing the value of water. To complete the model, bound constraints are introduced on the volumes V^r .

Some crucial observations are relevant at this point.

- The rate ρ_u^j of a given turbine does not depend on the volume in the reservoir feeding it; in other words: water elevation in the reservoirs is considered as a constant over the period $[0, T]$, thus allowing the simple expression (4). Due to the shortness of the optimized period (max. 2 days) and to the existence of bound constraints on the reservoir’s volumes, this is quite a valid approximation.
- A plant cannot discharge and pump simultaneously: there are constraints of the form $w_t^j w_t^{j'} = 0$. However, it can be shown that these constraints are automatically satisfied by an optimal solution – such a property relies heavily on the particular form of the Lagrangian, though.

¹Compare this simplification with that of thermal plants, where continuous variables were replaced by discrete ones. After the simplified model is solved, a smoothing “off-model” phase is performed, aimed at restoring feasible w^j ’s without changing the overall power produced.

– As will be seen in §2.3, dualization of the constraints results in a hydraulic Lagrangian of the form

$$c^i(w^i) - \lambda^\top \sum_j p^j(w^j).$$

If $\lambda \geq 0$ – quite a natural property, always true in practice –, the particular form (4) of $w \mapsto p(w)$ allows an LP formulation, via the addition of an extra variable for each unit u .

As a result “everything is linear”: minimizing the Lagrangian over the feasible set amounts to minimizing a linear function over a polyhedron, which we symbolize by $w^i \in W^i$. The resulting linear program is solved by an interior-point method, as described in [36] and implemented in [12]. In view of the form of the constraints, an LP solver tailored to flow problems would be more appropriate. However, some valleys might have additional constraints (for example limitations on flow variations); besides, quadratic Lagrangians will appear in §4. This motivates the choice of an interior-point methodology.

In what follows, we will use the notation ϑ [resp. \mathcal{H}] for the thermal [rep. hydraulic] mix. The total cost of a schedule is then

$$\sum_{i \in \vartheta} c^i(p^i) + \sum_{i \in \mathcal{H}} c^i(w^i). \quad (5)$$

2.3 Linking constraints

Demand Denote by $d^0 \in \mathbb{R}^T$ the predicted power demand vector. The mismatch between production and demand is

$$p_t^0 := d_t^0 - \sum_{i \in \vartheta \cup \mathcal{H}} p_t^i, \quad \text{for } t = 1, \dots, T. \quad (6)$$

This notation introduces a “fictitious power generator” $p^0 \in \mathbb{R}^T$, analogous to a slack variable. Imposing brutally the constraint $p^0 = 0$ – or even $p^0 \geq 0$ – would ignore various uncertainties present in the problem (the power demand vector d^0 is not 100% reliable, some production failures may occur, etc.). Rather, the model attaches a cost

$$c^0(p^0) = \sum_t c_t^0(p_t^0)$$

to the fictitious generation, where the elementary cost c_t^0 is depicted on Fig. 2. Then the balance constraint is just (6).

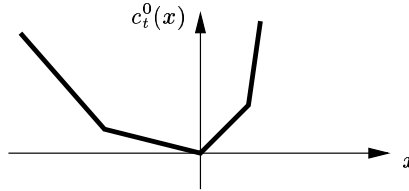


Figure 2: Cost of the fictitious power generator

Reserves In order to face unexpected events such as generator outages or sudden load variations in operational conditions, three types of *spinning reserves* impose constraints linking the production plants.

- The *primary* reserve continually adjusts (within a few seconds) small oscillations of the power load, simply by controlling automatically the angular speed of the alternators. This kind of reserve also aims at immediately compensating unexpected generator outages until the frequency of the network comes back to its nominal value. The total power thus available for adjustment must be at least some value $d^1 = (d_1^1, \dots, d_T^1)$.
- The *secondary* reserve also controls automatically the network’s frequency with a response time of minutes, and must be at least some value d^2 . The contribution of a given thermal unit i to the secondary reserve has a predefined value which depends whether p^i is at its minimal, maximal or intermediate value.
- The *tertiary* reserve is triggered manually and its time response is longer than 20 minutes. The contribution of a thermal unit i to this reserve is $\bar{p}^i - p^i$, but there are a number of conditions; for example, it is 0 if the unit is following a decreasing ramp, or during a startup sequence, and of course if it is off.

The total reserve mismatch is then

$$p^{-k} := d^k - \sum_{i \in \vartheta} R^{k,i}(p^i) - \sum_{i \in \mathcal{H}} R^{k,i} p^i \in \mathbb{R}^T, \quad k = 1, 2, 3. \quad (7)$$

Each $R^{k,i}$ is a more or less complex function of the power; if i indexes hydraulic production, $R^{k,i} p^i$ is actually linear. Analogously to the power constraint (6), the notation above introduces three ‘‘fictitious reserve generators’’ $p^{-k} \in \mathbb{R}^T$, which induce three costs, as displayed on Fig. 3. Note that $c^{-k}(x) = 0$ for $x \geq 0$: in fact, $p^k \geq 0$ is imposed as a hard constraint.

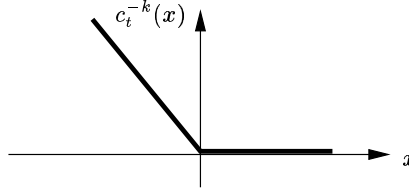


Figure 3: Cost of a fictitious reserve generator

In summary, our commitment problem is

$$\begin{cases} \min \sum_{k=0}^3 c^{-k}(p^{-k}) + \sum_{i \in \vartheta} c^i(p^i) + \sum_{i \in \mathcal{H}} c^i(w^i); \\ p^i \in P^i, i \in \vartheta; w^i \in W^i, i \in \mathcal{H}; p^{-k} \geq 0 \in \mathbb{R}^T, k = 1, 2, 3; \\ d^0 - p^0 - \sum_{i \in \vartheta} p^i - \sum_{i \in \mathcal{H}} p^i(w^i) = 0 \in \mathbb{R}^T; \\ d^k - p^{-k} - \sum_{i \in \vartheta} R^{k,i}(p^i) - \sum_{i \in \mathcal{H}} R^{k,i} p^i(w^i) = 0 \in \mathbb{R}^T, k = 1, 2, 3. \end{cases} \quad (8)$$

This problem is large-scale, heterogenous, mixed-integer, nonlinear; in a word: a nasty problem. However, dualizing the linking constraints is particularly attractive.

Remark 2.2 (Role of the fictitious costs) *Solving this problem is clearly impossible. More precisely, it is impossible to obtain schedules satisfying exactly the various linking constraints (or rather the associated complementary slackness). A common objective function is therefore required to assess a given schedule, in terms of its real cost and of its constraint mismatch. The fictitious costs c^{-k} do the job: the value of a ‘‘real’’ schedule $\{p^i\}_{i \in \vartheta \cup \mathcal{H}}$ is the sum of its real cost (5) and of the total fictitious cost*

$$c^0 \left(d^0 - \sum_{i \in \vartheta \cup \mathcal{H}} p^i \right) + \sum_{k=1}^3 c^{-k} \left(d^k - \sum_{i \in \vartheta \cup \mathcal{H}} R^{k,i}(p^i) \right)_-.$$

This remark will be important for Sections 4 and 5. □

3 The dual problem and its resolution

A first step for the resolution of (8) via Lagrangian relaxation is to solve the dual. This provides a lower bound on the optimal cost. More importantly, a solution of a convexified form of (8) is also provided; we call it a pseudo-schedule, and it will play a crucial role to recover a primal heuristic solution in §4. The reader may wish to consult [22, 23] for a more complete account of duality, following the same lines.

3.1 The dual problem

Taking dual variables $\lambda \in \mathbb{R}^T$ (for the balance constraint (6) at each time period) and $\mu \in \mathbb{R}_+^{3T}$ (for each reserve constraint (7) at each time period), we form the lagrangian

$$\begin{aligned} L(p, w, \lambda, \mu) &:= c^0(p^0) - \lambda^\top p^0 + \sum_{k=1}^3 (c^{-k}(p^{-k}) - \mu_k^\top p^{-k}) \\ &\quad + \sum_{i \in \vartheta} c^i(p^i) - \lambda^\top p^i - \mu^\top R^i(p^i) \\ &\quad + \sum_{i \in \mathcal{H}} c^i(w^i) - \lambda^\top p^i(w^i) - \mu^\top R^i p^i(w^i) \\ &\quad + \lambda^\top d^0 + \mu^\top d. \end{aligned}$$

Its minimization for given λ, μ amounts to solving the local problems

$$\begin{aligned}
& \min_{p^0 \in \mathbb{R}^T} c^0(p^0) - \lambda^\top p^0, \\
& \min_{p^{-k} \in \mathbb{R}_+^T} c^{-k}(p^{-k}) - \mu_k^\top p^{-k} \quad k = 1, 2, 3, \\
& \min_{p^i \in P^i} c^i(p^i) - \lambda^\top p^i - \mu^\top R^i(p^i) \quad i \in \vartheta, \\
& \min_{w^i \in W^i} c^i(w^i) - \lambda^\top p^i(w^i) - \mu^\top R^i p^i(w^i) \quad i \in \mathcal{H}.
\end{aligned} \tag{9}$$

The resulting optimal value is the dual function, which must be maximized over $(\lambda, \mu) \in \mathbb{R}^T \times \mathbb{R}_+^{3T}$ (note: the constraint $\lambda \geq 0$ could just be introduced as well, since negative values of λ would be economically absurd).

From now on, we will return to the condensed notation (1), (2) and neglect the sign-constraint on the dual variables; then we have $x = (p^{-k}, p^\vartheta, w^{\mathcal{H}})$ and $m = 4T$.

It is well known that any dual value $\theta(\lambda)$ is a lower bound for the optimal cost in (1)=(8), and that θ is a concave function. Calling $x(\lambda)$ an optimal solution in (3)=(9), it is also known that the vector $g := g(x(\lambda)) \in \mathbb{R}^m$ is a subgradient of θ (or rather a “supergradient”, θ being concave). Examining the local problems (9) shows that θ is a polyhedral function, with kinks at those λ such that $x(\lambda)$ is ambiguously defined. The dual problem is definitely a nonsmooth optimization problem.

The present dual approach is fairly classical, especially for dual commitment, and does not deserve much more comment. We just mention that the numerical illustrations to be reported below have been obtained with a bundle method. This is substantially more sophisticated than a mere subgradient method of the type $\lambda_+ = \lambda + sg(x(\lambda))$. The resulting dual algorithm therefore involves two (complicated but) very different worlds, sketched in Fig. 4: the primal world is in charge of solving (3)=(9), and the dual world iterates over the dual variable λ_k . For each of these two worlds, the other is a black box which delivers appropriate information upon request.

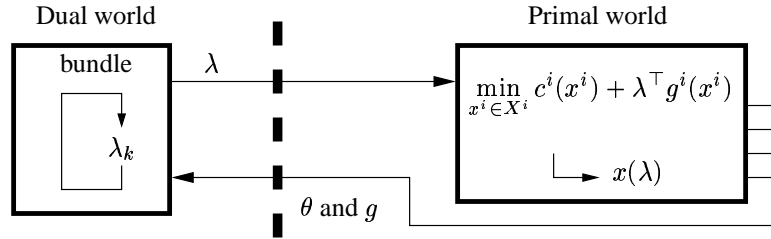


Figure 4: The primal-dual dialogue

3.2 Overview of bundle methods

A bundle method works in a way analogous to column generation.

Having solved the local problems (3) for a number of dual iterates λ_k , one has obtained corresponding optimal solutions $x_k = x(\lambda_k)$; by construction, each x_k lies in the set X of “technically feasible” schedules; note that none of them satisfies the linking constraints, though – otherwise (1) would be solved. Then one constructs the “restricted” dual function

$$\hat{\theta}(\lambda) := \min_k L(x_k, \lambda) = \min_k [c(x_k) + \lambda^\top g(x_k)], \tag{10}$$

which (over-)approximates the actual θ , insofar as the “restricted” set $\{x_k\}_k$ approximates the whole of X . Standard column generation finds the next iterate λ_+ by solving the LP restricted master $\max_\lambda \hat{\theta}(\lambda)$, i.e. $\max_\lambda \min_k L(x_k, \lambda)$, i.e.

$$\begin{cases} \max r & (\lambda, r) \in \mathbb{R}^m \times \mathbb{R}, \\ r \leq c(x_k) + \lambda^\top g(x_k) & \text{for each } k. \end{cases} \tag{11}$$

In nonsmooth optimization, this is known as the method of Kelley, or Cheney-Goldstein, or cutting planes [7, 18].

Bundle methods are stabilized versions of column generation. At the current iteration, one chooses a stability center $\hat{\lambda}$, a stepsize $s > 0$, and one solves the QP restricted master $\max \hat{\theta}(\lambda) - \|\lambda - \hat{\lambda}\|^2 / 2s$, i.e.

$$\begin{cases} \max r - \frac{1}{2s} \|\lambda - \hat{\lambda}\|^2 & (\lambda, r) \in \mathbb{R}^m \times \mathbb{R}, \\ r \leq c(x_k) + \lambda^\top g(x_k) & \text{for each } k. \end{cases} \tag{12}$$

Without entering details, the stepsize is adjusted at each call to the local problems; and the stability center is the best among the previous iterates: $\theta(\hat{\lambda}) = \max_k \theta(\lambda_k)$. It is $\hat{\lambda}$ which is supposed to converge to a dual solution, say $\hat{\lambda} \rightarrow \lambda^*$.

3.3 Primal interpretation and pseudo-schedules

The above algorithms have a primal interpretation, obtained by (bi-)dualization. This is well-known for pure cutting planes: the dual of (11) is

$$\min_{\alpha \in \Delta} \sum_k \alpha_k c(x_k), \quad \text{subject to} \quad \sum_k \alpha_k g(x_k) = 0, \quad (13)$$

where Δ is the unit simplex (of appropriate dimension, depending on the iteration number). In a bundle method, the QP restricted master (12) has also a dual, which turns out to be

$$\min_{\alpha \in \Delta} \sum_k \alpha_k c(x_k) + \frac{s}{2} \left\| \sum_k \alpha_k g(x_k) \right\|^2 + \hat{\lambda}^\top \sum_k \alpha_k g(x_k). \quad (14)$$

Call $\hat{\alpha} \in \Delta$ an optimal solution of (13) or (14). This (bi-)dualization reveals important primal objects:

$$\begin{array}{ll} \text{a pseudo-schedule} & \hat{x} := \sum_k \hat{\alpha}_k x_k \\ \text{with corresponding pseudo-cost} & \hat{c} := \sum_k \hat{\alpha}_k c(x_k) \\ \text{and pseudo-constraint} & \hat{g} := \sum_k \hat{\alpha}_k g(x_k). \end{array} \quad (15)$$

To help intuition, observe that

- pure cutting-planes imposes the constraint $\hat{g} = 0$: see (13);
- in a bundle method, this constraint is rather dualized (with multiplier $\hat{\lambda}$) and penalized (with penalty parameter $s/2$): see (14);
- if c is linear, then $\hat{c} = c(\hat{x})$: the pseudo-cost is the cost of the pseudo-schedule;
- if g is affine, then $\hat{g} = g(\hat{x})$: the pseudo-constraint is the constraint-value at the pseudo-schedule.

Observe also that \hat{x} does not normally lie in X : it is not a “technically feasible” schedule, hence our terminology.

Remark 3.1 *The solution of (12) is $\lambda_+ = \hat{\lambda} + s\hat{g}$, with \hat{g} given in (15) and $\hat{\alpha}$ solving (14). Knowing that \hat{g} is an “approximate subgradient” of θ at $\hat{\lambda}$, s appears as a stepsize and we see here that a bundle method resembles a subgradient method.*

Our use of simplified notation (1), (2) instead of (8), (9), eliminates the additional dual variable μ . A more serious simplification is that, if (12) had positivity constraints, say $\max_{\mu \geq 0} \hat{\theta}(\mu) - \frac{\|\mu - \hat{\mu}\|^2}{2s}$, its dual would be fairly more intricate than (14). Deriving this dual would be possible, though. Calculations are left to a motivated reader, they follow [32], [6, §14.3]. \square

Convergence of a bundle method means that

- the pseudo-cost tends to the dual optimal value: $\theta(\hat{\lambda}) \rightarrow \max \theta$ and $\hat{c} \rightarrow \max \theta$;
- the pseudo-constraint tends to 0: $\hat{g} \rightarrow 0$;
- besides, the stability center tends to a dual solution, if any: $\hat{\lambda} \rightarrow \lambda^*$.

Theory tells us that $\hat{c} \geq \max \theta$; besides, we have at each iteration

$$\theta(\lambda) \leq \theta(\hat{\lambda}) + \varepsilon + \hat{g}^\top (\lambda - \hat{\lambda}), \quad \text{for all } \lambda \in \mathbb{R}^m,$$

where we have set $\varepsilon := \hat{c} - \theta(\hat{\lambda}) \geq 0$. The algorithm can therefore be stopped when both ε and \hat{g} are small: $\hat{\lambda}$ is then approximately dual optimal.

In practice, ε tends to 0 much faster than \hat{g} . Altogether, we see that the crucial point in a bundle method is to obtain a small \hat{g} , i.e. to obtain a pseudo-schedule which is approximately “pseudo-feasible” with respect to the linking constraints (more simply, it is approximately “feasible” if these constraints are affine).

3.4 Numerical illustration

In EdF’s current implementation, the dual problem is solved by the bundle code described in [26], the restricted master problem being solved by K.C. Kiwiel’s code QPDF4 described in [19, 20]. Some numerical illustrations following this approach can be found in [24]; [29] gives comparative results with a mere subgradient algorithm. The use of bundle methods with positivity constraints is not too common, though; we therefore illustrate its behaviour in this setting.

We use three datasets, representative of various everyday situations, with different sets of available production units. In each of these sets, the primary reserve is neglected. There are 96 demand constraints, plus $2 \times 96 = 192$ reserve constraints, plus a number of daily constraints alluded to in Remark 2.1.

- Dataset A has 69 408 primal variables (power generation and reserves contributions for each unit at each time step) and 828 linking constraints (corresponding to 540 daily constraints); the maximal demand is 62 000 MW.
- Dataset B has 62 496 primal variables and 780 linking constraints (corresponding to 492 daily constraints); the maximal demand is 58 000 MW.
- Data C has 52 128 primal variables, 780 linking constraints, and a maximal demand of 52 000 MW.

Table 1 gives the statistics: number of iterations, and constraints violations by the final pseudo-schedule \hat{g} in (15). An iteration is one resolution of the quadratic master, followed by one minimization of the Lagrangian. In the last three columns of Table 1, the notation B [resp. R^k] stands for the *real mismatch* in the balance constraint (6) [resp. the reserve constraints (7)], excluding the fictitious generators. In other words we have

$$B_t = d_t^0 - \sum_{i \in \theta \cup \mathcal{H}} p_t^i \quad \text{and} \quad R_t^k = d_t^k - \sum_{i \in \theta \cup \mathcal{H}} R_t^{k,i}(p_t^i), \quad k = 1, 2, 3. \quad (16)$$

An overline means average value (over time); the powers are given in MW.

Data	m	# iters	$\overline{ B }$	$\overline{(R^2)_-}$	$\overline{(R^3)_-}$
A	828	200	6.70	5.59	2.18
B	780	189	2.21	2.99	0.0
C	780	174	1.70	2.86	0.0

Table 1: Statistics for three dual optimizations

It should be mentioned that excellent initializations are available for λ : between the initial and final iterates, the dual function increases by less than 5%. Figure 5 gives an idea of the behaviour of $\theta(\hat{\lambda})$ along the iterations. For industrial privacy, θ -values are normalized: K being the total iteration number, the picture displays the log of

$$\Theta_k := \frac{\theta(\hat{\lambda}_k) - \theta(\hat{\lambda}_K)}{\theta(\hat{\lambda}_1) - \theta(\hat{\lambda}_K)}, \quad \text{for } k = 1, \dots, K - 1,$$

a number varying from 1 to 0.

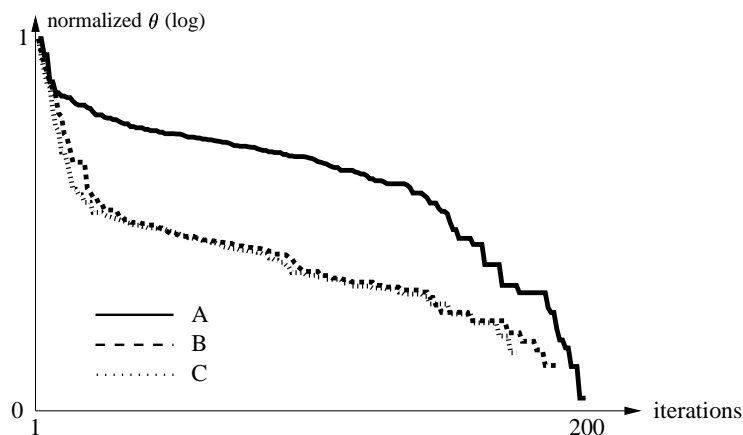


Figure 5: Evolution of the (normalized) dual function

Due to the logarithms, Fig. 5 is slightly deceiving, especially by the last iterations. Figure 6 zooms the values of Θ_k for dataset A, during the last 60 iterations.

As explained at the end of § 3.3, a crucial parameter in the algorithm is the pseudo-constraint \hat{g} in (15). Figure 7 shows the evolution of $\|\hat{g}\|$ for the three examples, also in logarithmic scale.

Remark 3.2 *Passing from one iteration to the next in the bundle algorithm does not change much the restricted master (12) or (14):*

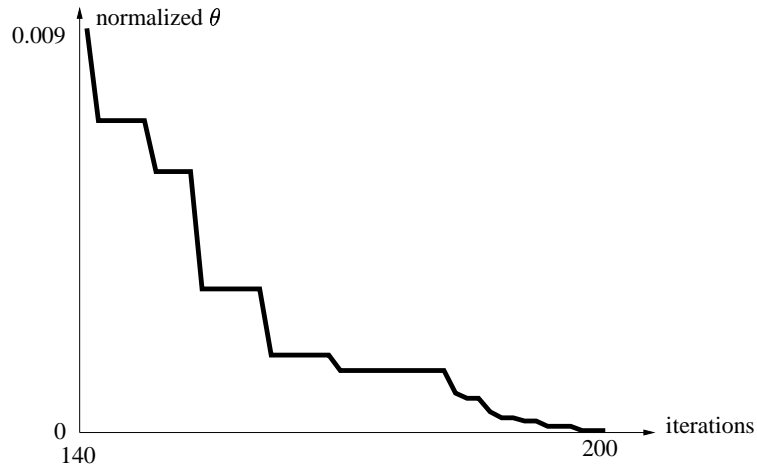


Figure 6: Tail of the (normalized) dual function, dataset A

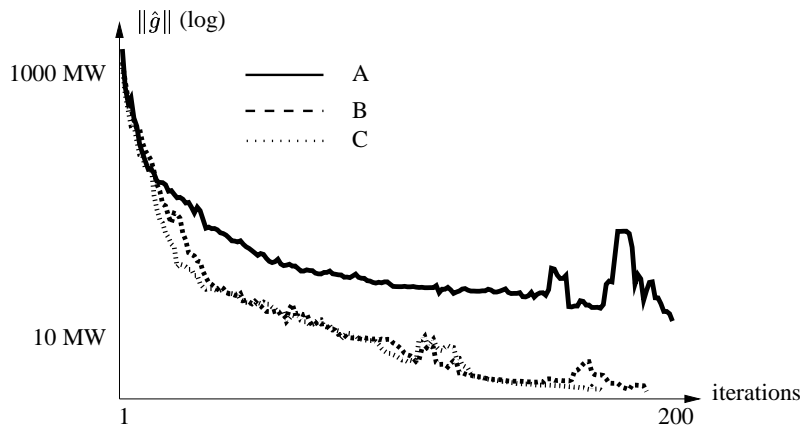


Figure 7: Evolution of the pseudo-constraint

- a new piece $L(x_+, \lambda)$ is introduced in $\hat{\theta}$ of (10) – resulting in an additional variable α_+ in (14),
- and/or the stepsize s is changed,
- and/or the stability center $\hat{\lambda}$ is updated.

Thus the restricted master lends itself to warm starts. However, the CPU time to solve (14) with QPDF4 is negligible compared with the time necessary for the local problems (9). In our reported experiments, the QP solver is actually used in a very primitive manner, where a cold start is performed at each iteration. This costs nothing in CPU and can but improve robustness of the QP solver: propagation of roundoff errors is avoided in the Choleski matrix associated with (14). \square

3.5 Reoptimization

A problem of raising importance in energy production is the following. Suppose (1)–(8) is solved, at least approximately. In particular the optimal cost is known, at least a good approximation of it (we will see in §5 that the dual bound *is* such a good approximation). Then comes a perturbation of the righthand side: d in (8) is replaced by $d + \delta d$, and one wishes to estimate *quickly* the corresponding perturbation of the optimal cost. A simple solution is to use duality: the sensitivity of the optimal cost (or of its dual approximation but this is good enough) with respect to the righthand side is λ^* . When the perturbation is large, however, this may result in substantial error; it is then safer to perform a reoptimization with the new righthand side.

Call $C(d)$ the dual optimal value of (8), considered as a function of the righthand side. Standard duality theory says that C is a convex function:

$$C(d + \delta d) \geq C(d) + \lambda^{*\top} \delta d.$$

Therefore $\lambda^{*\top} \delta d$ underestimates the change of the optimal cost, and is exact if *a posteriori*, λ^* is still a dual optimal solution for the new righthand side $d + \delta d$.

An obvious but important remark is that the righthand side has no influence whatsoever on the Lagrangian problems (9). Thus, each primal point x_k obtained during the first dual optimization would be obtained again during the second, if $L(x, \lambda_k)$ in (2) were minimized with the same λ_k . It follows that the function $\hat{\theta}(\lambda) + \lambda^\top \delta d$ – see (10) – is still a valid restricted dual function: it can be used right away for the second dual optimization, and this hot start may spare computing time.

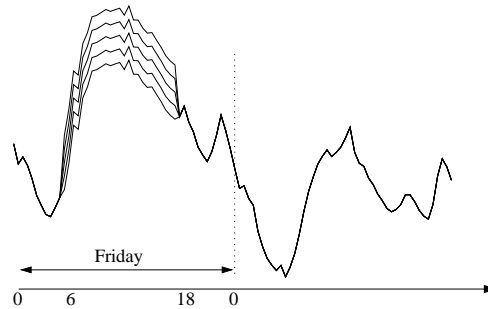


Figure 8: Typical perturbation of the demand

We illustrate this technique with four typical perturbations of the demand d^0 , given on Fig. 8. The data concern a Friday, with its dangling Saturday; the nominal demand oscillates between 37 000 and 54 000 MW; δd is a “stripe”, extending between two timesteps – here 6am and 6pm – representing an additional need of some industrial customer. We make 4 scenarios, with stripes of width ranking from 1 000 to 4 000 MW.

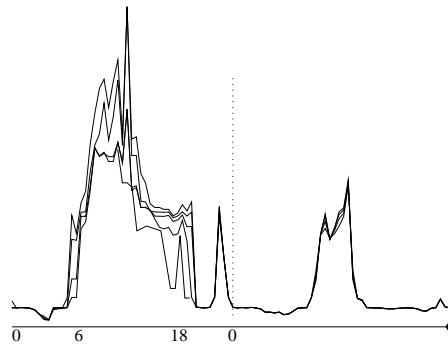


Figure 9: Corresponding perturbation of λ^*

Each of these perturbations does modify the optimal λ^* , as illustrated on Fig. 9.

Table 2 gives the statistics of the re-optimization, comparing the actual and linearized increases of the cost; increases are given in fractions of the nominal $\max \theta$ (normalized, as in §3.4).

δd^0 (MW)	# iters	CPU (s)	δC	instead of
0	123	406	0	0
1 000	89	294	1.7	1.4
2 000	68	224	3.6	2.8
3 000	130	429	5.6	4.2
4 000	111	366	7.9	5.6

Table 2: Statistics for reoptimization

4 Primal heuristics

As a rule, the dual algorithm of §3 does not produce any primal solution, not even an approximately feasible one. The aim of this section is to present and illustrate possible heuristic methods to obtain such an approximate solution. We will particularly be interested (§4.1) by the approach currently studied at EdF.

The situation is now as follows: we have solved the dual problem associated with (1) and we have on hand

- the last stability center $\hat{\lambda}$, which solves (approximately) the dual problem,
- the last pseudo-schedule \hat{x} , which solves (approximately) a convexified form of (1),
- its accompanying pseudo-cost \hat{c} and pseudo-constraint \hat{g} .

This data is available either from a standard column generation, or from a bundle algorithm².

4.1 The primal proximal idea

Consider an introductory example with one variable and one constraint:

$$\min x, \quad x \in X := \{-1, 0, +1\}, \quad g(x) := x = 0,$$

with solution $x^* = 0$. The Lagrangian is $L(x, \lambda) = (1 + \lambda)x$; assume that the Lagrangian problem (2) computes

$$x(\lambda) = \begin{cases} 1 & \text{if } \lambda \leq -1, \\ -1 & \text{if } \lambda > -1, \end{cases} \quad \text{and} \quad \theta(\lambda) = \begin{cases} 1 + \lambda & \text{if } \lambda \leq -1, \\ -1 - \lambda & \text{if } \lambda > -1. \end{cases}$$

Start the dual algorithm at $\lambda_1 = 1$, thus obtaining $x_1 = -1$ (so that $\theta(\lambda_1) = -2$ and $\hat{\theta}(\lambda) = -1 - \lambda$). Then come a number of iterations depending on the particular algorithm chosen. We may assume for simplicity that we are using a bundle method with $s = 2$; with $\hat{\lambda} = 1$, we obtain $\lambda_2 = -1$. Then $x_2 = 1$ is computed, with $\theta(\lambda_2) = 0$. At this stage, we detect that λ_2 is dual optimal: in fact $\hat{\theta}(\lambda) = \min\{1 + \lambda, -1 - \lambda\}$ is maximal at $\lambda = -1$ with value $\hat{\theta}(-1) = 0 = \theta(-1)$. At the same time, (13) or (14) (with $\hat{\lambda} = 0$) provides $\hat{x} = \frac{1}{2}(x_1 + x_2) = 0$, the required primal solution.

The above naive example reveals the irrelevance of *all possible* answers x_k from the Lagrange problem (2): none of them, taken individually, is any good in terms of the linking constraint. By contrast, their *convex combination* \hat{x} appears as “better than the others”. The remark can even be pushed further: X could be any set between $\{-1, +1\}$ and $[-1, +1]$, the situation would still be the same. All of the x_k ’s would look like random points, while \hat{x} would somehow be the best possible point computable via duality.

Therefore, duality would probably give better primal results if the x_k ’s could somehow be made closer together, namely closer to \hat{x} . This suggests a simple idea: to add a quadratic term $r\|x - \hat{x}\|^2$ to the cost function $c(x)$, thus forcing the Lagrangian problem (2) to answer better primal points, in terms of the linking constraints. Of course, note that the perturbed dual function

$$\begin{aligned} \theta_r(\lambda) &:= \min_{x \in X} L_r(x, \lambda), \quad \text{where} \\ L_r(x, \lambda) &:= c(x) + r\|x - \hat{x}\|^2 + \lambda^\top g(x) \end{aligned}$$

will no longer bound the primal cost from below. Besides, a dual algorithm will produce a new pair $(\hat{x}_r, \hat{\lambda}_r)$, biased by \hat{x} .

When specialized to the unit-commitment problem (8), the local problems (9) become

$$\begin{aligned} &\min_{p^0 \in \mathbb{R}^T} c^0(p^0) - \lambda^\top p^0 + r\|p^0 - \hat{p}^0\|^2, \\ &\min_{p^{-k} \in \mathbb{R}_+^T} c^{-k}(p^{-k}) - \mu_k^\top p^{-k} + r\|p^{-k} - \hat{p}^{-k}\|^2 && k = 1, 2, 3, \\ &\min_{p^i \in P^i} c^i(p^i) - \lambda^\top p^i - \mu^\top R^i(p^i) + r\|p^i - \hat{p}^i\|^2 && i \in \vartheta, \\ &\min_{w^i \in W^i} c^i(w^i) - \lambda^\top p^i(w^i) - \mu^\top R^i(p^i(w^i)) + r\|w^i - \hat{w}^i\|^2 && i \in \mathcal{H}. \end{aligned} \tag{17}$$

The decomposability property is preserved, only the form of the cost has changed; dynamic programming ($i \in \vartheta$) is not affected, and interior-point ($i \in \mathcal{H}$) accommodates the quadratic term easily; as for the fictitious problems, they are trivial anyway.

Remark 4.1 *Alternatively, one can penalize the deviation of the (local) constraint values. With the schematic form (1), (3), the penalized local problems become*

$$\min_{x^i \in X^i} c^i(x^i) + \lambda^\top g^i(x^i) + r\|g^i(x^i) - g^i(\hat{x}^i)\|^2.$$

²Incidentally, we also mention a lesser-known fact: \hat{x} would also be available from a standard subgradient algorithm, see [1, 21] or [35, p. 116].

This is probably more sensible, as it gets closer to the augmented Lagrangian idea of §4.2 below. Note also how this is reminiscent of primal decomposition [11].

With the specific problem (9), we obtain

$$\min_{p^i \in P^i} c^i(p^i) - \lambda^\top p^i - \mu^\top R^i(p^i) + r \|p^i - \hat{p}^i\|^2 + r \|R^i(p^i) - R^i(\hat{p}^i)\|^2, \quad i \in \vartheta$$

for the thermal units, and likewise for the others. However, a difficulty appears for the hydro-valleys: since the production is a nonlinear function of the control variable, the quadratic term $r \|p^i(w^i) - p^i(\hat{w}^i)\|^2$ results in a nonconvex quadratic program, which can no longer be conveniently solved.

In what follows, we will content ourselves with the formulation (17), which represents a handy notation for the actual implementation. \square

The schedules computed in (17) are *technically feasible* and *easy to compute*; and the hope is that some of them are qualitatively good, in terms of Remark 2.2. In summary, the resulting algorithm to solve (8) is composed of two phases:

Algorithm 4.2 (phase 1 - phase 2 for unit-commitment)

PHASE 1. Solve (by a bundle method) the dual problem associated with (9).

Obtain a dual solution $\hat{\lambda}$ and a pseudo-schedule \hat{x} .

Choose a penalty term $r > 0$.

PHASE 2. Solve (by a bundle method) the dual problem associated with (17).

During this resolution, record the best schedule computed in (17).

If no acceptable schedule has been obtained until convergence, repeat: form a new pseudo-schedule \hat{x}_r and execute

Phase 2 again. \square

Best schedules are of course those giving the best real + fictitious cost (remember Remark 2.2).

Note that a reasonable value for r can be computed: it must simply balance the primal costs and the (squared) constraint violations obtained during the first phase. Besides, the perturbed dual problem is again convex, so the dual algorithm can be started anywhere for the second phase – for example on $\hat{\lambda}$. However, in contrast with §3.5, the old x_k 's are of no use for the second phase.

Remark 4.3 One execution of Phase 2 in Algorithm 4.2 is one iteration – say ℓ – of the so-called proximal algorithm: to solve (1), this algorithm would define the sequence (x_ℓ) by

$$x_{\ell+1} = \operatorname{argmin} \{c(x) + r \|x - x_\ell\|^2 : x \in X, g(x) = 0\},$$

its motivation being that the above problem is “more convex” than (1). This motivates the terminology of primal proximal heuristic for Algorithm 4.2.

As for Phase 1, it can be viewed as an initialization of this proximal algorithm. \square

4.2 An approach via augmented Lagrangian

A known technique to cancel the duality gap is the augmented Lagrangian. Consider the equivalent form of (1)

$$\min c(x) + \pi \|g(x)\|^2, \quad x \in X, \quad g(x) = 0. \quad (18)$$

The corresponding dual function is

$$\begin{aligned} \theta^\pi(\lambda) &:= \min_{x \in X} L^\pi(x, \lambda), \quad \text{where} \\ L^\pi(x, \lambda) &:= c(x) + \lambda^\top g(x) + \pi \|g(x)\|^2. \end{aligned}$$

Under mild assumptions (certainly satisfied by the unit-commitment problem, see [16, Thm. XII.5.2.2]), there is no duality gap: for π large enough, $\sup \theta^\pi$ is the optimal value of (1). Unfortunately, the method is only theoretical, as $L^\pi(\cdot, \lambda)$ is “impossible” to minimize: cross-products of the type $x^i x^j$ destroy the decomposability property.

Remark 4.4 When (1) is convex, the above augmented dual function satisfies

$$\theta^\pi(\hat{\lambda}) = \max_{\lambda \in \mathbb{R}^m} \theta(\lambda) - \frac{1}{4\pi} \|\lambda - \hat{\lambda}\|^2$$

(see [17, Prop. XII.5.2.3]). With respect to Remark 4.3, a dual method using augmented Lagrangian appears as a proximal algorithm to maximize the dual function θ . Compare also with the bundle approach (12). \square

The method currently in operations at EdF is based on augmented Lagrangian; it is made implementable via a mechanism due to [8], which works as follows; see [31]. First linearize the trouble-making term $\|g(x)\|^2$; denoting by \tilde{x} the linearization point, the augmented Lagrangian becomes

$$\tilde{L}^\pi(x, \lambda, \tilde{x}) = c(x) + \lambda^\top g(x) + 2\pi g(\tilde{x}) \frac{\partial g}{\partial x}(\tilde{x})(x - \tilde{x}).$$

Then the local problems can cope with this Lagrangian, to compute the iterates x_k . However, the whole beneficial effect of augmented Lagrangian heavily relies on a *quadratic* term (to stabilize the x_k 's); \tilde{L}^π is therefore modified further to the “augmented-linearized-proximized” Lagrangian

$$\begin{aligned} \tilde{L}^{r\pi}(x, \lambda, \tilde{x}) &= \tilde{L}^\pi(x, \lambda, \tilde{x}) + r\|x - \tilde{x}\|^2 \\ &= c(x) + \lambda^\top g(x) + 2\pi g(\tilde{x}) \frac{\partial g}{\partial x}(\tilde{x})(x - \tilde{x}) + r\|x - \tilde{x}\|^2. \end{aligned}$$

Its minimization (for given $\lambda = \lambda_k$ and $\tilde{x} = \tilde{x}_k$) by the local problems gives the iterate x_k . Then λ is updated according to the standard scheme in augmented Lagrangian (see [4, §4.2] for example)

$$\lambda_{k+1} = \lambda_k + \pi \frac{\partial \tilde{L}^{r\pi}}{\partial \lambda}(x_k, \lambda_k, \tilde{x}_k) = \lambda_k + \pi g(x_k).$$

As for \tilde{x} , it is simply set to the current iterate: $\tilde{x}_{k+1} = x_k$.

The quadratic term $\|x - \tilde{x}\|^2$ introduces the proximal mechanism; the above approach resembles Algorithm 4.2, but with a number of differences:

- the insertion of the linear term $2\pi g(\tilde{x}) \frac{\partial g}{\partial x}(\tilde{x})x$,
- a “diagonalized” update of λ and of the primal stability center: here they are updated simultaneously, a strategy *à la* Arrow-Hurwicz [2]; by contrast, Algorithm 4.2 uses a strategy *à la* Uzawa [37], updating \hat{x} (or x_t , to follow the notation of Remark 4.3) only after having performed a full optimization on λ .

EdF’s implementation has two more differences:

- the stability center \tilde{x}_0 is initialized on the last x_k computed during Pase I (but our example at the beginning of §4 suggests that this is not a sound idea),
- the approach is actually applied to a Lagrangian decomposition form of (8), with a duplication of the primal variables (see [13, 14] and [25, §4.3]):

$$\min c(x), \quad x \in X \subset \mathbb{R}^n, \quad g(y) = 0 \in \mathbb{R}^m, \quad x = y$$

and it is the constraints $x = y$ that are dualized.

5 Comparative illustrations

This section presents a number of results to illustrate the recovery of a feasible schedule, after the dual problem is solved. They compare the behaviour of the two techniques described in §4: PS (pseudo-schedule) of §4.1 and AL (augmented Lagrangian) of §4.2.

5.1 Standard situations

We start our comparison with the three datasets of §3.4. For each example, we compare in Table 3 the schedules obtained, in terms of their cost and actual constraint violations. The column “cost” records the excess of the total cost (real + fictitious) over the lower bound $\max \theta$ obtained in Phase 1; and see (16) for the notation B, R^k .

The table also shows the iteration number where the best schedule has been obtained, and the number of iterations until termination (remember that AL has no reliable stopping criterion). An important advantages of PS is to find its best solution before AL – in addition to the availability of a stopping criterion.

5.2 A difficult case

In Case D, used in this section, the power load reaches a peak at 70 000 MW, thus providing a suitable test to illustrate the robustness of PS. Phase 2 becomes substantially harder to solve, as is suggested by the iteration numbers reported in Table 4; note also how expensive is the schedule found by AL (8.1% more than the lower bound). On the other hand, the average mismatches of the linking constraints seem acceptable. However, the maximal shortage in power, reached at

Data	Alg.	cost %	$\overline{ B }$	$\overline{(R^2)_-}$		$\overline{(R^3)_-}$	Iteration #	
				megawatts			best	stop
A	AL	1.27	44.0	12.7	0.8	155	200	
	PS	0.56	36.4	3.7	0.4	28	65	
B	AL	0.22	22.5	6.1	0.0	95	200	
	PS	0.14	21.6	6.4	0.0	8	33	
C	AL	0.26	17.6	11.8	0.0	163	200	
	PS	0.16	29.7	10.9	0.0	11	36	

Table 3: Statistics for three standard examples

Alg.	cost %	$\overline{ B }$	$\overline{(R^2)_-}$		$B_{\bar{t}}$	$R_{\bar{t}}^3$	Iter. #	
			megawatts				best	stop
AL	8.1	118.1	3.8	2.0	-1659	-5	200	200
PS	1.5	105.6	6.7	2.3	-786	-12	195	200

Table 4: Case D: an unusual peakload

a certain time \bar{t} , is unacceptable for AL (1 659 MW – remember from (16) $B < 0$ means an insufficient production; this schedule is totally useless) and very bad for PS. We note for future use that there is also a shortage of tertiary reserve at \bar{t} .

Obviously, both algorithms should have started an additional unit at \bar{t} , but they failed to finding such a solution. Actually, the blame is on the fictitious generator: at \bar{t} , the magnitude of λ becomes comparable to the slopes of c^0 (see Fig. 2). Then the algorithms use an unduly large amount of fictitious power: this is cheaper than starting an additional unit (probably most expensive, otherwise it would already be on!), and keeping it on for a possibly long time (because of technological constraints).

Alg.	cost %	cost (real)	$\overline{ B }$	$\overline{(R^2)_-}$		$B_{\bar{t}}$	$R_{\bar{t}}^3$	Iter. #	
				megawatts				best	stop
AL	33	-0.21	79.2	15.2	50.4	1301	-727	200	200
PS	12	-1.14	31.6	18.8	8.1	-42	-56	8	200

Table 5: Case D with increased c^0

Table 5 therefore reports on an experiment with a more brutal and simple $c_t^0(x) := 10^3|x|$. Because c^0 can hardly be given a sensible economic interpretation, the table also records real costs, excluding the fictitious ones. Observe the reasonable discrepancies for PS, which proposes a cheaper schedule than AL. Actually, the latter again failed to converge. Note in particular that there is now an (unacceptable) *overproduction* at \bar{t} ; and it is accompanied by an *underproduction* of tertiary reserve; this is nonsense: shifting some of the power to the tertiary reserve would do only good.

5.3 Inserting primary reserve

Finally, two cases – E and F – consider the three reserves altogether. Table 6 gives the statistics, as before. Only real costs are considered, the cost of AL being normalized to 1. Again, PS finds faster more economic schedules. The total numbers of iterations are not given.

Data	Alg.	cost	$\overline{ B }$	$\overline{(R^i)_-}$			Iteration #	
				megawatts			best	stop
E	AL	1	45.7	0.	15.3	0.	199	
	PS	0.98	29.8	0.2	9.0	0.	37	
F	AL	1	47.5	2.2	9.0	0.3	198	
	PS	0.99	46.7	1.5	9.0	0.1	72	

Table 6: Two cases with primary reserve

- As a general conclusion to all of these experiments, we summarize the advantages of the PS algorithm of §4.1.
- In normal situations, the best solution is generally found earlier in the iterative process.
 - It has an implementable stopping criterion; and it can be iteratively repeated if necessary.
 - Its solutions are statistically cheaper than the solutions obtained with the AL algorithm.
 - Mismatches of the linking constraints are smaller.
 - It is easier to tune and more robust, especially with respect to the fictitious cost.
 - It is less sensitive to the addition of new linking constraints.
 - More generally, it only uses standard duality, which makes it more satisfactory from a theoretical point of view.

References

- [1] K. Anstreicher and L.A. Wolsey. On dual solutions in subgradient optimization. Unpublished manuscript, CORE, Louvain-la-Neuve, Belgium, 1993.
- [2] K.J. Arrow and L. Hurwicz. Reduction of constrained maxima to saddle point problems. In J. Neyman, editor, *Proceedings of the 3rd Berkeley Symposium on Mathematical Statistics and Probability*, pages 1–26. University of California Press, Berkeley, 1956.
- [3] J. Batut and A. Renaud. Daily generation scheduling with transmission constraints: a new class of algorithms. *IEEE Transactions on Power Systems*, 7(3):982–989, 1992.
- [4] D.P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 1995.
- [5] D.P. Bertsekas, G.S. Lauer, N.R. Sandell, and T.A. Posberg. Optimal short-term scheduling of large-scale power systems. *IEEE Transactions on Automatic Control*, AC-28:1–11, 1983.
- [6] J.F. Bonnans, J.Ch. Gilbert, C. Lemaréchal, and C. Sagastizábal. *Numerical Optimization*. Springer Verlag, 2003.
- [7] E. Cheney and A. Goldstein. Newton’s method for convex programming and Tchebycheff approximations. *Numerische Mathematik*, 1:253–268, 1959.
- [8] G. Cohen. Auxiliary problem principle and decomposition of optimization problems. *Journal of Optimization Theory and Applications*, 32:277–305, 1980.
- [9] A. Daniilidis and C. Lemaréchal. Primal proximal convexification in discrete optimization ii: a theoretical study. In preparation. Also available as: Inria Research Report 4550, www.inria.fr/rrrt/rr-4550.html.
- [10] J.E. Falk. Lagrange multipliers and nonconvex programs. *SIAM Journal on Control*, 7(4):534–545, 1969.
- [11] A.M. Geoffrion. Primal resource-directive approaches for optimizing nonlinear decomposable systems. *Operations Research*, 18(3):375–403, 1970.
- [12] R. Gonzalez, M.P. Bongrain, and A. Renaud. Unit commitment handling transmission constraints with an interior point method. In *Proceedings of the 13th PSCC conference, Trondheim*, volume 2, pages 715–723, 1999.
- [13] M. Guignard and S. Kim. Lagrangean decomposition: a model yielding stronger Lagrangean bounds. *Mathematical Programming*, 39(2):215–228, 1987.
- [14] M. Guignard and S. Kim. Lagrangean decomposition for integer programming: theory and applications. *RAIRO Recherche Opérationnelle*, 21(4):307–323, 1987.
- [15] M.R. Hestenes. Multiplier and gradient methods. *Journal of Optimization Theory and Applications*, 4:303–320, 1969.
- [16] J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms*. Springer Verlag, Heidelberg, 1993. Two volumes.
- [17] J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms*. Springer Verlag, Heidelberg, 1993. Two volumes.
- [18] J. E. Kelley. The cutting plane method for solving convex programs. *J. Soc. Indust. Appl. Math.*, 8:703–712, 1960.

- [19] K. C. Kiwiel. A dual method for certain positive semidefinite quadratic programming problems. *SIAM Journal on Scientific and Statistical Computing*, 10(1):175–186, 1989.
- [20] K.C. Kiwiel. A Cholesky dual method for proximal piecewise linear programming. *Numerische Mathematik*, 68, 1994.
- [21] T. Larsson, M. Patriksson, and A.B. Strömberg. Ergodic, primal convergence in dual subgradient schemes for convex programming. *Mathematical Programming*, 86(2):283–312, 1999.
- [22] C. Lemaréchal. Lagrangian relaxation. In M. Jünger and D. Naddef, editors, *Computational Combinatorial Optimization*, pages 115–160. Springer Verlag, Heidelberg, 2001.
- [23] C. Lemaréchal. The omnipresence of Lagrange. *4OR*, 1(1):7,25, 2003.
- [24] C. Lemaréchal, F. Pellegrino, A. Renaud, and C. Sagastizábal. Bundle methods applied to the unit-commitment problem. In J. Doležal and J. Fidler, editors, *System Modelling and Optimization*, pages 395–402, 1996.
- [25] C. Lemaréchal and A. Renaud. A geometric study of duality gaps, with applications. *Mathematical Programming*, 90(3):399–427, 2001.
- [26] C. Lemaréchal and C. Sagastizábal. Variable metric bundle methods: from conceptual to implementable forms. *Mathematical Programming*, 76(3):393–410, 1997.
- [27] M. Madrigal and V.H. Quintana. An interior-point/cutting-plane method to solve unit commitment problems. *IEEE Transactions on Power Systems*, 15(3):1022–1027, 2000.
- [28] T.L. Magnanti, J.F. Shapiro, and M.H. Wagner. Generalized linear programming solves the dual. *Management Science*, 22(11):1195–1203, 1976.
- [29] F. Pellegrino, A. Renaud, and T. Socroun. Bundle method and augmented Lagrangian methods for short-term unit commitment. In *Proceedings of the 12th PSCC Conference, Dresden*, volume 2, pages 730–739, 1996.
- [30] M.J.D. Powell. A method for nonlinear constraints in minimization problems. In R. Fletcher, editor, *Optimization*. Academic Press, London, New York, 1969.
- [31] A. Renaud. Daily generation management at Electricité de France: from planning towards real time. *IEEE Transactions on Automatic Control*, 38(7):1080–1093, 1993.
- [32] R.T. Rockafellar. The multiplier method of Hestenes and Powell applied to convex programming. *Journal of Optimization Theory and Applications*, 6:555–562, 1973.
- [33] R.T. Rockafellar. Augmented Lagrange multiplier functions and duality in nonconvex programming. *SIAM Journal on Control*, 12:268–285, 1974.
- [34] G.B. Sheblé and G.N. Fahd. Unit commitment literature synopsis. *IEEE Transactions on Power Systems*, 9:128–135, 1994.
- [35] N.Z. Shor. *Minimization methods for non-differentiable functions*. Springer Verlag, Berlin, 1985.
- [36] T. Terlaky, editor. *Interior Point Methods of Mathematical Programming*. Kluwer Academic Press, Dordrecht, 1996.
- [37] H. Uzawa. Iterative methods for concave programming. In K. Arrow, L. Hurwicz, and H. Uzawa, editors, *Studies in Linear and Nonlinear Programming*, pages 154–165. Stanford University Press, 1959.
- [38] F. Vanderbeck. A generic view at the Dantzig-Wolfe decomposition approach in mixed integer programming: paving the way for a generic code. Working paper U0222, University of Bordeaux, Talence, France, 2002.



Unité de recherche INRIA Rhône-Alpes

655, avenue de l'Europe - 38330 Montbonnot-St-Martin (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique

615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur

INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)

<http://www.inria.fr>

ISSN 0249-6399