



HAL
open science

Prices stabilization for inexact unit-commitment problems

Sofia Zaourar, Jérôme Malick

► **To cite this version:**

Sofia Zaourar, Jérôme Malick. Prices stabilization for inexact unit-commitment problems. 2012. hal-00720352v1

HAL Id: hal-00720352

<https://hal.science/hal-00720352v1>

Preprint submitted on 24 Jul 2012 (v1), last revised 5 Jul 2013 (v3)

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.

Prices stabilization for inexact unit-commitment problems

Sofia Zaourar · Jérôme Malick

Received: date / Accepted: date

Abstract Unit-commitment problems in electricity production appeal to constraint decomposition techniques: by dualizing the linking constraints, the large-scale nonconvex problem decomposes into smaller independent subproblems. The dual problem consists then in finding the best Lagrangian multiplier (the optimal “price”); it is solved by a convex nonsmooth optimization method. Realistic modeling of technical production constraints makes the Lagrangian subproblems themselves difficult to solve. Nonsmooth optimization algorithms can cope with inexact solutions of the subproblems. In this case however, we observe that the computed optimal dual variables show a noisy and unstable behaviour, that could prevent their use as price indicators.

In this paper, we present a simple and easy-to-implement way to stabilize dual optimal solutions, by penalizing the noisy behaviour of the prices in the dual objective. After studying the impact of a general stabilization term on the model and the resolution scheme, we focus on the penalization by discrete total variation, showing the consistency of the approach. We illustrate our approach on a synthetic example, and real-life problems from EDF (the French Electricity Board).

Keywords Unit-commitment problems · Lagrangian duality · convex analysis · total variation regularization · inexact bundle method

1 Introduction

Let us consider n production units over T periods of time. Unit-commitment (UC) problem can be formulated as finding production schedules $p_i \in \mathbb{R}^T$ for each production unit, that minimize the total generation cost, satisfy operational constraints ($p_i \in \mathcal{P}_i$) and match the demand forecast $d \in \mathbb{R}^T$ (load and safety). The abstract form of the problem can be written as (see e.g. (Sheble and Fahd, 1994; Padhy, 2004; Dubost et al, 2005)):

$$\begin{cases} \min c(p) = \sum_{i=1}^n c_i(p_i) \\ \text{s.t. } Ap = d \\ p \in \mathcal{P} = \mathcal{P}_1 \times \dots \times \mathcal{P}_n, \end{cases}$$

A preliminary version of this work was presented at the EDF workshop “Advanced Optimization methods and their applications to unit commitment in energy management” (November 2011).

S. Zaourar
Univeristé Joseph Fourier - Inria Grenoble
655 avenue de l’Europe, Montbonnot, 38334 Saint Ismier, France
Tel.: +33(0)-47-6615461
E-mail: sofia.zaourar@inria.fr

J. Malick
CNRS, LJK
655 avenue de l’Europe, Montbonnot, 38334 Saint Ismier, France
Tel.: +33(0)-47-6615255
E-mail: jerome.malick@inria.fr

where A is a constant matrix such that each component $(Ap)_t$ represents the total production of period t . Since we only know a forecast of the demand, we do not impose the schedules to match exactly the demand, and instead, we penalize the mismatch with a function V . In this paper, we consider a general UC problem written as:

$$\begin{cases} \min c(p) + V(z) \\ \text{s.t. } z = d - Ap \\ p \in \mathcal{P}, z \in \mathbb{R}^T. \end{cases} \quad (1)$$

This optimization problem is easy to state in an abstract form as above, but is difficult to model and solve in practice. Operational constraints are numerous and often nonconvex or with a combinatorial nature. Furthermore, load and safety requirements couple all units together. In real situations, we end up with a large-scale, heterogeneous and nonconvex problem, with mixed integer variables (Hechme-Doukopoulos et al, 2010).

Example 1 (EDF daily UC problem) The EDF short-term electricity production management problem is described precisely in (Dubost et al, 2005; Hechme-Doukopoulos et al, 2010). We give here only its main characteristics. The model includes nearly $n = 200$ independent power plants of three types (nuclear, classical thermal and hydro-valleys). Every day, the state of the production park is known (available units and their operational constraints) as well as a demand forecast for the next 48 hours. Time being discretized in half-hourly time steps, the number of periods is $T = 96$. The UC problem is written as (1) and contains about 10^6 variables and 10^6 constraints. EDF provides us with 27 realistic instances of this problem, on which we will illustrate the interest of our approach. \square

Example 2 (A simple synthetic UC problem) In this paper, we also use a simple UC problem (provided to us by Claudia Sagastizabal) to precisely illustrate our results. In this problem, the costs are linear, the technical constraints are simply capacity constraints, and we impose the production to match the demand.

$$\begin{cases} \min c^\top p \\ Ap = d \\ 0 \leq p \leq p_{\max}. \end{cases} \quad (2)$$

This linear program can be solved by any linear programming solver for the instances size that we consider here. We will use the exact solution as a reference to illustrate our method. \square

Many exact methods and heuristics have been proposed to solve UC problems; see the surveys (Sheble and Fahd, 1994; Padhy, 2004). One of the most efficient and wide-spread approach is based on constraint (or price) decomposition and Lagrangian duality (Frangioni, 2010). By penalizing the supply-demand constraint in the objective function, using a Lagrangian multiplier $u \in \mathbb{R}^T$, the problem decomposes into smaller, homogeneous and independent subproblems:

$$\min_{p_i \in \mathcal{P}_i} \{c_i(p_i) - u^\top p_i\} =: \theta_i(u), \quad (3)$$

that represent the answer of the production park to the “price signal” $u \in \mathbb{R}^T$. As in the linear case, the optimal dual variables u^* can indeed also be interpreted in this (nonlinear, nonconvex) context, as marginal costs or prices (see e.g. (Boyd and Vandenberghe, 2004)). By denoting $\theta(u) := \sum_i \theta_i(u)$, the dual problem consists then (up to a change in signs) of finding the best prices:

$$\begin{cases} \min \Theta(u) := -\theta(u) - u^\top d + V^*(u) \\ \text{s.t. } u \in \mathbb{R}^T, \end{cases} \quad (4)$$

where V^* is the convex conjugate of V (Hiriart-Urruty and Lemaréchal, 2001). The dual problem is solved by a convex nonsmooth optimization method, bundle methods are methods of choice (Hiriart-Urruty and Lemaréchal, 1993, § XV). More and more realistic modeling of technical constraints makes the subproblems (3) themselves difficult to solve within the strict computational time limits. For example, in EDF daily UC problem, subproblems associated with hydro-valleys are large mixed integer linear programs, out of reach of current state of the art solvers. This implies that we only have an approximation of θ_i , and thus of the dual function. Bundle methods can handle inexact solutions of the subproblems and solve inexactly the dual problem (Kiwiel, 2006; Emiel and Sagastizábal, 2010; Oliveira et al, 2012).

In that case however, we observe (see following Section 2) that the dual optimal solutions show a noisy and unstable behaviour.

The objective of this paper is to stabilize the dual solutions of the UC problem, computed by the inexact bundle method. Our idea is simple: we add a stabilization term in the dual objective function in order to penalize the noisy behaviour of the solutions. We will see that this approach is versatile, easy to implement and gives good results. For example, numerical results of the final section show that we are able to compute dual solutions with 80% less variability on average, without losing in quality.

The structure of the paper is as follows. In Section 2, we recall the standard inexact resolution by Lagrangian duality and show the resulting bad behaviour of inexact prices. In Section 3, we present our general dual stabilization and its impact on the resolution scheme and on the primal UC problem. We give several possible stabilizations, and illustrate our approach on both a synthetic and real-life UC problems. In Section 4, we focus on the case of total variation regularization and we show in a particular example that this regularization enable to recover the exact prices. Finally, in Section 5, we present some practical issues and we validate our approach regarding the global problem.

2 Inexact dual resolution and noisy behaviour

This section gives more details about the inexact unit-commitment resolution and exhibits the noisy behaviour that we observe and want to reduce. Let us start with some notations.

We call *oracle* a procedure that solves the Lagrangian subproblems (3) and evaluates the dual function Θ : for a given multiplier $u \in \mathbb{R}^T$, it returns the value $\Theta(u)$ and a subgradient $g \in \partial\Theta(u)$, i.e. by definition, a vector $g \in \mathbb{R}^T$, such that $\Theta(v) \geq \Theta(u) + g^\top(v - u)$, for all $v \in \mathbb{R}^T$. In our case, some subproblems cannot be solved exactly within the time limit; this which leads to an *inexact oracle*, which produces an approximate value of the dual function and an approximate subgradient, in the following sense. We suppose that the oracle computes a feasible possibly non-optimal solution and that we know an upper bound ε on the distance to the optimal solution. Thus, for $u \in \mathbb{R}^T$, the inexact oracle returns a value $\tilde{\Theta}(u) \in \mathbb{R}$ such that

$$\Theta(u) - \varepsilon \leq \tilde{\Theta}(u) \leq \Theta(u), \quad (5)$$

and an ε -subgradient $\tilde{g} \in \mathbb{R}^T$ satisfying

$$\forall v \in \mathbb{R}^T, \quad \Theta(v) \geq \Theta(u) + \tilde{g}^\top(v - u) - \varepsilon.$$

In our synthetic example, we introduce and control the noise of the oracle. For the EDF problem, the mixed integer linear solver used to solve the hydro-valleys subproblems gives the best feasible solution found in a given time and an upper bound on the error (computed using the optimal value of the linear relaxation). On our 27 instances, the average upper bound on the error is about 3%/00.

The objective now is to solve (4), i.e. to minimize the dual function, which is convex by construction (as a maximum of affine functions), generally nonsmooth, and inexactly known via the oracle. Inexact bundle methods (Oliveira et al, 2012; Kiwiel, 2006; Emiel and Sagastizábal, 2010) are then the methods of choice to solve this problem. Indeed, one can show (Kiwiel, 2006) that the inexact bundle method converges to a price vector \tilde{u} such that:

$$\Theta(u^*) \leq \Theta(\tilde{u}) \leq \Theta(u^*) + \delta + \varepsilon, \quad (6)$$

where δ is the stopping tolerance of the bundle method and u^* the exact optimal prices. This previous guarantee is quite strong: since we know Θ with an error of ε , we inevitably undergo an error of ε on the “optimal” value that we compute and if we neglect the stopping tolerance δ , the approximation (6) says precisely that we minimize Θ with an error of at most ε . However, there is no such guarantee on how the inexact prices \tilde{u} are close to the exact one u^* . In practice, we observe on our two examples that they can be very different.

Example 3 (Inexact prices in EDF UC problem) Figure 1 represents the prices for an instance provided to us by EDF, with two models and two oracles. Red prices are obtained when the hydro-valleys subproblems are modeled as linear programs (LP) so that the oracle is exact. Blue prices are obtained when hydro-valleys subproblems are modeled as mixed integer linear programs (MILP) and then the oracle is inexact. For industrial privacy, price scales are not provided.

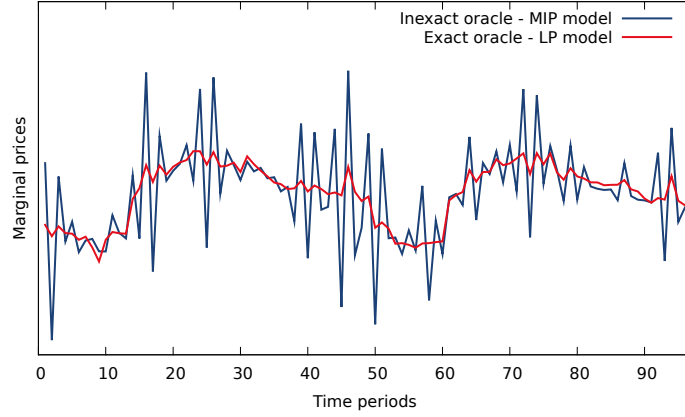


Fig. 1 Exact and inexact prices over time periods in EDF problem

The inexact prices (in blue) show a noisy behaviour: they have more variations and oscillations than the exact prices (in red), and more important peaks. There are two possible causes for this noisy behaviour:

1. the change in the model (LP to MILP),
2. the inexactness of the oracle.

We could distinguish the role of each cause by dropping the computing time constraint to give to the MILP solver the time to reach optimality. Unfortunately, the complexity of EDF subproblems modeling makes this impossible. To clearly highlight the effect of inexactness, we reproduce a similar phenomenon with the simple UC problem. \square

Example 4 (Inexact prices in the synthetic problem) Figure 2 shows the demand and the computed exact prices for a particular instance of the synthetic example (2) with $T = 14$, $n = 18$. In this example, exact prices are a few euros most of the time, with a peak to around 400 euros at $t = 4$ and $t = 5$ (which correspond to the highest demand periods). Figure 3 shows again the exact prices, together with inexact prices obtained with two different oracle errors. We see that inexact prices show more variations and in particular important peaks at $t = 7$ and $t = 9$, for oracle error $\varepsilon = 15\%$.




Fig. 2 Demand and exact prices over $T = 14$ time periods for the synthetic problem

\square

These examples suggest that the oracle inexactness gives undesirable behaviour to the dual solutions, which may trouble their use as price indicators. Our objective now is to get rid of the artificial oscillations caused by the noisy oracle to improve the quality of the computed prices.




Fig. 3 Prices corresponding to different oracle error $\varepsilon = 0, 10$, and 15% , over time periods

3 Prices stabilization by penalization

3.1 Dual penalization and primal interpretation

A first idea to limit the instability of the prices would be to smooth the curves a posteriori (by a running average for example). The drawback of this blind approach is that all the oscillations would be treated the same way: the “natural” ones (due to peak hours for example) as well as the “artificial” ones (due to oracle inexactness). For instance, for our simple problem (2), in figure 3, such approach will attenuate similarly the meaningful jump of periods $t = 4$ and $t = 5$, and the noisy peaks of periods $t = 7$ and $t = 9$. This example shows that we should incorporate the stabilization within the computation of the prices and not treat it as a separate task.

To reduce the instability of the prices without losing pertinent information, our idea is simple: we propose to add a parametrized penalization term in the dual objective function, that limits prices variations. Instead of (4), we consider a “stabilized dual” problem:

$$\begin{cases} \min \Theta(u) + S(u) \\ \text{s.t. } u \in \mathbb{R}^T, \end{cases} \quad (7)$$

where S is a well-chosen (closed) convex penalization function that aims to impose to the dual variables a desirable behaviour. We can think of different types of stabilization. For example, we can choose a function which is the combination of a norm and a function expressing the variation of the prices (as first or second order discrete derivatives):

$$S_1^1(u) = \text{TV}_\alpha(u) = \sum_{t=1}^{T-1} \alpha_t |u_{t+1} - u_t| \quad (\text{discrete total variation})$$

$$S_\infty^1(u) = \max_t \{\alpha_t |u_{t+1} - u_t|\}$$

$$S_2^1(u) = \sum_t \alpha_t |u_{t+1} - u_t|^2$$

$$S_1^2(u) = \sum_t \alpha_t |-u_t + 2u_{t+1} - u_{t+2}|$$

$$S_\infty^2(u) = \max_t \{\alpha_t |-u_t + 2u_{t+1} - u_{t+2}|\}$$

$$S_2^2(u) = \sum_t \alpha_t |-u_t + 2u_{t+1} - u_{t+2}|^2$$

where $\alpha = (\alpha_t) \in \mathbb{R}_+^{T-1}$ is a vector of nonnegative coefficients. We could also consider constrained versions of the above penalties. In Section 4, we will give a particular emphasis on the first penalty function, the discrete total variation, which is a popular regularization function, used successfully in many fields (for example in image processing for noise removal, see e.g. (Rudin et al, 1992; Chambolle, 2004)). Similar developments can be made for other stabilization functions.

In general, when $\alpha = 0$, the corresponding problem (7) is the initial problem (4). When α is very large, any constant $u \in \mathbb{R}^T$ is a solution of the problem. For well-chosen $\alpha \in \mathbb{R}_+^{T-1}$, we will see in forthcoming experiments that the solutions do not show the noisy behaviour anymore. Tuning the parameter α to get the desirable behaviour is obviously a delicate point. First, α should be adapted to make S have the same order of magnitude as Θ in (7). Second, α could also be chosen to more or less penalize the price variations according to the period. Typically, if we know peak periods where it is natural that the prices jump, we can choose the corresponding components α_t small or even null.

A primal view on the dual stabilization in (7) can also help in choosing a good α . The next proposition studies the case of a general stabilization. Forthcoming Corollary 1 specializes it in the case of total variation regularization.

Proposition 1 (Stabilized primal problem) *The dualization of the balance constraint in the following problem (8) leads to the stabilized dual problem (7).*

$$\begin{cases} \min c(p) + V(z) + S^*(\delta) \\ z = (d - \delta) - Ap \\ p \in \mathcal{P} \\ z, \delta \in \mathbb{R}^T. \end{cases} \quad (8)$$

Proof The dual function corresponding to the dualization of the linking constraint of (8) is for $u \in \mathbb{R}^T$: (with the same change of signs as in (4))

$$\begin{aligned} \Theta_S(u) &:= - \min_{p, z, \delta} \{c(p) + V(z) + S^*(\delta) + u^\top (d - \delta - Ap - z)\} \\ &= - \min_{p \in \mathcal{P}} \{c(p) + u^\top (d - Ap)\} - \min_{z \in \mathbb{R}^T} \{V(z) - u^\top z\} - \min_{\delta \in \mathbb{R}^T} \{S^*(\delta) - u^\top \delta\} \\ &= - \min_{p \in \mathcal{P}} \{c(p) + u^\top (d - Ap)\} + \max_{z \in \mathbb{R}^T} \{u^\top z - V(z)\} + \max_{\delta \in \mathbb{R}^T} \{u^\top \delta - S^*(\delta)\} \\ &= -\theta(u) + V^*(u) + S^{**}(u) \\ &= \Theta(u) + S(u) \quad (\text{because } S \text{ is closed and convex}). \end{aligned}$$

Therefore, the corresponding dual problem is (7). □

The effects of price stabilization on the primal problem interpret as follows: at each period of time t , demand forecast can be perturbed by a quantity δ_t , and this perturbation is penalized in the objective through the convex conjugate S^* of the stabilization function S .

3.2 Penalization in practice, and illustration

The additional term $S(u)$ in the dual objective does not change the nature of the dual problem (4). This is still an unconstrained convex nonsmooth problem, that can be solved by (inexact) bundle method. Actually, there are two ways to handle the additional term: (1) we can use the same bundle method and change the oracle, including S ; or (2) we can use a *disaggregated* bundle method (Bacaud et al, 2001) to exploit the simple structure of S . For simplicity, we consider in our computational tests the first strategy; so we deal with an oracle that returns, for u in \mathbb{R}^T

- the inexact dual objective value: $\tilde{\Theta}(u) + S(u)$, which has the same absolute error ε as $\tilde{\Theta}(u)$,
- the vector $\tilde{g} + g_S$, with $\tilde{g} \in \partial_\varepsilon \Theta(u)$ and $g_S \in \partial S(u)$, which is also an ε -subgradient of $(\Theta + S)$ at u .

We apply our stabilization approach on both our synthetic and real-life UC problems.

Example 5 (Price stabilization for the synthetic problem) Let us come back to the instance presented in Example 4. We set oracle error to 15%, and use as stabilization function the total variation, with constant coefficients ($\alpha_t = \alpha$, for all t) for simplicity. Which means that we apply the same penalization to all the jumps of the prices. Then, for different values of α , we obtain the range of prices represented in Figure 4. When $\alpha = 0$, we have the inexact prices of Figure 3. When α increases, we see that the oscillations are more and more attenuated, in particular on periods 7 and 9. Finally, for $\alpha \geq 300$, the peak between periods 4 and 5 is altered and it finally disappears. For α even larger, the prices are totally

flat. Note that for $\alpha \approx 50$, we almost recover the exact prices: artificial peaks of periods 7 and 9 are hardly visible whereas natural jumps of periods 4 and 5 are unaltered. It is remarkable that the total variation penalization erases first the artificial variations; we will give some theoretical insights into this fact in forthcoming Theorem 1. \square




Fig. 4 Variation of stabilized prices for different weights of total variation regularization in our simple problem

Example 6 (Price stabilization for EDF) We apply our total variation stabilization to a realistic instance provided by EDF (for industrial privacy, we only give the relative value of prices). Figure 5 represents the prices of the stabilized problem with the function TV_α , for different constant vectors α . When α is *small*, we observe that the stabilized prices have roughly the same look as the initial ones, with each peak attenuated. Some oscillations disappeared, for example between periods $t = 55$ to $t = 60$. For a *medium* α , prices are much more stable. We can notice though that they follow the main tendencies of the initial prices. When α is larger, the prices tend to be flat.




Fig. 5 Variation of stabilized prices for different weights of total variation regularization in EDF problem

Since rush hours are generally predictable, we can take them into account when choosing the coefficients α_t . For example, we tried:

$$\alpha_{\text{adaptive}}^t = \begin{cases} \alpha_{\text{small}}^t & \text{if } t \text{ in rush hours of the first day,} \\ \alpha_{\text{medium}}^t & \text{otherwise.} \end{cases}$$

Figure 6 gives the corresponding stabilized prices. We see clearly attenuated oscillations at the rush hours of the first day, and more stable prices on the other periods. \square




Fig. 6 Variation of stabilized prices with TV_α using a non constant vector α in EDF problem. The components corresponding to rush hours of the first day are less penalized than the regular periods

4 More on total variation regularization

In this section, we look closely to the stabilization by total variation to give some more theoretical insights. Given $\alpha \in \mathbb{R}_+^{T-1}$, we consider the (weighted) total variation function:

$$TV_\alpha(u) := \sum_{t=1}^{T-1} \alpha_t |u_{t+1} - u_t|. \quad (9)$$

Note that we can write $TV_\alpha = \|\cdot\|_1 \circ B_\alpha$, where B_α is the matrix of $\mathbb{R}^{T-1} \times \mathbb{R}^T$ defined by:

$$B_\alpha = \begin{pmatrix} -\alpha_1 & \alpha_1 & & & 0 \\ & -\alpha_2 & \alpha_2 & & \\ & & \ddots & \ddots & \\ 0 & & & -\alpha_{T-1} & \alpha_{T-1} \end{pmatrix}. \quad (10)$$

Obviously, TV_α is a (closed) convex function, as the composition of a norm and a linear function. Subdifferential calculus rules (see (Hiriart-Urruty and Lemaréchal, 2001)) give that $\partial TV_\alpha(u) = B_\alpha^\top \partial \|\cdot\|_1(B_\alpha u)$ and thus a subgradient of TV_α at u has the form:

$$g_t = \begin{cases} \alpha_1 s_1 & \text{if } t = 1, \\ \alpha_{t-1} s_{t-1} - \alpha_t s_t & \text{if } t \in \{2, \dots, T-1\}, \\ \alpha_{T-1} s_{T-1} & \text{if } t = T, \end{cases} \quad (11)$$

where $s_t = -1$ if $u_{t+1} < u_t$, $s_t = 1$ if $u_{t+1} > u_t$ and $s_t \in [-1; 1]$ if $u_{t+1} = u_t$.

4.1 Total variation stabilized primal problem

We give the impact of the total variation penalization on the primal UC problem. Using Proposition 1 with the explicit expression (9), we get the following corollary:

Corollary 1 (Primal interpretation of the regularization) *A primal problem giving the dual problem (7) stabilized with TV_α by dualizing the supply-demand constraint is the following*

$$\begin{cases} \min c(p) + V(z) \\ \text{s.t. } z = (d - \delta) - Ap \\ \sum_{i=1}^T \delta_i = 0 \\ \left| \sum_{i=1}^t \delta_i \right| \leq \alpha_t, \forall t = 1, \dots, T-1 \\ p \in \mathcal{P}, \delta \in \mathbb{R}^T. \end{cases} \quad (12)$$

Proof Using Proposition 1, we just have to compute the convex conjugate function TV^* . Theorem (Hiriart-Urruty and Lemaréchal, 2001, §E.2.2) gives:

$$\text{TV}_\alpha^*(x) = \inf\{\|y\|_1^* : y \in \mathbb{R}^{T-1}, B_\alpha^\top y = x\} = \inf\{i_{B_\infty(1)}(y) : y \in \mathbb{R}^{T-1}, B_\alpha^\top y = x\}$$

We can therefore write:

$$\text{TV}_\alpha^*(x) = \begin{cases} 0 & \text{if there exists } y \in B_\infty(1) \text{ such that } B_\alpha^\top y = x, \\ +\infty & \text{otherwise.} \end{cases}$$

Suppose now that x is in the image of B_α^\top , then simple computations give

$$\sum_{i=1}^T x_i = 0 \quad \text{and} \quad \text{for all } t \text{ such that } \alpha_t = 0, \quad \sum_{i=1}^t x_i = 0. \quad (13)$$

Let y be in the inverse image $(B_\alpha^\top)^{-1}(\{x\})$. We can show that two cases are possible for its components:

- For t such that $\alpha_t > 0$, $y_t = -\frac{1}{\alpha_t} \sum_{i=1}^t x_i$, and $|y_t| \leq 1$ is equivalent to $\left| \sum_{i=1}^t x_i \right| \leq \alpha_t$.
- For t such that $\alpha_t = 0$, y_t is free. Let us take $y_t = 0$ in this case.

By writing the second part of condition (13) as $\left| \sum_{i=1}^t x_i \right| \leq \alpha_t$, for all t such that $\alpha_t = 0$, the set of vectors where TV_α^* is zero is

$$\mathcal{C}_1 = \left\{ \delta \in \mathbb{R}^T : \left| \sum_{i=1}^t \delta_i \right| \leq \alpha_t, \forall t = 1, \dots, T-1, \sum_{i=1}^T \delta_i = 0 \right\}.$$

Finally, we can write $\text{TV}_\alpha^* = i_{\mathcal{C}_1}$ and plugging this expression in the problem (8), we get (12). \square

In this case, demand perturbation is constrained: its mean must be zero and the partial sums of its components are bounded by α . Note that the larger an α_t is (i.e. the more we penalize prices jump between periods t and $t+1$), the less constrained is the demand perturbation at t . Conversely, if $\alpha_t = 0$, i.e. prices jump between t and $t+1$ is free, then the total perturbation until period t must be zero. It means that the cumulative demand and perturbed demand until period t coincide. Note that if we know a limit under which a perturbation of the demand is acceptable (for example 1% of the forecast), we could use it to choose α .

4.2 Exact prices recovery in a particular case

In Example 5, we observed that for a certain choice of the stabilization parameter α , the stabilized prices are close to the exact ones. In this section, we prove that in a particular simple context, the stabilized prices are precisely the exact prices. We prove that we can indeed recover piecewise constant prices with a jump. We consider a general convex (dual) function Θ and a specific (artificial) oracle noise, which pushes the solutions toward a chaotic behaviour. Then, we show that there exists a constant vector α such that the solution of the stabilized problem using TV_α is precisely the exact prices.

The assumptions on the oracle noise are inspired from results on noise removal in image processing. In (Strong and Chan, 2003), it is shown that total variation regularization recovers exactly properties of the image in certain situations. Contrary to (Strong and Chan, 2003) though, we use directly convex analysis tools to prove our recovery result.

Theorem 1 *Let Θ be a convex function from $\mathbb{R}^T \rightarrow \mathbb{R}$. Assume that an optimal solution of the “initial” problem*

$$\min_{u \in \mathbb{R}^T} \Theta(u),$$

denoted by u^* , has the following form: $u_t^* = 1$ if $t \leq t_0$, $u_t^* = 0$ otherwise, for some $t_0 \in \{2, \dots, T-1\}$. Let v be a “noisy” vector of \mathbb{R}^T satisfying the following conditions:

$$\frac{1}{t_0} \sum_{i=1}^{t_0} v_i = 1 + \frac{\beta}{t_0}, \quad \frac{1}{T-t_0} \sum_{i=t_0+1}^T v_i = -\frac{\beta}{T-t_0}, \quad (14)$$

$$\min_{1 \leq i \leq t_0} v_i \geq 1, \quad \max_{t_0 < i \leq T} v_i \leq 0. \quad (15)$$

If $\alpha = \beta\sigma$, then u^* is also the optimal solution of the “stabilized” noisy problem

$$\min_{u \in \mathbb{R}^T} \Theta(u) + \frac{\sigma}{2} \|u - v\|^2 + \text{TV}_\alpha(u).$$

Proof Denote $f(u) = \Theta(u) + \frac{\sigma}{2} \|u - v\|^2 + \text{TV}_\alpha(u)$. As a sum of convex functions and a norm, f is strongly convex and u^* is its unique minimum if and only if $0 \in \partial f(u^*)$. We have:

$$\partial f(u^*) = \partial \Theta(u^*) + \sigma(u^* - v) + \partial \text{TV}_\alpha(u^*)$$

By construction, we have $0 \in \partial \Theta(u^*)$. So, it is sufficient to find a vector $g \in \partial \text{TV}_\alpha(u^*)$ such that

$$\sigma(u^* - v) + g = 0. \quad (16)$$

Let us write $\text{TV}_\alpha = \|\cdot\|_1 \circ B_\alpha$, where B_α is defined in (10), with the abuse of notation $\alpha_t = \alpha$, for all t . From the expression (11) of a subgradient of TV_α at $u = u^*$, it follows easily that finding $g \in \partial \text{TV}_\alpha(u^*)$ satisfying (16) amounts to finding $s_1, \dots, s_{t_0-1}, s_{t_0+1}, \dots, s_{T-1} \in [-1; 1]$ such that:

$$\begin{aligned} \sigma(u_1^* - v_1) & - \alpha s_1 & = 0 \\ \sigma(u_t^* - v_t) & - \alpha(s_{t-1} - s_t) & = 0 \quad \text{for } 2 \leq t \leq t_0 - 1 \\ \sigma(u_{t_0}^* - v_{t_0}) & - \alpha(s_{t_0-1} + 1) & = 0 \\ \sigma(u_{t_0+1}^* - v_{t_0+1}) & - \alpha(-1 - s_{t_0+1}) & = 0 \\ \sigma(u_t^* - v_t) & - \alpha(s_{t-1} - s_t) & = 0 \quad \text{for } t_0 + 2 \leq t \leq T - 1 \\ \sigma(u_T^* - v_T) & - \alpha s_{T-1} & = 0 \end{aligned}$$

Using assumptions (14), the previous system gives :

$$\begin{aligned} s_t &= \frac{\sigma}{\alpha} \sum_{i=1}^t (u_i^* - v_i) = \frac{\sigma}{\alpha} \sum_{i=1}^t (1 - v_i) \quad \text{if } 1 \leq t \leq t_0 - 1, \\ s_t &= -\frac{\sigma}{\alpha} \sum_{i=t+1}^T (u_i^* - v_i) = \frac{\sigma}{\alpha} \sum_{i=t+1}^T v_i \quad \text{if } t_0 + 1 \leq t \leq T - 1. \end{aligned} \quad (17)$$

It remains to check that these components s_t lie in $[-1; 1]$. Using assumptions (15), we see from expression (17) that for $1 \leq t \leq t_0 - 1$, s_t are negative, decreasing and bounded from below by $s_{t_0-1} = -1 + \frac{\sigma}{\alpha}(v_{t_0-1} - 1) \geq -1$. Similarly, for $t_0 + 1 \leq t \leq T - 1$, s_t are negative, increasing and bounded from below by $s_{t_0+1} = -1 + \frac{\sigma}{\alpha}(-v_{t_0-1}) \geq -1$. Therefore, equation (16) holds for the vector g , i.e. $0 \in \partial f(u^*)$. \square

5 Numerical experiments for EDF

In this section, we illustrate our approach on real-life UC problems, provided to us by EDF. We will show that we can drastically reduce the noisy behaviour of the prices, without losing pertinent information on the optimization problem.

Recall that approximation (6) essentially says that the computed prices \tilde{u} minimize the dual function Θ with the “minimal” error, which is ε the error of the oracle. We would like the same guarantee for \tilde{u}_s the optimal prices of the stabilized problem (7); that is, we would like that \tilde{u}_s also minimizes Θ up to an error of ε .

For EDF industrial problems, we cannot compute – and therefore compare – the exact values of Θ at the optimal prices u^* and the stabilized prices \tilde{u}_s . Fortunately, properties of inexact bundle methods recalled in Section 2 imply that the “exact” difference

$$\Delta_{\text{ex}} := \Theta(u^*) - \Theta(\tilde{u}_s)$$

is close to the observed difference

$$\Delta_{\text{obs}} := \tilde{\Theta}(\tilde{u}) - \tilde{\Theta}(\tilde{u}_s).$$

More precisely, we have the following lemma:

Lemma 1 *Let ε be the maximum error of the oracle and δ the stopping tolerance of the inexact bundle method. Then we have:*

$$|\Delta_{\text{obs}} - \Delta_{\text{ex}}| \leq \delta + 2\varepsilon$$

Proof First, using (5) and (6), we get that

$$\Theta(u^*) - \varepsilon \leq \tilde{\Theta}(\tilde{u}) \leq \Theta(u^*) + \delta + \varepsilon.$$

This implies that $\tilde{\Theta}(\tilde{u})$ approximates $\Theta(u^*)$ up to $\delta + \varepsilon$. Second, using (5) at \tilde{u}_s , we have

$$\Theta(\tilde{u}_s) - \varepsilon \leq \tilde{\Theta}(\tilde{u}_s) \leq \Theta(\tilde{u}_s),$$

so $\tilde{\Theta}(\tilde{u}_s)$ approximates $\Theta(\tilde{u}_s)$ with again an error of ε . Therefore, we have:

$$|\Delta_{\text{obs}} - \Delta_{\text{ex}}| \leq |\Theta(u^*) - \tilde{\Theta}(\tilde{u})| + |\tilde{\Theta}(\tilde{u}_s) - \Theta(\tilde{u}_s)| \leq \delta + \varepsilon + \varepsilon,$$

which is the desired approximation. □

Neglecting the stopping tolerance δ , we deduce that Δ_{obs} approximates Δ_{ex} with an error of 2ε . Thus, if Δ_{obs} is of the order of ε then, so is Δ_{ex} . In other words, the stabilized prices \tilde{u}_s have the same quality in terms of Θ as the initial prices \tilde{u} .

We analyze the impact of the stabilization on the 27 EDF instances. The (relative) error of the oracle is between 0.1‰ and 2%, with an average of about 3‰. Table 1 summarizes the average results between the stabilized prices \tilde{u}_s and the original ones \tilde{u} for the different settings of the stabilization parameter α of Section 3.2. Are reported:

- the average gain in total variation $(\text{TV}(\tilde{u}) - \text{TV}(\tilde{u}_s))/\text{TV}(\tilde{u})$,
- the average lost in $\tilde{\Theta}$: $\Delta_{\text{obs}}/\tilde{\Theta}(\tilde{u})$.

We observe that the total variation of the stabilized prices is considerably lower than the original one. Even for a small α , total variation decreases, in average, by half; and when α is larger the gain in total variation is more important.

We see that the loss in Θ also increases with α , but to a lesser extent. We note that for medium and adaptive α , the prices of the stabilized problem give an error on $\tilde{\Theta}$ smaller than the oracle error 3‰. Thus, these prices (for medium and adaptive α) minimize the dual function with an error of the same order as the oracle error, while being significantly more stable (more than 80% of total variation reduction).

Table 1 Average results on the 27 real-life UC instances for different values of α

Stabilization	Avg. gain in TV	Avg. lost in $\tilde{\Theta}$
Small α	52%	0.6 ‰
Medium α	88%	2.8 ‰
Large α	93%	3.3 ‰
Adaptive α	83%	2.6 ‰

6 Conclusions

In this paper, we considered unit-commitment resolution by Lagrangian duality, in the case where the subproblems are solved inexactly. We showed that the minimization of the resulting approximate dual function leads to unstable dual solutions. To limit this phenomenon, we proposed to add a stabilization in the dual objective, which penalizes the variations of the solutions. We showed that this can be interpreted as introducing a perturbation of the demand forecast in the primal problem.

We showed in a simple context that a certain choice of the stabilizer allows to actually recover the exact solutions. For a synthetic UC problem, we observed that we can indeed obtain stabilized prices close to the exact ones. For EDF industrial problems, we obtained a good compromise between the stabilization of the prices and the optimization of the dual function.

We finish with emphasizing that our approach is general, easy to implement and test, but requires adequate knowledge of the particular UC problem to fix the stabilization parameter. Efficient implementation can exploit the simple structure of the stabilization function in the optimization algorithm. Instead of including the stabilization term in the black-box oracle, one could consider explicitly the two terms of the dual function and use a disaggregated bundle method to accelerate the resolution.

Acknowledgements We would like to thank Grace Doukopoulos (EDF research) for providing us UC instances and the numerical results with EDF solver. The authors also acknowledge support of the “action strategique EDF/INRIA : Coûts marginaux en production journalière d’électricité”

Bibliography

- Bacaud L, Lemaréchal C, Renaud A, Sagastizábal C (2001) Bundle methods in stochastic optimal power management: A disaggregated approach using preconditioners. *Computational Optimization and Applications* 20:227 – 244
- Boyd S, Vandenberghe L (2004) *Convex Optimization*. Cambridge University Press, New York, NY, USA
- Chambolle A (2004) An algorithm for total variation minimization and applications. *Journal of Mathematical Imaging and Vision*
- Dubost L, Gonzalez R, Lemaréchal C (2005) A primal-proximal heuristic applied to the French unit-commitment problem. *Mathematical Programming* 104:129–151
- Emiel G, Sagastizábal C (2010) Incremental-like bundle methods with application to energy planning. *Computational Optimization and Applications* 46:305–332
- Frangioni A (2010) Unit commitment problems: A tale in Lagrangian optimization. *Optima* 84 104:6–7
- Hechme-Doukopoulos G, Brignol-Charousset S, Malick J, Lemaréchal C (2010) The short-term electricity production management problem at EDF. *Optima* 84 104:2–6
- Hiriart-Urruty JB, Lemaréchal C (1993) *Convex Analysis and Minimization Algorithms II*. Springer Verlag
- Hiriart-Urruty JB, Lemaréchal C (2001) *Fundamentals of convex analysis*. Springer Verlag
- Kiwiel K (2006) A proximal bundle method with approximate subgradient linearizations. *SIAM Journal on optimization* 4:1007–1023
- Oliveira W, Sagastizábal C, Lemaréchal C (2012) Inaccurate bundle methods in depth: unified analysis. (In preparation)
- Padhy N (2004) Unit commitment - a bibliographical survey. *IEEE Trans Power Syst* 19:1196 – 1205
- Rudin L, Osher S, Fatemi E (1992) Nonlinear total variation based noise removal algorithms. *Physica D* 60:259 – 268
- Sheble G, Fahd G (1994) Unit commitment literature synopsis. *IEEE Trans Power Syst* 9:128 –135
- Strong D, Chan T (2003) Edge-preserving and scale-dependent properties of total variation regularization. *Inverse Problems* 19(6):S165