CIRRELT

Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport

Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation

A Study of Auction Mechanisms for Multilateral Procurement Based on Subgradient and Bundle Methods

Jawad Abrache
Teodor Gabriel Crainic
Michel Gendreau
Tarik Aouam

August 2012

CIRRELT-2012-40

Bureaux de Montréal :

Université de Montréal
C.P. 6128, succ. Centre-ville
Montréal (Québec)
Canada H3C 3J7
Téléphone : 514 343-7575
Télécopie : 514 343-7121

Bureaux de Québec :

Université Laval
2325, de la Terrasse, bureau 2642
Québec (Québec)
Canada G1V 0A6
Téléphone : 418 656-2073
Télécopie : 418 656-2624

www.cirrelt.ca

A Study of Auction Mechanisms for Multilateral Procurement Based on Subgradient and Bundle Methods

Jawad Abrache[1], Teodor Gabriel Crainic[2,3,*], Michel Gendreau[2,4], Tarik Aouam[1]

[1] Al Akhawayn University, P.O. Box 104, Avenue Hassan II, Ifrane, 53000, Morocco

[2] Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT)

[3] Department of Management and Technology, Université du Québec à Montréal, P.O. Box 8888, Station Centre-Ville, Montréal, Canada H3C 3P8

[4] Department of Mathematics and Industrial Engineering, École Polytechnique de Montréal, P.O. Box 6079, Station Centre-ville, Montréal, Canada H3C 3A7

Abstract. The use of iterative auctions is very common in procurement processes, where the marketmaker often does not have access to complete and truthful information about the bidders' private valuations of the resources on sale. The literature on the design of iterative mechanisms for combinatorial auctions has addressed only the most basic cases and has been dominated by primal-dual approaches. In this paper, we consider a general production/consumption exchange of interdependent goods, for which we investigate iterative auction mechanisms based on mathematical programming dual decomposition methods. We focus on Lagrangian relaxation and the solution of the Lagrangian dual through subgradient algorithms and the bundle method. A case study of a simulated wood chip market is used to evaluate numerically the efficiency of the mechanisms.

Keywords: Procurement, iterative combinatorial auction mechanisms, mathematical programming, decomposition methods.

Acknowledgements. We would like to thank Benoît Bourbeau and Geneviève Hernu who helped in coding the numerical methods and in realizing the experimental study, as well as Prof. Frangioni who kindly provided his implementation of the bundle algorithm. This project has received financial support from the Bell Canada University Laboratories (LUB), the Natural Sciences and Engineering Research Council of Canada (NSERC) through their Discovery and Collaborative Industry-University Research grant programs, and the Fonds de recherche du Québec – Nature et technologies (FRQNT).

---

* Corresponding author: TeodorGabriel.Crainic@cirrelt.ca

Dépôt légal – Bibliothèque et Archives nationales du Québec
Bibliothèque et Archives Canada, 2012

1 Introduction and background

The spectacular growth of electronic commerce and information technology made Internet a place of choice for conducting business. The apparition of electronic marketplaces that followed has brought forth issues in relation to their structure and organization. Among these issues, the question of designing business rules governing the transactions in these marketplaces is of prime importance and has contributed heavily to the renewed interest in the discipline of market mechanism design.

A market mechanism can be defined as a set of deterministic rules that specify an allocation of the items traded in the market to the participants, as well as the corresponding payments the latter should make or receive. Market mechanisms can be classified into two classes (Mas-Colell et al., 1995): direct-revelation and indirect-revelation mechanisms. The first requires the participants to completely and truthfully report their types to the market-maker who "clears" the market such that an economic objective (e.g., overall social welfare) is optimized. For instance, in a production-consumption economy sellers need to disclose their production technologies and cost functions, while buyers need to reveal their preferences for the goods and their consumption constraints. Unfortunately, this strict requirement on information disclosure is very rarely realistic. Most of the time, the participants are unwilling and/or unable to disclose their private information to a third party. Indirect-revelation mechanisms helps circumvent some of these problems by not requiring systematic access by the market-maker to all the information. Iterative auctions, in particular, are an important family of indirect-revelation mechanisms that allow for progressive revelation of pertinent information. In each round of an iterative auction, sellers and buyers submit bids to sell or buy items, to which the market-maker responds by determining provisional allocations and payments, and by sending "signals" about the state of the auction; sellers and buyers then re-evaluate accordingly their bids, and so on. It is noteworthy that bids neither have to represent complete preferences - they only reflect participants' needs given the observed signals, nor to convey truthful information - unless the auction incentives prevents preference misrepresentation. The design of iterative auction mechanisms in optimized markets has attracted much attention recently. Yet, most of the auctions suggested in the literature are primal-dual algorithms that capitalize on linear programming formulations of the market clearing problem.

Mathematical programming decomposition approaches offer an interesting alternative. These methods have been used for decades to address large-scale structured optimization problems. Yet, their potential for decentralized decision making, if thoroughly analyzed, has seldom been exploited in the design of auction mechanisms. Hence, this paper aspires to contribute to a better understanding of the decomposition methods used as iterative auction mechanisms. The focus is on Lagrangian relaxation solved using basic subgradient algorithms and the bundle method. The methodology of the paper is as follows. We consider a general combinatorial exchange economy in which the participants trade heterogeneous divisible commodities. We assume that the participants are self-interested buyers and sellers who maximize their economic own surplus and react optimally to the prices announced on the market. We first formulate the allocation and payment rules of an ideal direct-revelation mechanism with the objective of maximizing the overall social efficiency of the market. Then we show that, under appropriate assumptions, the application of Lagrangian relaxation to the centralized allocation problem leads to indirect auction mechanisms that have the ability to achieve social efficiency without requiring complete information revelation from the participants. The efficiency of the different auctions derived is evaluated numerically on a simulated wood chip market.

The paper makes several important theoretical and practical contributions. While mathematical programming decomposition methods have been presented in the past as market mechanisms (notably Lagrangian relaxation, see de Vries and Vohra (2003)), the analysis has been limited to the one-sided case (that is, the market-maker selling several different items to many buyers). To the best of our knowledge, this is the first attempt to analyze these methods in a many-to-many exchange context. Moreover, the numerical results provide interesting insights into the potential and limitations of auction mechanisms based on decomposition approaches.

The paper is organized as follows. We formulate in Section 2 the problems of determining a socially efficient allocation and the corresponding equilibrium prices in a centralized many-to-many direct-revelation market. In Section 3, we present two relaxation-based methods, using the subgradient algorithm and the bundle method, respectively, and interpret them as iterative auctions. Finally, we devote Section 4 to the experimental study.

2 Centralized market-clearing

We consider a simplified economy with a set of divisible goods on sale and two categories of participants, sellers and buyers. Sellers have the capacity to produce the goods according to their own technology and production cost functions, while buyers consume goods either directly or as inputs to a transformation process. Hence, buyers have preferences for bundles of goods on sale and may also face technological requirements that constrain their consumption. The following notation is introduced:

- $\mathcal{L}$ (resp. $\mathcal{S}$, $\mathcal{J}$): set of goods (resp. sellers, buyers).
- $q_{s,l}$ (resp. $q_{j,l}$): quantity of good $l, l \in \mathcal{L}$ produced by seller s, $s \in \mathcal{S}$ (resp. consumed by buyer j, $j \in \mathcal{J}$).
- $\mathcal{D}_s$ (resp. $\mathcal{D}_j$): production (resp. consumption) feasibility set of seller s (resp. buyer j), containing all admissible quantities $q_s = \{q_{s,l}\}_{l\in\mathcal{L}}$ (resp. $q_j = \{q_{j,l}\}_{l\in\mathcal{L}}$)

that seller s (resp. buyer j) may produce (resp. consume). These sets are assumed to be convex and bounded.

- $C_s(.)$: production cost function of seller s, $s \in \mathcal{S}$; that is, $C_s(q_s)$ is the cost to seller s of producing q_s. This cost function is assumed to be continuous, convex, and monotone increasing.
- $V_j(.)$: valuation function of buyer j, $j \in \mathcal{J}$; similarly, $V_j(q_j)$ is buyer j's preference for consuming q_j. This valuation function is assumed to be continuous, concave, and monotone increasing.

In a direct-revelation market mechanism, sellers and buyers need to communicate to the market-maker their production and consumption feasibility sets and their cost and valuation functions, respectively. The mechanism's output is an allocation of goods and payments sellers (resp. buyers) need to make (resp. receive). The market-maker needs to determine a *socially-efficient* allocation, that is, a feasible allocation of goods that maximizes the overall welfare of all sellers and buyers. More precisely, a socially-efficient allocation is a solution of model (MC):

$$\max \quad \sum_{j \in \mathcal{J}} V_j(q_j) - \sum_{s \in \mathcal{S}} C_s(q_s) \tag{1}$$

$$s.t. \quad \sum_{j \in \mathcal{J}} q_{j,l} - \sum_{s \in \mathcal{S}} q_{s,l} = 0, \quad l \in \mathcal{L} \tag{2}$$

$$q_j \in \mathcal{D}_j, \quad j \in \mathcal{J}; q_s \in \mathcal{D}_s, \quad s \in \mathcal{S} \tag{3}$$

Model (MC) maximizes the market surplus, that is, the difference between the buyers' valuations and the sellers' production costs. Contraints (2) match the demand with the supply, while constraints (3) are buyer and seller quantity feasibility constraints.

With the classical assumptions on the buyers and sellers of having quasi-linear utility functions and being price-takers, the concept of Walrasian Equilibrium can be defined.

Definition 1 *The allocation $\tilde{q} = [\{\tilde{q}_j\}_{j \in \mathcal{J}}; \{\tilde{q}_s\}_{s \in \mathcal{S}}]$ and the price vector $p = \{p_l\}_{l \in \mathcal{L}}$ form a Walrasian Equilibrium if: (1) The allocation $\tilde{q}$ is feasible for Model (MC); (2) $V_j(\tilde{q}_j) - p.\tilde{q}_j = \max_{q_j \in \mathcal{D}_j}(V_j(q_j) - p.q_j)$, $\forall j \in \mathcal{J}$; and (3) $p.\tilde{q}_s - C_s(\tilde{q}_s) = \max_{q_s \in \mathcal{D}_s}(p.q_s - C_s(q_s))$, $\forall s \in \mathcal{S}$.*

Condition 1 above means that $\tilde{q}$ is an acceptable allocation for all sellers and buyers, that matches the total supply with the total demand. Conditions (2) and (3) point out the behavior of sellers and buyers as price-taking, utility-maximizing participants.

3 Auctions based on Lagrangian relaxation

The structure of the centralized market-clearing formulation (MC) naturally suggests Lagrangian relaxation as a decomposition approach. Hence, let us dualize (MC) by relaxing the supply/demand matching constraints (2). Let $\lambda = \{\lambda_l\}_{l\in\mathcal{L}}$ be the vector of Lagrangian multipliers associated with (2). The corresponding Lagrangian can be defined as:

$$\begin{aligned} L(q;\lambda) \quad &= \sum_{j\in\mathcal{J}} V_j(q_j) - \sum_{s\in\mathcal{S}} C_s(q_s) \\ &+ \sum_{l\in\mathcal{L}} \lambda_l (\sum_{s\in\mathcal{S}} q_{s,l} - \sum_{j\in\mathcal{J}} q_{j,l}); \\ &q_j \in \mathcal{D}_j, j \in \mathcal{J}; q_s \in \mathcal{D}_s, s \in \mathcal{S}; \lambda \in \mathbb{R}^{|\mathcal{L}|} \end{aligned} \tag{4}$$

Consider the Lagrangian dual function $\Theta(\lambda) = \max_q \{L(q;\lambda) : q_j \in \mathcal{D}_j, j \in \mathcal{J}, q_s \in \mathcal{D}_s, s \in \mathcal{S}\}$. The Lagrangian dual problem (LD) is $\min_\lambda \Theta(\lambda)$. It is noteworthy that the Lagrangian dual function can be formulated as (LR(λ)):

$$\max_q \quad \sum_{j\in\mathcal{J}} \left(V_j(q_j) - \sum_{l\in\mathcal{L}} \lambda_l q_{j,l} \right)$$

$$+ \sum_{s\in\mathcal{S}} \left(\sum_{l\in\mathcal{L}} \lambda_l q_{s,l} - C_s(q_s) \right) \tag{5}$$

$$s.t. \quad q_j \in \mathcal{D}_j, \quad j \in \mathcal{J} \tag{6}$$

$$q_s \in \mathcal{D}_s, \quad s \in \mathcal{S} \tag{7}$$

which is decomposable into $|\mathcal{J}| + |\mathcal{S}|$ independent sub-problems, one for each seller and buyer. Moreover, each one of these sub-problems consists of maximizing the surplus of a seller or buyer within the respective production or consumption feasibility sets.

Conditions under which an optimal solution of the Lagrangian dual problem and the corresponding optimal primal solutions correspond to an efficient allocation and form a Walrasian Equilibrium are stated in the following result.

Theorem 1 *Let $\lambda^\star$ be an optimal solution of the Lagrangian dual problem, and $q^\star$ an optimal primal solution of the corresponding Lagrangian relaxation $\Theta(\lambda^\star)$. Under conditions of: (1) convexity of seller production functions, concavity of buyer valuation functions, and convexity of the feasibility sets, (2) stability of problem (MC), and (3) feasibility of $q^\star$ for (MC), $q^\star$ is a socially-efficient allocation and $q^\star$,$\lambda^\star$ form a Walrasian Equilibrium.*

Proof 1 *The social-efficiency of the allocation $q^\star$ follows immediately from the strong Lagrangian duality theorem (Geoffrion, 1971), which precludes the existence of a duality* *gap between the primal problem and the Lagrangian dual problem under convexity and stability conditions*[1]*. With respect to the fact that* $q^\star$ *and* $\lambda^\star$ *constitute a Walrasian equilibrium, one may simply notice that each one of the* $|\mathcal{J}| + |\mathcal{S}|$ *sub-problems of the Lagrangian relaxation (*$LR(\lambda)$*) corresponds to the maximization of the utility of a seller or a buyer.*

3.1 The subgradient approach

The subgradient algorithm has been traditionally used to solve Lagrangian dual problems. In the basic version of the algorithm, a subgradient of the dual function Θ is computed at each iteration for the current vector of Lagrangian multipliers and the multipliers are updated along the direction of the subgradient. For the problem at hand, the algorithm can be stated as follows:

- **STEP 0**: Set $k = 0$, $\Theta_\star^{(0)} = +\infty$, $k_\star = 0$. Initialize the vector of Lagrangian multipliers, e.g., set $\lambda^{(0)} = 0$. Define a step size series ξ.
- **STEP 1**: Evaluate $\Theta(\lambda^{(k)})$. Let $q^{(k)}$ be an optimal solution of (LR($\lambda^{(k)}$)). Compute a subgradient $g^{(k)}$ of Θ at $\lambda^{(k)}$. For example: $g^{(k)} = \sum_{s\in\mathcal{S}} q_s^{(k)} - \sum_{j\in\mathcal{J}} q_j^{(k)}$ is a subgradient. Update $\Theta_\star^{(k)}$: if $\Theta(\lambda^{(k)}) < \Theta_\star^{(k)}$ then set $\Theta_\star^{(k)} = \Theta(\lambda^{(k)})$, and $k_\star = k$.
- **STEP 2: If $g^{(k)} = 0$ then $\lambda^{(k)}$ is an optimal solution of (LD). Return $\lambda^{(k)}$ and $q^{(k)}$.**
  Otherwise, adjust the Lagrangian multipliers according to $\lambda^{(k+1)} = \lambda^{(k)} - \xi^{(k)} g^{(k)}$.
- **STEP 3: If an appropriate stopping criterion is satisfied, return $\lambda^{(k)}$ and $q^{(k)}$.**
  Otherwise, set $k = k + 1$ and return to STEP 1.

The choice of the step size series $\xi^{(k)}$ is critical to the convergence of the algorithm. In practice, the following two schemes are the most commonly used.

1. The series $\{\xi^{(k)}\}_{k\in\mathbb{N}}$ is such that $\lim_{k\to\infty} \xi^{(k)} = 0^+$ and $\sum_{k=0}^{\infty} \xi^{(k)} = +\infty$.

[1]The *stability* of problem (MC) is a key condition for having no "duality gap" between the primal problem (MC) and its Lagrangian dual. Stability is guaranteed when the *filling property* (Hiriart-Urruty and Lemaréchal, 1991), which requires the subsets $\mathcal{D}_j, j \in \mathcal{J}$ and $\mathcal{D}_s, s \in \mathcal{S}$ to be compact and the functions $V_j, j \in \mathcal{J}$ and $C_s, s \in \mathcal{S}$ to be continuous. This condition is satisfied in particular, when the objective and the constraints of the primal problem are affine or quadratic.

2. Suppose an "estimate" $\bar{\Theta}^{(k)}$ on the optimal value of (MC) is available at each iteration k. Consider the series $\{\xi^{(k)}\}_{k\in\mathbb{N}}$ such that $\xi^{(k)} = \rho^{(k)}\frac{\Theta(\lambda^{(k)})-\bar{\Theta}^{(k)}}{\|g^{(k)}\|^2}$, where $\rho^{(k)}$ is a scaling factor. Typically, the sequence $\{\rho^{(k)}\}_{k\in\mathbb{N}}$ is such that $\rho^{(0)} \in [0,1]$ and $\rho^{(k)}$ is halved each time $\Theta_\star^{(k)}$ has not been updated for n (generally equal to 10 or 20) consecutive iterations.

Concerning the algorithm's stopping criterion, a reasonable condition would be $\|g^{(k)}\| \leq \epsilon$ (the rationale of this criterion is Everett's theorem (Everett, 1963), which states that, as far as the supply/demand constraints 2 are concerned, the corresponding primal solution $q^{(k)}$ would be $\epsilon-$feasible). However, the condition $\|g^{(k)}\| \leq \epsilon$ may never be satisfied due to the fact that the subgradient algorithm only requires one subgradient of the dual function, and not the entire subdifferential, to be evaluated at each iteration. Thus, this test is often combined to other "heuristic" criteria, such as: "stop if $\Theta_\star^{(k)}$ has not been improved in the last N last iterations", or "stop if the gap $\Theta(\lambda^{(k)}) - \bar{\Theta}^{(k)}$ is smaller than a threshold ϵ".

The subgradient algorithm has an interesting interpretation as an iterative auction mechanism, driven by prices. The market-maker arbitrarily sets an initial vector $\lambda^{(0)}$ of single-product prices. At a given round k of the process, each seller $s, s \in \mathcal{S}$, determines a production level $q_s^{(k)}$ that maximizes its surplus given the current prices of the goods and formulates a bid $B_s^{(k)} = \{q_s^{(k)}\}$ that specifies this production level. Similarly, each buyer $j, j \in \mathcal{J}$, formulates a unique bid $B_j^{(k)} = \{q_j^{(k)}\}$, where $q_j^{(k)}$ is a surplus-maximizing consumption level for buyer j at the given prices. The market-maker then revises the prices of the goods along a steepest descent direction given by the excess vector $\sum_{s\in\mathcal{S}} q_s^{(k)} - \sum_{j\in\mathcal{J}} q_j^{(k)}$. The auction continues until there is a marginal difference between the supply and the demand in the market.

Given that the subgradient-based auction may stop before an "implementable" outcome is reached (one that satisfies - at least approximately - the balance of supply and demand), a process for recovering feasible primal solutions is needed. In many practical cases, specialized heuristics can often be designed for that purpose. The general nature of the problem at hand, however, only allows for equally general procedures to recover primal feasibility. Hence, we adapt a very simple approach due to Larsson, Patriksson, and Strömberg (Larrson et al., 1999), which consists in *projecting* upon the feasible domain of (MC) the elements of an ergodic sequence of primal solutions converging to an optimal solution of (MC). More specifically, it can be shown that the sequence $\{\bar{q}^{(k)}\}_{k\in\mathbb{N}}$ such that $\bar{q}^{(k)} = \frac{\sum_{r=1}^{k-1}\xi^{(r)}q^{(r)}}{\sum_{r=1}^{k-1}\xi^{(r)}}$, $k \in \mathbb{N}$ converges to the set of optimal solutions of (MC). Let K be the last iteration of the auction. The (Euclidean) projection of $\bar{q}^{(K)}$ on the feasible domain of (MC) corresponds to allocation vectors q that solve:

$$\min_q \quad \|q - \bar{q}^{(K)}\|^2 \tag{8}$$

$$s.t. \quad \sum_{j\in\mathcal{J}} q_{j,l} - \sum_{s\in\mathcal{S}} q_{s,l} = 0, \quad l \in \mathcal{L} \tag{9}$$

$$q_j \in \mathcal{D}_j, \quad j \in \mathcal{J} \tag{10}$$

$$q_s \in \mathcal{D}_s, \quad s \in \mathcal{S} \tag{11}$$

In order to implement this recovery procedure, the market-maker nevertheless needs complete knowledge of feasibility sets $\mathcal{D}_j, j \in \mathcal{J}$ and $\mathcal{D}_s, s \in \mathcal{S}$. In the absence of this knowledge, bids submitted by sellers and buyers in previous iterations of the auction can be used to "shape" approximations of the actual feasibility sets. Consider the convex hulls: $\hat{\mathcal{D}}_j = \{q_j = \sum_{k=0}^K q_j^{(k)}\alpha_j^k : \sum_{k=0}^K \alpha_j^k = 1; \alpha_j^k \geq 0, k = 0,\ldots,K\}$, $j \in \mathcal{J}$, and $\hat{\mathcal{D}}_s = \{q_s = \sum_{k=0}^K q_s^{(k)}\beta_s^k : \sum_{k=0}^K \beta_s^k = 1; \beta_s^k \geq 0, k = 0,\ldots,K\}$, $s \in \mathcal{S}$. Thanks to the convexity of the feasibility sets, $\hat{\mathcal{D}}_j, j \in \mathcal{J}$ and $\hat{\mathcal{D}}_s, s \in \mathcal{S}$ are inner-approximations of $\mathcal{D}_j, j \in \mathcal{J}$ and $\mathcal{D}_s, s \in \mathcal{S}$, respectively. The projection of $\bar{q}^{(K)}$ on the approximated feasible sets yields the following quadratic problem:

$$\min_q \quad \|q - \bar{q}^{(K)}\|^2 \tag{12}$$

$$s.t. \quad \sum_{j\in\mathcal{J}} q_{j,l} - \sum_{s\in\mathcal{S}} q_{s,l} = 0, \quad l \in \mathcal{L} \tag{13}$$

$$q_j \in \hat{\mathcal{D}}_j, \quad j \in \mathcal{J} \tag{14}$$

$$q_s \in \hat{\mathcal{D}}_s, \quad s \in \mathcal{S} \tag{15}$$

One should be aware of the heuristic nature of this recovery procedure. Problem (12-15) is a restriction of (8-11) and, thus, is not necessarily feasible and may not be successful in providing a feasible outcome when a small number of bids are used to approximate the feasibility sets, or when the feasible domain of the allocation problem is originally tight.

3.2 The bundle approach

Bundle methods, originally developed for nonsmooth optimization (Wolfe, 1975; Lemaréchal, 1989), may equally be suggested for solving the Lagrangian dual problem. These methods rely basically on the concept of *bundle of information*, which is used to build "good" approximation models of the dual function Θ - at least in the vicinity of an optimal solution. Hence, let the bundle $\mathcal{B} = \{(\lambda^{(k)}; \Theta(\lambda^{(k)}); g^{(k)})\}_{k=1,\ldots,K}$ represent the information gathered at a given time, where $g^{(k)} \in \partial\Theta(\lambda^{(k)})$, $\forall k = 1,\ldots,K$ is a subgradient of Θ

at $\lambda^{(k)}$. The first-order approximation of Θ with the information in bundle $\mathcal{B}$ yields the *cutting-plane model* of Θ: $\Theta_{\text{cp}}(\lambda) = \max_{1 \leq k \leq K}\{\Theta(\lambda^{(k)}) + {g^{(k)}}^T(\lambda - \lambda^{(k)})\}$.

The early cutting-plane algorithm (Kelley, 1960) is an iterative procedure that consists in minimizing the approximate model Θ_{cp} and using the optimal solution $\lambda^{(K+1)}$ obtained at iteration K to enrich Θ_{cp} with a new cutting plane. Practical experience with the cutting-plane algorithm has nonetheless revealed its *instability*: the iterate $\lambda^{(K+1)}$ is often very remote from $\lambda^{(K)}$, even if the latter is very close to an optimal dual solution. A significant number of the cutting planes generated by the algorithm are consequently of little help in closing the gap between Θ and Θ_{cp} in the neighborhood of an optimal solution. Bundle methods address the instability issue by defining a *stability center* $\bar{\lambda}$ and requiring that the approximate model produces an iterate $\lambda^{(K+1)}$ "not too far" from $\bar{\lambda}$. This is done by the introduction of a stabilizing term $\frac{1}{2t^K}\|\lambda - \bar{\lambda}\|^2$ in the expression of the cutting-plane model Θ_{cp}, where t^K is a parameter that can be interpreted both as a step size and a *trust-region* parameter. The new approximation model of Θ is thus $\Theta_B(\lambda) = \Theta_{\text{cp}}(\lambda) + \frac{1}{2t^K}\|\lambda - \bar{\lambda}\|^2$. The minimization of Θ_B at iteration K corresponds to the quadratic problem $(Q_{\mathcal{B}})$:

$$
\begin{aligned}
\min_{\nu,\lambda} \quad & \nu + \frac{1}{2t^K}\|\lambda - \bar{\lambda}\|^2 \\
s.t. \quad & \nu \geq \Theta(\lambda^{(k)}) + {g^{(k)}}^T(\lambda - \lambda^{(k)}) \\
& \quad , k = 1, \ldots, K
\end{aligned}
\tag{16}
$$

Let $(\nu^{(K+1)}; \lambda^{(K+1)})$ be an optimal solution of this problem, and let $\Delta^K = \Theta(\bar{\lambda}) - \Theta(\lambda^{(K+1)})$ and $\tilde{\Delta}^K = \Theta_B(\bar{\lambda}) - \Theta_B(\lambda^{(K+1)})$ denote the *actual* and the *predicted* (by model Θ_B) decrease of Θ, respectively. If $\Delta^K \geq m\tilde{\Delta}^K$ (m is a pre-specified parameter such that $0 < m < 1$), i.e., the value of Θ has actually been "sufficiently" decreased with respect to the predicted value, the bundle method performs a *serious-step*: accept $\lambda^{(K+1)}$ as the new stability center. Otherwise, a *null-step*, which consists in leaving the stability center unchanged but adding $(\lambda^{(K+1)}; \Theta(\lambda^{(K+1)}); g^{(K+1)})$ to the bundle for a more refined approximation of Θ, is made.

The setting of the parameter sequence $\{t_k\}_{k\in\mathbb{N}}$ is extremely important in defining the behavior of a bundle algorithm and its numerical efficiency (Hiriart-Urruty and Lemaréchal, 1991). Small values of t tend to drive the bundle algorithm into making relatively few null-steps and also "small" serious-steps resulting only in marginal improvement of the dual functions (when $t \to 0$, the bundle method is nothing else than the subgradient algorithm). On the other hand, with large values of t the bundle algorithm tends to performs few serious-steps when large values of t are considered, moving toward the cutting-plane algorithm as $t \to +\infty$. The design of variable sequences $\{t_k\}_{k\in\mathbb{N}}$ is indeed a complex issue and the literature is clearly lacking in theoretical results on provably "good" sequences. To date, heuristic approaches that consist in increasing t_k after a serious-step and decreasing it after a null-step (Kiwiel, 1990; Schramm and Zowe,

1992) seem to provide the best results.

It is interesting to compare the way bundle methods manage prices with the simpler price update scheme of the subgradient algorithm. Basically, two fundamental observations can be made: (a) the bundle approach relies on a collection of information representing a "history" of the market, that is, a set of prices and the corresponding bidder reactions (the desired production and consumption levels at these prices); (b) a specific price vector (the stability center), in the neighborhood of which the approximate cutting-plane model can be reasonably "trusted", is given a special status. In that regard, the dual viewpoint provides additional insight. Let us consider $\Pi_{\mathcal{B}}$, the problem obtained by dualizing constraints (16) of $Q_{\mathcal{B}}$:

$$\max_{\delta} \quad \sum_{k=1}^{K} \delta_k \{\Theta(\lambda^{(k)}) - {g^{(k)}}^T \lambda^{(k)}\}$$

$$-\frac{1}{2} t^K \|{z_\delta}^{(K)}\|^2 + {z_\delta}^{(K)T} \bar{\lambda} \tag{17}$$

$$s.t. \quad \sum_{k=1}^{K} \delta_k = 1 \tag{18}$$

$$\delta_k \geq 0, \quad k = 1, \ldots, K \tag{19}$$

where ${z_\delta}^{(K)} = \sum_{k=1}^{K} \delta_k g^{(k)}$.

By writing optimality conditions of $Q_{\mathcal{B}}$ and $\Pi_{\mathcal{B}}$, we obtain that, for an optimal solution $(\nu^\star; \lambda^\star)$ of $Q_{\mathcal{B}}$, there exists an optimal solution $\{\delta_k^\star\}_{k=1,\ldots,K}$ of $\Pi_{\mathcal{B}}$ such that:

(i) $\lambda^\star = \bar{\lambda} - t^K \sum_{k=1}^{K} \delta_k^\star g^{(k)}$;

(ii) $\sum_{k=1}^{K} \delta_k^\star = 1$ and $\delta_k^\star \geq 0, k = 1, \ldots, K$;

which indicates that the bundle algorithm actually constructs *aggregated subgradients* $z^{(K)} = \sum_{k=1}^{K} \delta_k^\star g^{(k)}$ as convex combinations of the subgradients available in the bundle, and moves (in the case of a serious-step) in the opposite direction of the aggregated subgradient, to an extent given by step size t^K. By pushing the analysis a little bit further, it is possible to establish (Lemaréchal, 2001) that the convex combination $\sum_{k=1}^{K} \delta_k^\star g^{(k)}$ tends towards 0 as $K \to +\infty$. This result is fundamental for the bundle-based auction, as it implies asymptotic *feasibility* of $\sum_{k=1}^{K} \delta_k^\star q^{(k)}$ for the centralized market model (MC), and by consequence, optimality for (MC). In a nutshell, the bundle-based auction does not need an ad-hoc feasibility recovery procedure: convex combinations $\sum_{k=1}^{K} \delta_k^\star q^{(k)}$ are asymptotically social-efficient.

4 Application to a procurement case study

4.1 The experimental setting

Our proposed auction mechanisms are illustrated on a more detailed model of multilateral multi-commodity markets presented in Bourbeau et al. (2005). This model has the advantage of being closer to actual applications in procurement, especially in the context of regulated marketplaces for the trade of natural resources. We briefly present in the following the notation and the important elements of the model. We refer the reader interested in more details about the model to Bourbeau et al. (2005).

Participants in the market seek to trade a set of products. A product is a basic commodity with a specific physical denomination (e.g., a wood specie). Products are generally not available in a "pure" state and come rather as part of lots that are "mixtures" of several products. Hence, let $\mathcal{K}$ be the set of basic products, $\mathcal{L}$ the set of lots, and b_l^k be the proportion of product k in lot l, $k \in \mathcal{K}$, $l \in \mathcal{L}$.

It is assumed for simplicity (but with no loss of generality) that each seller may only offer a single lot. Thus, a lot $l \in \mathcal{L}$ is attached to seller l and Q^l denotes the maximum quantity produced of that lot. The production cost function $C_l(.)$ of lot l is assumed to have a continuous, piecewise-linear, and strictly increasing marginal cost function $C_l'(.)$. On the buyer side, Bourbeau *et al.*'s model accounts for the differences in quality among the various lots by considering: (i) a multiplicative adjustment coefficient r_j^l, which indicates that one unit of lot l is equivalent for buyer j to r_j^l units of a standard lot; and (ii) an additive coefficient s_j^l, which denotes how much more or less buyer j values, in absolute terms, a unit of lot l with respect to a unit of the standard lot. Furthermore, the model considers a unit transportation cost t_j^l between the seller producing lot l and buyer j. The latter's preference for a bundle $q_j = \{q_{j,l}\}_{l \in \mathcal{L}}$ can accordingly be expressed as $V_j(q_j) = U_j(\sum_{l \in \mathcal{L}} r_j^l q_{j,l}) + \sum_{l \in \mathcal{L}} (s_j^l - t_j^l) q_{j,l}$, where $U_j(.)$ is a utility function such that $U_j'(.)$ is continuous, piecewise-linear, and strictly decreasing. Buyers need also to express requirements regarding the composition of the lots they purchase. More specifically, let M_j^k and m_j^k denote respectively the maximum and minimum proportions of product k that buyer j may tolerate in the acquired lots, and Q^j the maximum total volume - expressed in terms of the standard lot - buyer j requires.

With the notation above, the market-clearing problem corresponds to the following formulation

$$\max \quad \sum_{j\in\mathcal{J}} U_j(\sum_{l\in\mathcal{L}} r_j^l q_{j,l}) + \sum_{l\in\mathcal{L}} (s_j^l - t_j^l) q_{j,l} - \sum_{l\in\mathcal{L}} C_l(q_l)$$

$$s.t. \quad \sum_{j\in\mathcal{J}} q_{j,l} - q_l = 0, \quad l \in \mathcal{L} \tag{20}$$

$$\sum_{l\in\mathcal{L}} r_j^l q_{j,l} \leq Q^j, \quad j \in \mathcal{J} \tag{21}$$

$$m_j^k \sum_{l\in\mathcal{L}} r_j^l q_{j,l} \leq \sum_{l\in\mathcal{L}} b_l^k r_j^l q_{j,l}, \quad j \in \mathcal{J}, k \in \mathcal{K} \tag{22}$$

$$\sum_{l\in\mathcal{L}} b_l^k r_j^l q_{j,l} \leq M_j^k \sum_{l\in\mathcal{L}} r_j^l q_{j,l}, \quad j \in \mathcal{J}, k \in \mathcal{K} \tag{23}$$

$$0 \leq q_l \leq Q^l, q_{j,l} \geq 0, \quad j \in \mathcal{J}, l \in \mathcal{L} \tag{24}$$

where $q_{j,l}$ denotes the quantity of lot l purchased by buyer j and q_l the total quantity of lot l procured by the corresponding seller.

The computational experiments aim to compare auction processes based on various decomposition methods from the perspective of achieved economic efficiency. For that purpose, the "benchmark" used is a mechanism based on the centralized market-clearing formulation (20-24), which assumes that the market-maker has access to complete information about the sellers and buyers valuation and cost functions, as well as their technological constraints.

We have performed tests on several problem series made of instances obtained from a custom problem generator we have developed. Given values for the numbers of buyers, sellers (lots), and basic products, volumes $\{Q_j\}_{j\in\mathcal{J}}$ and $\{Q_l\}_{l\in\mathcal{L}}$, proportions $\{b_l^k\}_{l\in\mathcal{L},k\in\mathcal{K}}$, and tolerances $M_j^k, m_j^k, j \in \mathcal{J}, k \in \mathcal{K}$ are randomly generated according to continuous uniform distributions over pre-specified intervals. For the sake of simplicity, we considered *purely quadratic* buyer utility functions $U_j(.), j \in \mathcal{J}$ and seller cost functions $C_l(.), l \in \mathcal{L}$. This implies no loss of generality, since a simple transformation suggested in Bourbeau et al. (2005) allows to deal with a general piecewise-quadratic formulation as a purely quadratic one. Furthermore, our instances involved no transportation costs t_l^j or additive adjustment coefficients s_l^j, $j \in \mathcal{J}, l \in \mathcal{L}$. Table 1 displays the characteristics of the problem series considered in the study. The series are subdivided into two categories according to (1) $|\mathcal{K}|$, the number of basic products; and (2) Δ_m, the minimum difference between tolerances M_j^k, m_j^k ($\Delta_m = \underline{M} - \overline{m}$, where $\underline{M}$ designates the minimum value M_j^k can take, and $\overline{m}$ the maximum value of m_j^k). These two parameters are important since they directly impact the number and forcefulness of constraints (23 and 22) in the market-clearing formulation. Each series of problems consisted of 10 randomly generated instances.

We have set up four auction processes. The first three are based on different variants

Table 1: Characteristics of Problem Instances

Problem series	Problem description			
	# buyers	# lots	# products	$\Delta_m = \underline{M} - \overline{m}$ (%)
$S-01$	50	100	3	30
$S-02$	50	250	3	30
$S-03$	100	50	3	30
$S-04$	100	200	3	30
$S-05$	100	500	3	30
$S-06$	50	100	10	10
$S-07$	50	250	10	10
$S-08$	100	50	10	10
$S-09$	100	200	10	10
$S-10$	100	500	10	10

of the subgradient algorithm, and the fourth on Frangioni's implementation of the bundle method (Carraresi et al. (1996)). More specifically, the subgradient variants used are:

1. The basic subgradient method: $\lambda^{(k+1)} = \lambda^{(k)} - \xi^{(k)} g^{(k)}$, with step size formula $\xi^{(k)} = \rho^{(k)} \frac{\Theta(\lambda^{(k)}) - \bar{\Theta}^{(k)}}{\|g^{(k)}\|^2}$. We have used the simple estimate $\bar{\Theta}^{(k)} = 0.5.\Theta_\star^{(k)}$, where $\Theta_\star^{(k)}$ denotes the best value of the Lagrangian dual function $\Theta(.)$ found so far. Parameter $\rho^{(0)}$ has been calibrated in the set $\{0.1, 0.3, 0.5, 0.7, 1.0\}$ for each problem series.

2. The subgradient method with the Camerini-Fratta-Maffioli rule (Cam). This variant relies on an elementary aggregation of the subgradients to compute a direction along which to move at each iteration. Thus, $\lambda^{(k+1)} = \lambda^{(k)} - \xi^{(k)} d^{(k)}$, where $d^{(k)} = g^{(k)} + \sigma^{(k)} d^{(k-1)}$ and $\sigma^{(k)}$ is such that

$$\sigma^{(k)} = \begin{cases} -\mu \frac{g^{(k)} d^{(k-1)}}{\|d^{(k)}\|^2} & \text{if} g^{(k)} d^{(k-1)} < 0, \\ 0 & \text{otherwise}, \end{cases}$$

   where parameter μ is set to 1.5. The step size formula used is similar to that of the basic subgradient, that is $\xi^{(k)} = \rho^{(k)} \frac{\Theta(\lambda^{(k)}) - \bar{\Theta}^{(k)}}{g^{(k)} d^{(k)}}$.

3. The subgradient method with the modified Camerini-Fratta-Maffioli rule: $\mu = -\frac{\|g^{(k)}\| \|d^{(k-1)}\|}{g^{(k)} d^{(k-1)}}$.

The bundle algorithm requires the calibration of many parameters, the most important of which are: (1) The maximum size of the bundle, with the following values: 10, 20, 50, 100, and 200; (2) the parameter m controlling the serious step condition ($\Delta^K \geq m\tilde{\Delta}^K$, see section 3.2), for which we tested the two values that are known to work best in practice, 0.1 and 0; and (3) the relative accuracy $\epsilon = 10^{-6}$ of the bundle algorithm's stopping criterion $t^\star\|z^{(K)}\|^2 - \omega^{(K)} \leq \Theta(\bar\lambda)$, where $z^{(K)} = \sum_{k=1}^{K}\delta_k^\star g^{(k)}$ is the aggregated subgradient, $\omega^{(K)} = \sum_{k=1}^{K}\delta_k^\star[\Theta(\bar\lambda) - (\Theta(\lambda^{(k)}) + g^{(k)}(\bar\lambda - \lambda^{(k)}))]$ the aggregated linearization error, and $t^\star$ a fixed step size, generally one order of magnitude larger than $t^{(0)}$.

In order to compare the different auction mechanisms, we fixed the maximum number of rounds to 1000 for the subgradient-based auctions and to 2000 for the bundle-based auction. The following metrics are used:

(a) For the auction processes based on the subgradient and its variants (CFM and modified CFM), we measured:

1. the gap $GAP_{\text{lr}} = (Z_{\text{lr}} - Z_{\text{cent}})/Z_{\text{cent}}$, where Z_{lr} is the best upper bound obtained by the corresponding subgradient methods;
2. the gap $GAP^P_{\text{lr}} = (Z_{\text{cent}} - \tilde Z)/Z_{\text{cent}}$, where $\tilde Z$ is the economic surplus achieved by the "closest" feasible allocation to the primal solution $q^{(k_\star)}$, obtained by projecting the latter on the feasible domain of model (20-24);
3. the gap GAP^E_{lr} corresponding to the allocation q^E obtained through the projection of the last term $q^{(K)}$ of the ergodic sequence $\{\bar q^{(k)}\}_{k\in\mathbb{N}}$ defined in Section 3.1 on the domain of feasible allocations, that is $GAP^E_{\text{lr}} = (Z_{\text{cent}} - \bar Z)/Z_{\text{cent}}$, where $\bar Z$ is the economic surplus achieved by q^E;
4. the Euclidean norms $\|g^{(k_\star)}\|$ and $\|g^{(K)}\|$ of the trivial subgradients corresponding to allocations $q^{(k_\star)}$ and $q^{(K)}$, respectively.

(b) For the bundle process, the following quantities have been measured:

1. the gap $GAP_{\text{b}} = (Z_{\text{b}} - Z_{\text{cent}})/Z_{\text{cent}}$, where Z_{b} is the best upper bound obtained by the bundle algorithm;
2. the Euclidean norm $\|z_\delta{}^{(K)}\|$ of the aggregated subgradients.

Finally, the experiments were carried out on a 64-processor, 64 Gigabytes of RAM Sun Enterprise 10000 operated under SunOS 5.8, with versions 8.0 and 1.2 of the CPLEX solver and the Concert library, respectively.

4.2 Numerical results

Table 2 displays the results obtained by the basic subgradient, the Camerini-Fratta-Maffioli (CFM), and the modified CFM methods. The column (GAP_{lr}) indicates the average gaps corresponding to the best upper bound achieved at the last bidding round (1000th in this case). We only retained the best gaps with respect to the five possible values of the initial scaling factor $\rho^{(0)}$) (the corresponding value of $\rho^{(0)}$ is listed in the table). The second and the third entries of the table are the average gaps of the allocations obtained by projecting $q^{(k_\star)}$ and $q^{(K)}$, respectively, on the feasible allocation domain, while the average norms of the subgradients corresponding to $q^{(k_\star)}$ and $q^{(K)}$ are indicated in the last two columns.

Several observations can be made regarding these results. First, both the basic subgradient and the CFM method displayed relatively small average gaps, consistently converging to within 4% of the optimal solution of the centralized market-clearing formulation. By comparison, the convergence of the modified CFM was not as uniform, and the gaps obtained on some series ($S-03$, $S-07$, and $S-08$ in particular) were much larger. In that regard, the poor performance of the modified CFM method seems likely to be a consequence of its relatively greater sensitivity to the choice of the initial scaling factor $\rho^{(0)}$, rather than an inherent lack of effectiveness.

Yet, feasibility of the primal solutions obtained by the three methods is a major source of concern. Fairly large subgradients (taking into consideration the magnitude of the randomly generated quantities $Q^j, j \in \mathcal{J}$ and $Q^l, l \in \mathcal{L}$) were obtained. The results in the table also indicate that the heuristic approaches suggested do not satisfactorily resolve this issue, as (1) the projection of the primal solution on the feasible domain produces very poor allocations, and (2) despite significantly reducing infeasibility, the ergodic sequence of Larrson et al. (1999) seems to suffer from its notoriously slow convergence rate.

In order to gain more insight into the behavior of the three subgradient-based auctions, we have taken the one instance of series $S-01$ and we mapped out in Figure 1 the best upper bounds as each auction progresses. The figure shows quite clearly the relative superiority of the two CFM methods in terms of speed of convergence: 189 and 183 rounds of bidding were enough for the CFM and the modified CFM auctions, respectively, to attain a less-than 1% gap, while the basic subgradient method needed 463 rounds to achieve comparable gap levels. Another interesting question, raised in our analysis of the results in Table 2, is the sensitivity of the subgradient variants to the initial scaling factor $\rho^{(0)}$. We have thus plotted the upper bounds obtained by the three methods with different values of $\rho^{(0)}$ (see Figures 2,3, and 4). The results corroborate our previous observation, that the modified CFM method is much more sensitive to the value of parameter $\rho^{(0)}$: therefore it might perform very poorly and fail to converge within a reasonable number of rounds if a wrong value of the initial scaling factor is chosen (e.g., with $\rho^{(0)} = 1.0$).

Finally, we present in Table 3 the results of the bundle-based auction process. The table shows for each problem series the largest, smallest, and average number of rounds (up to a limit of 2000 rounds) needed for convergence, as well as the average of the corresponding gap GAP_{b} and of the norm of the aggregated subgradient $z^{(K)}$. In order to compare the bundle results with those of the three subgradient variants, we also listed these quantities up to the 1000^{th} round, which was the limit for the subgradient auctions. The results show the superiority of the bundle process regarding dual con-

Table 2: Behavior of the Subgradient, the CFM, and the Modified CFM Auctions.

Series	Basic subgradient					
	GAP_{lr} (%)	GAP^P_{lr} (%)	GAP^E_{lr} (%)	$\lVert g^{(k_\star)}\rVert$	$\lVert g^{(K)}\rVert$	$\rho^{(0)}_\star$
$S-01$	0.11	2.84	17.68	954.51	260.38	0.5
$S-02$	0.69	5.15	57.41	1480.18	243.71	0.5
$S-03$	1.29	12.07	76.80	1663.27	185.89	0.5
$S-04$	0.05	3.17	27.09	2118.53	442.49	0.7
$S-05$	1.76	10.47	68.44	3247.60	378.24	0.5
$S-06$	0.23	3.78	24.85	1337.48	328.63	0.7
$S-07$	2.99	15.37	70.25	2131.56	276.35	0.5
$S-08$	1.70	14.69	78.53	2342.54	190.53	0.7
$S-09$	1.21	6.28	43.84	3071.38	578.84	0.7
$S-10$	4.05	20.22	73.10	4323.10	401.54	0.7
	CFM					
	GAP_{lr} (%)	GAP^P_{lr} (%)	GAP^E_{lr} (%)	$\lVert g^{(k_\star)}\rVert$	$\lVert g^{(K)}\rVert$	$\rho^{(0)}_\star$
$S-01$	0.10	2.58	24.81	925.06	311.12	0.3
$S-02$	0.71	5.46	58.97	1509.12	247.88	0.5
$S-03$	1.22	11.32	76.56	1691.69	185.73	0.5
$S-04$	0.12	3.14	27.07	1999.62	441.17	0.7
$S-05$	1.57	9.29	67.85	3115.90	377.95	0.5
$S-06$	0.39	3.93	26.61	1302.91	342.53	0.7
$S-07$	2.49	13.40	70.05	2055.00	275.99	0.5
$S-08$	1.88	16.43	79.72	2347.58	192.34	0.7
$S-09$	1.54	7.08	44.06	3055.72	580.74	0.7
$S-10$	4.03	20.32	74.18	4314.83	404.95	0.7
	Modified CFM					
	GAP_{lr} (%)	GAP^P_{lr} (%)	GAP^E_{lr} (%)	$\lVert g^{(k_\star)}\rVert$	$\lVert g^{(K)}\rVert$	$\rho^{(0)}_\star$
$S-01$	0.08	2.40	39.92	910.22	412.35	0.1
$S-02$	0.78	5.35	62.19	1552.66	262.85	0.3
$S-03$	5.61	35.75	85.58	1788.30	199.65	0.3
$S-04$	0.16	3.22	39.15	2036.43	545.61	0.3
$S-05$	2.11	10.33	74.04	3234.97	409.98	0.3
$S-06$	0.26	3.55	39.14	1287.81	430.49	0.3
$S-07$	11.27	41.38	84.87	2318.16	311.49	0.3
$S-08$	23.30	76.02	91.69	2566.29	211.83	0.5
$S-09$	3.18	8.84	62.98	3209.39	720.68	0.3
$S-10$	28.32	67.58	92.59	4896.05	455.18	0.3

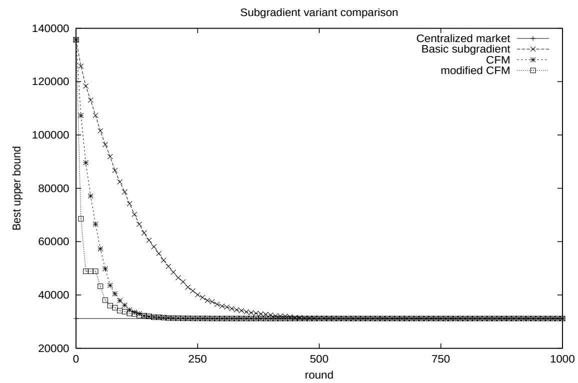

Figure 1: Evolution of the Best Upper Bound for the Three Subgradient Auctions.

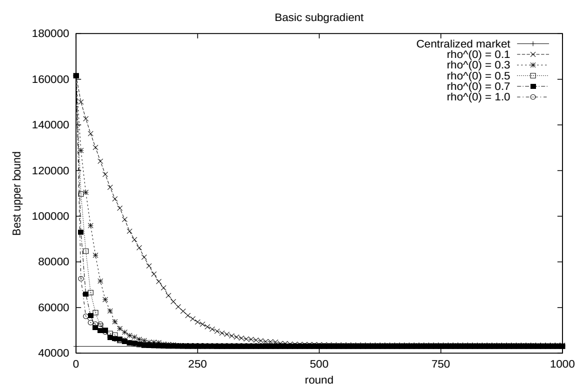

Figure 2: Evolution of the Best Upper Bound for the Basic Subgradient Auction.

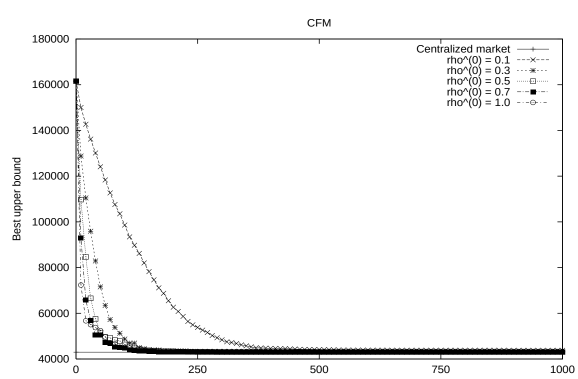

Figure 3: Evolution of the Best Upper Bound for the CFM Auction.

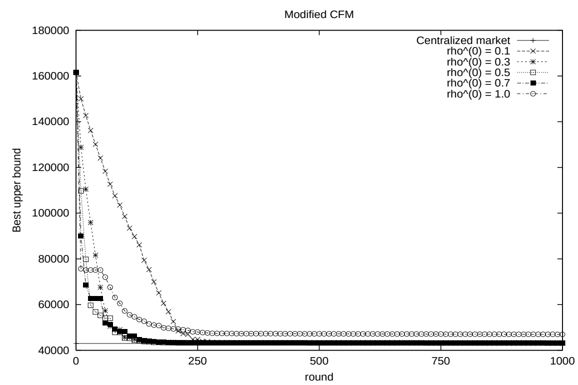

Figure 4: Evolution of the Best Upper Bound for the modified CFM Auction.

Table 3: Behavior of the Bundle-based Auction.

Series	**Total # rounds**			GAP_{b}	$\|z^{(K)}\|$	$GAP_{\text{b}}^{(1000^+)}$	$\|z^{(K)}\|^{(1000^+)}$
	min.	max.	avg.				
$S-01$	379	591	458.7	2.02E-5	0.05	2.02E-5	0.05
$S-02$	662	851	737.4	5.25E-5	0.01	5.25E-5	0.01
$S-03$	266	346	297.9	6.08E-5	0.01	6.08E-5	0.01
$S-04$	654	1438	998.6	6.59E-6	0.09	1.74E-5	0.24
$S-05$	1413	>2000	1643.4	3.30E-5	0.07	0.556	3.71
$S-06$	388	548	461.3	2.07E-5	0.05	2.07E-5	0.05
$S-07$	671	960	825.1	1.85E-5	0.05	1.85E-5	0.05
$S-08$	280	352	300.7	5.5E-5	0.01	5.5E-5	0.01
$S-09$	844	1103	990.2	3.9E-5	0.05	8.8E-5	0.21
$S-10$	1323	>2000	1720.5	5.6E-6	0.11	0.565	17.20

vergence to the optimal objective, as most problem series displayed average gaps in the order of 10^{-5}. By comparison, the two series with the largest instances, $S-05$ and $S-10$, have large gaps up to the 1000^{th} round; but remarkably, extending this limit to 2000 rounds has brought down the gaps dramatically, and the stopping criterion of the bundle algorithm was actually met in 98% of the instances before the 2000^{th} round. As for the aggregated subgradients, they are several orders of magnitude smaller than the subgradients displayed in Table 2, which tends to confirm the asymptotic convergence of the bundle-based auctions to primal feasible allocations.

5 Conclusions

This paper has presented a new perspective of mathematical decomposition methods as iterative auctions in combinatorial exchanges of interdependent goods. We have focused on the Lagrangian relaxation, and we have shown that auctions inspired from subgradient and bundle algorithms could be interpreted as iterative mechanisms in which the participants progressively reveal their preferences to the market-maker. Under certain conditions, these auctions yield outcomes that reconcile the overall welfare-maximization market objective with the individual views of participants seeking to maximize their surplus. Numerical results obtained on a wood chip market case study show that the different variants of the subgradient method often converge in the dual space to the optimal market surplus but generally fail to produce feasible allocations. The bundle-based auction, with its more sophisticated price update rules, resolves this primal feasibility issue.

Several interesting research issues could be addressed in the future. First, the convexity assumption we made about feasibility sets of participants is a key one. In particular, when indivisible goods are considered, duality gaps prevent the interpretation of dual multipliers as prices. Two avenues that seem to be attractive are: (a) the exploration of extended formulations of the market-clearing allocation problem; and (b) pricing schemes based on approximated linear prices, which sacrifices either dual feasibility or complementary slackness (e.g., DeMartini et al. (1999)). Second, we have only considered a minimal set of constraints on the market side (demand and supply balance). Real world markets would typically add other constraints derived from specific business rules, such as buyers requiring to be matched with a few "qualified" sellers, and the decomposition approaches need to be adequately adapted to deal with the additional constraints. Third, the basic iterative auctions designed in this paper do not provide mechanisms to differentiate between identical bids, or to stop before the subgradient and the bundle stopping criteria are met. These refinements will be the subject of further research. Finally, incentive compatibility of the auction mechanisms associated with the decomposition approaches is an important and challenging issue we plan to explore.

Acknowledgments

We would like to thank Benoît Bourbeau and Geneviève Hernu who helped in coding the numerical methods and in realizing the experimental study, as well as Prof. Frangioni who kindly provided his implementation of the bundle algorithm. This project has received financial support from the Bell Canada University Laboratories (LUB), the Natural Sciences and Engineering Research Council of Canada through their Discovery and Collaborative Industry-University Research grant programs, and the FRQ_NT research development fund of the Province of Québec.

References

B. Bourbeau, T. G. Crainic, M. Gendreau, and J. Robert. Design for Optimized Multilateral Multi-commodity Markets. *European Journal of Operations Research*, 163(2): 503–529, 2005.

P. Carraresi, A. Frangioni, and M. Nonato. Applying Bundle Methods to the Optimization of Polyhedral Functions: An Applications-Oriented Development. Technical Report TR-96-17, Dipartimento di Informatica, Univesita di Pisa, Corso Italia 40, 56125 Pisa, Italy, 1996.

S. de Vries and S. Vohra. Combinatorial Auctions: A Survey. *INFORMS Journal on Computing*, 15(3):284–309, 2003.

C. DeMartini, A. M. Kwasnica, J. O. Ledyard, and D. P. Porter. A New and Improved Design for Multi-object Iterative Auctions. Social Science Working Paper SSWP 1054, Caltech, 1999.

H. Everett. Generalized Lagrange Multiplier Method for Solving Problems of Optimum Allocation of Resources. *Operations Research*, 11:399–417, 1963.

A. M. Geoffrion. Duality in Nonlinear Programming: A Simplified Applications-Oriented Development. *SIAM Review*, 13(1):1–37, 1971.

J. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms II: Advanced Theory and Bundle Methods*. A Series of Comprehensive Studies in Mathematics. Springer, New York, 1991.

J. E. Kelley. The cutting-plane method for solving convex programs. *J. SIAM*, 8:703–712, 1960.

K. C. Kiwiel. Proximity Control in Bundle Methods for Convex Nondifferentiable Optimization. *Mathematical Programming*, 46:105–122, 1990.

T. Larrson, M. Patriksson, and A.-B. Strömberg. Ergodic Primal Convergence in Dual Subgradient Schemes for Convex Programming. *Mathematical Programming*, 86:283–312, 1999.

C. Lemaréchal. Nondifferentiable Optimization. In Nemhauser, G.L., Rinnoy Kan, A.H.G, and Todd, M.J., editors, *Handbooks in Operations Research and Management Science, Vol.1*. North-Holland, Amsterdam, 1989.

C. Lemaréchal. Lagrangian Relaxation. In Jünger, M. and Naddef, D., editors, *Computational Combinatorial Optimization*. Springer-Verlag, Berlin, 2001.

A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

H. Schramm and J. Zowe. A Version of the Bundle Idea for Minimizing a Nonsmooth Function: Conceptual Idea, Convergence Analysis, Numerical Results. *SIAM Journal on Optimization*, 2(1):121–152, 1992.

P. Wolfe. A Method of Conjugate Subgradients for Minimizing Nondifferentiable Convex Functions. *Mathematical Programming Study*, 3:145–173, 1975.