

Université de Montréal

**Modèles et algorithmes
pour les enchères combinatoires**

par

Jawad Abrache

Département d'informatique et de recherche opérationnelle

Faculté des arts et des sciences

Thèse présentée à la Faculté des études supérieures en vue de
l'obtention du grade de Philosophiæ Doctor (Ph.D.) en informatique

Octobre, 2003

©Jawad Abrache, 2003



QA

76

U54

2004

V.035

AVIS

L'auteur a autorisé l'Université de Montréal à reproduire et diffuser, en totalité ou en partie, par quelque moyen que ce soit et sur quelque support que ce soit, et exclusivement à des fins non lucratives d'enseignement et de recherche, des copies de ce mémoire ou de cette thèse.

L'auteur et les coauteurs le cas échéant conservent la propriété du droit d'auteur et des droits moraux qui protègent ce document. Ni la thèse ou le mémoire, ni des extraits substantiels de ce document, ne doivent être imprimés ou autrement reproduits sans l'autorisation de l'auteur.

Afin de se conformer à la Loi canadienne sur la protection des renseignements personnels, quelques formulaires secondaires, coordonnées ou signatures intégrées au texte ont pu être enlevés de ce document. Bien que cela ait pu affecter la pagination, il n'y a aucun contenu manquant.

NOTICE

The author of this thesis or dissertation has granted a nonexclusive license allowing Université de Montréal to reproduce and publish the document, in part or in whole, and in any format, solely for noncommercial educational and research purposes.

The author and co-authors if applicable retain copyright ownership and moral rights in this document. Neither the whole thesis or dissertation, nor substantial extracts from it, may be printed or otherwise reproduced without the author's permission.

In compliance with the Canadian Privacy Act some supporting forms, contact information or signatures may have been removed from the document. While this may affect the document page count, it does not represent any loss of content from the document.

Université de Montréal
Faculté des études supérieures

Cette thèse intitulée :

“Modèles et algorithmes
pour les enchères combinatoires”

présentée par :
Jawad Abrache

a été évaluée par un jury composé des personnes suivantes :

Patrice Marcotte président du jury

Teodor Gabriel Crainic directeur de recherche

Michel Gendreau codirecteur de recherche

Bernard Gendron membre du jury

Patrick Jaillet examinateur externe

Jacques Robert représentant du doyen

Thèse acceptée le :

01\03\04

Sommaire

Dans un grand nombre de contextes de marché, plusieurs objets hétérogènes doivent être simultanément transigés. Pour de tels marchés, les enchères combinatoires, où les participants peuvent placer des mises portant sur des “paquets” d’objets interdépendants, constituent une classe importante de mécanismes. L’attrait principal de ce type d’enchères est de permettre aux participants d’exprimer leurs préférences exactes pour les différents paquets d’objets désirés, ce qui favorise très souvent l’efficacité économique de l’enchère. Toutefois, et en raison du nombre exponentiel de paquets d’objets possibles, les enchères combinatoires sont des mécanismes de marché “complexes”, tant pour l’encanteur que pour les participants.

Cette thèse porte sur quelques aspects fondamentaux de la conception de mécanismes d’enchère combinatoire. Nous présentons dans un premier temps une revue critique de la littérature consacrée aux enchères combinatoires. Nous identifions en particulier quatre facettes de la problématique générale, que nous discutons plus en détail. Plus précisément, ces dernières sont reliées à : (i) la classification des marchés combinatoires et les différentes formulations correspondantes du problème d’allocation, (ii) le besoin de langages expressifs de mise, (iii) la conception d’enchères combinatoires itératives où les participants révèlent progressivement leurs préférences, et (iv) le besoin d’outils d’aide à la décision (“aviseurs”) destinés à assister les participants dans l’élaboration de stratégies de mise profitables.

Le second volet de la thèse présente une contribution importante de la thèse à la conception de langages de mise pour les enchères combinatoires. À cet effet, nous avons pu noter que les langages jusqu’ici suggérés dans la littérature ne s’adressaient qu’aux enchères d’objets indivisibles. Nous développons ainsi un nouveau cadre formel unificateur pour la définition de langages de mises pour des enchères de biens divisibles et indivisibles. Sur le plan méthodologique, le nouveau cadre repose sur une définition à deux niveaux d’une mise combinée : à un niveau interne, des mises “atomiques”, désignant des ordres simples d’achat ou de vente, sont combinées de façon à exprimer des contraintes sur les proportions d’exécution de ces ordres ; la définition de mises combinées est complétée à un niveau externe par application récursive d’opérateurs

“logiques” de mise. L’analyse formelle de langages de mises dérivés au sein de ce cadre démontre la remarquable expressivité de ces langages.

Nous consacrons le troisième volet de la thèse à une application des enchères combinatoires aux marchés financiers. À cet égard, la possibilité de soumettre des ordres composites formés de mises simples d’achat et de vente de valeurs exécutées à proportions égales, est particulièrement utile pour des gestionnaires d’actifs désirant rééquilibrer leurs portefeuilles dans un contexte de “fin de séance”. Nous suggérons une nouvelle formulation des problèmes d’allocation et de détermination des prix dans une enchère d’actifs financiers avec ordres composites. Cette nouvelle formulation, comparativement aux modèles de la littérature, permet notamment aux participants de spécifier des proportions d’exécution minimales de leurs ordres, ainsi que d’exiger qu’une sélection soit effectuée entre ordres jugés “équivalents”. Nous développons également deux procédures permettant la discrimination des allocations et des prix optimaux sur la base d’un critère “éthique” simple. L’étude expérimentale réalisée a permis d’évaluer l’efficacité économique et la complexité numérique des modèles proposés.

La littérature sur la conception de mécanismes itératifs d’enchère combinatoire est dominée par des processus basés sur le tâtonnement walrasien classique et l’approche primaire/duale. Dans le quatrième volet de la thèse, nous explorons une nouvelle avenue de recherche prometteuse qui fait appel aux méthodes de décomposition en programmation mathématique. Ainsi, nous montrons qu’aussi bien des approches “duales” basées sur la relaxation lagrangienne (notamment le sous-gradient et les méthodes de “faisceaux”) que la décomposition de Dantzig-Wolfe correspondent bien à des processus itératifs d’enchère. Nous établissons également certaines différences importantes entre ces deux types d’approches en ce qui a trait à l’utilisation de l’information contenue dans les mises des participants et aux hypothèses faites sur le comportement stratégique de ces derniers.

Mots-clés : Commerce électronique, mécanisme de marché, enchères combinatoires, langage de mise, marchés financiers, enchères itératives, méthodes de décomposition.

Abstract

In many market contexts, several heterogeneous goods need to be simultaneously traded. In such contexts, combinatorial auctions, in which participants submit combined bids on bundles of items, constitute an important class of market mechanisms. By allowing participants to express directly their preferences for combinations of interrelated items, combinatorial auctions alleviate economic inefficiency encountered in simultaneous ascending auctions due to the participants' aversion to bid aggressively on individual items in a desired bundle. On the other hand, combinatorial auctions are complex market mechanisms that often require the auctioneer and the participants to solve difficult decision problems.

This thesis deals with some fundamental aspects of combinatorial auction design. Despite being relatively new, this area of research has received increased attention in recent years. We therefore present a comprehensive survey of the literature on combinatorial auctions. We put the emphasis on four important design issues : (i) the classification of combinatorial markets and the corresponding formulations of the allocation problem ; (ii) the need for expressive languages to formulate combined bids ; (iii) the design of iterative auction mechanisms characterized by progressive disclosure of the participants' preferences ; and (iv) the need for optimization-based decision support tools ("advisors") to assist participants work out profitable bidding strategies.

Our second paper presents an important contribution of the thesis to the design of combinatorial bidding languages. One may notice that the bidding languages previously proposed in the literature have all been formalized for combinatorial auctions of indivisible items. We thus suggest a novel unified framework for bidding in combinatorial auctions of divisible or indivisible items. The framework relies on a two-level representation of a combined bid : at the inner level, "atomic" bids that designate single-item sell or buy orders are combined so that requirements on execution proportions of these orders are expressed ; the "partial" bids defined this way are recursively combined, at the outer level, through logical bidding operators. Formal analysis of bidding languages derived from the framework proves the remarkable expressiveness

of these languages.

The ability to submit consolidated orders to sell and buy various financial assets in pre-specified proportions is particularly appealing for traders involved in “end-of-day” portfolio balancing operations. We propose in our third paper a new market-clearing formulation that extends the previous models of the literature through a more detailed portfolio representation and the formulation of new bidding requirements. These concern the ability to set minimal execution proportions and to define exclusive OR relations between “equivalent” orders. We equally suggest tie-breaking procedures that allow to discriminate between optimal allocations and prices on the basis of submission times. Experimental results seek to analyze bundle trading impacts and to evaluate the consequences of the bidding requirements introduced in our model from the perspective of allocation efficiency and computational complexity.

Multi-round auction processes based on classical Walrasian tâtonnement and the primal/dual approach prevail in the literature on iterative combinatorial auctions. A promising research avenue we explore in our fourth paper is that of mathematical programming decomposition methods. In that regard, we show that dual approaches based on Lagrangian relaxation, as well as Dantzig-Wolfe decomposition, can be interpreted as iterative market mechanisms in which the participants’ own interests are reconciled with the global objective of the market. We furthermore establish some important distinctions between the suggested auction processes regarding their use of the information disclosed in the participants’ bids and the various assumptions they make on the strategic behavior of the participants.

Keywords : E-commerce, market mechanism, combinatorial auctions, bidding languages, financial markets, iterative auctions, decomposition approaches.

Table des matières

Sommaire	i
Abstract	iii
Table des matières	v
Table des figures	x
Liste des tableaux	x
Dédicace	xii
Remerciements	xiii
1 Introduction	1
2 Revue de littérature	11
2.1 Premiers travaux	12
2.2 Enchères multi-unités	14
2.3 Enchères de plusieurs objets interdépendants : modèles de base et applications choisies	15
2.3.1 Modèles d'équilibre de prix	17
2.3.2 Modèles d'enchères combinatoires	21
2.3.3 Applications choisies	33
2.4 Mécanismes progressifs d'enchère combinatoire	37
2.4.1 Mécanismes basés sur des approches primales/duales	38
2.4.2 Mécanismes expérimentaux	44
2.5 Expression de la préférence	46
2.5.1 Langages de mises	47

2.5.2	Révélation incrémentale de la préférence	50
2.6	Mécanismes d'enchère combinatoire incitatifs	51
2.6.1	Le mécanisme de Vickrey-Clarke-Groves	52
2.6.2	Approximations du mécanisme VCG	55
2.6.3	Implantation indirecte du mécanisme VCG	57
2.7	Conclusion	60
3	Design Issues for Combinatorial Auctions	61
3.1	Introduction	63
3.2	Basic formulations	67
3.2.1	The one-to-many indivisible case	69
3.2.2	Many-to-one combinatorial auctions	72
3.2.3	A network formulation	73
3.2.4	Combinatorial exchanges	76
3.2.5	Conclusion	78
3.3	Expression of combined bids	78
3.3.1	Motivation and state of the art	79
3.3.2	A new bidding framework	81
3.3.3	Impact on the allocation problem	86
3.4	Iterative combinatorial auctions	88
3.4.1	Design of auction rules	91
3.4.2	Pricing	92
3.4.3	Incentive-compatibility issues	95
3.5	Participant decision problems	97
3.6	Conclusion	99
4	A Framework for Combinatorial Auctions	101
4.1	Introduction	104
4.2	Purpose and methodology	109
4.3	The inner level	113
4.3.1	Basic notation and definitions	113
4.3.2	Inner-level bidding operators	115
4.4	The outer level	119
4.5	Price consistency	121

4.6	Analysis of the framework	125
4.6.1	The indivisible case	125
4.6.2	The divisible case	126
4.7	Impact on the allocation problem	129
4.7.1	Lower bound constraints	131
4.7.2	Partial bid execution	132
4.7.3	The ORDERING operator	132
4.7.4	The EQUAL operator	132
4.7.5	The SIMPLEX operator	132
4.7.6	Quantity operators	133
4.7.7	The BUDGET operator	133
4.7.8	The SELECT-INNER operator	134
4.7.9	Hybrid operators	134
4.7.10	Trivial combined bids	134
4.7.11	The SELECT-OUTER operator	135
4.8	An application to portfolio bundle trading	135
4.9	Conclusion and perspectives for future work	138
5	Bundle Trading in Financial Markets	140
5.1	Introduction	142
5.2	Portfolio bundle trading market mechanisms	146
5.2.1	The allocation problem	147
5.2.2	The pricing problem	151
5.3	Discrimination procedures	153
5.4	Experimental analysis	157
5.4.1	Basic bundle-based problems	162
5.4.2	Lower bound problems	164
5.4.3	XOR problems	166
5.5	Concluding remarks	171
6	Decomposition and Combinatorial Auctions	174
6.1	Introduction	177
6.2	Prior work	180

6.3	Centralized market-clearing for a combinatorial exchange economy	183
6.4	Market-clearing based on Lagrangian relaxation	186
6.4.1	The subgradient approach	187
6.4.2	The bundle approach	192
6.5	An auction scheme based on column-generation	194
6.6	Computational study	198
6.6.1	The experimental setting	198
6.6.2	Numerical results	202
6.7	Conclusions	207
7	Conclusion	209
7.1	Principales contributions	210
7.2	Avenues de recherche	212
	Bibliographie	215

Table des figures

1.1	Enchères : mécanismes de marché indirects	3
2.1	Tâtonnement à base d'ajustement des prix.	20
2.2	Une représentation arborescente des allocations réalisables des objets 1, . . . , 5 aux mises b_1, \dots, b_{10}	27
2.3	Illustration du mécanisme AUSM.	45
3.1	Combinatorial bids on capacity	75
3.2	Direct revelation mechanisms vs multi-round auctions	90
4.1	Trade execution constraints in the divisible case.	110
4.2	The two-level bidding framework.	112
4.3	A graph representation of the bidding structure.	113
4.4	Condition subsets corresponding to instances of the EQUAL, SELECT- INNER, and SELECT-INNER + EQUAL operators.	119
4.5	Bidding Example 1.	122
4.6	Bidding Example 2.	123
4.7	Price consistency in the graph representation of the bidding structure.	124
5.1	A model for bundle price formation	159
5.2	DATASET-1 : Market surplus	163
5.3	DATASET-1 : Cumulative value aggregation	164
5.4	Integrality gaps and CPU times for DATASET-2 test problems . . .	167
5.5	Economic gaps for DATASET-3 test problems	169
5.6	Integrality gaps and CPU times for DATASET-3 problems	170
6.1	A direct-revelation market mechanism.	184
6.2	The subgradient algorithm as an iterative auction.	190

6.3	Evolution of the best upper bound for the three subgradient methods.	205
6.4	Sensitivity of the subgradient methods to $\rho^{(0)}$.	206
6.5	Evolution of the gap associated with the DW-based auction for instances of $S - 01$, $S - 04$, and $S - 09$.	207

Liste des tableaux

2.1	Exemple de non-existence d'équilibre de prix	22
2.2	Exemple d'allocation combinatoire : une allocation fractionnaire est strictement meilleure qu'une allocation entière	41
2.3	Prix des ordres de transport	47
5.1	Example of portfolio bundle trading	144
5.2	DATASET-1 - Basic bundle trading allocation problems	162
5.3	DATASET-2 - Lower bound allocation problems	165
5.4	DATASET-3 - XOR allocation problems	168
6.1	Characteristics of problem instances	201
6.2	Behavior of the subgradient, the CFM, and the modified CFM methods	204

À mes parents

Là où il n'y aura rien, lisez que je vous aime. *Denis Diderot*

Remerciements

J'aimerais tout d'abord exprimer mes remerciements à mes directeurs de recherche, les professeurs Teodor Gabriel Crainic et Michel Gendreau, pour m'avoir proposé le sujet de cette thèse et m'avoir soutenu moralement et financièrement afin de la mener à bien. Je leur suis également très reconnaissant pour la grande confiance qu'ils m'ont témoignée depuis notre première rencontre en 1995, et par la suite durant mes études graduées. Je puis les assurer que le fait de travailler sous leur direction a été une expérience très enrichissante et une introduction gratifiante au monde de la recherche scientifique. J'ose espérer que nous aurons l'occasion de collaborer souvent à l'avenir.

Je désire remercier les membres du jury pour avoir si aimablement accepté d'évaluer cette thèse et pour l'intérêt qu'ils ont bien voulu témoigner à mes travaux de recherche.

Je remercie également les professeurs, chercheurs et professionnels de recherche que j'ai pu côtoyer au Département d'informatique et de recherche opérationnelle de l'Université de Montréal, ainsi qu'au Centre de recherche sur les transports et au CIRANO, pour leur contribution directe ou indirecte à ce travail de recherche, et plus généralement pour le rôle qu'ils ont joué dans ma formation. Je tiens tout particulièrement à rendre hommage aux professeurs Jacques A. Ferland, Patrice Marcotte, Bernard Gendron, Michael Florian, Felisa J. Vázquez-Abad, Sang Nguyen, Jacques Robert et Gilbert Laporte. Je désire par ailleurs saluer messieurs Zine Khélil et Benoît Bourbeau dont l'aide a été très appréciable.

Je profite de l'occasion pour remercier mesdames Lucie l'Heureux, Lucie-Nathalie Cournoyer et Josée Vignola, pour les nombreux services rendus, leur disponibilité de tous les jours et leur extrême gentillesse. Je dois également remercier les responsables des ressources informatiques du Centre de recherche sur les transports, messieurs Daniel Charbonneau et Luc Rocheleau pour leur soutien technique.

J'adresse un clin d'œil amical à tous mes collègues étudiants et amis au Centre de recherche sur les transports et ailleurs, qui ont insufflé de la vie dans les quelques années passées à travailler sur la thèse. Ils sont trop nombreux pour que je les cite par leur nom à chacun. Qu'ils soient néanmoins assurés que j'ai en ce moment une pensée pour eux.

J'aimerais souligner l'appui des organismes suivants, qui m'a permis de mener à bien mon projet de recherche : le Conseil de recherche en sciences naturelles et en génie du Canada (CRSNG), le fonds FQRNT de la province du Québec, le fonds Valorisation-Recherche Québec et les Laboratoires universitaires Bell (LUB)

Je ne saurais conclure ce mot sans rendre un hommage tout particulier à mes parents, Monsieur M'barek Abrache et Madame Halima Abrache, dont le dévouement et l'abnégation ont été mes principales sources d'inspiration. Mes très chers parents, vous avez été une des raisons majeures qui m'ont incité à entreprendre des études doctorales et à persévéérer dans les moments difficiles. Cette thèse vous est ultimement dédiée.

Chapitre 1

Introduction

L'essor des technologies de l'information et la popularité grandissante de l'Internet ont radicalement transformé bon nombre de nos habitudes, à un tel point qu'il est parfois difficile de s'en rendre pleinement compte sans un regard rétrospectif. Un des épisodes marquants de cette révolution numérique a sans doute été l'avènement du *commerce électronique*, terme qui désigne l'ensemble des échanges de biens et de services conduits à travers des médias électroniques, en particulier sur le Web.

Même si le commerce électronique est davantage connu du grand public sous ses formes consommateur-à-consommateur (C2C) et entreprise-à-consommateur (B2C) (en partie grâce au succès de certaines "dotcom" vedettes telles que eBay et Amazon), le volet inter-firmes (B2B) du commerce électronique constitue, sur le plan du volume et de l'importance stratégique des échanges, la forme actuellement prédominante et, selon toute vraisemblance, celle qui possède le plus de potentiel de croissance. Le développement du commerce électronique inter-firmes, ainsi que les vagues récentes de déréglementation dans plusieurs industries (énergie, télécommunications, etc.) et les campagnes de privatisation de services publics conduites par les instances gouvernementales, ont ainsi grandement contribué à l'apparition de *places de marché électroniques*.

Par rapport aux marchés traditionnels, les places de marché électroniques présentent plusieurs avantages. D'abord, leur caractère "virtuel" fait en sorte qu'elles s'affranchissent des contraintes d'espace physique et leur permet d'atteindre et de rassembler des participants géographiquement éloignés les uns des autres. Cette ouverture à des communautés plus importantes de participants favorise les opportu-

nités d'affaire et augmente les possibilités que des liens d'échange potentiel puissent s'établir entre acheteurs et vendeurs de biens et de services. D'autre part, l'automatisation du processus d'échange permet de réduire les procédures administratives et par conséquent les coûts de transaction. Enfin, un gain de compétitivité et d'efficacité économique découle naturellement de la liquidité et de la fluidité accrues du marché.

Les places de marchés électroniques s'organisent très souvent autour de *mécanismes de marché*. Un mécanisme de marché désigne un ensemble de règles formelles qui, à partir d'*informations* sur les préférences des participants pour les objets transigés sur le marché, spécifient une *allocation* des objets aux participants et des *paitements* que ces derniers doivent effectuer ou recevoir. Dans plusieurs cas de marchés, dit *optimisés*, les règles déterminant l'allocation et les paiements sont telles qu'un objectif de marché doit être atteint. Typiquement, cet objectif découle d'un impératif d'efficacité sociale et consiste à produire une allocation maximisant la somme des préférences des participants, ou représente les intérêts d'un participant ou d'un groupe de participants particuliers, en maximisant les surplus de ces derniers.

Une *enchère* est un mécanisme de marché particulier caractérisé par : (i) les participants expriment leurs préférences pour les objets transigés en plaçant des *mises* indiquant les quantités des différents objets que le participant désire vendre ou acheter, et les prix correspondants que ce dernier est prêt à payer ou à recevoir; et (ii) un agent particulier, l'*encanleur*, sert d'intermédiaire entre les participants et prend en charge l'implantation du mécanisme. Compte tenu du fait que les préférences d'un participant pour les différents objets transigés constituent généralement une information privée de ce dernier, une enchère peut être considérée comme étant un mécanisme de marché *indirect* (Fig. 1.1) : les participants mettent au point des *stratégies* de mises qui spécifient, à partir des préférences privées et de l'information disponible sur l'état du marché et sur les compétiteurs, la structure et la dynamique des mises à placer sur le marché.

L'utilisation des enchères est bien entendu antérieure à l'avènement du commerce électronique et à l'étude systématique des mécanismes de marché. Leur importance provient du nombre considérable d'activités économiques qui y ont recours, notamment dans des situations où l'équité et la visibilité sont des critères importants (Roth-

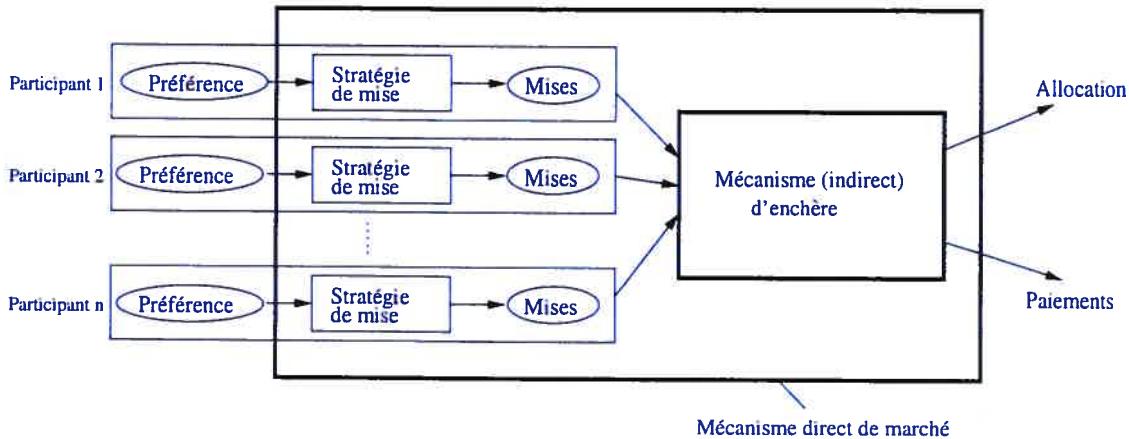


FIG. 1.1 – Enchères : mécanismes de marché indirects

kopf et Park [124]). Les privatisations de services publics, la vente de droits d'accès ou d'exploitation de certaines ressources ou infrastructures (champs pétroliers, exploitations minières, lignes de chemins de fer, etc.) et les appels d'offres pour la propreté d'équipements ou de services ne sont que quelques uns des nombreux contextes où des mécanismes d'encheres sont sollicités. Les encheres sont également le modèle de référence pour transiger des biens qui ne peuvent l'être à travers les canaux de distribution traditionnels, tels qu'objets rares, produits reconditionnés ou en rupture de stock, ou encore certains produits à durée de vie très limitée (billets d'avion, par exemple). D'un autre point de vue, un mécanisme d'encheres peut être utilisé comme processus distribué pour le partage équitable de ressources (temps CPU dans un système multi-tâches, utilisation des machines dans un atelier de production, etc.). Les avantages des encheres sont nombreux : étant très souvent des mécanismes de marché optimisés, elles peuvent se prévaloir de leur efficacité économique par rapport à des procédés tels que loteries, audit de propositions, et autres procédures administratives. Par ailleurs, le fait qu'elles soient des mécanismes de marché indirects signifie qu'elles préservent l'autonomie et la confidentialité de l'information des participants.

Dans les situations de marché les plus simples, un objet unique disponible en une unité est mis en vente. Pour ces marchés, des mécanismes d'encheres classiques tels que l'encheres anglaise ascendante, l'encheres hollandaise descendante et les encheres à enveloppe fermée "plus-haut-prix" et "second-prix" peuvent être suggérés. Ces mécanismes ont depuis longtemps été étudiés (Vickrey [140]) et leurs propriétés

théoriques sont bien connues. Néanmoins, dans un grand nombre de marchés importants, plusieurs objets hétérogènes doivent être simultanément transigés, et ces objets sont *interdépendants* pour les participants, dans la mesure où la préférence d'un participant pour l'achat ou la vente d'un objet donné dépend de l'achat ou de la vente d'autres objets. La vente de droits d'utilisation de la bande de fréquences pour les télécommunications sans fil constituent un exemple probant de marché d'objets interdépendants. Typiquement, les objets transigés dans ces marchés sont des licences pour l'utilisation de différentes bandes de fréquences couvrant différentes régions. D'une part, la proximité géographique et la possibilité de consolider les équipements font en sorte que les opérateurs télécom ont tendance à considérer que des licences sur des régions connexes sont complémentaires les unes des autres. D'autre part, le fait que deux bandes de fréquences différentes puissent être indifféremment utilisées pour assurer un service donné les rend substituables les unes des autres.

L'*enchère simultanée ascendante* (SAA), qui revient tout simplement à conduire en parallèle, mais de façon indépendante, autant d'enchères (anglaises) que le marché compte d'objets différents, est un mécanisme classique qui peut être proposé pour la vente de plusieurs objets hétérogènes. Une série d'enchères de ce type a d'ailleurs été conduite par la Commission fédérale américaine des Communications (FCC) pour la cession des premières licences d'utilisation de la bande de fréquences (Milgrom [100], Cramton [34]). L'analyse de la SAA dans Milgrom [102] indique cependant le handicap suivant : en présence de complémentarité, un participant désirant obtenir un paquet d'objets donné peut décider de continuer de miser sur les objets individuels du paquet même si le prix de certains de ces objets sur le marché dépasse leur valeur pour le participant. Ce faisant, ce dernier court le risque de se retrouver avec un sous-ensemble d'objets du paquet désiré qu'il aura payé plus que sa valeur. Afin d'éviter ce risque, les participants adoptent très souvent un comportement stratégique qui peut potentiellement conduire à des enchères économiquement inefficaces. Les *enchères combinatoires*, où les participants peuvent placer des mises combinées portant sur des "paquets" d'objets, contournent cette difficulté en rendant possible l'expression explicite des préférences des participants pour les différents ensembles d'objets.

La conception de mécanismes d'enchère combinatoire est une problématique *fon-*

damentalement complexe. Bien que les premiers travaux portant sur les enchères combinatoires aient été consacrés à un problème particulier, en l'occurrence, la détermination des mises gagnantes dans une enchère combinatoire à enveloppe fermée d'objets indivisibles visant à maximiser le revenu de l'enchanteur (Rothkopf, Pekeč et Harstad [125]), la situation a radicalement changé depuis, et la littérature des enchères combinatoires s'est considérablement enrichie et compte désormais des contributions importantes de plusieurs disciplines (économie, recherche opérationnelle, intelligence artificielle, etc.). Le chapitre 2 donne un aperçu de cet état de l'art. En particulier, l'apport de la recherche opérationnelle dans la modélisation de problèmes de décision reliés à la conception de mécanismes d'enchère et la mise au point d'approches de résolution efficaces est primordial. Ces problèmes se situent aussi bien dans le périmètre strict du mécanisme (règles d'allocation et de paiements), qu'à sa périphérie, c'est-à-dire au niveau de l'élaboration des stratégies de mises des participants (quand, sur quoi, et combien miser ?).

Nous abordons dans le chapitre 3 la problématique de la conception de mécanismes d'enchère combinatoire. Ce chapitre comporte quatre volets, présentant chacun une facette de la problématique générale. Dans le premier volet, nous procédons à une classification multidimensionnelle de l'espace des marchés et des enchères, qui complète les efforts de catégorisation semblables de Engelbrecht-Wiggans [44] et de Wurman, Wellman et Walsh [150], et qui en même temps nous permet d'uniformiser la terminologie que nous utilisons tout au long de la thèse. Nous dérivons naturellement plusieurs formulations de base du problème de la détermination des mises gagnantes correspondant à des contextes différents : enchères combinatoires “directes” (un vendeur, plusieurs acheteurs), “renversées” (un acheteur, plusieurs vendeurs), “doubles” (plusieurs acheteurs, plusieurs vendeurs), ainsi que des enchères combinatoires de ressources dites “réseau”, qui peuvent être assimilées à de la capacité dans une structure de réseau. Le second volet discute le besoin de *langages de mise*, qui sont des formalismes permettant aux participants d'exprimer de manière succincte leurs préférences pour les objets transigés. Le troisième volet est consacré aux mécanismes itératifs d'enchère combinatoire qui procèdent en plusieurs rondes de mise, et dans lesquels les participants ne sont pas astreints de révéler intégralement leurs préférences privées

à l'enchanteur, mais se contentent de miser en fonction de "signaux" sur l'état de l'enchère reçus de l'enchanteur au fur et à mesure que cette dernière progresse. Finalement, le quatrième volet discute brièvement des problèmes de décision des participants et du besoin d'outils d'aide à la décision ("aviseurs") pour assister ces derniers dans l'élaboration de leurs stratégies de mise en fonction de leur contexte économique propre, de la connaissance qu'ils ont de la compétition, et de l'évolution de l'état du marché.

L'ingénierie des langages de mise pour les enchères combinatoires est un axe de recherche qui a été particulièrement fécond en réalisations durant les dernières années. Les langages proposés dans la littérature (Sandholm [130], Fujushima, Leyton-Brown et Shoham [55], Nisan [105], Boutilier et Hoos [23]) ont cependant en commun le fait qu'ils ne s'adressent qu'à des enchères d'objets indivisibles. Notre contribution à cet égard, qui fait l'objet du chapitre 4, consiste en un nouveau cadre formel unificateur pour la définition de langages de mise pour les enchères combinatoires de biens divisibles et indivisibles. Une remarque fondamentale est que dans le cas divisible, un langage de mise doit disposer des outils nécessaires permettant aux participants d'exprimer non seulement des conditions logiques (reliées à l'exécution ou non d'une mise), mais également de conditions relatives aux proportions d'exécution de ces mises. Sur le plan méthodologique, cela se traduit par une extension de la notion d'*opérateur de mise*, et une structure à deux niveaux d'une mise combinée : un niveau interne, où des mises "atomiques", désignant des ordres simples d'achat ou de vente, sont combinées au besoin pour exprimer des conditions sur les proportions d'exécution, et un niveau externe où la définition d'une mise combinée est complétée par application récursive d'opérateurs logiques. L'analyse de langages de mise particuliers dérivés de ce cadre méthodologique montre que, dans le cas indivisible, son expressivité est au moins égale à celle du langage "Formules OR/XOR" de Nisan. Quant au cas divisible, nous établissons sa capacité à supporter une classe très générale de fonctions de préférence, ainsi que son potentiel d'expression *concise* de quelques fonctions de préférence particulières.

Les marchés financiers constituent un domaine d'application de choix des mécanismes d'enchère combinatoire. Plusieurs auteurs (Fan, Stallaert et Whinston [50],

Bossaerts, Fine et Ledyard [21]) ont récemment considéré la possibilité de permettre des ordres *composites* formés de plusieurs mises simples d'achat et de vente de valeurs financières, avec la condition que ces mises soient exécutées dans les mêmes proportions. Les ordres composites sont particulièrement utiles dans un contexte de "fin de séance", où les gestionnaires d'actifs financiers doivent souvent "rééquilibrer" leurs portefeuilles afin d'atteindre une composition cible, dans la mesure où la consolidation des opérations d'achat et de vente diminue le risque de terminer la séance avec un portefeuille non équilibré. Nous présentons dans le chapitre 5 la formulation de base du problème d'allocation dans une enchère mono-ronde de valeurs financières avec ordres composites. Les extensions que nous suggérons permettent notamment aux participants de spécifier des bornes inférieures sur les proportions d'exécution de leurs ordres, ainsi que d'exiger qu'une sélection soit faite entre ordres jugés "équivalents" et par conséquent substituables les uns aux autres.

L'éventualité que le mécanisme d'enchère puisse fournir une multitude de solutions optimales soulève la question du choix de l'allocation et des prix à implanter. Dans un contexte comme celui des marchés financiers, il est important qu'une procédure de discrimination des solutions sur la base d'un critère "éthique" soit mise en place. Quoique ce problème ait été discuté dans la littérature, à notre connaissance aucune solution pratique satisfaisante n'a été proposée. Notre approche au problème est comme suit. Nous suggérons le temps de soumission des ordres comme premier critère de sélection possible, sur la base duquel nous définissons deux relations d'ordre lexicographique sur les espaces des solutions primales et duales optimales du problème d'allocation. Ces relations d'ordre formalisent des stratégies de choix des solutions qui consistent à privilégier (aux niveaux des proportions d'exécution et des paiements à effectuer), les participants ayant soumis les premiers ordres. Nous montrons ensuite qu'en procédant par améliorations successives des solutions optimales, il est possible d'aboutir aux "meilleurs" allocations et prix selon le critère du temps de soumission.

Dans un grand nombre de contexthes de marché, l'hypothèse que le mécanisme de marché dispose *directement* de toute l'information pertinente nécessaire à la détermination des allocations aux participants et des paiements effectués par ces derniers est trop contraignante. En effet, quand l'objectif du marché est "global" (par

exemple, maximiser le bien-être social des participants), il n'est pas acquis, a priori, que des participants aspirant chacun à maximiser leur propre surplus acceptent de dévoiler leurs préférences privées pour les objets transigés sur le marché. En outre, et c'est particulièrement tangible dans les marchés combinatoires, l'évaluation *complète* et *précise* de ces préférences constitue souvent une tâche plus complexe que ce dont les ressources limitées des participants sont en mesure d'accomplir. Ce problème constitue la motivation principale pour la conception de mécanismes de marché à *révélation progressive* de l'information : sur la base de *mises* effectuées par les participants (qui ne représentent pas les demandes "définitives" des participants, mais reflètent plutôt leurs besoins étant donné l'état actuel du marché), l'encanteur détermine une allocation et des paiements provisoires et "signale" le nouvel état temporaire du marché aux participants, de manière à ce que ces derniers puissent mettre à jour leurs mises en conséquence.

Les enchères combinatoires *itératives* (ou multi-rondes), dans lesquelles l'activité de mise et la détermination des allocations et des paiements provisoires sont synchronisées par des événements prédéfinis (débuts et fins de ronde), forment une classe importante de mécanismes de marché progressifs. À la lumière de l'état de l'art actuel, la littérature des enchères combinatoires itératives est dominée par deux types d'approches. Pour les enchères d'objets divisibles, des résultats classiques de la théorie de l'équilibre général dans les économies d'échange (Mas-Colell, Whinston et Green [95]), notamment ceux concernant l'existence et la stabilité d'équilibres de prix, sont mis à profit dans la conception de mécanismes itératifs d'ajustement de prix selon le principe du tâtonnement walrasien (les prix de biens excédentaires diminuent et ceux de biens déficitaires augmentent). En ce qui concerne les marchés d'objets indivisibles, l'emphase a été placée sur le cas particulier d'enchères combinatoires faisant intervenir un vendeur et plusieurs acheteurs potentiels d'objets disponibles chacun en une unité. En effet, l'existence d'une formulation de programmation linéaire du problème consistant à déterminer l'allocation maximisant le bien-être social des participants (Bikhchandani et Ostroy [20]) a rendu possible la mise en oeuvre de l'approche *principale/duale* dans le développement d'un certain nombre de mécanismes itératifs. Une avenue de recherche jusqu'ici peu explorée est celle des méthodes de décomposition en

programmation mathématique et la possibilité d'en dériver des mécanismes de marché itératifs. Utilisées depuis les débuts de la recherche opérationnelle dans l'optimisation de problèmes structurés de grande taille, les approches de décomposition ont à cet effet un potentiel de prise de décision décentralisée et une interprétation économique remarquables. Ainsi, nous considérons dans le chapitre 6 une économie très générale composée de producteurs et consommateurs de biens, et nous établissons que, aussi bien des approches de décomposition "duales" basées sur la relaxation lagrangienne (notamment l'algorithme du sous-gradient et les méthodes dites de "faisceaux") que la génération de colonnes selon le principe de décomposition de Dantzig-Wolfe peuvent être associées à des processus itératifs d'enchère. Nous discutons ensuite certaines différences importantes qui existent entre les deux types d'approches, notamment en ce qui a trait à leur utilisation de l'information contenue dans les mises des participants et aux hypothèses qu'elles font sur le comportement stratégique de ces derniers. Enfin, l'étude expérimentale entreprise dans ce chapitre a pour objectif l'évaluation des mécanismes d'enchère suggérés du point de vue de leur efficacité économique et du volume d'informations divulguées par les participants.

Les chapitres 3 à 7 consistent en des articles rédigés en collaboration avec Teodor Gabriel Crainic et Michel Gendreau (ainsi que Benoît Bourbeau en ce qui concerne l'article du chapitre 4). Les trois premiers articles ont été soumis à des revues scientifiques de langue anglaise et ont été acceptés pour publication. Quant au quatrième article, nous publions dans cette thèse la toute dernière version dont nous disposons, un manuscrit de l'article devant être soumis incessamment. Disposant chacun de sa propre revue de littérature, les articles peuvent, dans une grande mesure, être lus indépendamment les uns des autres.

Cette thèse a été conduite dans le cadre du projet TEM ("Towards Electronic Marketplaces"), un projet de recherche qui a rassemblé sur une période de quatre années des chercheurs de plusieurs disciplines (sciences économiques, recherche opérationnelle, génie logiciel, télé-informatique) autour de problématiques diverses reliées à la conception d'infrastructures, de mécanismes de marché et d'outils de décision pour les places de marché électroniques. À cet effet, nous tenons à souligner le support financier du Conseil de recherche en sciences naturelles et en génie du Canada (CRSNG), du Fonds

québécois de la recherche sur la nature et les technologies, de Valorisation-Recherche Québec et des Laboratoires universitaires Bell (LUB).

Chapitre 2

Revue de littérature

Entreprendre une revue de littérature des enchères combinatoires est une tâche qui s'avère être d'un réel intérêt intellectuel mais qui comporte bon nombre de difficultés. En effet, l'étendue et la diversité des problèmes que l'on peut placer sous la bannière des "enchères combinatoires", ainsi que le caractère multidisciplinaire des contributions sont probablement les premières constatations qui interpellent le chercheur désireux de dresser l'état de l'art. D'autre part, la relative nouveauté du domaine de recherche fait en sorte que beaucoup de travaux et de résultats importants continuent d'affluer, ce qui exige un effort constant de suivi.

Il est possible d'envisager deux façons différentes d'aborder une telle revue de littérature. La première consiste tout simplement à suivre le fil chronologique des principaux travaux (problématique, modèles, formalismes, mécanismes, etc.) entrepris par les chercheurs. La seconde approche, que nous avons retenue pour notre revue, se veut plutôt "thématische" et se base sur une sélection de certains axes de recherche importants. Nous avons jugé que ce parti était plus à même de faciliter la compréhension des développements, au prix d'une présentation parfois anachronique des résultats. Notons également que cette revue n'a pas la prétention d'être un exposé exhaustif des travaux réalisés dans le domaine. À cet égard, nous référons le lecteur intéressé aux revues de De Vries et Vohra [39], de Kalagnanam et Parkes [70], et de Xia, Koehler et Whinston [151] qui apportent un point de vue complémentaire.

Ce chapitre est organisé comme suit. Dans les sections 2.1 et 2.2, nous présentons quelques résultats fondamentaux concernant les enchères d'un objet unique et les enchères multi-unités. La section 2.3 est consacrée au cadre qui nous intéresse, ce-

lui des enchères de plusieurs objets hétérogènes interdépendants. Dans cette section, nous exposons les développements reliés aux modèles d'équilibre de prix pour les objets divisibles, les principales formulations d'encheré combinatoire pour les biens indivisibles, ainsi que quelques domaines d'application choisis. Nous dédions la section 2.4 aux mécanismes progressifs d'encheré combinatoire. La section 2.5 traite de l'expression de la préférence et met l'emphase sur les langages de mise combinatoire. Enfin, nous présentons dans la section 2.6 les travaux reliés au développement de mécanismes d'encheré combinatoire ayant la propriété d'inciter les participants à dévoiler leurs véritables préférences.

2.1 Premiers travaux

Bien que les enchères aient vraisemblablement été utilisées depuis l'aube de l'humanité, leur étude formelle et systématique est relativement récente et date du début des années 60. Dans un article célèbre, Vickrey [140] étudie le cas d'encheres d'un objet unique, et compare deux mécanismes d'encheré classiques : (i) l'encheré anglaise ascendante, où les participants misent à tour de rôle des prix de plus en plus élevés, jusqu'à ce qu'aucun d'eux n'ait plus envie de continuer, l'objet allant alors au participant ayant placé la plus haute mise et ce dernier paie la valeur de sa mise ; et (ii) l'encheré hollandaise descendante où des prix de moins en moins élevés sont annoncés par l'encanteur jusqu'à ce qu'un seul participant soit prêt à payer le prix annoncé. Vickrey montre que dans le cas symétrique, c'est-à-dire lorsque tous les participants ont des fonctions de préférence qui suivent la même distribution de probabilité, ces deux mécanismes d'encheré sont équivalents dans le sens où les stratégies Pareto-optimales des participants, ainsi que les valeurs espérées des paiements, sont les mêmes dans les deux enchères. Vickrey note également que le placement de nouvelles mises dans l'encheré anglaise s'arrête approximativement quand le prix de l'objet atteint la *seconde plus haute* valeur de l'objet pour les participants, ce qui l'amène à prouver que l'encheré anglaise ascendante est en fait une implantation itérative d'un mécanisme d'encheré "second-prix" (dans lequel le participant qui obtient l'objet ne paie pas le prix déclaré dans sa propre mise, mais plutôt celui de la plus haute mise perdante), "à enveloppe fermée" (qui se déroule en une seule ronde, et où les mises sont confi-

dentielles). Néanmoins, l'importance de l'article de Vickrey provient surtout du fait qu'il démontre que l'enchère “second-prix” possède la propriété remarquable d'inciter les participants à miser leur vraie préférence pour l'objet mis en vente. Enfin, Vickrey propose une extension du mécanisme “second-prix” aux enchères multi-unités, où les participants sont en compétition pour l'achat de plusieurs unités d'un même objet.

L'article de Myerson [103] est un apport significatif à la théorie des enchères. Myerson considère le problème de la conception d'une enchère *optimale* (maximisant le revenu du vendeur), où plusieurs participants entrent en compétition pour l'achat d'un objet unique. Selon le modèle de l'information utilisé, dit des *valeurs indépendantes privées*, chaque participant connaît de façon confidentielle la valeur qu'il attribue à l'objet, alors que le vendeur ne dispose que de variables aléatoires mutuellement indépendantes représentant les valeurs de l'objet pour les différents participants. L'auteur définit des conditions de “réalisabilité” de l'enchère (rationalité des participants et enchère incitative au dévoilement de la valeur de l'objet), puis exprime le problème de déterminer l'enchère réalisable optimale sous forme d'un problème de programmation mathématique. L'étude qui en découle fournit un certain nombre de résultats importants. Parmi ceux-ci, le plus fondamental est sans doute une preuve très générale du principe dit de *l'équivalence des revenus* (the “Revenue Equivalence Theorem”), qui stipule que les mécanismes d'enchère dans lesquels sont vérifiées les deux conditions suivantes : (i) l'objet mis en vente va à un des participants ayant la plus haute valeur, et (ii) tout participant reçoit une utilité espérée nulle si sa valeur privée est la plus petite, ont tous le même revenu espéré. Indépendamment de Myerson, Riley et Samuelson [122] démontrent le même résultat et proposent différents exemples d'enchères illustrant le principe de l'équivalence des revenus. Un traitement plus général des enchères optimales, où sont relaxées certaines hypothèses du modèle d'information de Myerson, est réalisé par Harris et Raviv [64] et Maskin et Riley [96].

Dans le modèle dit de la *valeur commune*, la valeur intrinsèque de l'objet mis en vente est *la même* pour tous les participants. Cependant, cette valeur ne sera rendue publique - au gagnant ou à tous les participants à l'enchère - qu'après la dernière terminée. Chacun des participants doit donc se contenter, pour élaborer sa stratégie de mise, d'une estimation propre de la valeur de l'objet. Par conséquent, le

modèle de la valeur commune convient particulièrement bien à des enchères reliées aux privatisations de biens publics et d'exploitation d'infrastructures au profit de participants partiellement informés sur la valeur des biens (contrats d'exploitation de champs pétroliers, de terrains forestiers, etc.). Parmi les contributions importantes réalisées dans ce cadre, Wilson [146] étudie le processus de mise d'une enchère à enveloppe fermée dans un environnement à valeur commune, et détermine des stratégies d'équilibre du jeu correspondant. Plusieurs autres auteurs, dont Ortega-Reichert [110], Rothkopf [123], Reece [121] et Milgrom [101] analysent plus en détail le modèle de la valeur commune et le phénomène de la surestimation de l'objet par le gagnant (“the winner's curse”) qui l'accompagne.

2.2 Enchères multi-unités

Nous commençons par présenter un modèle simplifié des enchères multi-unités. Une version légèrement différente est proposée dans Ausubel et Cramton [10].

Soit K unités d'un objet et n participants. Le participant j désire acquérir jusqu'à m_j unités de l'objet, avec $\sum_{j=1}^n m_j > K$. L'objet a pour le participant j une valeur unitaire v_j , connue avec certitude uniquement par ce dernier. Sans perte de généralité, nous supposons que pour le vendeur et les autres participants, v_j est une variable aléatoire qui suit une certaine loi de probabilité F_j sur $[0, 1]$.

Chaque participant j , $1 \leq j \leq n$, déclare une fonction de mise $b_j : \{0, \dots, m_j\} \rightarrow [0, 1]$, où $b_j(q)$ est le prix maximal que le participant j est prêt à payer pour acquérir la q -ième unité de l'objet. Nous supposons que b_j est monotone décroissante et que $b_j(0) = 1$ et $b_j(m_j) = 0$.

Une fonction de demande $q_j : [0, 1] \rightarrow \{0, \dots, m_j\}$, qui est en quelque sorte la réciproque de b_j , est définie pour tout prix b comme étant le nombre maximal d'unités que le participant j voudrait acquérir à ce prix, $q_j(b) = \max\{q \in \{0, \dots, m_j\} : b_j(q) \geq b\}$.

Un prix d'équilibre du marché peut alors être calculé, sur la base des fonctions q_j , $1 \leq j \leq n$, comme suit : $p_{bal} = \inf\{b \in [0, 1] : \sum_{j=1}^n q_j(b) \leq K\}$. Les mises

supérieures au prix de balance p_{bal} sont acceptées et la quantité $q_j(p_{bal})$ est attribuée au participant j , $1 \leq j \leq n$.

Dans le contexte multi-unités, les différents mécanismes d'enchère se distinguent par les paiements effectués par les participants. Le plus souvent, on a recours à l'un des deux modes de paiement suivants :

- Enchère à *prix uniforme*, avec un paiement total du participant j ,
 $P_j = p_{bal}q_j(p_{bal})$, $1 \leq j \leq n$.
- Enchère “*pay-your-bid*”, avec un paiement total du participant j ,
 $P_j = \sum_{q=0}^{q_j(p_{bal})} b_j(q)$, $1 \leq j \leq n$.

L'usage de ces deux mécanismes d'enchère présente toutefois des difficultés. Ainsi, Noussair [108] étudie le cas particulier de la vente par l'enchère uniforme de deux objets identiques et établit qu'une stratégie dominante consiste pour un participant à déclarer sa véritable préférence sur le premier objet, mais à offrir un prix moindre pour le second. Engelbrecht-Wiggans et Kahn [45] et Ausubel et Cramton [10] généralisent ce résultat et montrent, autant pour l'enchère à prix uniforme que pour l'enchère “*pay-your-bid*”, que ce comportement stratégique des participants peut mener à une réduction de la demande et à des enchères inefficaces où certaines unités ne vont pas aux agents qui les évaluent le plus. Ausubel et Cramton montrent également que l'enchère de Vickrey généralisée au cas multi-unités, en étant incitative à la déclaration des vraies préférences, évite la réduction de la demande. Ausubel [8] propose une implantation de l'enchère de Vickrey sous forme d'un mécanisme d'enchère progressive ascendante.

2.3 Enchères de plusieurs objets interdépendants : modèles de base et applications choisies

Un certain nombre de contextes de marché importants ne rentrent pas dans le cadre simple des enchères d'un objet unique, ni dans celui des enchères multi-unités. Dans ces marchés, mettant en vente des objets différents, le fait que la valeur d'un objet pour un participant puisse dépendre de l'obtention ou non d'autres objets joue un rôle déterminant. Pour fixer les idées, considérons un ensemble S d'objets indivi-

sibles et désignons par $v_j(S)$ la valeur de S pour un participant j . L'interdépendance des objets prend principalement deux formes. Deux sous-ensembles disjoints d'objets S_1 et S_2 de S peuvent être *complémentaires* si $v_j(S_1 \cup S_2) > v_j(S_1) + v_j(S_2)$ ou *substituables* si $v_j(S_1 \cup S_2) < v_j(S_1) + v_j(S_2)$. La complémentarité et la substitutabilité des objets est illustrée par les deux exemples suivants.

Exemple - 1 : Fenêtres temporelles de décollage et d'atterrissement aux aéroports (Rasenti, Smith et Bulfin [120])

Une fenêtre de décollage d'un aéroport correspondant à l'origine d'un vol est complémentaire à une fenêtre précise d'atterrissement à l'aéroport correspondant à la destination du vol, car les deux vont de pair et aucune d'elles n'a de valeur en soi pour une compagnie aérienne. Par contre, deux paquets de fenêtres de décollage/atterrissement, correspondant tous deux aux mêmes aéroports origine et destination, mais à des heures différentes de la journée, peuvent se substituer l'un à l'autre si la compagnie n'est intéressée à offrir qu'un seul service quotidien entre les deux aéroports.

Exemple - 2 : Droits d'utilisation de la bande de fréquences en télécommunications

Les licences d'utilisation de la bande de fréquences pour les télécommunications sans fil sont souvent fortement interdépendantes les unes des autres pour les opérateurs offrant les services de télécommunications. Cette interdépendance est surtout due à des raisons de couverture géographique (économies d'échelle réalisées avec des équipements adéquatement placés) et de convenance technique des fréquences aux besoins de l'opérateur. À titre d'exemple, considérons l'enchère no.26 conduite par la Commission fédérale américaine des télécommunications (FCC), où des licences d'utilisation de fréquences destinées à un service de pagette ("Paging") ont été mises en vente. Ces fréquences font partie des bandes 929 MHz (12 fréquences) et 931 MHz (37 fréquences) et couvrent 51 régions économiques principales (MEA) du territoire des États-Unis. Pour un opérateur typique présent sur le marché de la côte-est américaine, les licences CZMEA001AA, CZMEA002AA et CZMEA004AA, toutes reliées à la fréquence 931.0125 Mhz, mais couvrant respectivement les régions de Boston, New York et Philadelphie, seront très probablement considérées comme étant complémentaires.

taires les unes des autres. Par contre, les licences CZMEA001AA (931.0125 Mhz) et CZMEA001AB (931.0375 Mhz) aux fréquences très proches pourraient selon toute vraisemblance se substituer l'une à l'autre.

La présente section est consacrée à l'essentiel de la littérature concernant les enchères de plusieurs objets interdépendants. Cette littérature se divise en deux groupes de travaux : le premier concerne les enchères de biens divisibles et est surtout dominé par la théorie économique de l'équilibre. Le second groupe concerne quant à lui les objets indivisibles.

2.3.1 Modèles d'équilibre de prix

La théorie de l'équilibre général en micro-économie fournit un cadre théorique adéquat pour l'étude des enchères d'objets divisibles différents. Afin de clarifier les notions importantes, nous introduisons dans un premier temps le modèle relativement simple de l'équilibre de prix dans une économie *à échange pur*.

Considérons un ensemble de L produits et n participants. À chaque participant j , $1 \leq j \leq n$, sont associés un vecteur $w_j \in \mathbf{R}_+^L$ de quantités initiales des L produits dont dispose le participant et une fonction d'utilité $u_j : \mathbf{R}_+^L \rightarrow \mathbf{R}$ telle que $u_j(x_j)$ est la "valeur" (préférence) de l'allocation $x_j = (x_j^1, \dots, x_j^L)$ pour le participant j . L'objectif de l'économie est d'organiser des échanges de produits entre les participants de telle manière à réaliser une allocation *socialement efficace*, qui maximise la somme des préférences des participants. Pour des participants réagissant de manière compétitive à un système de prix, la notion d'équilibre walrasien est fondamentale.

Définition 2.3.1.1 Une allocation $\{x_j^*\}_{1 \leq j \leq n} \in \mathbf{R}_+^L$ et un vecteur de prix $p^* \in \mathbf{R}_+^L$ constituent un équilibre walrasien si :

- (i) $\forall j, 1 \leq j \leq n, x_j^* = x_j(p^*, p^* \cdot w_j)$ maximise u_j sur l'ensemble $\{x \in \mathbf{R}_+^L : p^* \cdot x_j \leq p^* \cdot w_j\}$
- (ii) $\sum_{j=1}^n x_j^* - \sum_{j=1}^n w_j \leq 0$
- (iii) $p^* \cdot \sum_{j=1}^n (x_j^* - w_j) = 0$

La condition (i) exprime le fait que chaque participant détermine, en fonction des prix des produits, l'allocation qui maximise son utilité tout en respectant sa contrainte de budget limité. Le fait que les quantités échangées ne peuvent dépasser, en volume total, les quantités initiales dont disposent les participants est exprimé par la condition (ii). La condition (iii), quant à elle, est une conséquence directe de la loi de Walras, qui stipule que dans une économie d'échange pur, il n'y a pas d'interaction avec l'extérieur du système et ainsi la somme des dépenses engagées dans l'obtention des quantités $\{x_j^*\}_{1 \leq j \leq n}$ est égale à la somme des ressources découlant de la possession de quantités initiales $\{w_j\}_{1 \leq j \leq n}$.

On définit la fonction d'excès de demande $z_j : \mathbf{R}^L \rightarrow \mathbf{R}$ d'un participant j par $z_j(p) = x_j(p^*, p^* \cdot w_j) - w_j$, et la fonction d'excès cumulé de demande $z(p) = \sum_{j=1}^n z_j(p)$. On peut, grâce à ces fonctions, caractériser de façon plus compacte un prix d'équilibre. En effet, $p^* \in \mathbf{R}_+^L$ est un prix d'équilibre si et seulement si $z(p^*) \leq 0$.

Quand les fonctions d'utilité u_j , $1 \leq j \leq n$ sont *continues, strictement concaves* et *fortement monotones* (Mas-Colell, Whinston et Green [95]), un équilibre walrasien existe, et possède des propriétés économiques intéressantes. Parmi ces propriétés, les deux résultats suivants découlent des théorèmes fondamentaux du bien-être, et permettent de voir l'allocation découlant d'un équilibre walrasien comme étant une solution économiquement acceptable : chaque équilibre walrasien $(\{x_j^*\}_{1 \leq j \leq n}, p^*)$ correspond à une allocation $\{x_j^*\}_{1 \leq j \leq n}$ Pareto-optimale, et inversement, pour chaque allocation $\{x_j^*\}_{1 \leq j \leq n}$ Pareto-optimale, il existe un vecteur prix p^* tel que $(\{x_j^*\}_{1 \leq j \leq n}, p^*)$ est un équilibre walrasien.

La détermination de l'équilibre walrasien a suscité énormément d'intérêt. Scarf [134] a été le premier à développer des algorithmes élémentaires d'approximation basés sur les théorèmes du point fixe de Brouwer et de Kakutani. Mathiesen [97] propose de résoudre le problème de complémentarité non linéaire correspondant à l'équilibre walrasien par une série d'approximations linéaires. Nagurney [104] dérive une formulation en inégalité variationnelle du modèle de l'équilibre walrasien et propose deux algorithmes itératifs de résolution : une méthode projective et un algorithme d'approximation basé sur la diagonalisation.

Une extension naturelle de l'économie à échange pur est une économie qui distingue les participants selon qu'ils sont *producteurs* ou *consommateurs* de biens. Afin d'alléger la présentation, nous nous contenterons de définir sommairement les grandes lignes d'une économie de production-consommation. Chaque producteur possède une capacité et un coût de production donnés des différents biens, et obéit à un impératif économique de maximiser son profit, qui correspond à la différence entre le paiement qu'il reçoit de la part des consommateurs et son coût de production. Les consommateurs, comme dans l'économie de l'échange pur, cherchent à maximiser leur utilité tout en respectant la contrainte sur le budget disponible. La notion d'équilibre reste sensiblement la même, en tenant bien sûr compte de ces nouveaux éléments.

Le mécanisme classique du *tâtonnement* basé sur l'ajustement des prix ("Price-tâtonnement"), suggéré initialement par Léon Walras [141], est vraisemblablement l'une des toutes premières tentatives de déterminer un équilibre compétitif dans une économie de production-consommation. Le mécanisme de tâtonnement procède de la manière suivante : les producteurs et les consommateurs déterminent leurs fonctions d'offre et de demande pour chaque produit en fonction des prix courants et les envoient à l'encanteur. Si pour un produit donné, l'offre excède la demande, l'encanteur révise le prix à la baisse. Si, au contraire, il y a davantage de demande que d'offre, le prix est révisé à la hausse. Ce mécanisme est illustré par la figure (2.1) dans le cas d'un seul produit. Quand l'équilibre de prix est *stable*, le processus d'ajustement converge vers un équilibre walrasien. Une condition suffisante pour la stabilité est ce qu'on appelle la condition de substitutabilité, "gross substitutability" : l'augmentation du prix d'un produit ne fait pas diminuer la demande pour les autres produits (Arrow et Hahn [7]).

Wellman [143] présente la *programmation orientée-marchés* ("Market-Oriented Programming") comme étant un paradigme générique pour le développement de systèmes de contrôle distribués, permettant la modélisation et l'implantation de ces systèmes sous forme de communautés d'agents autonomes réagissant à un système de prix. En profitant de la synergie de la théorie économique de l'équilibre et des techniques multi-agents, la programmation orientée-marchés se fixe comme objectif principal la mise au point de systèmes distribués au design robuste et qui, de sur-

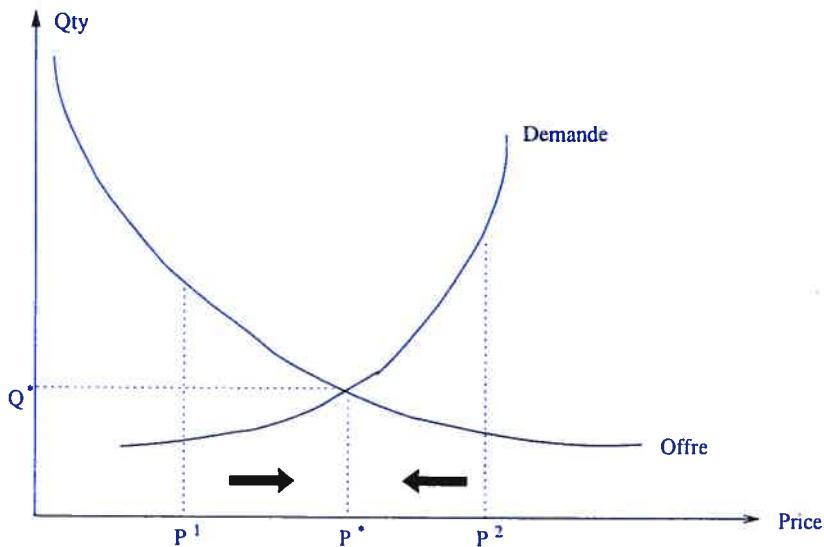


FIG. 2.1 – Tâtonnement à base d'ajustement des prix.

croît, fournissent des allocations de ressources économiquement efficaces. Parmi les réalisations à inscrire au chapitre de la programmation orientée-marchés, citons le système WALRAS (Wellman [143], Cheng et Wellman [30]), qui est une implantation quasi-directe du mécanisme d'ajustement de prix original de Walras. La programmation orientée-marchés a été mise à contribution dans plusieurs applications concrètes, dont une implantation distribuée du problème de transport multi-commodités (Wellman [144]), l'allocation de la qualité de service (QoS) dans les applications multimédia (Yamaki, Wellman et Ishida [152]), et l'ordonnancement dans un environnement distribué (Wellman *et al.* [145]).

Les mécanismes d'enchère basés sur l'ajustement des quantités “Quantity-tâtonnement” forment une seconde classe de processus de tâtonnement, appellés également processus de tâtonnement orientés-ressources. Bien qu'ayant sollicité moins d'effort d'étude de la part des chercheurs, comparativement au tâtonnement walrasien, cette classe de mécanismes est intéressante sur le plan méthodologique, et il convient de la présenter brièvement (nous référerons le lecteur intéressé à la thèse doctorale de Ygge [153] qui fournit plus de détails à cet effet). Un mécanisme de tâtonnement orientés-ressources est un processus itératif dans lequel l'encanteur détermine des allocations provisoires des produits aux participants, et ces derniers soumettent les prix qu'ils sont prêts à payer ou à recevoir pour produire ou consommer

à ce niveau d'allocation. Ces prix prennent souvent la forme de préférences *marginales*, c'est-à-dire de prix associés à une unité supplémentaire produite ou consommée. D'une itération à l'autre, l'enchanteur ajuste les allocations selon le principe suivant : la quantité échangée d'un bien donné augmente si le prix marginal proposé par un consommateur est plus élevé que celui proposé par un producteur, et diminue dans le cas contraire. Le processus d'ajustement s'arrête quand tous les participants sont prêts à payer ou à recevoir le même prix marginal pour tous les produits. Parmi les auteurs qui ont utilisés des processus de tâtonnement basés sur l'ajustement des quantités, citons Kurose et Simha [79] qui proposent une approche orientée-ressources pour l'allocation distribuée de ressources système dans un réseau informatique. Les algorithmes suggérés par Kurose et Simha se basent sur une recherche multi-dimensionnelle de l'équilibre à pas fixes et variables. Ygge et Akkermans [154] utilisent, quant à eux, un algorithme de recherche plus efficace basé sur la méthode de Newton-Raphson.

2.3.2 Modèles d'enchères combinatoires

L'analyse faite par Milgrom [102] de l'*enchère simultanée ascendante* (SAA) mise en pratique lors de la mise en vente, par la FCC, des premières licences d'utilisation des bandes de fréquences aux États-Unis montre que la non-divisibilité des objets d'une part, et la complémentarité entre ces derniers de l'autre, peuvent être source de non-existence d'équilibre de prix. Dans une enchère simultanée ascendante, les différents objets sont mis en vente simultanément par l'enchanteur, et les participants peuvent placer autant de mises qu'ils le désirent, mais sur des objets *individuels* uniquement. Milgrom présente l'exemple d'une vente de deux objets *A* et *B* à deux participants 1 et 2. Les valeurs des objets pour les deux participants sont spécifiées dans le tableau 2.1. On peut noter que les objets *A* et *B* sont complémentaires pour le participant 1. Il est facile de conclure à la non-existence d'un prix d'équilibre. En effet, une solution efficace sur le plan économique consiste à allouer les deux objets au participant 1. Pour empêcher le participant 2 de demander les deux objets, leur prix respectifs p_A et p_B doivent toutefois vérifier : $p_A > 7$, $p_B > 12$. D'autre part, à ces prix, le participant 1 ne pourra acquérir individuellement aucun des deux objets. De fait, une enchère simultanée ascendante finira par allouer les deux objets au participant 2.

	<i>A</i>	<i>B</i>	<i>AB</i>
Participant 1	5	10	20
Participant 2	7	12	19

TAB. 2.1 – Exemple de non-existence d'équilibre de prix

L'exemple précédent dévoile une limitation importante de tout mécanisme d'enchère multi-produits faisant appel à des mises sur des objets individuels : les mises des participants sur des objets individuels ne permettent pas d'exprimer directement les valeurs résultant de la synergie des objets. Dans la pratique, un participant intéressé à obtenir un ensemble donné d'objets continuera à placer des mises individuelles sur ces objets, jusqu'à ce que la valeur de réserve de l'ensemble soit atteinte. Ce faisant, il est possible que le participant n'obtienne qu'un sous-ensemble de l'ensemble cible, qui devra être payé plus que sa valeur intrinsèque. Ce phénomène, désigné dans la littérature sous le nom de problème du risque d'exposition (“exposure problem”), est à l'origine d'un comportement stratégique de la part des participants aux enchères simultanées ascendantes, et peut par conséquent être source d'inefficacité économique.

Les *enchères combinatoires* désignent communément ces mécanismes d'enchère où il est possible de placer des mises consolidées sur des *ensembles* d'objets. En permettant aux participants de spécifier, à même leurs mises, leurs préférences exactes pour des paquets d'objets plutôt que pour des objets individuels uniquement, les enchères combinatoires éliminent le risque d'exposition et atténuent en conséquence le comportement stratégique des participants.

Nous présentons dans ce qui suit un certain nombre de formulations d'enchères combinatoires. Ces formulations constituent des modèles “génériques” dans la mesure où elles ne sont pas reliées à une application en particulier. Nous accompagnons chaque formulation d'une revue des approches de résolution les plus importantes qui ont été suggérées dans la littérature.

Le problème classique de l'allocation combinatoire

Soit un vendeur unique (l'encanteur) mettant en vente un ensemble G de m objets non-divisibles différents, et un ensemble N de n acheteurs potentiels. De chaque objet

$i \in G$, le vendeur dispose d'une seule unité. Sans perte de généralité, nous supposons également que ce dernier n'a pas de prix de réserve sur les objets mis en vente.

Une *mise combinée* placée par un acheteur $j \in N$ s'exprime sous la forme d'un couple $(S, p_{j,S})$ où S est un sous-ensemble d'objets de G , et $p_{j,S}$ est le prix que l'acheteur j serait prêt à payer pour obtenir S . Quitte à ajouter des mises artificielles, nous supposons que les acheteurs misent sur tous les sous-ensembles d'objets de G .

Définissons les variables de décision suivantes :

$$\forall j \in N, \forall S \subseteq G, \quad x_{j,S} = \begin{cases} 1 & \text{si } S \text{ est alloué à l'acheteur } j, \\ 0 & \text{sinon.} \end{cases}$$

Le problème de la détermination des mises gagnantes (“the winner-determination problem”) correspond au modèle (CAP-WD) suivant :

$$\max \quad \sum_{j \in N} \sum_{S \subseteq G} p_{j,S} x_{j,S} \tag{2.1}$$

$$\text{s. à} \quad \sum_{S \subseteq G} x_{j,S} \leq 1, j \in N \tag{2.2}$$

$$\sum_{j \in N} \sum_{S \subseteq G} \delta_{i,S} x_{j,S} \leq 1, i \in G \tag{2.3}$$

$$x_{j,S} \in \{0, 1\}, S \subseteq G, j \in N \tag{2.4}$$

Dans le modèle (CAP-WD), l'objectif est de maximiser le revenu du vendeur. Les relations (2.2) établissent la contrainte qu'un acheteur donné ne peut obtenir plus d'un sous-ensemble (ce qui, précisons-le, ne compromet nullement la généralité du modèle puisque les acheteurs misent sur tous les sous-ensembles possibles d'objets). Quant aux contraintes (2.2), elles expriment le fait qu'un objet donné ne peut être alloué à plus d'un acheteur.

Certains auteurs font référence au modèle (CAP-WD) comme étant le problème de l'allocation combinatoire (“the Combinatorial Allocation Problem”), mais confondent parfois l'objectif de maximisation du revenu avec l'objectif, fondamentalement différent, qui consiste à optimiser l'efficacité économique de l'allocation des objets aux

acheteurs ou, en d'autres termes, à maximiser le bien-être social total des acheteurs. Afin de définir ce second objectif, dénotons par $v_{j,S}$ la préférence de l'acheteur $j \in N$ pour l'obtention du sous-ensemble $S \subseteq G$. Ici, $v_{j,S}$ est plutôt la valeur (*privée*) de S pour l'acheteur j , et non un prix misé par l'acheteur. Le problème de la détermination de l'allocation efficace des objets correspond alors au modèle (CAP-SE) suivant :

$$\max_{j \in N} \sum_{S \subseteq G} v_{j,S} x_{j,S} \quad (2.5)$$

$$\text{s. à } \sum_{S \subseteq G} x_{j,S} \leq 1, j \in N \quad (2.6)$$

$$\sum_{j \in N} \sum_{S \subseteq G} \delta_{i,S} x_{j,S} \leq 1, i \in G \quad (2.7)$$

$$x_{j,S} \in \{0, 1\}, S \subseteq G, j \in N \quad (2.8)$$

Dans le reste de cette section, nous nous limiterons aux travaux dédiés au modèle (CAP-WD), ainsi qu'à ses différentes variantes et extensions. Dans la mesure où l'analyse de l'efficacité économique fait intervenir les valeurs privées des participants, cette analyse est fortement liée à la capacité du mécanisme d'enchère à induire ces derniers à dévoiler leur préférences. Nous remettions donc la présentation de la littérature relative au modèle (CAP-SE) au sections dédiées aux enchères itératives et aux propriétés d'incitation des enchères combinatoires.

Le modèle (CAP-WD) correspond au problème d'empaquetage maximal ("Set Packing") de l'ensemble G . Pour ce problème, Sandholm [130] note qu'au pire cas, c'est-à-dire lorsque les mises sont formées à partir de toutes les combinaisons possibles des objets de G , le nombre d'allocations réalisables est dans $O(m^m)$ et $\Omega(m^{m/2})$. Rothkopf, Pekec et Harstad [125] proposent un algorithme basé sur la programmation dynamique pour déterminer une allocation optimale. L'idée de leur algorithme est simple et se base sur la remarque que pour tout sous-ensemble S d'objets, le revenu maximal que l'on peut tirer des objets de S provient d'une mise $(S, p_{j,S})$ sur le paquet S lui-même, ou bien c'est la somme des utilités maximales obtenues à partir de deux sous-ensembles S_1 et S_2 partitionnant strictement S . Ainsi, pour des sous-ensembles de cardinalité de plus en plus grande, l'algorithme détermine la partition

qui procure la plus grande utilité totale. Les auteurs montrent que l'algorithme trouve systématiquement une allocation efficace en un temps dans $O(3^m)$, ce qui est significativement meilleur qu'une énumération exhaustive des allocations réalisables mais demeure toujours trop coûteux.

La complexité du problème a incité certains auteurs, notamment Rothkopf, Pekeč et Harstad [125], à suggérer des *restrictions* sur les combinaisons d'objets sur lesquelles les agents peuvent placer des mises, espérant ainsi concevoir des algorithmes spécialisés permettant de déterminer une allocation optimale en un temps polynômial. Trois catégories de restrictions sont proposées, correspondant chacune à une structure particulière de l'ensemble des mises admissibles :

1. **Structures hiérarchiques.** L'ensemble des mises possède une structure arborescente, ce qui veut dire que quelles que soient les mises $(S^{(1)}, p_{j^{(1)}, S^{(1)}})$ et $(S^{(2)}, p_{j^{(2)}, S^{(2)}})$, une des deux conditions suivantes est vérifiée : (i) $S^{(1)}$ et $S^{(2)}$ sont disjoints, et (ii) $S^{(1)} \subseteq S^{(2)}$ ou $S^{(2)} \subseteq S^{(1)}$. Un algorithme itératif simple, construisant une solution en parcourant la structure arborescente à partir de ses feuilles, permet de déterminer une allocation optimale dans ce cas en un temps dans $O(m^2)$.
2. **Structures basées sur des mises de cardinalité limitées.** Ces structures correspondent aux mises $(S, p_{j, S})$ telles que $|S| \leq k$ où k est un entier tel que $1 \leq k \leq m$. Rothkopf, Pekeč et Harstad montrent qu'il est possible de déterminer une allocation optimale en un temps polynômial lorsque $k \leq 2$. Notamment, le cas $k = 2$ correspond à un problème de couplage pondéré maximal (“maximum weighted matching”) que l'on peut résoudre en $O(m^3)$. Malheureusement, les auteurs montrent également que quand $k > 2$, le problème correspondant est NP-complet.
3. **Structures ordonnées.** On suppose qu'un ordre total \preceq peut être défini sur l'ensemble G des objets, et que seules les mises portant sur des produits *successifs* selon cet ordre sont acceptées ; en d'autres termes, $\exists a, b, a \preceq b$, tels que $S = \{x \in G : a \preceq x \preceq b\}$. Remarquons que les structures hiérarchiques constituent un cas particulier de structures ordonnées. Il est alors possible de montrer

que l'algorithme dynamique proposé par Rothkopf, Pekeč et Harstad dans le cas général peut exploiter la structure en “intervalles” des mises pour trouver une allocation optimale en un temps dans $O(m^2)$. Cependant, les auteurs démontrent que dans le cas plus intéressant correspondant au produit de deux ordres linéaires (et donc à une localisation géométrique des mises dans un espace euclidien à deux dimensions), déterminer l'allocation optimale est encore un problème NP-complet.

Une caractérisation différente des structures ordonnées a été proposée par Nissan [105] qui démontre qu'en présence de telles structures, la relaxation continue du modèle (CAP-WD) possède la propriété d'intégralité, c'est-à-dire qu'il existe une allocation fractionnaire optimale qui est entière.

Sandholm [130] adopte dans son traitement du problème une approche différente de l'analyse structurelle précédente. Tout en n'imposant a priori aucune restriction sur les combinaisons d'objets sur lesquelles il est permis aux participants de placer des mises, il exploite essentiellement le fait que pour un assez grand nombre d'objets à vendre, les mises *réellement* reçues de la part des participants à l'enchère forment nécessairement un sous-ensemble de mises de taille très réduite relativement à l'ensemble de toutes les mises possibles. Il est donc tout à fait envisageable de procéder par une recherche énumérative de la meilleure solution, en autant que l'algorithme utilisé examine chaque allocation réalisable formée à partir des mises reçues exactement une fois. Plus précisément, l'algorithme proposé construit, à mesure que les mises sont reçues, un arbre où chacun des nœuds (sauf la racine), désigne un sous-ensemble d'objets correspondant à une mise combinée, et où tout chemin de la racine vers un autre nœud représente une allocation réalisable d'un sous-ensemble d'objets. Ainsi, si une mise $b = (S, p)$ est reçue, afin de garantir que toutes les allocations réalisables contenant la mise b seront représentées dans l'arbre, un nœud représentant b doit a priori être ajouté au bout de chaque chemin C tel que le sous-ensemble S' des objets faisant partie des mises déjà représentées dans C vérifie $S \cap S' = \emptyset$. Toutefois, une telle construction de l'arbre peut entraîner qu'une même allocation soit représentée par plusieurs chemins. Une procédure plus astucieuse permet en fait de garantir la

représentation de chaque allocation une et une seule fois dans l'arbre. Ainsi, Sandholm montre que si les objets de G sont indexés, et si, en plus de la condition de disjonction évoquée par la procédure ci-dessus, un nœud représentant une nouvelle mise est ajouté au bout d'un chemin \mathcal{C} uniquement si la mise inclut l'objet *de plus petit indice* parmi ceux qui ne figurent pas déjà dans \mathcal{C} , la représentation qui en résulte inclut toutes les allocations réalisables des objets aux mises reçues et est minimale. Un exemple d'une telle construction, tiré de l'article de Sandholm [130] est illustré par la figure 2.2 ci-dessous.

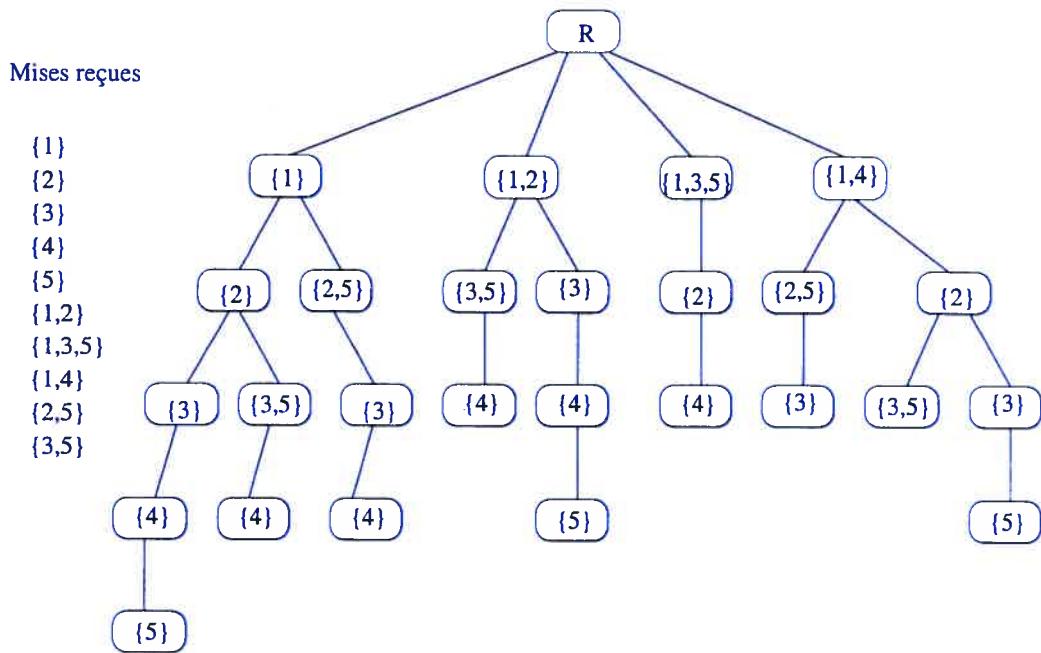


FIG. 2.2 – Une représentation arborescente des allocations réalisables des objets $1, \dots, 5$ aux mises b_1, \dots, b_{10} .

Une fois l'arbre construit, une recherche en profondeur (préférée dans ce cas car elle permet généralement d'obtenir des allocations réalisables plus rapidement qu'une recherche par largeur) peut être effectuée. Afin d'améliorer la recherche, Sandholm propose une phase de prétraitement incluant notamment une recherche préliminaire visant à éliminer des mises dominées par d'autres ensembles de mises, ainsi qu'une procédure de recherche par approfondissement itératif (“Iterative Deepening A*”).

Fujushima, Leyton-Brown et Shoham [55] proposent une procédure de recherche

proche de celle de Sandholm, en l'occurrence, l'algorithme CASS (Combinatorial Allocation Structured Search). L'algorithme CASS exploite deux idées fondamentales qui permettent de réduire de façon significative le nombre d'allocations non réalisables considérées lors d'une recherche. La première est l'organisation des mises en sous-ensembles de mises mutuellement incompatibles ("Bins"), qui, en raison d'un conflit sur un objet donné, ne peuvent être simultanément gagnantes dans une allocation. Formellement, à chaque objet $i \in G$, est associé l'ensemble de mises $D_i = \{(S, p) : i \in S, k \notin S, \forall k < i\}$. Ainsi, durant une passe d'une recherche en profondeur, (i) un seul élément de chaque D_i est considéré à la fois et (ii) si D_i est considéré, et $k > i$ est alloué dans la solution provisoire correspondante, alors D_k ne peut plus être considéré dans la même passe. La seconde idée, issue de la programmation dynamique, est l'exploitation des résultats intermédiaires ("caching") dans la réduction de l'espace de recherche. Ainsi, à une étape donnée de l'algorithme de recherche, soit r_C^* le revenu maximal (supposé connu) que l'on peut tirer d'un sous-ensemble d'objets C , et soit r_F le revenu tiré d'une allocation partielle $F \subseteq G$ telle que $G - F \subseteq C$. Si $r_C^* + r_F < r_{max}$, où r_{max} désigne le revenu de la meilleure allocation trouvée jusqu'à présent par l'algorithme de recherche, alors il est inutile de poursuivre la recherche dans la branche de l'arbre correspondant à F .

L'étude menée par Andersson, Tenhunen et Ygge [5] est intéressante à plusieurs égards. Ainsi, les auteurs passent en revue les techniques mises en œuvre dans les procédures de recherche précédentes et établissent ce qu'elles ont de commun avec les techniques classiques de résolution du problème de partitionnement d'ensemble (Balas et Padberg [12]), en particulier avec un algorithme classique dû à Garfinkel et Nemhauser [56]. D'autre part, Andersson, Tenhunen et Ygge comparent numériquement les performances des algorithmes de Sandholm et de Fujushima, Leyton-Brown et Shoham, avec celles d'un solveur commercial (CPLEX 6.5). Les résultats obtenus montrent qu'en moyenne, les performances de CPLEX se comparent avantageusement aux algorithmes de recherche. L'étude ne tient toutefois pas compte des tous derniers développements, notamment de l'algorithme CABOB proposé par Sandholm *et al.* [132], qui utilise des techniques telles que la décomposition de l'espace de recherche et l'utilisation d'heuristiques de branchement élaborées.

Bien que des algorithmes approximatifs pour le problème général d'empaquetage d'ensemble aient déjà été proposés dans la littérature (Chandra et Halldórson [27]), l'algorithme Casanova de Hoos et Boutilier [66] est, à notre connaissance, la seule application au problème de l'allocation combinatoire. Casanova est un algorithme de recherche locale de l'espace des allocations réalisables utilisant une notion simple de *voisinage* (une seule mise non satisfaite est choisie pour intégrer l'allocation courante). Le choix des mises est tributaire de leur *score*, qui désigne le rapport du revenu tiré de la mise au nombre d'objets de la mise. L'algorithme a été comparé à CASS [55] et les résultats obtenus sont fort encourageants.

Enchères combinatoires renversées

Intimement lié au problème classique de l'allocation combinatoire est le modèle d'enquête combinatoire *renversée* (“Reverse Combinatorial Auction”). Par opposition au modèle d'allocation, une enquête renversée fait intervenir un acheteur et plusieurs vendeurs. De ce fait, elle s'avère particulièrement utile dans la modélisation des situations de marché où l'objectif est la fourniture de biens et de services. Nous présentons dans ce qui suit une formulation de base du modèle d'enquête combinatoire renversée. Soit un acheteur désirant obtenir un ensemble G de biens différents. Ces biens sont fournis par plusieurs vendeurs sous forme de paquets d'objets (appelés communément *lots* dans ce cas). Une offre d'un vendeur s'exprime ainsi sous la forme d'une mise combinée $b = (S_b, p_b)$, où $S_b \subseteq G$ est un sous-ensemble de biens et p_b le prix demandé par le vendeur pour la fourniture de S_b . Soit B l'ensemble des mises soumises par les vendeurs. Les variables de décision du problème sont :

$$\forall b \in B, \quad x_b = \begin{cases} 1 & \text{si la mise } b \text{ est gagnante,} \\ 0 & \text{sinon.} \end{cases}$$

La détermination des mises gagnantes correspond ainsi au modèle (RevCA) suivant :

$$\min \quad \sum_{b \in B} p_b x_b \quad (2.9)$$

$$\text{s. à } \sum_{b \in B} \delta_{i,S_b} x_b \geq 1, i \in G \quad (2.10)$$

$$x_b \in \{0, 1\}, b \in B \quad (2.11)$$

L'objectif de l'acheteur est de sélectionner un sous-ensemble de lots qui recouvre l'ensemble G au plus bas coût. Ce problème est connu dans la littérature sous le nom de recouvrement minimal (“Set Covering”) de l'ensemble G . Sandholm *et al.* [133] notent l'importance de l'hypothèse implicite du modèle (RevCA) que l'acheteur tolère des unités supplémentaires de chaque bien (contraintes 2.10). En effet, quand cette hypothèse de libre disposition des biens n'est pas vérifiée par le contexte de l'enchère, la détermination des mises gagnantes devient plutôt un problème de partitionnement d'ensemble, qui est notoirement plus difficile à résoudre. Pour une version multi-unités du modèle (RevCA), Davenport et Kalagnanam [37] considèrent des contraintes additionnelles sur le nombre de fournisseurs gagnants et les quantités des différents biens obtenus de chaque fournisseur, et évaluent l'impact numérique de ces contraintes.

Allocation combinatoire multi-unités

Le modèle de l'enchère combinatoire multi-unités (MUCA) est une extension directe du modèle (CAP-WD) qui tient compte de la disponibilité, chez le vendeur, de plusieurs unités de chaque produit. Ainsi, soit M_i le nombre d'unités disponibles de l'objet $i, i \in G$. Une mise combinée placée par un acheteur peut être exprimée, dans ce cas, sous la forme $b = (\{a_{b,i}\}_{i \in G}, p_b)$, où $a_{b,i}$ désigne le nombre d'unités de l'objet i requises par l'acheteur dans la mise b . Soit B l'ensemble des mises combinées formulées par les acheteurs. Définissons les variables de décision suivantes :

$$\forall b \in B, \quad x_b = \begin{cases} 1 & \text{si la mise } b \text{ est gagnante,} \\ 0 & \text{sinon.} \end{cases}$$

Le problème de la détermination des mises gagnantes dans le cas d'une enchère combinatoire multi-unités correspond au modèle (MUCA) suivant :

$$\max \quad \sum_{b \in B} p_b x_b \quad (2.12)$$

$$\text{s. à} \quad \sum_{b \in B} a_{b,i} x_b \leq M_i, i \in G \quad (2.13)$$

$$x_b \in \{0, 1\}, b \in B \quad (2.14)$$

Le modèle (MUCA) correspond à un problème de sac à dos multi-dimensionnel. Pour cette catégorie de problèmes, l'ouvrage de référence de Martello et Toth [94] est une revue détaillée de l'essentiel des méthodes de résolution exactes et heuristiques. Parmi les contributions récentes, l'algorithme CAMUS (“Combinatorial Auction Multi-Unit Search”) de Leyton-Brown, Shoham et Tennenholz [89] utilise, grossièrement, les mêmes techniques mises à profit dans l'algorithme CASS [55], généralisées aux enchères multi-unités. Dans Gonen et Lehmann [59] et Lehmann et Gonen [84], les auteurs proposent l'utilisation conjointe, au sein d'une procédure de Branch-and-Bound, de bornes supérieures et inférieures sur la valeur de l'allocation optimale. Ces dernières sont dérivées, respectivement, de relaxations simples du problème de sac-à-dos multi-dimensionnel (dont la relaxation linéaire) et d'algorithmes d'allocation vorace basés sur différentes heuristiques de choix de mises. Mansini et Grazia Sprenzana [92] montrent qu'une borne supérieure sur le *nombre* de mises gagnantes dans une allocation optimale peut être calculée en résolvant $|G|$ problèmes auxiliaires de sac à dos. Ce résultat est d'une importance fondamentale car il permet d'adoindre des inégalités valides à la formulation (MUCA) et d'améliorer la qualité des bornes supérieures. Les résultats des tests préliminaires menés par Mansini et Grazia Sprenzana montrent un gain de performance appréciable par rapport à CPLEX 7.0.

Enchères combinatoires doubles

Les enchères combinatoires *doubles* (“Combinatorial Exchanges”) font référence à des mécanismes de marché basés sur des mises combinées d'achat et de vente formulées

par plusieurs acheteurs et vendeurs potentiels. Formellement, une mise combinée b formulée par un participant à une enchère combinatoire double s'exprime sous la forme $b = (\{q_{b,i}\}_{i \in G}, p_b)$ où :

- $q_{b,i} > 0$ signifie que le participant désire vendre $q_{b,i}$ unités du produit $i, i \in G$ dans la mise b et $q_{b,i} < 0$ qu'il veut acquérir $-q_{b,i}$ unités du produit i ;
- $p_b > 0$ (resp. $p_b < 0$) est un prix que le participant est prêt à payer (resp. recevoir) pour l'exécution de la mise b .

Soit B l'ensemble des mises combinées placées par les participants. Considérons les variables de décision suivantes :

$$\forall b \in B, \quad x_b = \begin{cases} 1 & \text{si la mise } b \text{ est gagnante,} \\ 0 & \text{sinon.} \end{cases}$$

Le problème de la détermination des mises gagnantes correspond au modèle (CE) :

$$\max_{b \in B} \sum_{b \in B} p_b x_b \tag{2.15}$$

$$\text{s. à } \sum_{b \in B} q_{b,i} x_b \leq 0, i \in G \tag{2.16}$$

$$x_b \in \{0, 1\}, b \in B \tag{2.17}$$

Dans le modèle (CE), les contraintes (2.16) indiquent que l'offre est suffisante pour satisfaire la demande. Notons que ces contraintes doivent être remplacées par des contraintes d'égalité dans le cas où l'enchanteur ne peut disposer librement de volumes de produits offerts en trop. L'objectif est de maximiser le surplus économique du marché.

Étant donné que les enchères combinatoires multi-unités constituent un cas particulier d'enchères combinatoires doubles, le problème de la détermination des mises gagnantes dans ces dernières est, a fortiori, NP-complet. Sandholm et Suri [129] présentent l'algorithme BOB, dans lequel bon nombre de techniques de recherche de solutions exactes proposées pour le problème classique de l'allocation combinatoire sont adaptées aux enchères doubles. Kothari, Sandholm et Suri [77] étudient la complexité d'enchères combinatoires doubles dans le cadre plus général des enchères *hy-*

brides, dans lesquelles certaines mises combinées peuvent être *partiellement* exécutées. Les auteurs établissent que, dans le cas où le nombre de mises dont on tolère l'exécution partielle ne dépasse pas $|G|$, le problème de la détermination des mises gagnantes peut être réduit à un problème linéaire.

Kalagnanam, Davenport et Lee [69] considèrent un modèle d'enchère double standard adapté aux besoins des industries de transformation de produits semi-finis (papiers, métaux, etc.). Le modèle est quelque peu différent de (CE) dans le sens que les mises portent sur l'offre ou la demande d'un seul produit. Cependant, les produits échangés ont différents niveaux de qualité, et il est possible de satisfaire la demande pour un produit de qualité Q par l'exécution d'une ou de plusieurs mises de vente de produits similaires de qualités supérieures à Q . La substitution entre produits donne ainsi lieu à des contraintes additionnelles établissant des associations d'exécution possible entre mises d'achat et mises de vente. Les auteurs notent que, quand plusieurs mises de vente peuvent être consolidées pour satisfaire une mise d'achat, la détermination des mises gagnantes revient à résoudre un problème de flot maximal. Toutefois, dès que la consolidation n'est plus possible, déterminer une allocation optimale requiert la résolution d'un problème d'affectation généralisée, qui est NP-complet.

2.3.3 Applications choisies

Depuis quelques années, les enchères combinatoires connaissent un franc succès comme mécanismes d'allocation efficace de ressources. À défaut de pouvoir dénombrer tous les domaines où une forme ou une autre d'enchère combinatoire a pu être mise en pratique, nous présentons une sélection d'applications dans une variété de contextes.

Droits d'utilisation de l'infrastructure ferroviaire

Le mécanisme BICAP (BInary Conflict Ascending Price) a été développé par Bremer et Plott [24] pour l'allocation des droits d'utilisation de l'infrastructure ferroviaire suédoise. Le mécanisme BICAP est essentiellement une enchère combinatoire progressive où des mises croissantes sur différents trains sont placées en temps continu par des agents. Chaque fois qu'une nouvelle mise est reçue par l'encanteur, ce dernier

détermine une allocation réalisable optimale du système de voies ferrées (ensemble de trains ne générant pas de conflit d'utilisation de la voie et maximisant le revenu de l'enchanteur). L'enchère se termine lorsqu'aucune nouvelle mise n'est reçue pendant une période de temps prédéterminée.

Une enchère combinatoire pour le choix de cours

Graves, Schrage et Sankaran [60] présentent un système automatisé pour le choix et l'enregistrement aux cours destiné aux étudiants de la Graduate School of Business de l'Université de Chicago. Le système se présente sous la forme d'un mécanisme d'enchère combinatoire où les étudiants, disposant d'un capital de "points" fictifs, misent sur des *programmes* formés de paquets de cours. Dans une première phase du mécanisme, l'enchanteur détermine une allocation optimale de la capacité d'enseignement disponible (salles d'enseignement et corps enseignant) compte tenu des mises reçues. Une seconde phase permet aux étudiants, à travers plusieurs rondes de mise, d'ajuster cette allocation en ajoutant, supprimant, ou échangeant des cours individuels dans leur programme. Bien que les auteurs n'aient pas effectué de mesures d'efficacité, le système est utilisé avec succès par l'administration de la faculté depuis 1981 et semble satisfaire les besoins des étudiants et des professeurs.

Problèmes de tournées de véhicules

Davis et Smith [38] proposent un protocole simple de négociation et de coordination d'actions entre les agents d'un système distribué. Le protocole en question, dénommé "Contract Net Protocol" ou CNP, identifie les tâches du système (par exemple des ordres à exécuter dans un problème de tournées de véhicules) et partage les agents en deux catégories : les coordinateurs, responsables de la supervision de l'exécution d'une tâche, et les contractants, exécutant la tâche sur le terrain. Un processus de négociation, qui peut être vu comme une série d'enchères simples à enveloppe fermée, se déroule entre coordinateurs et contractants comme suit. Un message annonçant une tâche est diffusé par un coordinateur vers les contractants qui sont sous sa tutelle. Ceux-ci font l'évaluation de la tâche en fonction des informations qu'ils possèdent et des ressources qui sont à leur disposition, placent des mises

pour l'exécution totale ou partielle de la tâche, et les envoient au coordinateur. Une fois toutes les mises reçues, le coordinateur désigne par des messages de décernement les contractants qui exécuteront la tâche, scellant ainsi le contrat la concernant. Ce protocole a été appliqué par Davis et Smith à un système distribué de détection de véhicules. Sandholm [127] étend le protocole en permettant des mises sur des paquets d'ordres. Fischer, Müller et Pischel [54] développent un système multi-agents pour un problème de tournées de véhicules avec fenêtres horaires, qui inclut en plus d'un mécanisme basé sur le protocole CNP, une phase de négociation entre coordinateurs et un mécanisme supplémentaire d'enchère double entre contractants ("Simulated Trading", voir Bachem, Hochstättler et Malich [11]) permettant l'amélioration de la qualité des allocations d'ordres faites par le protocole CNP et la réallocation dynamique des ordres en cas de l'impossibilité de les exécuter.

Fourniture de matériaux d'emballage

En février 2001, la compagnie Volvo a organisé une enchère combinatoire pour la fourniture de produits d'emballage. L'enchère a porté sur plus de 600 produits différents répartis sur 14 segments. Une phase préliminaire de présélection des fournisseurs et de négociation des prix de réserve sur les produits individuels précède la phase principale de mise. Les résultats obtenus ont pu confirmer les principaux avantages des enchères combinatoires. À titre d'exemple, notons que la compétition accrue résultant du caractère "tout-ourien" des mises combinées a permis à Volvo de réaliser une économie moyenne de 4% par rapport au processus d'approvisionnement traditionnel de Volvo.

Enchère combinatoire de temps d'antenne pour "spots" publicitaires

Jones et Koehler [68] présentent un mécanisme d'enchère combinatoire destiné à l'allocation de temps d'antenne pour la diffusion de spots publicitaires. L'originalité du mécanisme réside dans le fait que les annonceurs peuvent spécifier, dans un *langage de mise* élaboré, des besoins spécifiques complexes quant aux spots obtenus. Ainsi, l'annonceur peut exiger des spots dans des émissions précises ou demander à ce qu'ils apparaissent dans des émissions faisant partie d'un ensemble cible. En outre, les

annonceurs précisent très souvent un niveau minimal de “couverture”, exprimant le niveau de l’auditoire (classifié en plusieurs catégories selon la tranche d’âge, le sexe, les centres d’intérêt, etc.) rejoint par leurs spots publicitaires. Enfin, il est également possible à un annonceur de demander l’agrégation de plusieurs spots afin de permettre la mise en place de “campagnes” publicitaires. Dans la formulation du problème d’allocation des spots aux mises des annonceurs, l’encanteur doit tenir compte, en plus des besoins de ces derniers, de contraintes additionnelles reliées à la réglementation en cours, de la nature des émissions, et des objectifs commerciaux propres de la chaîne.

Allocation de contrats de restauration pour le système scolaire chilien

Depuis 1999, le gouvernement chilien organise, sur une base annuelle, des enchères pour l’allocation de contrats d’approvisionnement en denrées alimentaires du système scolaire national. Le mécanisme correspondant, tel que décrit dans Epstein *et al.* [46], est basé sur une enchère combinatoire à enveloppe fermée. Dans une phase préliminaire de qualification des participants, l’organisme gouvernemental responsable de l’allocation des contrats organise un appel d’offres initial puis pré-selectionne et classe les contractants potentiels sur la base de critères techniques, légaux et financiers. Ensuite, l’autorité publie les détails techniques concernant la structure des repas, les divers standards de qualité requis, etc. Les contractants sélectionnés peuvent soumettre des mises combinées portant sur l’approvisionnement couvrant plusieurs régions du pays. Un ensemble de mises minimisant le *coût* de l’allocation, obtenu en agrégant les prix proposés par les participants et des critères de qualité associés à chaque contractant, est alors déterminé par l’organisme. Grâce principalement aux économies d’échelle réalisées par les contractants en desservant plusieurs régions de manière consolidée, ainsi qu’au niveau accru de compétitivité, le mécanisme d’enchère a permis au gouvernement chilien de réduire le coût total de l’approvisionnement de 22% par rapport au processus informel de négociation des contrats jusque-là en place.

2.4 Mécanismes progressifs d'enchère combinatoire

Tous les modèles d'enchère combinatoire considérés jusqu'à présent font implicitement référence à un mécanisme de marché à "enveloppe fermée", avec des déclarations *complètes* et *définitives*, par les participants, de leurs préférences à l'enchanteur. C'est le cas notamment du modèle (CAP-SE) où l'enchanteur doit a priori disposer des valeurs $v_{j,S}$ de tous les paquets d'objets $S \subseteq G$ pour tous les participants $j \in N$. Toutefois, plusieurs considérations font en sorte que l'hypothèse selon laquelle les participants révèlent de manière intégrale leurs préférences est très peu réaliste dans la pratique. Tout d'abord, les valeurs $v_{j,S}, j \in N, S \subseteq G$, font généralement partie de l'information confidentielle des participants, et il est tout à fait naturel que ces derniers soient réticents à divulger ces valeurs sans incitatif approprié. D'autre part, un mécanisme à révélation complète suppose que les participants évaluent les préférences, mais ne tient nullement compte de la *complexité* d'une telle évaluation, qui peut être considérable étant donnés le nombre exponentiel de paquets et la possibilité que le problème de décision correspondant à l'évaluation de chaque paquet puisse être intrinsèquement difficile.

Les mécanismes progressifs d'enchère combinatoire répondent ainsi au souci de permettre aux participants de révéler, au fil de l'avancement de l'enchère, suffisamment d'information pertinente sur leurs préférences pour permettre à l'enchanteur de déterminer une allocation optimale des objets. Au fur et à mesure que l'enchère progresse, les participants reçoivent de la part de l'enchanteur un "signal" portant généralement sur une allocation et des prix *temporaires*, et dont le but est de fournir aux participants des indications sur l'état de l'enchère et de guider leur implication future dans cette dernière.

Les approches de tâtonnement basées sur l'équilibre walrasien des prix présentées à la section 2.3.1 peuvent être considérées comme des mécanismes progressifs d'enchère combinatoire dans le cadre d'économies de production/consommation de biens divisibles. Dans le cas d'objets indivisibles, l'essentiel de l'effort de recherche a été consacré au problème classique de l'allocation combinatoire. Sur le plan méthodologique, les mécanismes proposés peuvent être classifiés selon deux grandes catégories d'approches.

2.4.1 Mécanismes basés sur des approches primales/duales

Les approches primales/duales constituent un cadre méthodologique intéressant pour la conception de mécanismes d'enchère à révélation progressive de l'information. Considérons le problème classique de l'allocation combinatoire (CAP-SE), et supposons que ce problème puisse être exprimé sous forme d'un programme linéaire (P). Le principe d'un algorithme primal/dual typique pour déterminer une solution optimale de (P) est le suivant. Une solution duale réalisable y de (P), qui peut être assimilée à un vecteur de prix, est maintenue par l'algorithme. Dans le but de déterminer une solution primaire réalisable (allocation admissible des objets) x , telle que le couple primal/dual (x, y) vérifie la condition d'optimalité des écarts complémentaires, un *problème primal relaxé* (RP) est formulé. Si la solution optimale \hat{x} de ce dernier est réalisable pour (P), alors \hat{x} est également une solution optimale de (P). Sinon, la solution du dual de (RP) fournit suffisamment d'informations pour mettre à jour le vecteur de prix y de telle façon à ce que la réalisabilité duale soit maintenue, tout en progressant vers une solution réalisable de (P).

Le problème d'affectation a donné lieu aux premières enchères itératives d'objets multiples hétérogènes basées sur des approches primales/duales d'ajustement des allocations et des prix. Le problème d'affectation peut être vu comme étant un cas particulier de (CAP-SE) où chaque participant désire obtenir, au plus, un seul objet ($v_{j,S} = 0, \forall S \subseteq G : |S| \neq 1$). Ainsi, notons $v_{ji} = v_{j,\{i\}}, \forall i \in G$ et supposons, sans perte de généralité, que $|G| \geq |N|$. Soit $x_{ji} = 1$ si l'objet $i \in G$ est affecté au participant $j \in N$, $x_{ji} = 0$ autrement. Une affectation efficace des objets aux participants est une solution de :

$$\max \sum_{j \in N} \sum_{i \in G} v_{ji} x_{ji} \quad (2.18)$$

$$\text{s. à } \sum_{i \in G} x_{ji} \leq 1, j \in N \quad (2.19)$$

$$\sum_{j \in N} x_{ji} \leq 1, i \in G \quad (2.20)$$

$$x_{ji} \geq 0, i \in G, j \in N \quad (2.21)$$

Étant totalement unimodulaire, le problème d'affectation possède la propriété d'intégralité. Son dual s'écrit :

$$\min \sum_{i \in G} p_i + \sum_{j \in N} s_j \quad (2.22)$$

$$\text{s. à } p_i + s_j \geq v_{ji}, i \in G, j \in N \quad (2.23)$$

$$p_i, s_j \geq 0, i \in G, j \in N \quad (2.24)$$

Si on interprète $\{p_i\}_{i \in G}$ comme un vecteur de prix, il n'est pas difficile de voir que la condition des écarts complémentaires correspond à un comportement stratégique particulier des participants. En effet, on aurait $s_j = v_{ji} - p_i$, si $x_{ji} = 1, \forall j \in N$. En tenant compte également de la contrainte de réalisabilité duale (2.23), ceci indique qu'aux prix $p_i, i \in G$, un participant $j \in N$ ne peut obtenir qu'un objet i qui maximise son surplus s_j (différence entre sa préférence pour l'obtention de l'objet et le prix payé pour l'obtenir). Le problème primal relaxé formulé à une itération donnée d'une approche primale/duale consiste donc à allouer le maximum d'objets à des participants qui réagissent de façon *myope* aux prix courants en plaçant des mises uniquement sur les objets maximisant leur surplus.

Ce principe a été mis en œuvre par Bertsekas [15] dans son algorithme AUCTION pour le problème d'affectation, et a été généralisé par la suite (Bertsekas [16], Bertsekas [17]) à d'autres problèmes de flot dans les réseaux, notamment les problèmes du plus court chemin et du flot minimum. Des considérations algorithmiques (développer des approches de résolution distribuées, donc potentiellement parallélisables, pour les problèmes de flots dans les réseaux) sont toutefois les principales motivations de l'auteur. La détermination de prix dans le cadre du problème d'affectation a également été l'objet d'une étude de Leonard [88], qui démontre qu'à travers des variables duales optimales maximisant $\sum_{i \in G} p_i$, l'enchanteur obtient des prix tels qu'aucun participant

n'a d'intérêt à miser des prix autres que ses vraies préférences pour les objets. Demange, Gale et Sotomayor [40] proposent un mécanisme itératif d'enchère pour le problème d'affectation, auquel Bikhchandani *et al.* [18] donnent une interprétation primale/duale. L'importance de ce mécanisme provient du fait qu'il possède la propriété d'inciter les participants à une déclaration véridique de leurs préférences. Plus précisément, Demange, Gale et Sotomayor définissent la notion d'*ensemble à “excès de demande”* (“overdemanded sets”) pour lesquels le nombre d'objets disponibles est strictement inférieur au nombre de participants misant sur des objets de l'ensemble. Les auteurs montrent qu'en ajustant, à une itération donnée du processus primal/dual, les prix des objets faisant partie d'ensembles à “excès de demande” de cardinalité minimale, le processus est induit à déterminer des paiements équivalents à ceux d'un mécanisme “second-prix” de Vickrey (ou, plus précisément, du mécanisme de Vickrey-Clarke-Groves ; voir la section 2.6).

En ce qui concerne le problème plus général de l'allocation combinatoire, la formulation (CAP-SE) ne possède malheureusement pas la propriété d'intégralité. En effet, si l'on considère l'exemple du tableau 2.2, l'allocation optimale procure un bien-être social de 1 et correspond, par exemple, à allouer les trois objets au participant 1. D'autre part, on peut noter que l'allocation fractionnaire ($x_{1,\{A,B\}} = 0.5$, $x_{2,\{B,C\}} = 0.5$, et $x_{3,\{A,C\}} = 0.5$) est réalisable et procure un bien-être supérieur de 1.5. En fait, Bikhchandani et Mamer [19] montrent que l'absence d'une marge duale entre la formulation (CAP-SE) et sa relaxation linéaire est une condition nécessaire et suffisante pour l'existence d'un *équilibre walrasien* : un vecteur de prix $\{p_i\}_{i \in G}$ supportant une allocation efficace x^* , c'est-à-dire tel que chaque participant reçoit dans x^* un paquet d'objets maximisant son surplus étant donné les prix. D'autre part, Gul et Stacchetti [62] introduisent une propriété de “substitutabilité” entre produits (“Gross Substitutes”), qui exprime le fait que, étant donnés deux vecteurs de prix $\{p_i\}_{i \in G}$ et $\{p'_i\}_{i \in G}$ tels que $p \leq p'$, si le participant demande le paquet d'objets S aux prix p (S maximise le surplus du participant à ces prix), il continue de demander les objets i' de S qui n'ont pas changé de prix (tels que $p_{i'} = p'_{i'}$). Les auteurs montrent que cette propriété de substitutabilité est une condition suffisante pour l'existence

d'un équilibre walrasien.

	{A}	{B}	{C}	{A, B}	{B, C}	{A, C}	{A, B, C}
1	0	0	0	1	0	0	1
2	0	0	0	0	1	0	1
3	0	0	0	0	0	1	1

TAB. 2.2 – Exemple d'allocation combinatoire : une allocation fractionnaire est strictement meilleure qu'une allocation entière

En ajoutant des variables auxiliaires au modèle (CAP-SE), Bikhchandani et Ostroy [20] obtiennent deux autres formulations, plus fortes que (CAP-SE), du problème de l'allocation combinatoire. Soit \mathcal{G} l'ensemble des partitions de G . Définissons les variables de décision suivantes : soit $z_\alpha = 1$ si la partition α est retenue pour l'allocation aux participants, et $z_\alpha = 0$ sinon, $\forall \alpha \in \mathcal{G}$. La formulation (CAP-SE-2) s'écrit alors :

$$\max \sum_{j \in N} \sum_{S \subseteq G} v_{j,S} x_{j,S} \quad (2.25)$$

$$\text{s. à } \sum_{S \subseteq G} x_{j,S} \leq 1, j \in N \quad (2.26)$$

$$\sum_{\alpha \in \mathcal{G}} z_\alpha \leq 1 \quad (2.27)$$

$$\sum_{j \in N} x_{j,S} \leq \sum_{\alpha \in \mathcal{G}: S \in \alpha} z_\alpha, S \subseteq G \quad (2.28)$$

$$x_{j,S}, z_\alpha \in \{0, 1\}, S \subseteq G, j \in N, \alpha \in \mathcal{G} \quad (2.29)$$

La contrainte (2.27) indique qu'une seule partition des objets doit être retenue pour l'allocation, tandis que les contraintes (2.28) font en sorte que seuls les sous-ensembles de la partition retenue peuvent être alloués aux participants.

Une alternative pour déterminer une allocation réalisable des objets est donc de choisir une partition de G , puis d'allouer chaque sous-ensemble d'objets de la partition à un participant. Dénotons par $\beta = (\alpha, A)$ où α est une partition de G et A une allocation donnée des sous-ensembles de α aux participants. Ainsi, $(S, j) \in \beta$ signifie que sous β , le sous-ensemble d'objets S est alloué au participant j . Soit Γ l'ensemble

de toutes les valeurs possibles de β . Définissons la variable de décision $z_\beta = 1$ si “l’allocation” β est retenue, et $z_\beta = 0$ autrement. La formulation (CAP-SE-3) correspond alors au problème :

$$\max \sum_{j \in N} \sum_{S \subseteq G} v_{j,S} x_{j,S} \quad (2.30)$$

$$\text{s. à } \sum_{S \subseteq G} x_{j,S} \leq 1, j \in N \quad (2.31)$$

$$\sum_{\beta \in \Gamma} z_\beta \leq 1 \quad (2.32)$$

$$x_{j,S} \leq \sum_{\beta \in \Gamma: (S,j) \in \beta} z_\beta, S \subseteq G, j \in N \quad (2.33)$$

$$x_{j,S}, z_\beta \in \{0, 1\}, S \subseteq G, j \in N, \beta \in \Gamma \quad (2.34)$$

L’intérêt de la formulation (CAP-SE-3) est qu’elle possède la propriété d’intégralité (Bikhchandani et Ostroy [20]) et se prête par conséquent bien à une analyse primaire/duale. D’autre part, les variables duales de (CAP-SE-3) ont une interprétation intéressante. En particulier, les variables duales $p_{S,j}, S \subseteq G, j \in N$ associées aux contraintes (2.33) correspondent à *des prix de paquet non-anonymes* : lors d’un processus primal/dual, une variable duale réalisable $p_{S,j}$ désigne un prix du paquet S , personnalisé au participant j , que ce dernier doit “battre” d’un certain incrément pour espérer obtenir S . Quant aux variables duales $s_j, j \in N$, et r , correspondant aux contraintes (2.31) et (2.32), elles désignent respectivement les surplus réalisés par les participants et le revenu dégagé par l’enchanteur, étant donnés les prix $p_{S,j}, S \subseteq G, j \in N$.

Parkes [111] et Parkes et Ungar [115] présentent le mécanisme iBundle qui, à notre connaissance, est le premier mécanisme itératif d’enchère pour le problème général de l’allocation combinatoire basé sur un processus primal/dual d’ajustement des allocations et des prix. iBundle est un mécanisme générique qui se décline en trois variantes. La première, iBundle(2), fait appel à la formulation (CAP-SE-2) et utilise des prix de paquets anonymes. Toutefois, elle requiert que les mises placées

par les participants vérifient une condition dite de “*sûreté*” (“safety”) : l’ensemble de tous les paquets d’objets assurant à un participant le surplus maximal compte tenu des prix courants ne contient aucune paire de paquets disjoints. La seconde variante, iBundle(3) s’affranchit de la condition de “*sûreté*”, au prix de devoir faire appel à la formulation (CAP-SE-3). Enfin, iBundle(d) met à jour les prix de façon dynamique, utilisant par défaut des prix anonymes et ne recourant à des prix personnalisés que quand celà devient nécessaire (quand la condition de “*sûreté*” n’est plus vérifiée par les mises des participants). Dans les trois cas, l’hypothèse que les participants s’en tiennent à la stratégie “*myope*”, consistant à miser en tenant uniquement compte des prix courants, demeure toutefois fondamentale pour que les enchères correspondantes puissent aboutir, en fin de compte, à des allocations efficaces des objets.

Wurman et Wellman [148] proposent une alternative pour la détermination d’un équilibre de prix qui se base sur le principe suivant. Soit une allocation efficace x^* des objets, solution de (CAP-SE). Si l’on considère uniquement les paquets d’objets *alloués* dans x^* , le problème d’allocation peut être réduit à un problème d’affectation de ces paquets aux participants (en faisant appel au *besoin* à des paquets d’objets fictifs). Des prix d’équilibre, portant sur les paquets d’objets, peuvent alors être obtenus à partir du dual du problème d’affectation. Les solutions optimales de ce dual n’étant en général pas uniques, Wurman et Wellman formulent plutôt deux problèmes auxiliaires du dual qui fournissent les prix optimaux minimisant respectivement le revenu de l’encanteur et le surplus total des participants. Par la suite, les prix des paquets non alloués dans x^* sont trivialement calculés de façon à ce que l’équilibre de prix soit préservé (plus précisément, de manière à ce qu’aucun participant ne soit intéressé à acquérir ces paquets). Étant anonymes, les prix déterminés de cette manière n’assurent toutefois pas l’encanteur de réaliser le revenu maximal pour l’allocation x^* . Cette approche est à la base du mécanisme itératif AkBA, qui généralise en quelque sorte l’enchère simultanée ascendante aux enchères combinatoires : à une itération donnée du mécanisme, des prix d’équilibre sont déterminés selon le principe précédent sur la base des mises gagnantes (provisoires) et des nouvelles mises reçues durant l’itération courante, et doivent être améliorés d’un certain incrément par les nouvelles mises admissibles à l’itération suivante ; le mécanisme s’arrête lorsqu’aucune

nouvelle mise n'est reçue à une itération donnée.

2.4.2 Mécanismes expérimentaux

Conjointement aux enchères combinatoires progressives mettant en œuvre un processus d'ajustement primal/dual, un certain nombre de mécanismes progressifs ont été proposés dans un cadre plus expérimental. Nous présentons brièvement trois représentants importants de cette classe de mécanismes.

La procédure AUSM

Un des tout premiers mécanismes d'enchère combinatoire progressive est la procédure AUSM (“Adaptive User Selection Mechanism”) développée par Banks, Ledyard et Porter [13]. AUSM est une enchère en temps continu où des mises combinées peuvent être placées de manière asynchrone par les participants. L'enchère fonctionne selon un principe très simple. Dès qu'une nouvelle mise combinée est reçue, l'enchanteur consulte l'ensemble des mises faisant partie de l'allocation courante et qui entrent en conflit, sur un ou plusieurs objets, avec la nouvelle mise. Si la valeur cumulée de cet ensemble de mises (total des prix associés aux mises combinées) dépasse la nouvelle mise, celle-ci est rejetée ; sinon, c'est la nouvelle mise qui déplace l'ensemble de mise dans l'allocation courante. Bien entendu, cette procédure d'allocation “vorace” ne donne pas, en général, une allocation économiquement efficace. Afin de renforcer le mécanisme, les auteurs proposent d'utiliser une file d'attente qui accueille les mises rejetées par l'enchanteur. En tout temps, les participants peuvent ainsi constituer des *coalitions* de mises rejetées et déplacer des mises gagnantes dans l'allocation courante. L'exemple de la figure 2.3 permet d'illustrer le mécanisme AUSM. L'analyse expérimentale du mécanisme par Bykowsky, Cull et Ledyard [25] montre des gains intéressants en termes d'efficacité économique par rapport à une enchère simultanée ascendante.

Le mécanisme RAD

DeMartini *et al.* [41] présentent le mécanisme RAD comme étant une tentative de

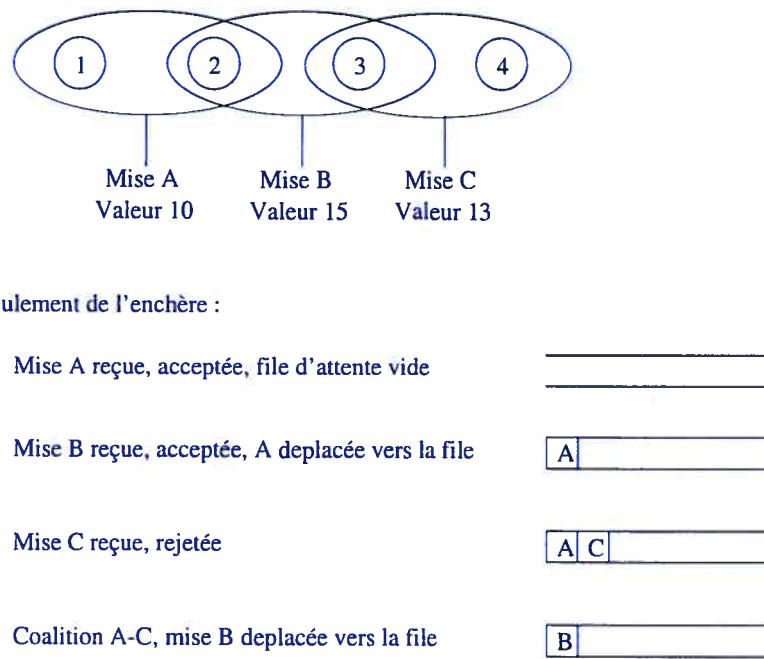


FIG. 2.3 – Illustration du mécanisme AUSM.

combiner les meilleurs aspects de l'enchère simultanée ascendante et du mécanisme AUSM. Plus précisément, RAD est une enchère multi-rondes qui permet aux participants de soumettre des mises combinées sur des paquets d'objets. Durant une ronde de l'enchère, une allocation provisoire des objets aux mises combinées est déterminée. À l'instar d'une enchère simultanée ascendante, le mécanisme RAD signale des prix sur les objets individuels et impose aux nouvelles mises reçues une amélioration relative minimale - calculée sur la base des prix annoncés - par rapport aux mises gagnantes provisoires. Le caractère combinatoire de l'enchère excluant la possibilité d'utiliser des variables duales du problème d'allocation comme prix d'équilibre, le mécanisme RAD calcule des prix "approchés" qui minimisent la violation des écarts complémentaires. Par ailleurs, le mécanisme RAD implante une règle d'*activité simple* qui conditionne le nombre de nouvelles mises qu'un participant peut soumettre dans une ronde donnée et oblige ce dernier à rester actif tout au long de l'enchère.

Le mécanisme PAUSE

Le mécanisme PAUSE de Kelly et Steinberg [72] est une enchère combinatoire progressive adaptée aux besoins de l'industrie des télécommunications, qui innove en

comprenant deux phases distinctes. Dans la première, que l'on peut assimiler à une phase d'“apprentissage” du marché, les opérateurs misent sur des licences individuelles comme dans une enchère simultanée ascendante classique. Un processus combinatoire similaire au mécanisme AUSM suit dans la seconde phase. Dans les deux phases du mécanisme, l'enchère adopte des règles d'activité relativement complexes portant sur la population totale couverte par les mises des opérateurs.

2.5 Expression de la préférence

En soumettant des mises sous forme de prix associés à des paquets d'objets, les participants à une enchère combinatoire peuvent en principe refléter de façon précise leurs fonctions de préférence pour les objets, quelles que soient les structures de ces dernières. Toutefois, miser *explicitement* sur les différents sous-ensembles d'objets intéressants peut s'avérer très complexe en pratique. Ainsi, pour un grand nombre de fonctions de préférence, le nombre de mises nécessaires est exponentiel par rapport au nombre d'objets sur le marché. Ceci est évidemment le cas de fonctions de préférences générales, sans structure particulière : si m objets sont mis en vente, un acheteur potentiel pourrait placer jusqu'à $2^m - 1$ mises sur les différents paquets d'objets. Par contre, une énumération explicite peut être évitée quand la fonction de préférence d'un participant possède une structure *logique* particulière. L'exemple suivant permet d'illustrer ce dernier point.

Exemple : Mises combinées sur des ordres de transport

Soit une bourse de fret dans laquelle des transporteurs misent sur des ordres de transport de marchandises. Dans le mode de transport dit “à charges pleines” (“Truckload”), les transporteurs doivent allouer un véhicule à chaque charge et transporter cette dernière de son point de départ jusqu'à sa destination. L'aspect combinatoire du marché vient du fait que les transporteurs peuvent placer des mises consolidées sur des paquets d'ordres qu'il se proposent d'exécuter comme un tout.

Considérons un transporteur donné disposant d'un véhicule qu'il est possible d'allouer à un, et un seul, parmi les quatre ordres de transport A , B , C et D . Les prix demandés par le transporteur sont donnés par le tableau 2.3 ci-dessous.

Ordre	A	B	C	D
Prix	5	5	7	8

TAB. 2.3 – Prix des ordres de transport

Sans un moyen d'exprimer implicitement la relation d'exclusion dans l'exécution des quatre ordres de transport, le transporteur serait obligé de formuler les mises suivantes : $(\{A\}, 5)$, $(\{B\}, 5)$, $(\{C\}, 7)$, $(\{D\}, 8)$, $(\{A, B\}, 0)$, $(\{A, C\}, 0)$, $(\{A, D\}, 0)$, $(\{B, C\}, 0)$, $(\{B, D\}, 0)$, $(\{C, D\}, 0)$, $(\{A, B, C\}, 0)$, $(\{A, B, D\}, 0)$, $(\{A, C, D\}, 0)$, $(\{B, C, D\}, 0)$, $(\{A, B, C, D\}, 0)$. Notons qu'une énumération partielle des mises aux prix non nuls uniquement n'est pas suffisante dans ce cas car elle ne prévient pas le marché d'allouer plus d'un ordre de transport au transporteur.

Supposons maintenant que le marché permet aux participants d'utiliser un *opérateur de mise* OU-exclusif (XOR) défini comme suit. Soit b_1, \dots, b_l des mises combinées usuelles ; une mise $b = XOR\{b_1; \dots; b_l\}$ placée par un participant signifie que ce dernier requiert qu'une, et une seule, parmi les mises b_1, \dots, b_l soit exécutée. L'utilisation de cet opérateur permettrait alors au transporteur d'exprimer implicitement sa préférence en plaçant la mise $XOR\{(\{A\}, 5); (\{B\}, 5); (\{C\}, 7); (\{D\}, 8)\}$.

Dans cette section, nous passons en revue les différents travaux reliés à l'expression de la préférence dans les enchères combinatoires. Ces travaux s'inscrivent dans le cadre de deux types d'approches : les langages de mises et la révélation incrémentale de la préférence.

2.5.1 Langages de mises

Les articles précurseurs dédiés au problème classique de l'allocation combinatoire contiennent presque tous des prémisses de langages de mises. Ainsi, Rothkopf, Pekeč et Harstad [125] présentent une formulation du problème de détermination des mises gagnantes qui, bien qu'équivalente à (CAP-WD), ne fait pas l'hypothèse que le participant place des mises sur tous les sous-ensembles possibles d'objets. Cette formulation exploite une logique OR implicite : chaque acheteur potentiel peut soumettre plusieurs mises combinées et accepte qu'un nombre quelconque de ces mises soient

exécutées. D'autres auteurs se sont plutôt tournés vers la logique XOR, présentée dans l'exemple précédent. C'est le cas notamment de Sandholm [130] et de Fujushima, Leyton-Brown et Shoham [55]. Ces derniers, en particulier, suggèrent l'utilisation d'objets "fictifs" pour simuler une relation d'exclusion. Par exemple, afin d'exprimer la mise $XOR\{(\{A\}, 5); (\{B\}, 5); (\{C\}, 7); (\{D\}, 8)\}$, il suffit de définir un objet fictif E et de soumettre les mises $b_1 = (\{A, E\}, 5)$, $b_2 = (\{B, E\}, 5)$, $b_3 = (\{C, E\}, 7)$ et $b_4 = (\{D, E\}, 8)$, la disponibilité d'une seule "unité" de E faisant en sorte qu'au plus une seule mise parmi b_1, \dots, b_4 peut être exécutée.

Hoos et Boutilier [66] proposent les premiers langages de mise supportant des constructions logiques complexes basées sur la combinaison de plusieurs opérateurs de mise. Étant donné un ensemble d'objets G , les auteurs définissent une *clause* C comme étant un prédicat logique relié à un sous-ensemble S_C de G . La satisfaction de C par une allocation des objets signifie que l'acheteur obtient *un ou plusieurs* objets de S_C . Une mise combinée B est alors exprimée sous la forme $B = < C, p >$, où C est un ensemble de clauses avec la condition implicite que *toutes* les clauses de C soient satisfaites par une allocation des objets, et p un prix que l'acheteur serait prêt à payer le cas échéant. Conceptuellement parlant, le langage de mise \mathcal{L}_{CA}^{cnf} qui en résulte combine donc un opérateur logique OR au niveau de clauses et un opérateur AND au niveau des mises combinées. Hoos et Boutilier présentent également le langage \mathcal{L}_{CA}^{k-of} qui fait aussi appel, au niveau des clauses, à un opérateur plus général de sélection ("Sélectionner k objets parmi n ").

L'article de Nisan [105] constitue, à plusieurs égards, une contribution majeure au développement de langages de mise pour les enchères combinatoires. Sur le plan purement méthodologique, Nisan définit deux critères d'évaluation d'un langage de mise :

1. *Expressivité*. L'expressivité d'un langage de mise mesure, d'une part, la capacité brute du langage à représenter des fonctions de préférence. Ainsi, plus large est l'éventail de fonctions supportées, plus expressif est le langage. D'autre part, il y a une corrélation entre l'expressivité d'un langage et sa concision. Par conséquent, un langage expressif, tout en supportant des fonctions de préférence

importantes, permet l'expression des mises de la manière la plus compacte possible.

2. *Simplicité*. La simplicité d'un langage de mise mesure sa commodité d'utilisation. Un langage simple doit donc être facilement compréhensible et utilisable par les participants, en plus de nécessiter un traitement peu complexe de la part de l'encanteur.

Nisan définit formellement plusieurs langages de mises qu'il est possible de classifier comme suit.

- Mises “atomiques”. Dans ce langage, le plus simple que l'on puisse concevoir, un participant soumet uniquement une *seule* mise combinée de la forme $b = (S, p)$, où S est un sous-ensemble d'objets et p est un prix que le participant serait prêt à payer pour obtenir S . Bien entendu, ce langage est très peu expressif puisque même une fonction de préférence additive n'y est pas représentable.
- Mises-OR, Mises-XOR, OR-de-XORs, XOR-de-ORs, Formules-OR/XOR. Ces langages correspondent respectivement à l'application simple, croisée et récursive des logiques OR et XOR à des mises atomiques. Une analyse empirique de l'expressivité de ces langages, qui fait appel à un ensemble représentatif de fonctions de préférence, fournit nombre d'enseignements, dont le fait que les langages “OR-de-XORs” et “XOR-de-ORs” ne sont pas comparables au sens de l'expressivité.
- Mises-OR*. Ce langage est tout simplement une variante du langage “mises-OR” utilisant des objets fictifs comme moyen de simuler la logique XOR, ce qui rejoint le concept déjà introduit dans Fujushima, Leyton-Brown, et Shoham [55]. Nisan démontre que le langage “mises-OR*” est strictement plus expressif que les langages “OR-de-XORs” et “XOR-de-ORs”, et qu'il permet de simuler, de manière compacte, le langage “Formules-OR/XOR”.

Boutilier et Hoos [23] suggèrent le langage \mathcal{L}_{GB} qui généralise les langages de mise précédents. Les auteurs font la remarque que ces langages peuvent être classifiés, du point de vue de la sémantique de prix, en deux groupes. Dans les langages du premier groupe (Sandholm [130], Fujushima, Leyton-Brown, et Shoham [55], Nisan [105]), les prix sont spécifiés par les participants au niveau des mises atomiques, tandis que les

prix se présentent dans les langages \mathcal{L}_{CA}^{cnf} et \mathcal{L}_{CA}^{k-of} (Hoos and Boutilier [66]) au niveau le plus haut, celui de la mise combinée finale. L'idée de permettre la spécification de prix à un niveau intermédiaire de la formule logique correspondant à une mise combinée, s'impose donc naturellement et est mise en œuvre dans le langage \mathcal{L}_{GB} . Boutilier et Hoos démontrent que, pour certaines fonctions de préférence, cette nouvelle sémantique de prix permet d'exprimer des mises de manière strictement plus concise que les autres langages de la littérature.

2.5.2 Révélation incrémentale de la préférence

Dans tous les langages de mise de la section précédente, les participants à l'enchère disposent de cadres formels plus ou moins sophistiqués leur permettant d'exprimer, selon une syntaxe et une sémantique précises, des mises reflétant leurs fonctions de préférence privées. Les approches basées sur la révélation incrémentale de la préférence (“Preference Elicitation”), quant à elles, délèguent la charge de découvrir les préférences des participants à l'encanteur. Plus précisément, l'encanteur formule des *requêtes* d'information sur ces préférences et les transmet aux participants. Les réponses fournies par ces derniers sont exploitées dans la recherche d'allocations efficaces, ainsi que dans la formulation de requêtes plus précises. Un processus de révélation bien conçu doit permettre à l'encanteur d'inférer des allocations efficaces exactes ou des approximations avec un minimum de requêtes.

Conen et Sandholm [32] proposent un processus d'enchère basé sur la révélation incrémentale de la préférence pour le problème classique de l'allocation combinatoire. Étant donné un ensemble G de m objets différents mis en vente et n acheteurs potentiels. Pour chaque acheteur, une relation d'ordre est définie sur l'ensemble des 2^m paquets d'objets (incluant \emptyset). Cet ordre de préférence, relié à la valeur des sous-ensembles pour l'acheteur, classe ces derniers du plus désirable au moins désirable. Il est alors possible de noter qu'à une allocation $X = \{X_1, \dots, X_n\}$ ($X_i \subseteq G$, $i = 1, \dots, n$ désigne le sous-ensemble d'objets alloué à l'acheteur i) correspond un vecteur unique $R(X) = [R_1(X_1), \dots, R_n(X_n)]$ où $R_i(X_i)$ est le rang de X_i dans l'ordre de préférence de l'acheteur i . Forts de ce constat, les auteurs développent des algorithmes de recherche d'allocations efficaces opérant dans l'espace des vecteurs “rang” et non

pas directement dans l'espace des allocations. Au cours du processus de recherche, l'enchanteur dispose d'informations partielles sur les préférences des acheteurs, qu'il enrichit au fur et à mesure en formulant deux types de requêtes :

- Requêtes portant sur le rang. Exemples : Quel est le paquet préféré, A ou B ? Quel est le rang d'un paquet A ? Quel paquet est au k -ième rang?
- Requêtes portant sur les valeurs de paquets. Notons qu'un acheteur peut répondre à ces questions avec des valeurs exactes ou encore avec des bornes (inférieures ou supérieures) sur les valeurs.

Hudson et Sandholm [67] évaluent numériquement l'efficacité, en ce qui a trait au volume d'informations révélées à l'enchanteur, des algorithmes de révélation de préférence de Conen et Sandholm. Les résultats obtenus confirment qu'en général, un faible pourcentage de l'information complète disponible chez les participants est réellement nécessaire pour déterminer une allocation efficace. Un bémol doit toutefois être apporté à ce résultat : pour que l'allocation obtenue maximise le bien-être social total des participants, il est primordial que ceux-ci répondent sincèrement aux requêtes de l'enchanteur. Bien que Conen et Sandholm suggèrent une variante de leur processus d'enchère qui possède les propriétés économiques requises pour inciter les participants à des réponses sincères, cette variante est basée sur un parallèle trivial avec le processus classique de Vickrey-Clarke-Groves, ce qui signifie que la détermination des paiements nécessite de mettre en place n processus d'enchère secondaires parallèlement au processus principal.

2.6 Mécanismes d'enchère combinatoire incitatifs

Dans plusieurs contextes de marché, maximiser l'efficacité économique de l'allocation est le but ultime de l'enchère. Dans ce cas, il faut tenir compte du fait que la préférence pour l'obtention des paquets d'objets constitue une information privée de chaque participant. L'enchanteur est alors confronté au problème fondamental suivant : comment déterminer une allocation efficace sans avoir accès aux préférences des participants ? Il s'avère qu'une approche élégante permet de résoudre ce problème de manière radicale : il suffit que l'enchanteur fasse en sorte que les règles d'allocation et de détermination des paiements du mécanisme d'enchère soient conçues de telle

façon à ce que la *déclaration véridique* des préférences soit une stratégie dominante d'un participant dans le jeu correspondant à l'enchère. Les participants n'auraient alors aucun intérêt à placer des mises qui reflètent autre chose que leur véritables préférences. Dans ce qui suit, nous désignerons par *mécanismes incitatifs* ("Incentive Compatible Mechanisms") des mécanismes d'enchère vérifiant cette propriété.

2.6.1 Le mécanisme de Vickrey-Clarke-Groves

Nous avons vu que, dans le cas d'enchères de plusieurs objets identiques, l'enchère "second-prix" de Vickrey peut être généralisée en un mécanisme où miser sa vraie valeur est une stratégie dominante pour chaque participant. Les travaux de Clarke [31] et Groves [61] fournissent une généralisation de l'enchère de Vickrey au problème classique de l'allocation combinatoire. Nous présentons le mécanisme en question, largement connu sous le nom de mécanisme de Vickrey-Clarke-Groves (ou VCG).

Commençons d'abord par introduire la notation suivante. Soit un encanteur disposant d'un ensemble G de m objets différents à allouer à un ensemble N de n participants. On désigne par $x = \{x_j\}_{j \in N}$ une allocation des objets aux participants, où $x_{j,i} = 1$ si l'objet $i \in G$ est alloué au participant $j \in N$, $x_{j,i} = 0$ sinon. Soit $v_j(\cdot)$ la fonction de préférence du participant j telle que $v_j(x_j)$ désigne l'utilité pour le participant j d'obtenir l'allocation correspondant au vecteur x_j . Afin d'exprimer le fait que cette utilité est une information privée, nous définissons le concept de *type* t_j du participant j défini sur l'ensemble de types possibles T_j , $j \in N$, et nous introduisons le paramètre type dans la fonction de préférence de chaque participant. Ainsi, $v_j(x_j, t_j)$ désigne l'utilité pour le participant j d'obtenir l'allocation correspondant à x_j étant donné que le type du participant est t_j .

Le mécanisme VCG s'énonce comme suit :

1. Chaque participant j , $j \in N$, déclare un type t_j . Soit $t = \{t_j\}_{j \in N}$ le vecteur des types déclarés par les participants. Une allocation réalisable $x^*(t)$ qui maximise le revenu de l'encanteur étant donnés les types déclarés est calculée comme étant une solution du problème :

$$\max \sum_{j=1}^n v_j(x_j, t_j) \quad (2.35)$$

$$\text{s. à } \sum_{j=1}^n x_{j,i} \leq 1, i \in G \quad (2.36)$$

$$x_{j,i} \in \{0, 1\}, j \in N, i \in G \quad (2.37)$$

2. Soit $h_j(t)$ le paiement du participant j , $j \in N$, étant donné le vecteur t des types déclarés par les participants. $h_j(t)$ est déterminé par :

$$h_j(t) = \sum_{k \neq j} v_k(x_k^*(t_{-j}; 0), t_k) - \sum_{k \neq j} v_k(x_k^*(t), t_k) \quad (2.38)$$

où t_{-j} désigne le vecteur formé des composantes $(t_1, \dots, t_{j-1}, t_{j+1}, \dots, t_n)$ du vecteur t .

Notons que le premier terme de la formule de paiement (2.38) est la somme des préférences des autres participants pour une solution efficace lorsque le participant j annonce une valeur de type nulle, ce qui désigne, par convention, que ce dernier est absent de l'enchère. Quant au second terme, il désigne cette même somme en la présence du participant j . Le paiement d'un participant, dans le mécanisme VCG, doit donc être interprété comme étant l'impact de sa présence sur la valeur d'une allocation efficace des objets aux autres participants. Il est également possible de noter que le paiement selon le mécanisme correspond à une "remise" par rapport à la mise du participant. En effet, le paiement $h_j(t)$ peut être reformulé comme suit :

$$h_j(t) = v_j(x_j^*(t), t_j) - [\sum_{j \in N} v_k(x_k^*(t), t_k) - \sum_{k \neq j} v_k(x_k^*(t_{-j}; 0), t_k)] \quad (2.39)$$

Il est très simple de voir que ce mécanisme incite les participants à révéler leur véritable type. En effet, si t_j désigne maintenant le véritable type du participant j , l'optimalité de $x^*(t)$ pour le problème (2.35-2.37) donne :

$$\begin{aligned}
\sum_{j=1}^n v_j(x_j^*(t), t_j) &= v_j(x_j^*(t_{-j}, t_j), t_j) + \sum_{k \neq j} v_k(x_k^*(t_{-j}, t_j), t_k) \\
&\geq v_j(x_j^*(t_{-j}, \tilde{t}_j), t_j) + \sum_{k \neq j} v_k(x_k^*(t_{-j}, \tilde{t}_j), t_k) \\
&\forall \tilde{t}_j \in T_j, \forall t_{-j} \in \prod_{k \neq j} T_k
\end{aligned}$$

En notant que :

$$\sum_{k \neq j} v_k(x_k^*(t_{-j}, t_j), t_k) = \sum_{k \neq j} v_k(x_k^*(t_{-j}, 0), t_k) - h_j(t) \quad (2.40)$$

on obtient :

$$\begin{aligned}
v_j(x_j^*(t_{-j}, t_j), t_j) - h_j(t_{-j}, t_j) &\geq v_j(x_j^*(t_{-j}, \tilde{t}_j), t_j) - h_j(t_{-j}, \tilde{t}_j) \\
&\forall \tilde{t}_j \in T_j, \forall t_{-j} \in \prod_{k \neq j} T_k
\end{aligned}$$

Ce qui signifie que, pour le participant j , $j \in N$, annoncer sa véritable préférence est une stratégie dominante.

Le mécanisme de Vickrey-Clarke-Groves possède donc la propriété remarquable d'inciter les participants à une enchère combinatoire à dévoiler leur véritable fonction de préférence pour l'obtention des paquets d'objets. De plus, et dans ce qui peut être perçu comme une généralisation des travaux de Myerson sur les enchères optimales d'un objet unique, Krishna et Perry [78] montrent qu'une version plus générale du mécanisme VCG maximise le revenu espéré du vendeur parmi tous les mécanismes *efficaces*, qui satisfont les conditions énoncées par Myerson, à savoir la rationalité individuelle des participants et l'incitation aux déclarations véridiques des préférences. Toutefois, au vu du problème (2.35-2.37) et de l'équation (2.38) qui expriment respectivement les problèmes de la détermination des mises gagnantes, et du calcul du paiement de chaque participant, une implantation directe du mécanisme VCG nécessiterait la résolution de $n + 1$ problèmes similaires à (2.35-2.37), ce qui peut être

extrêmement coûteux. De plus, les participants se trouvent dans l'obligation de communiquer à l'encanteur des fonctions de préférence complètes (c'est-à-dire en précisant leur préférence pour chaque paquet d'objets qu'il leur est possible d'obtenir), afin de permettre à ce dernier d'effectuer l'allocation et de déterminer les paiements.

2.6.2 Approximations du mécanisme VCG

Une manière possible de contourner la complexité du mécanisme VCG est de résoudre le problème d'allocation (2.35-2.37) de manière approximative. Ainsi, on remplacerait la solution optimale $x^*(t)$ de ce problème par une solution approchée $\tilde{x}(t)$ obtenue à partir d'une procédure heuristique. La question suivante s'impose alors immédiatement : la propriété d'incitation à la déclaration vérifique des préférences est-elle maintenue ? Malheureusement, la réponse à cette question est négative en général (voir Parkes [113], par exemple).

Plusieurs contributions récentes ont néanmoins permis d'apporter des éclaircissements additionnels au sujet de la pertinence de mécanismes VCG approximatifs. Ainsi, Kfir-Dahav, Monderer et Tennenholtz [73] présentent trois axiomes qui assurent qu'un mécanisme VCG, utilisant une heuristique dans la détermination des mises gagnantes par le problème (2.35-2.37), garde la propriété d'incitation. En particulier, le second axiome correspond à la condition qu'un participant ne peut améliorer la valeur de l'allocation, étant donné un ensemble de types reportés, en modifiant unilatéralement son propre type. Toutefois, les auteurs ne fournissent aucune indication quant à la conception de telles heuristiques, ni des garanties sur leurs performances. Un autre type d'approche est suggéré par Lehmann, O'Callaghan et Shoham [85], qui considèrent - dans le même esprit que les restrictions de Rothkopf, Pekeč et Harstad [125] - le cas particulier d'enchères combinatoires où les participants sont intéressés à un paquet d'objets unique. Bien que le problème d'allocation demeure NP-complet, Lehmann, O'Callaghan et Shoham montrent que des règles d'allocation et de détermination des paiements vérifiant des propriétés très naturelles donnent lieu à des mécanismes incitatifs et efficaces sur le plan numérique. Plus précisément, ces règles doivent faire en sorte que (i) l'allocation est *exacte* : le participant obtient le

paquet demandé, ou rien du tout; (ii) l'allocation est *monotone* : si le participant obtient le paquet S en misant (S, p) , il aurait tout aussi bien obtenu S' s'il avait misé (S', p') , où $S' \subseteq S$ et $p' \geq p$; (iii) le paiement est *critique*, c'est-à-dire qu'un participant avec une mise gagnante paie le prix minimal qu'il aurait pu miser et demeurer gagnant. Les auteurs proposent, d'autre part, une procédure d'allocation vorace et des règles de détermination de paiements vérifiant ces propriétés et possédant par conséquent la propriété d'incitation.

Une approche d'approximation originale est proposée par Nisan et Ronen [107]. Ces derniers définissent la classe de mécanismes de marché dits “basés” sur le mécanisme VCG comme étant l'ensemble des mécanismes faisant appel à une procédure d'allocation *optimale* ou *approximative* des objets, mais utilisant une formule semblable à (2.38) pour le calcul des paiements. Pour de tels mécanismes, il n'est pas difficile d'observer que l'unique raison pour qu'un participant déclare faussement ses préférences est de permettre à l'encanteur d'améliorer son allocation des objets. Ceci amène les auteurs à proposer une variante du mécanisme VCG, qu'ils appellent *mécanisme de la seconde chance*. Dans ce mécanisme, chaque participant j , $j \in N$, place, comme dans un mécanisme VCG classique, une déclaration t_j de son type. En outre, le participant soumet à l'encanteur une fonction de recours (“appeal function”) $l_j : \prod_{k \in N} T_k \rightarrow T_j$. La signification de cette fonction est la suivante : “Si les types déclarés par les participants sont t_1, \dots, t_n , alors je juge que l'algorithme d'allocation produira une meilleure allocation si $l_j(t_1, \dots, t_n)$ est substitué à t_j ”. L'encanteur détermine alors la meilleure allocation en considérant respectivement les types originaux et les fonctions de recours, et calcule les paiements de la manière habituelle (formule (2.38)). Les auteurs spécifient enfin des conditions nécessaires pour que le mécanisme de la seconde chance incite des participants disposant de connaissances partielles ou de moyens de calcul limités (ne leur permettant, par exemple, que de dériver des stratégies “approximatives”) à déclarer leurs vraies préférences.

2.6.3 Implantation indirecte du mécanisme VCG

Une autre limitation majeure du mécanisme Vickrey-Clarke-Groves, dans sa version originale, est son caractère “enveloppe fermée”, c'est-à-dire que les participants doivent transmettre leur mises à l'enchanteur d'un seul trait et de manière intégrale. On serait alors en droit de se poser la question suivante : existe-t-il des mécanismes d'enchères implantant *indirectement* le mécanisme VCG ? Plus précisément, ces mécanismes indirects doivent procéder en permettant aux participants de dévoiler progressivement leurs préférences, et aboutir en fin de compte au résultat escompté du mécanisme VCG, à savoir une allocation efficace et les paiements “second-prix” donnés par la formule (2.38).

L'approche primale/duale présentée à la section 2.4 est encore une fois à la base des mécanismes itératifs implantant indirectement l'enchère VCG. Un certain nombre de notions et de résultats théoriques additionnels sont toutefois nécessaires. Parmi ceux-ci, les plus importants concernent le concept de *produit marginal* d'un participant et le lien entre le produit marginal et les équilibres compétitifs de prix. Bikhchandani *et al.* [18] définissent le produit marginal du participant j comme étant la quantité $V^* - V_{-j}^*$, où V^* est la valeur d'une allocation efficace et V_{-j}^* désigne la valeur d'une allocation efficace ne tenant pas compte de la présence du participant j . Il convient de noter que le produit marginal d'un participant n'est autre que le surplus réalisé par le participant dans une enchère VCG (c'est précisément la “remise” monétaire de l'équation (2.39)). Considérons maintenant la formulation (CAP-SE-3) (qui, rappelons-le, possède la propriété d'intégralité) et les variables duales de sa relaxation linéaire $\{s_j\}_{j \in J}$, r et $\{ps_{j,i}\}_{S \subseteq G, j \in N}$ associées respectivement aux contraintes (2.31), (2.32) et (2.33). En faisant appel à l'analyse de sensibilité en programmation linéaire, il est possible d'interpréter une variable duale optimale s_j^* comme étant l'impact d'une perturbation (qui doit toutefois être infinitésimale !) du terme de droite de l'inégalité correspondante (2.31) sur l'objectif de (CAP-SE-3). Si de plus, on note qu'un changement de ce terme de 1 à 0 revient à écarter le participant j de l'enchère, il est tentant de voir sous quelles conditions $s_j^* = V^* - V_{-j}^*, \forall j \in N$.

Bikhchandani et Ostroy [20] définissent une condition nécessaire et suffisante

pour qu'une des variables duales optimales $[\{s_j^*\}_{j \in J}, r^*, \{p_{S,j}^*\}_{S \subseteq G, j \in N}]$ vérifie $s_j^* = V^* - V_{-j}^*, \forall j \in N$. Cette condition, que les auteurs appellent "substituabilité" des participants ("agents-are-substitutes") s'écrit :

$$V^* - V_{-\tilde{N}}^* \geq \sum_{j \in N \setminus \tilde{N}} V^* - V_{-j}^*, \forall \tilde{N} \subseteq N, \quad (2.41)$$

où $V_{-\tilde{N}}^*$ désigne la valeur de l'allocation efficace en écartant tous les participants de \tilde{N} . Intuitivement, la "substituabilité" des participants signifie que l'impact d'une "coalition" de ces derniers est toujours supérieur à l'impact cumulé des membres de la coalition. Bikhchandani et Ostroy montrent que, quand cette condition est satisfaite, les solutions optimales duales $[\{s_j^*\}_{j \in N}, r^*, \{p_{S,j}^*\}_{S \subseteq G, j \in N}]$ vérifiant $s_j^* = V^* - V_{-j}^*, \forall j \in N$ sont précisément celles qui correspondent au niveau de revenu r^* le plus faible parmi toutes les solutions duales optimales.

La possibilité de concevoir des implantations itératives du mécanisme VCG a également été envisagée pour des fonctions de préférence possédant la propriété de substituabilité entre produits ("Gross Substitutes"), qui se résume à dire que la demande d'un participant pour un objet donné dont le prix demeure inchangé ne diminue pas si les prix d'autres objets augmentent. Dans ce cadre, Gul et Stacchetti [63] proposent un mécanisme d'enchère ascendante s'inspirant de l'enchère anglaise classique. L'enchère proposée procède de la même façon que le mécanisme itératif suggéré par Demange, Gale et Sotomayor [40] pour le problème d'affectation, en ajustant à la hausse les prix des objets faisant partie des ensembles d'objets "à excès de demande" de cardinalité minimale. La convergence, en un nombre fini d'itérations, vers le vecteur de prix walrasiens *minimaux* est prouvée. Gul et Stacchetti vont néanmoins plus loin, et montrent que : (a) ces prix walrasiens ne correspondent pas toujours aux paiements du mécanisme VCG, et (b) aucun mécanisme itératif "simple", c'est-à-dire basé sur une trajectoire ascendante de prix linéaires, ne permet de calculer ces paiements pour toutes les fonctions de préférence vérifiant la propriété de substituabilité entre produits.

Ce résultat négatif est quelque peu tempéré par un mécanisme d'enchère innovateur mis au point récemment par Ausubel [9]. Le mécanisme suggéré comporte une

enchère principale et $|N|$ enchères secondaires. Dans une enchère secondaire j , correspondant au participant $j \in N$, ce dernier est supposé *absent*, et un vecteur de prix d'équilibre walrasien p_{-j}^* est déterminé grâce à un processus de tâtonnement classique (l'existence de ces prix est garantie par la condition de substitutabilité entre produits). L'enchère principale, quant à elle, fait intervenir tous les participants et converge vers un vecteur de prix d'équilibre p^* . Ausubel montre qu'en démarrant l'enchère principale avec les prix p_{-j}^* , la trajectoire des allocations et des prix de l'enchère contient suffisamment d'informations pour déterminer le paiement du participant j dans le mécanisme VCG. Il convient de noter la remarquable généralité du cadre théorique dans lequel l'auteur évolue, et qui permet de concevoir autant des variantes itératives que des versions progressives continues du mécanisme.

À notre connaissance, le mécanisme “iBundle Extend & Adjust”, ou iBEA (Parkes [112], Parkes et Ungar [116]) est la seule implantation itérative du mécanisme VCG pour des fonctions de préférence générales (ne vérifiant pas forcément les propriétés de substituabilité entre produits ou entre participants). À ce propos, Parkes et Ungar [116] notent que les approches primales/duales classiques, qui déterminent les paiements VCG à l'aide des prix d'équilibre compétitif minimisant le revenu de l'encanteur, requièrent au moins la substituabilité entre participants et sont donc vouées à l'échec dans le cas général. Aussi le mécanisme iBEA se démarque-t-il de ces approches par un processus à deux phases. Dans la première phase, que l'on peut assimiler au mécanisme iBundle(3) (Parkes [111]), une enchère ascendante détermine une allocation efficace des objets à des participants réagissant aux prix annoncés par l'encanteur en plaçant des mises sur les paquets générant le surplus maximal. La seconde phase a pour objectif de déterminer des remises aux participants par rapport aux prix d'équilibre compétitif de la première phase de façon à ce que les paiements révisés correspondent aux paiements VCG. Pour cela, cette phase incite les participants ayant obtenu des paquets dans l'allocation efficace à continuer à miser sur ces paquets en introduisant des participants “fictifs” pour maintenir la compétition. Les auteurs définissent le *test de Vickrey*, qui se résume à ce qui suit : aucun des participants *dépendant* d'un participant j faisant partie des gagnants dans l'allocation

efficace (dans le sens où ces participants voient leur allocation affectée par l'absence de j) n'est plus actif dans la seconde phase du mécanisme. En fait, Parkes et Ungar démontrent que ce test permet d'arrêter le mécanisme car il indique que suffisamment d'informations ont été rassemblées pour le calcul des paiements VCG.

2.7 Conclusion

Les enchères combinatoires ont été à l'origine d'un important volume de travaux durant les dix dernières années. Dans ce chapitre, nous avons présenté une revue de littérature des mécanismes d'enchère combinatoire. Après un bref aperçu des premiers travaux sur les enchères d'un objet unique et multi-unités, nous avons exposé les principaux modèles d'enchère combinatoire et des exemples d'applications, les mécanismes progressifs d'enchère combinatoire, les langages de mise combinatoire, ainsi que les mécanismes "second-prix", incitatifs à une révélation véridique des préférences.

Chapitre 3

Design Issues for Combinatorial Auctions

La référence de cet article est :

J. Abrache, T.G. Crainic, and M. Gendreau, "Design Issues for Combinatorial Auctions", à paraître dans *4OR*.

Design Issues for Combinatorial Auctions

Jawad Abrache^{1,3}, Teodor Gabriel Crainic^{2,3}, and Michel Gendreau^{1,3}

¹ Département d'informatique et de recherche opérationnelle, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada, H3C 3J7.

² Département management et technologie, Université du Québec à Montréal, C.P. 6192, succursale Centre-ville, Montréal, Canada H3C 4R2.

³ Centre de recherche sur les transports, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada, H3C 3J7.

Abstract. Combinatorial auctions are an important class of market mechanisms in which participants are allowed to bid on bundles of multiple heterogeneous items. In this paper, we discuss several complex issues that are encountered in the design of combinatorial auctions. These issues are related to the formulation of the winner determination problem, the expression of combined bids, the design of progressive combinatorial auctions that require less information revelation, and the need for decision support tools to help participants make profitable bidding decisions. For each issue, we survey the existing literature and propose avenues for further research.

Keywords : E-commerce, Mechanism Design, Combinatorial Auctions, Bidding Languages, Iterative Auctions, Advisors

3.1 Introduction

Auctions are important market mechanisms, used since the earliest of times for the allocation of goods and services. Public and private institutions generally prefer them to other common trading processes (lotteries, price-posting, etc.) because they are open, quite fair, generally easy to understand by participants, and often lead to economically efficient outcomes. However, a real surge in their popularity has only been observed during the last decade, due in part to the emergence of e-commerce and the increasing tendency to shift important business activities to the Internet, as well as to a deregulation wave that led to the privatization of several industries. While the well-publicized Federal Commission for Communications (FCC) auctions of spectrum licenses (McMillan [99]) remain the most striking example, auctions have been used for a variety of other purposes. These include the allocation of airport take-off and landing time slots (Rassenti, Smith, and Bulfin [120]), course registration (Graves, Schrage, and Sankaran [60]), private and public procurement (Davenport and Kalagnanam [37]), sale of online seats (Eso [47]), distribution routes (Ledyard *et al.* [83]), job shop scheduling (Wellman *et al.* [145]), etc.

Whether they involve spectrum rights, transportation routes, or computer hardware parts, many markets of interest have one thing in common : they all trade items of different nature that are *interrelated* from the perspective of the participants to the market. Item interrelation means that, independently of the way items are traded in the market, the value of a given item to a participant depends on whether or not that participant has been able to trade some other items as well. Items may in that regard be *complementary* or *substitutable* to each other. More precisely, if A and B are two items, and $v(\cdot)$ denotes the participant's (supposed to be a buyer) function of preference, A and B are said to be complementary if $v(\{A, B\}) > v(\{A\}) + v(\{B\})$, and substitutable if $v(\{A, B\}) < v(\{A\}) + v(\{B\})$. Consider time slots in airports as an example. A take-off slot associated to the origin airport of a flight and a corresponding landing slot in the destination airport indeed complement each other. On the other hand, two pairs of take-off and landing time slots that correspond to the same origin and destination airports within the same period of time (e.g., from 8h00 to

8h30) are likely to be substitutable for an airline company operating one daily service between the two airports.

The way item interrelation impacts the trading strategies of a participant depends primarily on how items are traded in the market. For example, if the market maker has several different items to sell and decides to do so by running several parallel auctions (one for each item), the participant could of course submit several simultaneous bids on all the items in which it is interested. It may continue to bid on complementary items that constitute a desirable collection of items till the total value of its bids reach its preference for the collection. But since the auctions of the different items run independently of each other, the participant may find himself stuck with a subset of the desired collection, which it would have paid more than its value. This *exposure* problem often leads in practice to strategic bidding and therefore to economically inefficient auctions.

Combinatorial auctions are increasingly considered as an alternative to simultaneous single-item auctions. Combinatorial auctions commonly refer to auction mechanisms in which participants are allowed to bid on combinations, or bundles of items. Being able to bid on bundles clearly mitigates the exposure problem, since it gives the participants the option to bid their precise valuations for any collection of items they desire. On the other hand, combinatorial auctions often require the market maker, the participants, or both, to solve complex decision problems. Hence, consider what might arguably be the simplest setting for a combinatorial auction : an auctioneer selling n different items to several potential buyers, which are allowed to submit sealed bids on bundles of items. On the basis of the bids it receives, the auctioneer must decide which bids win and which ones lose, under the condition that no single item is allocated to more than one bid, and such that its revenue from the sale of the items is maximized. This *winner determination problem* is well-known to be NP-hard (Rothkopf, Pekeč, and Harstad [125]) and even difficult to approximate (see Sandholm [130], for example).

The challenge of *mechanism design* (MasColell, Whinston, and Green [95]) for combinatorial auctions is much broader. In the context of auctions, a mechanism can be defined as the specification of all possible *bidding strategies* available to the par-

ticipants, and of an outcome function that maps these strategies to an *allocation* of items (who gets what ?) and corresponding *payments* the participants need to make or receive. The mapping is generally done with respect to an objective that can be the maximization of the revenue of the sellers, the maximization of the overall social efficiency of the allocation, or any other objective. Market designers trying to implement auction mechanisms therefore find themselves faced with many complicated issues to address. While some of these issues, such as deciding on bidder qualification, entry fees, or scoring rules, call mainly upon the experience of the designer and its knowledge of the context of the auction (Rothkopf and Park [124]), some others are indeed fundamental. These issues concern, for instance, the decision to make bidding in the auction one-shot or progressive, the nature and timing of the information to be revealed to the participants in intermediary stages of the auction. More importantly, the designer often needs to ensure, by properly setting the rules of the auction, that the objectives are always achieved, even in environments involving self-interested participants and characterized by incomplete information (i.e., participants keeping private their preferences). Since Myerson's seminal paper on optimal auction design (Myerson [103]), tackling these questions has greatly motivated the overall research effort on mechanism design, and a significant part of this effort has been devoted to combinatorial auction mechanisms. Important examples of these include the AUSM auction (Banks, Ledyard, and Porter [13]), the RAD mechanism (DeMartini *et al.* [41]), the PAUSE mechanism (Kelly and Steinberg [72]), the AkBA family of auctions (Wurman and Wellman [148]), and the iBundle mechanism (Parkes [111]).

As pointed out in Rothkopf and Park [124], market design is a multidisciplinary effort made of contributions from economics, operations research, computer science, and many other disciplines. Economists, in particular, have played a decisive role in the exploration of the theoretical properties of auctions (Klemperer [76]) : by putting game theory into application, they have built models that describe the strategic behavior of the participants in the many auction types ; they have shaped powerful theories for economic efficiency, pricing, incentives, and collusive behavior, among other issues relevant to auctions ; and last but not least, they have set in experimental economics the scientific foundations for testing their theories. The contribution of computer

science lies mostly in (a) the development of appropriate software architectures and tools for the deployment of auctions; (b) the design of software *agents* capable of interacting competitively or cooperatively in an “intelligent” way; and (c) the design and implementation of the simulation platform for the evaluation of auction mechanisms in controlled artificial environments. As for operations research, it will play in our opinion an increasingly important role in the modeling of the many decision problems encountered by the auctioneer and the participants during the course of an auction mechanism. Being closer to the actual applications, it has tendency to develop more detailed models of the reality than, say, economics, and thus may be particularly appealing to engineers and practitioners. Furthermore, when these models are “hard” - as this is the case in combinatorial auctions, optimization techniques can be extremely valuable in the design of efficient exact and heuristic solution approaches to the proposed formulations, and may even be impossible to circumvent.

The goal of this paper is to identify and discuss some of the complex issues related to the design of combinatorial auctions. We put the emphasis on four issues : a classification of combinatorial auctions and the associated formulations, the expression of combinatorial bids, the design of multi-round mechanisms intended to determine allocations and prices in situations where complete information about the participants’ preferences is not available, and the decision problems faced by participants in combinatorial auctions.

The paper is organized as follows. In Section 3.2, we present an elementary taxonomy of auctions and survey several important formulations of the winner determination problem. In Section 3.3, we tackle the important issue of the expression of combined bids and give evidence of the need for a bidding framework that goes beyond what the basic languages currently permit participants to express. In Section 3.4, we discuss progressive auction mechanisms that approximate the behavior of an “ideal” complete information market, when only incomplete information about participants’ valuation functions is available to the auctioneer. For these mechanisms, pricing schemes and the design of auction rules are interesting but generally challenging issues that need to be studied more extensively. We conclude in Section 3.5 with a general discussion about the role of advisors to participants in combinatorial

auctions.

3.2 Basic formulations

In order to make the presentation as uniform as possible, we present a taxonomy of auctions we use throughout the paper. It is not our intention, however, to realize an exhaustive parameterization of auctions. For a fuller treatment of auction classification, we refer the reader to Engelbrecht-Wiggans [44] and Wurman, Wellman, and Walsh [150]. Hence, we limit ourselves to the following dimensions of the auction space.

– **What is traded ?** Items that are traded can be :

1. *Indivisible* goods versus *divisible* ones. Capacity in telecommunication networks is divisible, but rail right-of-way is not. It is noteworthy that, when multiple units of items are traded, *item* divisibility should be clearly distinguished from *bid* divisibility, in the sense that the former depends intrinsically on the physical nature of goods while the latter refers to bidders' tolerance to obtain partial execution of their bids. Note also that auctions of divisible items sometimes provide acceptable models for the sale of physically indivisible goods, especially when large volumes are involved in the trade (assets in financial markets, for example).
2. *Pure* commodities that have no special structure versus *network* commodities which refer to capacity or services that belong to systems with network structure.

– **What roles do the participants play in the auction ?** It is possible to distinguish between *one-sided* auctions and *multilateral* ones. One-sided auctions correspond to trading situations in which there is (a) one seller and multiple buyers (one-to-many), or (b) many sellers and one buyer (many-to-one). Multilateral auctions, often designated by the name *exchanges*, involve many sellers and many buyers (many-to-many). It is noteworthy that a participant in an exchange can be only a seller, only a buyer, or both.

– **What are the objectives of the auction ?** Auctions can be *optimized* or not. In optimized auctions, the market mechanism ensures that a given goal is

achieved when the auction *clears*, i.e., when (provisional or final) allocations and payments are determined. Hereby, we may separately consider :

1. The *allocation rule*, which induces : (a) *locally efficient* outcomes (Wurman, Wellman, and Walsh [150]) when the revenue of the seller in a one-to-many configuration, the cost to the buyer in a many-to-one configuration, or the surplus of the auction in multilateral cases are optimized given the bids of the participants ; or (b) *socially efficient* outcomes when the overall social welfare of the participants is optimized.

2. The *pricing rule*, which indicates what participants should pay or receive. For example, a participant whose bids win may have to pay a *uniform* price corresponding to an “equilibrium” state of the market, the exact amount of money specified in the bids (first-price auctions), or the price of the “second-best” bid (Vickrey-based payments).

- **How “complex” are the participants’ bids ?** If we limit ourselves to combinatorial auctions, we would have to decide whether the participants bid *simply*, i.e., make *unrelated* bids in which they specify only the composition of the bid and a corresponding price, or are allowed to use sophisticated *bidding languages*, in terms of which they may express more complex bidding requirements. We will elaborate further on this issue in Section 3.3.
- **How is the auction organized ?** An auction may be :

1. *Single-round* if it clears only once, or *progressive* if provisional outcomes are determined during the course of the auction and participants are allowed to update their bids. Progressive auctions can be *iterative* (multi-round) if there are pre-specified events that schedule bidding and clearing in the auction, or *continuous* if clearing may occur asynchronously (for instance, whenever new bids are submitted by participants, no bidding activity is observed during a given period of time, etc.).

2. Based on an *ascending* price update scheme (English-like auction), *descending* price update scheme (Dutch-like auction), or non-monotone price updates (e.g., Walrasian tâtonnement).

- 3. *Sequential* when items are traded one at a time (e.g., art auctions), or *parallel* when they are traded simultaneously.
- **What information is revealed to participants?** We distinguish between *sealed-bid* auctions in which no information is disclosed to the participants, and *open* auctions that provide them with “signals” about the state of the auction. Very often, the information handed over to participants consists of anonymous or personalized price quotes new bids need to beat in order to be eligible to be provisional winners.

By giving specific values to parameters in each one of these dimensions, one may derive different combinatorial auction situations and mechanisms. In order to illustrate the modeling challenges, we limit ourselves in this section to the first two dimensions (the nature of items and the roles of participants), and only consider local efficiency as objective. The basic winner determination formulations have already been studied in the literature on combinatorial auctions. In this survey, we connect them to classical optimization problems to help gain useful insights into the complexity of tackling winner determination and integrating it into complex auction mechanisms.

3.2.1 The one-to-many indivisible case

In this configuration, one seller has a set G of m indivisible items to sell to n potential buyers. Let us suppose first that items are available in single units. A bid made by buyer j , $1 \leq j \leq n$ is defined as a tuple $(S, p_{j,S})$ where $S \subseteq G$ and $p_{j,S}$ is the amount of money buyer j is ready to pay to obtain bundle S . Define $x_{j,S} = 1$ if S is allocated to buyer j , and 0 otherwise. The winner determination problem can be formulated as model (M1) :

$$\max \sum_{1 \leq j \leq n} \sum_{S \subseteq G} p_{j,S} x_{j,S} \quad (3.1)$$

$$\text{s.t. } \sum_{1 \leq j \leq n} \sum_{S \subseteq G} \delta_{i,S} x_{j,S} \leq 1, \forall i \in G \quad (3.2)$$

$$\sum_{S \subseteq G} x_{j,S} \leq 1, \forall j, 1 \leq j \leq n \quad (3.3)$$

$$x_{j,S} \in \{0, 1\}, \forall S \subseteq G, \forall j, 1 \leq j \leq n \quad (3.4)$$

where $\delta_{i,S} = 1$ if $i \in S$, and 0 otherwise. Constraints (3.2) establish that no single item is allocated to more than one buyer, while constraints (3.3) ensure that no buyer obtains more than one bundle. The objective is to maximize the revenue of the seller given the bids made by buyers.

Model (M1) corresponds to a set-packing problem (de Vries and Vohra [39]). The classical account of the problem by Rothkopf, Pekeč, and Harstad [125] first proposes a dynamic programming algorithm that can determine a revenue-maximizing allocation in $O(3^m)$ iterations. The algorithm is based on the straightforward remark that, given a subset S of items, the maximum revenue that can be achieved from the sale of S comes from a bid $(S, p_{j,S})$ on S itself, or from the sale of two subsets S_1 and S_2 that form a partition of S . The authors also consider several restrictions on allowable bids that make the problem computationally manageable. Hence, they show that winner determination can be solved in polynomial time if bids have nested structure (any two bundles are either disjoint or one of them is a subset of the other), some cardinality-based restrictions are imposed (e.g., allow only bundles of two items or less), or bids have some inherent geometric structure (notably when items can be linearly ordered and bundles may only contain items that are adjacent to each other).

By opposition to the worst case analysis of Rothkopf, Pekeč, and Harstad, the search algorithms that have been proposed in the literature (Fujushima, Leyton-Brown, and Shoham [55]; Sandholm [130]; Sandholm *et al.* [132]; Hoos and Boutilier [66], for instance) capitalize on the observation that, when the number of items is large, bidders are likely to formulate bids on only a small subset of all possible bundles. In particular, Sandholm [130] proposes a tree representation of the solution space in which items are judiciously indexed so that a feasible allocation can be represented only once. Fujushima, Leyton-Brown, and Shoham [55] suggest in their CASS algorithm a structured depth-first search procedure in which two fundamental ideas are put forward to avoid unnecessary computation. The first is the identification of subsets of mutually incompatible bids (“bins”), i.e., which cannot be simultaneously

executed due to a conflict on one item, so that the exploration of a solution can be interrupted as soon as two items in a same “bin” are encountered. The second idea, inspired by dynamic programming, is the use of intermediate results to prune the search tree : suppose we already know the maximum revenue r_C^* that can be achieved from the sale of $C \subseteq G$, and we consider an partial feasible allocation of the subset $F \subseteq G$ at a given step of the search such that $G \setminus F \subseteq C$; if $r_C^* + r_F$ is lower than the revenue of the best feasible allocation found up to that point, then there is no need to explore the tree beyond F . The CABOB algorithm of Sandholm *et al.* [132] calls upon additional techniques that include pruning with upper and lower bounds, decomposition of the bid graph, and dynamic branching heuristics. Andersson, Tenhunen, and Ygge [5] have recently made insightful computational comparisons between some of the search algorithms (namely the CASS and Sandholm’s algorithm) and standard MIP techniques used in commercial solvers (CPLEX 6.5). Although it has not taken the most recent developments into account, their study comes to the conclusion that the overall performance of CPLEX is actually very good compared to that of the search algorithms.

While many approximate methods for the general set packing problem have been suggested in the literature (Chandra and Halldörsson [27]), the Casanova algorithm by Hoos and Boutilier [66] is, to the best of our knowledge, the only representative of this class of methods in the context of combinatorial auctions. Casanova is a stochastic search algorithm using in its exploration of the allocation space a simple concept of neighborhood. More specifically, a single non executed bid in the solution corresponding to the current feasible allocation of items is chosen for execution in the next feasible allocation. The choice is done according to the bid’s “score”, which designates the ratio of the bid’s price to the number of items in the bid. Numerical comparison with the CASS algorithm seems to indicate promising results.

The multi-unit combinatorial auction (MUCA) extends model (*M1*). Here, the seller has M_i available units of item i to sell. A bid submitted by a buyer takes the form $b = (\{a_{b,i}\}_{i \in G}, p_b)$, where $a_{b,i}$ is the number of units of item i that are requested by bid b , and p_b is the price the buyer offers for the collection $\{a_{b,i}\}_{i \in G}$. Let B denote the set of all bids made by buyers, and $x_b = 1$ is bid b wins, and 0 if it loses, $\forall b \in B$.

The winner determination problem can be written in this case as model $(M2)$:

$$\max \quad \sum_{b \in B} p_b x_b \quad (3.5)$$

$$s.t. \quad \sum_{b \in B} a_{b,i} x_b \leq M_i, \forall i \in G \quad (3.6)$$

$$x_b \in \{0, 1\}, \forall b \in B \quad (3.7)$$

Model $(M2)$ is a 0-1 multidimensional knapsack problem for which exact and heuristic solution methods have been designed and implemented (Martello and Toth [94]). In the context of combinatorial auctions, many search algorithms have been recently proposed. Among the important contributions, Leyton-Brown, Shoham, and Tennenholtz [89] present the CAMUS (“Combinatorial Auction Multi-Unit Search”) algorithm in which the main techniques introduced by former search algorithms are generalized to deal with the multi-unit model. Various bounding techniques (using notably linear relaxation of the multidimensional knapsack problem and greedy allocation procedures) are suggested in the Branch-and-Bound algorithms of Gonen and Lehmamn [59] and Lehmamn and Gonen [84]. Finally, the observation in Mansini and Speranza [92] that a lower bound on the *number* of winning bids in an optimal allocation can be derived from the solution of $|G|$ auxiliary knapsack problems is instrumental in allowing the formulation of good valid inequalities for the $(M2)$ formulation and in improving the upper bounds on the optimal solution. Preliminary results obtained by the authors suggest significant improvement in performance in comparison with CPLEX 7.0, particularly for large problems.

3.2.2 Many-to-one combinatorial auctions

In a many-to-one configuration of combinatorial auctions (sometimes called *reverse* combinatorial auctions), one buyer needs to obtain a set G of items, supplied by several potential sellers. A bid b made by a seller can be defined as $b = (S_b, p_b)$, where S_b is a subset of items, and p_b an ask price the seller requires to be paid for S_b to be supplied. Again consider the set B of all bids, and define binary decision

variables $x_b = 1$ if bid b wins, and 0 if it loses, $\forall b \in B$. The winner determination problem is to find the less expensive set of bids that provide the buyer with all items in G , and corresponds to model $(M3)$:

$$\min \quad \sum_{b \in B} p_b x_b \quad (3.8)$$

$$s.t. \quad \sum_{b \in B} \delta_{i,S_b} x_b \geq 1, \forall i \in G \quad (3.9)$$

$$x_b \in \{0, 1\}, \forall b \in B \quad (3.10)$$

where $\delta_{i,S_b} = 1$ if $i \in S_b$, and 0 otherwise. Model $(M3)$ is a set-covering problem, which is also NP-hard. In practice, it is particularly useful in modeling procurement of goods and services. It is important to note that an implicit *free disposal* assumption is made in model $(M3)$; that is, the buyer tolerate more than one unit of each item to be supplied. If this tolerance to extra units cannot be assumed in a particular market context, constraints (3.9) need to be changed to equalities. The corresponding set partitioning problem proves to be relatively more difficult to tackle (Sandholm *et al.* [133]). Two recent applications of reverse combinatorial auctions are trucking service acquisition for Sears Logistical Services (Ledyard *et al.* [83]) and procurement of direct inputs for a food manufacturer (Davenport and Kalagnanam [37]). In the latter, side constraints that enforce the buyer's tolerance to the number of winners and volumes of goods received from each of them have been added to the basic model, and their impact on solution times has been investigated.

3.2.3 A network formulation

All previous models have dealt with pure items with no special structure. We claim that, when the traded commodities correspond to network resources (e.g., capacity in telecommunication networks), complex bidding requirements related to flow conservation, required offer and demand, etc., can be directly represented on network structures, and specialized network flow algorithms can help in finding the optimal allocations more efficiently than plain LP solution methods. By way of illustration,

and in order to give an empirical support to our claim, we present in this section a basic formulation of the winner determination problem in a combinatorial auction for selling network capacity.

Let $G = (V, A)$ be a network, where V is a set of vertices and A a set of links. To each link $a \in A$ is associated a capacity v_a . It is assumed that the capacity is owned by a single seller and that there are several buyers. The combinatorial aspect of the problem ensues from the fact that buyers desire to obtain capacity between pairs of vertices, rather than on individual links. To simplify the presentation, we define a bid b_j submitted by buyer $j \in N$ as $b_j = (\{O_j, D_j\}, c_j, G_j, p_j)$ (we suppose, with no loss of generality, that a buyer submits a single bid), where

1. $(\{O_j, D_j\})$ is an origin-destination pair of vertices specifying that buyer j needs capacity between O_j and D_j ;
2. c_j is the required capacity between O_j et D_j ;
3. $G_j \subseteq G$ is a subnetwork such that $O_j, D_j \in G_j$, with the condition that all the capacity required between O_j and D_j must be within paths in G_j ;
4. p_j is the price offer of participant j for the bundle.

Figure 3.1 illustrates such capacity bidding. Two bids have been submitted. b_1 is a \$100 bid for a capacity of 20 contained in the sub-network G_1 between O_1 and D_1 , while b_2 is a \$80 bid for a capacity of 10 on path $O_2 - I_2 - D_2$.

Let K_j be the set of paths between O_j and D_j that are in G_j . We define the decision variables $x_j, \forall j \in N$ and $h_k, \forall k \in K_j, \forall j \in N$, as follows :

$$x_j = \begin{cases} 1, & \text{if bid } b_j \text{ is winning,} \\ 0, & \text{otherwise;} \end{cases}$$

and h_k is the capacity allocated to participant j on path $k \in K_j$.

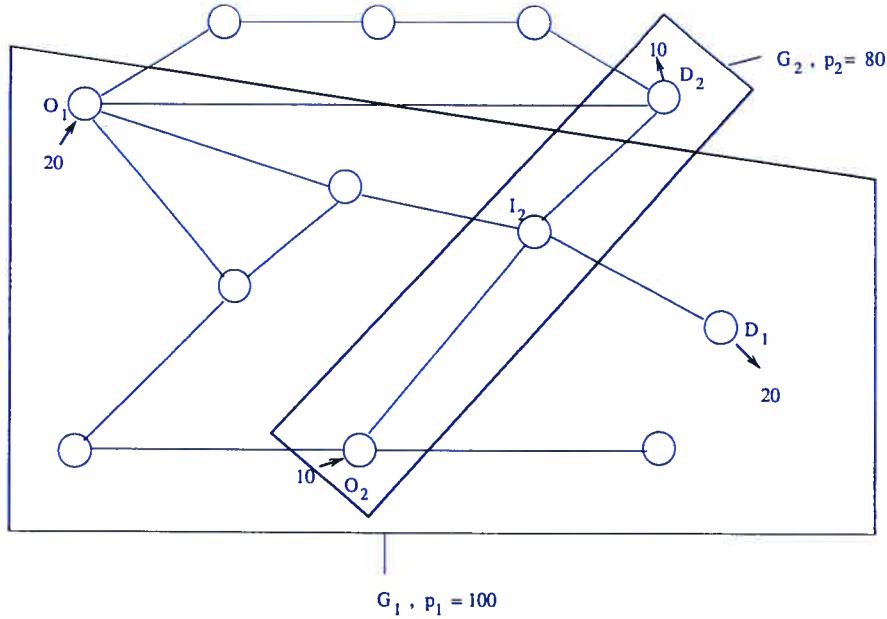


FIG. 3.1 – Combinatorial bids on capacity

The winner determination problem can be formulated as model (M4) :

$$\max \quad \sum_{j \in N} x_j p_j \quad (3.11)$$

$$s.t. \quad \sum_{k \in K_j} h_k = c_j x_j, \forall j \in N \quad (3.12)$$

$$\sum_{j \in N} \sum_{k \in K_j} \delta_{a,k} h_k \leq v_a, \forall a \in A \quad (3.13)$$

$$x_j \in \{0, 1\}, \forall k \in K_j, \forall j \in N \quad (3.14)$$

$$h_k \geq 0, \forall k \in K_j, \forall j \in N \quad (3.15)$$

where $\delta_{a,k} = 1$ if $a \in k$, 0 otherwise. Constraints (3.12) state that the capacity allocated to a winning bid must be within the bid's sub-network, while constraints (3.13) correspond to capacity availability on links.

When paths are completely specified by buyers (G_j is limited to a single path between O_j and D_j , $\forall j \in N$), and single units of capacity are available on links ($v_a = 1, \forall a \in A$) and requested by buyers ($c_j = 1, \forall j \in N$), model (M4) is equivalent to model (M1), in which items are links and bundles are paths. The particularity of

model ($M4$) lies in the fact that buyers do not need in general to indicate a specific path along which the capacity should be allocated. It is up to the auctioneer to assume the additional task of routing the requested capacities between the origins and the destinations in order to determine the winning bids. Model ($M4$) could of course be solved directly by a commercial MIP solver. However, significant gains in computational efficiency may probably be obtained if one exploits the remark that the LP relaxation of ($M4$) can be formulated as a multicommodity network flow problem. Efficient specialized algorithms (Ahuja, Magnanti, and Orlin [4]) can therefore be used instead of plain simplex in a Branch-and-Bound procedure, for example.

3.2.4 Combinatorial exchanges

Combinatorial exchanges refer to many-to-many combinatorial auctions, in which there are many sellers and many buyers. A participant in this category of auctions may submit bids $b = (\{q_{b,i}\}_{i \in G}, p_b)$ where $q_{b,i}$ is a quantity of item i to trade in bid b ($q_{b,i} > 0$ in case of a buy, and $q_{b,i} < 0$ in case of a sell), and p_b is a price that the participant is ready to pay $p_b > 0$ or asks to receive $p_b < 0$. If bids are indivisible, i.e., the whole bundles $\{q_{b,i}\}_{i \in G}$ are traded, or nothing at all, then denote by G the set of all bids and define $x_b = 1$ if bid b wins, 0 otherwise. The winner determination problem is formulated as model ($M5 - a$) :

$$\max \quad \sum_{b \in B} p_b x_b \quad (3.16)$$

$$s.t. \quad \sum_{b \in B} q_{b,i} x_b \leq 0, \forall i \in G \quad (3.17)$$

$$x_b \in \{0, 1\}, \forall b \in B \quad (3.18)$$

Model ($M5 - a$) maximizes the total surplus of the market under the constraint that sales should cover buys. Notice that inequalities in constraints (3.17) assume free disposal by the market maker of any extra quantity of items supplied in the market, and must be changed to equalities if that assumption cannot be made.

When bids are divisible, let decision variable x_b , $b \in B$ designate the *execution*

proportion of bid b , and $p_b(x_b)$ the price the participant is ready to pay or receive if proportion x_b of bid b is executed. The allocation problem can be formulated as model $(M5 - b)$:

$$\max \quad \sum_{b \in B} p_b(x_b) \quad (3.19)$$

$$\text{s.t.} \quad \sum_{b \in B} q_{b,i} x_b \leq 0, \forall i \in G \quad (3.20)$$

$$0 \leq x_b \leq 1, \forall b \in B \quad (3.21)$$

While model $(M5 - b)$ is generally easy to solve, especially when the price mappings $p_b(\cdot)$ are linear, model $(M5 - a)$ remains NP-complete, since the one-to-many indivisible case corresponding to model $(M1)$ may be seen as a particular instance of combinatorial exchanges with indivisible bids. Sandholm and Suri [129] suggest the BOB algorithm, in which they adapt various search techniques previously suggested for the one-to-many allocation model $(M1)$. Hybrid clearing models for exchanges, in which some bids can be subdivided while others cannot, have also been considered in the literature. Thus, Kothari, Sandholm, and Suri [77] consider combinatorial exchanges with bundle bids including only sell or buy components, and show that when a few bids (no more than the number of commodities traded) may be partially executed, there is no integrality gap between the corresponding market clearing formulation and its LP relaxation.

Applications of combinatorial exchanges have also been suggested for trading assets in financial markets (e.g., Fan, Stallaert, and Whinston [52]), supply chain formation and coordination (Walsh, Wellman, and Ygge [142]), and market clearing in process industries (Kalagnanam, Davenport, and Lee [69]). The latter is particularly interesting, as the model considered by the authors considers supply and demand bids on single products, but with various levels of quality. The fact that the model tolerates substitution between products having different levels of quality give rise to additional constraints on possible matchings between sell and buy orders. The authors note also that whether or not it is possible to consolidate several sell orders to satisfy a buy order is crucial, as the allocation problem can be modeled as a maximum

flow problem when consolidation is tolerated, whereas it corresponds to an NP-hard generalized assignment problem otherwise.

3.2.5 Conclusion

In this section, we have presented a few basic formulations of the winner determination problem in combinatorial auctions. These formulations are important from a mechanism design perspective because they may serve as starting points for modeling more complex settings. Moreover, they provide insights on the computational complexity of more elaborate market clearing algorithms.

Real-world markets often require, however, that designers of combinatorial auction mechanisms extend the basic formulations by addressing a certain number of additional issues. Thus, in many important markets, participants do not limit their bid definition to desired bundles of items and prices to pay or receive, but may also bid on other attributes, such as quality of service, delivery times, requirements on technology, etc. Handling these requirements may sometimes be achieved through bid re-weighting schemes that take into account the additional attributes in the winner determination objective (Sandholm and Suri [128]). The most common approach nevertheless consists in adding side constraints to the basic formulations. Side constraints present the advantage of encompassing both market requirements derived from business practices (e.g., guarantee a minimal market share to a given group of participants) and constraints formulated by participants when complex bidding languages are used to express bids. They may, however, significantly increase the complexity of the corresponding market clearing formulations. A comprehensive compilation of generic classes of side constraints for combinatorial markets, and examination of their impacts on the complexity of winner determination formulations can be found in (Sandholm and Suri [128]).

3.3 Expression of combined bids

The framework that describes how bids are defined in a combinatorial market should be sufficiently powerful to allow the representation of the preferences and objectives of the various participants. From a market design perspective, it should

be also flexible and general so that one does not need to invent a new formalism for every new application. In this section, we survey the existing literature on bidding languages, and briefly present a new unified bidding framework for combinatorial auctions of divisible and indivisible items recently introduced.

The definition of bidding languages is also closely related to issues relative to the user interfaces and how easy it is for auction participants to enter their bids. The study of these questions is, however, outside the scope of the current paper.

3.3.1 Motivation and state of the art

By submitting combined bids that consist of the specification of a bundle of items and an associated price, a participant actually could, at least in theory, reflect accurately its preference for any subset of items. Yet, this can be difficult and costly in practice. Consider for instance a combinatorial freight exchange in which shippers submit orders to move loads between different locations and carriers bid from the execution of these orders. In order to permit an optimal usage by of the transportation resources available to the carriers, the exchange allows the latter to consolidate several individual loads and submit package bids on complete routes. Suppose now that, at a given stage of the auction, a (small) carrier with only one available truck is interested in (and able to) service loads in five different bundles A, B, C, D , and E . With no mean to express succinctly the condition that it could serve *anyone*, but *only one* of these bundles, the carrier will have to enumerate explicitly all subsets of $\{A, B, C, D, E\}$, evaluate them, then bid accordingly.

Concise expression of such queries through an appropriate “logic” has thus naturally motivated the first bidding languages proposed in the literature. Hence, in Fujushima, Leyton-Brown, and Shoham [55], as well as in Sandholm [130], one may find the expression of exclusive OR (XOR) conditions through the usage *dummy goods*. These are items with no value to bidders, and intended only to enforce exclusion in the execution of the corresponding bids. For instance, our carrier may define a dummy load l , construct bundles $A \cup \{l\}, B \cup \{l\}, \dots, E \cup \{l\}$, and submit five unrelated bids on these new bundles. Hoos and Boutilier [66] suggest one of the first combinatorial bidding languages in which different logics are combined. More specifically, Hoos and

Boutilier define (1) *clauses* as subsets of items such that a bidder formulating a clause expresses its willingness to obtain any number of items in the clause ; and (2) *bids* as sets of clauses along with a price, such that the bidder requires all the clauses of a bid to be satisfied by an allocation of the items and declares its willingness to pay the associated price in that case. The resulting \mathcal{L}_{CA}^{cnf} bidding language can thus be seen as a two-level logical formalism in which an OR logic governs the clause level, while an AND logic applies at the bid level. The authors introduce also a slightly more general language (\mathcal{L}_{CA}^{k-of}), in which a selection operator takes place of the conjunctive logic.

As far as we know, Nisan [105] is the first successful effort to systematically *analyze* a bidding language. Hence, the author defines two important concepts : (1) the *expressiveness* of a bidding language, which is a measure of the language's ability to express concisely bids that are consistent with (support) a certain family of bidder valuation functions ; and (2) its *simplicity*, which indicates how easy it is, for the bidders and the auctioneer, to understand and use the language. Additionally, Nisan formally defines and analyzes seven bidding languages :

- Atomic bids. In this language, the simplest possible in the combinatorial bidding world, a bidder may only submit a single bid $b = (S, p)$, where $S \subseteq G$ and p is the price the bidder is willing to pay for S . Obviously, this language provides very little expressiveness since even additive preferences are not supported.
- OR-bids, XOR-bids, OR-of-XORs, XOR-of-ORs, OR/XOR-formulae. These languages correspond to the application of the OR and XOR logics on the atomic bids.
- The OR* language. This language is simply a variation of the OR-bids language in which “dummy” bids can be used to express disjunction (in basically the same way that Fujushima *et al.* and Sandholm previously suggested). Surprisingly, the OR* language is provably more expressive than both the OR-of-XORs and the XOR-of-ORs languages.

Boutilier and Hoos [23] is an attempt to generalize the prior combinatorial bidding languages by focusing on the semantics of prices. Thus, while in Nisan's language the emphasis is on logical conditions (in the sense that prices are only relevant at the

atomic bid level), the bidding framework suggested by Boutilier and Hoos allows to associate prices at any level of the logical formulae associated with a combined bid. More specifically, three bidding operators are introduced : \wedge , \vee , and \oplus . The semantics of the language can be summarized as follows. A bid can basically take the form $b = \langle \{i\}, p \rangle$, where $i \in G$ is a single item and p is a price the bidder is willing to pay if it obtains item i . Otherwise, if b_1 and b_2 denote two combined bids formulated in the language, with respective price valuations p_1 and p_2 , and $\mathcal{X} \in \{\wedge, \vee, \oplus\}$. A combined bid $b = \langle b_1 \mathcal{X} b_2, p \rangle$ has the following interpretation, dependent of operator \mathcal{X} .

1. If $\mathcal{X} = \wedge$, the bidder expresses its willingness to execute bids b_1 and b_2 for their corresponding price valuations, *and* to pay a “premium” of p if both bids are executed.
2. If $\mathcal{X} = \vee$, the bidder requires that (i) b_1 is executed for $p_1 + p$; (ii) b_1 is executed for $p_2 + p$; or (iii) b_1 and b_2 are executed for $p_1 + p_2 + p$.
3. If $\mathcal{X} = \oplus$, the bidder expresses that it is willing to pay $\max(p_1, p_2) + p$ if b_1 , b_2 , or both of them are executed.

Obviously, the operator \wedge is intended to express bid complementarity, while \vee and \oplus reflect two different forms of bid substitutability. Actually, Boutilier and Hoos argue that this language has the potential to represent any utility function and to express certain bids more succinctly than prior bidding languages (in particular, Nisan’s OR* language).

3.3.2 A new bidding framework

A major limitation of the bidding languages proposed so far is that they apply only to combinatorial auctions of indivisible goods. It is legitimate to think that, as important markets trading commodities that are intrinsically divisible (e.g., electricity power, telecommunication capacity) or can be safely be considered as divisible (assets in financial markets), a comprehensive and unified bidding framework, which would encompass both the divisible and the indivisible cases, would prove much more appropriate.

The bidding framework we propose relies on a two-level representation of a bid. Physical items traded in the market constitute the framework's elementary ingredients. At the lower, *inner* level, we define the *atomic bid* as a sell or buy request of a quantity q of a given item, along with a price valuation p . In the divisible case, the atomic bid can be “subdivided” into arbitrarily small fractions and its *execution* within a trade that is acceptable to the participant means essentially that a positive proportion of the quantity q is traded; otherwise, in the indivisible case, the whole quantity q should be traded for the atomic bid to be executed.

Partial bids are then introduced at the inner level to formalize the combination of atomic bids and, in the divisible case, the expression of conditions related to their traded proportions. Hence, a partial bid refers to a collection of atomic bids and relies on a *bidding operator* that contains information on the execution conditions. For instance, a partial bid can be used to express the following request : “I desire to sell up to 40 units of item r_1 at \$100 and to buy up to 20 units of item r_3 at \$90. Moreover, I want *equal proportions* of these orders to be traded”, by combining atomic bids corresponding to the buy and sell orders with an **EQUAL** bidding operator. A partial bid is *executed* if all the conditions included in its associated bidding operator are satisfied.

The *outer* level of the framework is mainly concerned with providing means to define and express *logical* conditions related to the execution of partial bids. At the outer level, the most important concept is that of the *combined bid*, which is basically a collection of partial bids that are combined with the help of a logical bidding operator. Hence, it would be possible, for instance, to formulate a bidding requirement such as “Execute Partial Bid 1 *or* Partial Bid 2, *but not both of them*” by the means of a combined bid containing references to Partial Bid 1 and Partial Bid 2, and a bidding operator **XOR** representing the exclusive OR execution condition. It is of course possible, just like in other combinatorial languages for indivisible items, to define more complex bidding requirements that involve logical expressions, or formulae, by allowing for the recursive application of a few basic logical operators in the expression of the combined bid. Thus, a final combined bid that carries all the relevant bidding information should be *submitted* by each participant in the auction.

The full description of the bidding framework can be found in Abrache *et al.* [1]. In

the following, we present a brief survey of the important concepts of the framework.

The inner level

Let G be the set of items traded in the market and L the set of participants.

Definition 3.1 (Atomic bid) An atomic bid is a 4-tuple $\delta = (r, q, b, p)$ where

- $r \in G$ is a reference to an item ;
- q is the maximum quantity of item r to be traded in δ ;
- b is a lower bound on the execution proportion x of atomic bid δ , which means that the participant asks for the execution of at least the proportion b of the maximum quantity q ;
- p is a price valuation related to δ .

The interpretation of the quantity q and the price valuation p depends on the divisibility of the atomic bid. In the indivisible case, the whole quantity q of item r should be traded in δ , or nothing at all, and p indicates a price the participant may pay or receive if the q units are traded. Whereas in the divisible case, an atomic bid can be subdivided into arbitrarily small portions. Quantity q is therefore interpreted as the maximum quantity of item r to be traded in δ , and an *execution proportion* $x \in [0, 1]$ may be associated to atomic bid δ to indicate that a quantity xq of item r is traded in δ . Accordingly, price valuation p corresponds in this case to a price mapping defined on $[0, b]$ such that $p(x)$ is the price the participant may pay or receive if a proportion x of atomic bid δ is executed. We say that atomic bid $\delta = (r, q, b, p)$ is *executed* in a trade if the lower bound condition $x \geq b$ is satisfied by the outcome of the trade.

Partial bids combine atomic bids and formulate conditions related to their execution proportions. A partial bid may be defined as follows :

Definition 3.2 (Partial bid) Let \mathcal{A}_l be the set of atomic bids of participant l . A partial bid θ_i formulated by participant l may take one of the two following forms :

1. $\theta_i = \delta_h, h \in \mathcal{A}_l$ (θ_i is an atomic bid) ;
2. $\theta_i = (\Delta_i, \mathcal{X}_i, p_i)$ where
 - $\Delta_i = \{\delta_k\}_{k \in K_i}, K_i \subseteq \mathcal{A}_l$ is a subset of atomic bids defined by participant l ;
 - \mathcal{X}_i is a **bidding operator** of the inner level applied to Δ_i ;

- p_i is a price valuation related to θ_i .

A bidding operator \mathcal{X}_i of the inner level can be associated to a *condition subset* $E_{\mathcal{X}_i} \subseteq [0, 1]^{|K_i|}$ which represents analytically, as a mathematical set of constraints, the execution conditions of the operator. Partial bid θ_i is *executed* if the vector $x^i = \{x_k\}_{k \in K_i}$ of execution proportions of atomic bids in Δ_i is in the condition subset $E_{\mathcal{X}_i}$. If partial bid θ_i is not executed, then no atomic bid in Δ_i should be executed.

Providing participants in actual combinatorial auctions with an adequate set of operators that have a well-defined “meaning”, equally understood by the participants and the auctioneer, constitutes an extremely important design step. In Abrache *et al.* [1], we introduce some important classes of inner-level bidding operators. We may classify these operators into three categories :

- *Composition* operators express directly conditions on the execution proportions of atomic bids. An example of a composition operator is the **EQUAL** operator, which expresses the requirement that equal proportions of atomic bids in a partial bid θ_i are executed when θ_i is executed, and corresponds to the condition subset

$$E_{\mathcal{X}_i} = \{\{x_k\}_{k \in K_i} \in [0, 1]^{|K_i|} : x_{k_1} = x_{k_2}, \forall k_1, k_2 \in K_i\}.$$

- The *selection* operator specifies constraints on the *number* of atomic bids to be executed. Let $\theta_i = (\Delta_i, \mathcal{X}_i, p_i)$ be a partial bid, where $\Delta_i = \{\delta_k\}_{k \in K_i}$, and denote by $\Pi_i = \{k \in K_i : \delta_k \text{ is executed}\}$ the set of atomic bids in Δ_i that are executed in the trade. Consider the logical operator

$$S_{k^l, k^u}(\Delta_i) = \begin{cases} 1 & \text{if } k^l \leq |\Pi_i| \leq k^u, \\ 0 & \text{otherwise,} \end{cases}$$

where k^l and k^u are integer parameters such that $0 \leq k^l \leq k^u \leq |K_i|$. The associated selection operator **SELECT-INNTER** expresses the condition that no less than k_l and no more than k_u atomic bids should be executed when θ_i is executed, and corresponds to the condition subset

$$E_{\mathcal{X}_i} = \{\{x_k\}_{k \in K_i} \in [0, 1]^{|K_i|} : S_{k^l, k^u}(\Delta_i) = 1\}.$$

- *Hybrid* operators combine functions of composition and selection operators. More precisely, a hybrid operator consists of composition constraints that should be applied only to the atomic bids that are selected to be executed by a selection constraint. We may then define, for example, the **SELECT-INNER + EQUAL** operator as follows :

$$E_{\mathcal{X}_i} = \{\{x_k\}_{k \in K_i} \in [0, 1]^{|K_i|} : S_{k^l, k^u}(\Delta_i) = 1; x_{k_1} = x_{k_2}, \forall k_1, k_2 \in \Pi_i\};$$

The outer level

The following recursive definition of a combined bid allows for the definition of bid execution constraints that correspond to complex logical formulae.

Definition 3.3 (*Combined bid*) Let I_l be the set of partial bids defined by participant $l, l \in L$.

A combined bid Θ_j that participant l formulates can take one of the two following forms :

1. $\Theta_j = \theta_i, i \in I_l$ (Θ_j is a partial bid);

2. $\Theta_j = (\Omega_j, \mathcal{X}_j, p_j)$ where

- $\Omega_j = \{\Theta_j\}_{j \in J_j}$ is a subset of other previously defined **combined bids** formulated by participant l , J_j being the index set of these combined bids;
- \mathcal{X}_j is a **logical bidding operator** of the outer level applied to Θ_j ;
- p_j is a price valuation related to Θ_j .

We suggest the selection operator as our bidding operator of choice at the outer level. Let us consider combined bid $\Theta_j = (\Omega_j, \mathcal{X}_j, p_j)$, where $\Omega_j = \{\Theta_j\}_{j \in J_j}$. Denote by $\Psi_j = \{\bar{j} \in J_j : \Theta_{\bar{j}} \text{ is executed}\}$ the set of combined bids in the expression of Θ_j that are executed when Θ_j is executed in the trade. The outer level selection operator **SELECT-OUTER** corresponds to the following logical operator

$$S_{N^l, N^u}(\Omega_j) = \begin{cases} 1 & \text{if } N^l \leq |\Psi_j| \leq N^u, \\ 0 & \text{otherwise.} \end{cases}$$

Here N^l and N^u are integer parameters such that $0 \leq N^l \leq N^u \leq |J_j|$. In this case, the **SELECT-OUTER** operator indicates that no less than N^l and no more than N^u combined bids in Ω_j have to be executed, should combined bid Θ_j be executed. Otherwise, if the selection condition is not satisfied, then no combined bid in Ω_j should be executed.

It is noteworthy that the usual logical operators AND, OR, and XOR are in fact special cases of the **SELECT-OUTER** operator : if Θ_1 and Θ_2 are two combined bids, then $\Theta_1 \text{ AND } \Theta_2 \equiv S_{2,2}(\{\Theta_1, \Theta_2\})$, $\Theta_1 \text{ OR } \Theta_2 \equiv S_{1,2}(\{\Theta_1, \Theta_2\})$, and $\Theta_1 \text{ XOR } \Theta_2 \equiv S_{1,1}(\{\Theta_1, \Theta_2\})$.

In a complex bidding framework, price semantics have considerable importance and should be clarified. Among the important questions related to prices that need to be addressed are the following : what do prices specified at the atomic, partial and combined bid level mean ? which ones of these values are relevant ? can we have conflicting prices ? In Abrache *et al.* [1], we precise the meaning of prices and propose general-purpose (and in that sense minimal) conditions that need to be verified to ensure that the pricing information submitted by a participant is *complete* (i.e., the auctioneer would always be able to determine, whatever the allocation of items, the payment a bidder is ready to make or receive) and *coherent* (i.e., prices specified by a bidder in its bids are not conflicting with each other).

3.3.3 Impact on the allocation problem

We illustrate the impact of bidding languages on market clearing formulations by considering a simple application in financial markets. In many contexts, traders need to submit *bundle orders* to simultaneously sell and buy different assets, along with prices they are willing to pay or receive if the orders are executed. This is notably the case when they rebalance their portfolios at the end of a trading session. After receiving all the trade orders, the market maker determines the executed proportions of each order and payments the traders should make or receive such that total surplus of the market is maximized. A bundle order j defined by trader l is basically a vector $O_{lj} = (\{q_{ljr}\}_{r \in G}, p_{lj})$ where :

- q_{ljr} is the maximum number of units of asset r that may be traded in order j ;
 $q_{ljr} > 0$ corresponds to a buy, $q_{ljr} < 0$ to a sell, and $q_{ljr} = 0$ if asset r is not traded in order j ;
- p_{lj} is the bundle price the trader is willing to make or receive if order j is entirely executed.

Let J_l be the set of bundle orders formulated by trader l , $l \in L$, and define the primary decision variables :

- x_{lj} = the traded proportion of bundle order j formulated by trader l .

The basic formulation of the market clearing problem can be expressed as the following LP model :

$$\max \quad \sum_{l \in L} \sum_{j \in J_l} p_{lj} x_{lj} \quad (3.22)$$

$$s.t. \quad \sum_{l \in L} \sum_{j \in J_l} q_{ljr} x_{lj} = 0, \quad r \in G \quad (3.23)$$

$$0 \leq x_{lj} \leq 1, \quad l \in L, j \in J_l \quad (3.24)$$

Among the many additional bidding requirements traders may formulate in such markets, we focus on : a) lower bounds on the executed proportion of orders, which indicate that traders prefer an order not to be executed at all unless a minimal execution proportion is guaranteed (trading small volumes may sometimes be non profitable if there are transaction fees to pay) ; and b) XOR relations between certain orders, such that at most one of these orders may be executed (may indicate, for example, that the trader consider the orders as “equivalent”, but is averse to the fragmentation of its portfolio).

Let us associate a lower bound b_{lj} to a bundle order j defined by trader l and define \mathcal{X}_l as the set of all XOR relations defined by trader l , where $\mathcal{X} \in \mathcal{X}_l$ is a subset bundles in J_l such that at most one order in \mathcal{X} should be executed. Consider the auxiliary decision variables :

- $y_{lj} = 1$ if order j formulated by trader l is executed, = 0 otherwise.

When lower bounds constraints and XOR relations are taken into account, market clearing corresponds to the following MIP formulation :

$$\max \quad \sum_{l \in L} \sum_{j \in J_l} p_{lj} x_{lj} \quad (3.25)$$

$$s.t. \quad \sum_{l \in L} \sum_{j \in J_l} q_{ljr} x_{lj} = 0, \quad r \in G \quad (3.26)$$

$$b_{lj} y_{lj} \leq x_{lj} \leq y_{lj}, \quad l \in L, j \in J_l \quad (3.27)$$

$$\sum_{j \in \mathcal{X}} y_{lj} \leq 1, \quad \mathcal{X} \in \mathcal{X}_l, l \in L \quad (3.28)$$

Hence, even the most elementary bidding operators can significantly increase the complexity of market clearing formulations. A numerical investigation of the impact on economic surplus and computational complexity of lower bound and XOR operators in the context of bundle trading of financial assets is presented in Abrache, Crainic, and Gendreau [3].

3.4 Iterative combinatorial auctions

In a number of settings, knowing how to write bids and determine the winning allocation and prices is sufficient. Single-round, sealed-bid auctions are a case in point. Simply put, participants submit all their bids more or less simultaneously, and the auctioneer determines the winning bid by simply identifying the “best” one with respect to pre-defined rules. The fact that bids are sealed implies that, in general, a bidder will have no information about other participants’ behavior and, consequently, will derive its bidding strategy from incomplete and abstract (i.e., not related to the current auction) assessment of competition, as well as from its own valuation of the items on the market. Such a market is inefficient (in an economic sense of the term) due, in particular, to a lack of information concerning the cost and utility functions of the market participants.

Assuming such information is available, one could build a model to determine optimal allocations and prices. To illustrate, consider an idealized multi-commodity, multi-lateral market (Bourbeau *et al.* [22]). Participants, which are sellers and buyers

of products, communicate all the relevant information about their production costs and demand functions, respectively. The market maker also requires the participants to reveal complete information about the transportation costs between sellers and buyers, as well as all the technological constraints related to the production and consumption of the products. It then solves a large non-linear market clearing model to identify an allocation and a set of equilibrium prices such that the total social efficiency is maximized.

Situations in which bidders hand over to market makers complete and truthful information are very rare, however. The participants are generally *unwilling*, and sometimes even *unable*, to disclose all the relevant information required by the auction mechanism to optimize the market. In this, they may be motivated by several considerations :

- Information confidentiality. Given that participants are often self-interested agents, they are generally reluctant to disclose proprietary data, even to an electronic agent representing the market maker.
- Uncertainty in the valuation of items. In some contexts, the value of items or bundles of items is not known with certainty to the participants and only estimates of the actual valuations can be communicated to the auctioneer (e.g., oil and gas lease auctions ; see Oren and Williams [109]). In some other contexts, that information is imprecise (art auctions, for instance) and needs to be adjusted according to what competitors actually bid.
- Complexity of evaluating and communicating preferences. Especially when the number of items on the market is large and the bidding requirements of participants are complex, evaluation by participants of their own preferences can become a hard task. Moreover, a communication “bottleneck” exists (Nisan [106]) and implies that optimal outcomes in combinatorial auctions cannot be achieved, in the general case, with sub-exponential data communication.

Iterative auctions (Cramton [35]) alleviate some of these concerns since they require significantly less *a priori* information and allow participants to progressively reveal their private information by altering their selling or buying offers in light of the market information and their own assessment of the market. The idea of iterative

auctions (Figure 3.2) is the following. In each auction round, participants submit bids on bundle of items. These bids do not need to represent the complete and definitive needs of participants, nor to convey, *a priori*, truthful information. Given the bids, the auctioneer uses a market-clearing process to determine a set of provisional allocations and payments. Information – signals – related to the temporary state of the market, and intended to incite the participants to commit themselves further in the auction, is then returned to them. Consequently, in the following rounds, the participants may alter their bids or make new ones, according to the signals received from the market and to their bidding strategies. The process continues until a stopping criterion is met (e.g., no new bids or bid updates are submitted in a given round), and the outcome of the auction becomes a final one. Bid changes from one round to the next are often governed by activity rules whose function is to give impetus to the market by prompting participants to be active and reveal their real needs as early as possible.

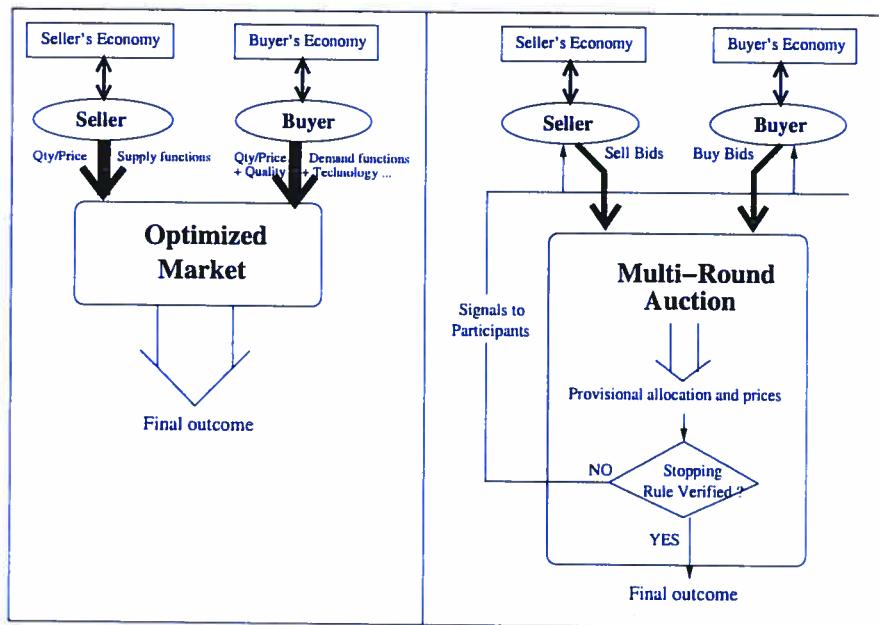


FIG. 3.2 – Direct revelation mechanisms vs multi-round auctions

Many “classical” auction mechanisms, such as the ascending English and the descending Dutch auctions, are actually iterative auctions. These are well known and understood. The design of iterative combinatorial auctions gives rise to many complex problems, however. In the following, we focus on three particularly important

aspects : the design of auction rules, the pricing schemes in price-directed auctions, and the incentive-compatibility properties of an iterative combinatorial auction.

3.4.1 Design of auction rules

A primary objective of the market maker is to move the auction at a steady pace, ensuring that the participants progressively, but actively, commit themselves, and that the whole process eventually converges to allocations and prices close to what would have been obtained in the “ideal” case of an optimized, centralized market. Several rules must be set in order to achieve this goal :

- Admissibility rules. These rules govern the way participants update their bids as the auction goes on. For the most part, they consist of constraints on the composition of bids (a bidder, for instance, may only bid on increasingly larger packages), or on price offers (bidders should bid incrementally on each bundle of items, for example).
- Activity rules. In order to give impetus to the market, participants should express their real needs reasonably early in the auction. Hence, the mechanism should prevent, for instance, participants from simply observing the market, or making infinitesimal modifications to current bids, and waiting for the final stages of the auction to submit “jump” bids in an attempt to throw everybody else out of the auction.
- Stopping rules. These rules specify criteria according to which the process ends. Examples of these are predefined numbers of rounds, predefined auction times, and the absence of significant bidding activity. Note that a stopping rule can be rather complex and may consist, for instance, in the combination of several simpler stopping rules.
can be made of the composition of several such stopping criteria.

It is noteworthy that the complexity of real-world applications often requires the auctioneer to do a finer subdivision of an iterative auction mechanism, first in *phases*, then in rounds. A phase can be defined as a sequence of rounds intended to reach an important intermediary step of the auction. Each phase may consequently be characterized by its own rules and “mechanism”, and produce a provisional outcome

supplied in input to the next phase. To illustrate the concept of multi-phase auctions, let us briefly consider an iterative procurement auction for transportation services designed for a large mining company. In the auction, the mining company acts as the buyer, while sellers are carriers that provide the transportation services. Carriers bid on long-term contracts on transportation *lots* (a lot designates a required transportation capacity between two locations). A two-phase prototype mechanism has been suggested. The first phase is composed of a single round, in which the carriers submit bids on *individual* lots. The purpose of this preliminary phase is to “heat” the market and to gently introduce the carriers into the bidding process (bidding decisions are relatively simple here since no combination of lots is permitted). The second phase is a multi-round process, in which the carriers are allowed to bid on *bundles* of lots (routes). The auctioneer maintains prices on individual lots that are disclosed at the end of each round. Provisional winners are notified individually. To be admissible, a new bundle bid needs to beat a provisional winner by at least a given threshold.

3.4.2 Pricing

The nature of the information disclosed to participants at intermediary stages of the auction is a central issue in the design of multi-round auction mechanisms. An important class of such mechanisms are *price-directed* iterative auctions, in which that information is primarily related to prices of items or bundles of items. While price-directed iterative auctions can easily be designed and implemented in the simple case of single-item bidding auctions, deriving prices in combinatorial settings is much more challenging.

The divisible case (commodities or bids are divisible) is encompassed by the theory of general equilibrium in exchange economies. To ease the presentation, we limit ourselves to the one-sided, one-to-many case. Let G be a set of m divisible goods, and J a set of buyers. The seller has an endowment $M = [M_1, \dots, M_m]$ of the goods. Each buyer j has a preference $v_j(x_j)$ for bundle $x_j = [x_{j,1}, \dots, x_{j,m}]$. A *socially-efficient* allocation x^* is an allocation that solves $\max_{x \in \mathcal{D}} \sum_{j \in J} v_j(x_j)$, where \mathcal{D} is the set of all feasible allocations of goods to buyers. *Walrasian* equilibrium prices that

support the efficient allocation are single-item prices $\{p_i\}_{i \in G}$, such that x^* maximizes the payoff of each buyer ; that is

$$v_j(x_j^*) - p_j \cdot x_j^* = \max_{x \in \mathcal{D}} \{v_j(x_j) - p_j \cdot x_j\}$$

with the usual quasi-linear utilities assumption.

A classical result of the general equilibrium theory (Arrow and Debreu [6]) establishes that Walrasian equilibrium prices exist under conditions of continuity, monotony, and concavity of preference functions $v_j(\cdot)$. Reaching an equilibrium, when it exists, through a Walrasian *tâtonnement* process is however dependent of whether or not that equilibrium is *stable*. In that regard, the economic literature notoriously identifies the *gross substitutes* (GS) property, which essentially states that the demand for a given good does not decrease when individual prices of other goods increase, as a sufficient condition for the stability of equilibria (see Arrow and Hahn [7], for example).

The literature of iterative auction design in presence of indivisibilities has mainly focused on the Combinatorial Allocation Problem (CAP) (de Vries and Vohra [39]). The CAP has the same settings as the basic winner determination formulation (*M1*), but seeks to maximize the overall social efficiency of the market, rather than the revenue of the seller given buyer bids. So, with the notation of subsection 3.2.1 and $v_j(S)$ defined as the preference of buyer j for getting bundle $S \subseteq G$, a basic formulation of the CAP can be written as model (CAP) :

$$\max \quad \sum_{1 \leq j \leq n} \sum_{S \subseteq G} v_j(S) x_{j,S} \quad (3.29)$$

$$s.t. \quad \sum_{1 \leq j \leq n} \sum_{S \subseteq G} \delta_{i,S} x_{j,S} \leq 1, \forall i \in G \quad (3.30)$$

$$\sum_{S \subseteq G} x_{j,S} \leq 1, \forall j, 1 \leq j \leq n \quad (3.31)$$

$$x_{j,S} \in \{0, 1\}, \forall S \subseteq G, \forall j, 1 \leq j \leq n \quad (3.32)$$

If one does not take into account the integrality gap that may exist between model (CAP) and its LP relaxation, equilibrium single-item prices can be derived from the

dual of the LP relaxation (Bikhchandani and Mamer [19]) : optimal solutions $\{p^*_i\}_{i \in G}$ and $\{\tau^*_j\}_{1 \leq j \leq n}$ of

$$\min \quad \sum_{1 \leq j \leq n} \tau_j + \sum_{i \in G} p_i \quad (3.33)$$

$$s.t. \quad \sum_{i \in S} p_i + \tau_j \geq v_j(S), \forall j \in N, \forall S \subseteq G \quad (3.34)$$

$$p_i \geq 0, \forall i \in G \quad (3.35)$$

$$\tau_j \geq 0, \forall j, 1 \leq j \leq n \quad (3.36)$$

can be interpreted in this case as Walrasian equilibrium prices and optimal payoffs of participants, respectively.

The existence of Walrasian equilibrium prices (and therefore the integrality of the LP relaxation of (CAP)) requires stronger conditions in the indivisible case. Notably, Gul and Stacchetti [62] show that the GS property is a sufficient one, which implies, essentially, that linear (single-item) prices may not exist when there are complementarities between items. Bikhchandani and Ostroy [20] propose two extended formulations of the CAP. While these formulations have many more variables and constraints than model (CAP), they are generally stronger and duals of their LP relaxations provide, respectively, *anonymous* bundle prices (all buyers pay the same price for a bundle), and *discriminatory* bundle prices (what a buyer pays for a bundle depends on its identity). Interestingly, the strongest formulation has an integral LP relaxation, which means it is always possible to compute discriminatory bundle prices that support an efficient allocation of the CAP.

Researchers have recently considered embedding Bikhchandani and Ostroy's formulations in primal-dual frameworks. The iBundle family of ascending-price auctions (Parkes [111]) is an example of such approach. The iBundle mechanisms assume that participants are self-interested price-taker buyers that react myopically to prices by bidding on bundles giving them the most payoff at these prices, and manage to reach a competitive equilibrium (an equilibrium that maximizes also the revenue of the seller) by carefully increasing prices on over-demanded bundles.

Many other iterative auctions based on different price-adjustment schemes have been suggested. For instance, the experimental RAD mechanism of DeMartini *et al.* [41] announces, at the end of each round, single-item prices to the participants. These prices are “approximated” equilibrium prices that minimize the violation of complementary slackness. To be admissible in the next round, new bids by the participants need to beat the current provisional prices by a certain increment. Wurman and Wellman [149] prove the existence of anonymous bundle equilibrium prices and give a procedure to compute them. Their proof is constructive and proceeds in two steps. Once a provisional optimal allocation has been determined on the basis of bids submitted by the participants, prices for the assigned bundles are computed using the dual of the corresponding assignment problem (see Leonard [88]). Given that this problem has generally multiple solutions, the authors suggest rather two auxiliary problems that provide optimal prices minimizing the revenue of the auctioneer and maximizing the surplus of the participants, respectively. Then, prices of unassigned bundles are set such that no participant is distracted from the provisional optimal allocation. Equilibrium prices computed this way are not competitive equilibrium prices, though, in the sense that they do not guarantee the auctioneer of getting the highest possible revenue.

A somewhat different but promising line of research consists in the adaptation of mathematical decomposition approaches, especially price-driven ones (Dantzig-Wolfe, Lagrangian relaxation and decomposition). These methods have been used for decades to tackle large-scale optimization of problems with special structure, but there have been very few efforts to take profit of their potential for decentralized decision making to design corresponding iterative auction mechanisms (see for instance Kutanoglu and Wu [80] for an application to distributed job shop scheduling).

3.4.3 Incentive-compatibility issues

Up to this point, the reader may judiciously ask : how could the auctioneer determine the socially-efficient allocation if the participants do not accept to reveal (progressively or in one shot) their valuations without misreporting them ? Actually, this question brings out a focal aspect of an auction mechanism, which is its ability to

provide the right incentives for participants to bid *truthfully*. Regarding this issue, a mechanism is said to be *strategy-proof* if it is a dominant strategy for any participant to report its true valuations, whatever the strategies adopted by the other participants. Strategy-proofness, when it can be achieved, is indeed a very powerful property since it means that participants will confine themselves to the simplest strategy available to them (which is to report truthfully their private types), being assured that doing so is in their best interest.

The Vickrey-Clarke-Groves auction (VCG) (Vickrey [140]; Clarke [31]; Groves [61]) is known to be an economically efficient, strategy-proof mechanism. Given preferences $\{\tilde{v}_j(\cdot)\}_{j \in J}$ reported by participants, the VCG's allocation and payment rules are :

- Return an allocation $x^* \in \arg \max_{x \in \mathcal{D}} \sum_{j \in J} \tilde{v}_j(x_j)$ that maximizes total value given the reported valuations.
- Participant j pays $V_{-j}^* - (V^* - \tilde{v}_j(x_j^*))$, where $V^* = \max_{x \in \mathcal{D}} \sum_{j \in J} \tilde{v}_j(x_j)$ and $V_{-j}^* = \max_{x \in \mathcal{D}} \sum_{k \in J - \{j\}} \tilde{v}_k(x_k)$ (a participant receives a “discount” on its reported value equal to the economic impact of its presence in the market).

A serious limitation of the VCG auction lies in the fact that it is a sealed-bid mechanism that requires complete information about participants' preferences to be revealed to the auctioneer. This fact has motivated the design of iterative incentive-compatible auctions, that would end up with the same outcome as the direct-revelation VCG mechanism. Among the most important developments recently reported, we may cite Gul and Stacchetti [63] who show that, under the GS condition, a simple tâtonnement process that generalizes the English auction leads to the smallest Walrasian prices, which in turn correspond to the Vickrey-Clarke-Groves payments with further restrictions on the GS preferences. Bikhchandani *et al.* [18] give a primal-dual interpretation to Gul and Stacchetti's auction. Ausubel [9] suggests an iterative implementation of the VCG with GS preferences. Ausubel's mechanism requires however to run $|J| + 1$ parallel auctions in order to compute the Vickrey-Clarke-Groves payments. The Extend & Adjust iterative mechanism (Parkes [112]; Parkes and Ungar [116]) computes the Vickrey-Clarke-Groves payments through a two-phase process. In the first phase, an iBundle ascending-price auction determines

an efficient allocation and competitive equilibrium prices. The second phase collects “just enough” additional information from participants to compute Vickrey discounts.

The simple fact of being able to derive Vickrey-Clarke-Groves payments in iterative auction mechanisms does not necessarily mean, however, that it is always *desirable* to do so. Indeed, Vickrey auctions suffer from many other shortcomings (Rothkopf, Teisberg, and Kahn [126]), such as their sensitivity to collusion and cheating, and the fact that they do not guarantee the budget-balance of the market and may give a seller a marginally small revenue. The latter stands out since it can be shown that for the important case of exchanges (even non combinatorial ones), budget-balance may not be achieved. The design of alternative auction mechanisms in this case is an interesting avenue of research recently investigated by Parkes, Kalagnanam, and Eso [114].

3.5 Participant decision problems

Auction participants need, of course, to construct initial bids and to modify them (both their composition and the associated price offers) as the multi-round process goes on. Yet, this is not the only problem they face. To be able to decide on profitable bid strategies, participants have to analyze complex information disclosed by the market mechanism and combine it with their business processes : internal cost policies, current operations and activities, knowledge of the economic sector and competition, etc. Thus, there is need to develop optimization-based decision support tools - *advisors* - to help participants tackle these decisions.

To illustrate the types and role of advisors, consider applications to electronic freight marketplaces (Chang, Crainic, and Gendreau [28]; Figliozi, Mahmassani, and Jaillet [53]). Participants are shippers (production firms, freight forwarders, etc.) that need commodities (for simplicity, assume full load trailers or containers) moved between various locations, and motor carriers bidding for the loads. In designing their bidding strategies, carriers are faced with several questions : (1) on which loads to bid ? (2) when to bid ? and (3) at what prices ? Decisions have to be coherent with the current and forecast fleet deployment and demand. It is also noteworthy that

particular groups of loads may present a special interest for a given carrier when, for example, they may be blocked into a route performed by a single driver or they allow to bring home a driver and its empty vehicle. In this context, advisors are software agents that assist carriers in making “profitable” bidding decisions, by processing the information available in the market, and realize its integration into the dynamic planning of transportation operations.

A major difference between advisors and classical decision support systems is that while the latter have to interact only with the planning methods and data of the respective firm, the former have also to deal with the many forms of marketplaces encountered on the Internet. Consequently, other than the particular transportation sector in which they will evolve (truckload, less-than-truckload, container, a combination of the three, etc.) advisors may be classified according to their response to the following characteristics :

1. Market type. Advisors may be developed for *single* or *multiple* marketplaces. In the latter, carriers are interested in loads appearing on several different marketplaces. Indeed, while marketplaces are independent of each other, loads are often interdependent for carriers (e.g., to form an interesting route loads have to be negotiated on different marketplaces). In this case, advisors have the additional burden of coordinating the carriers’ bidding activities in the different marketplaces. Advisors of this type have been proposed in the literature for very simple multi-market negotiations Benyoucef *et al.* [14], but no known multi-market advisors exist for more complex settings.
2. Auction type. Auctions can be *single* or *multi-round*, *continuous* or *periodic*, and may involve bidding on *independent single* loads, or on bundles (*combinatorial bidding*).
3. Integration with the planning process. Advisors can be *remotely coupled* to the (dynamic) planning of operations, or *tightly coupled* to it. In the first case, the advisor rely generally on predetermined lists of available vehicles. On the other hand, tightly coupled advisors need to interact, at regular intervals with the (dynamic) fleet management process in order to evaluate loads to bid on. The length of these intervals depends on the response time of the model, as

well as on the delays tolerated by the auction and the carrier trade-off between profitability and risk of loosing loads.

In computer science terms, the advisor (or agent) *ménagerie* is even more diverse and complex. We may mention, for example, that while *planning* advisors, such as the ones described above, may be used to select loads on which to bid and determine the corresponding pricing data, *negotiators* are required to actually conduct the bidding. The complexity of the negotiation strategy, as well as its call to planning advisors during the auction, depend largely on the market characteristics and the time available for computations.

The development of dynamic advisors, for freight as for other types of markets, is thus a key design issue, especially critical for a successful deployment of combinatorial market designs.

3.6 Conclusion

Since the very first attempts to use combinatorial auctions for the allocation of heterogeneous commodities, there has been increasing awareness that the design of this class of auctions is a complex and multi-faceted problem. While the early literature has naturally started by addressing the winner determination problem, that proved to be only the beginning. Thus, a remarkable multidisciplinary effort has been initiated to investigate original issues raised by combinatorial bidding (bidding languages, for instance), as well as to rethink some other well-known problems in auction theory and practice (such as incentives) that become particularly difficult when package bidding is allowed.

In this paper, we gathered and discussed a few interesting issues in the design of combinatorial auctions. The first of these is the classification of combinatorial auctions and the corresponding formulations of the winner determination problem. We put the emphasis on five important models representing direct and reverse one-sided combinatorial auctions, auctions of network resources, and combinatorial exchanges. These generic models provides preliminary insights on the complexity of clearing the market. However, the auction designer should be concerned with additional attributes and side constraints that follow from real-world applications, and which may

complicate significantly the market-clearing formulations.

We next discussed the need for high-level bidding languages that give participants the means to express succinctly their bidding requirements. We presented a novel formalism that goes a step beyond the existing languages of the literature by allowing combined bid formulation for divisible and indivisible commodities. The impact of using a bidding language on the formulation of the allocation problem is illustrated through a simple application in finance. This analysis suggests that dealing with expressive bidding formalisms can be challenging for both the auctioneer, which has to handle potentially large market-clearing MIP formulations, and the participants, which need to figure out how to construct and update bids that are coherent with their business processes and their knowledge of the market and their competitors. It also points out the importance, in practice, of reaching an acceptable trade-off between the expressiveness of the bidding language and its simplicity of use.

The design of iterative incentive-compatible combinatorial auctions, which give participants the impetus to progressively reveal truthful information about their preferences, has been a key objective of recent research on mechanism design. Although important breakthroughs have been reported for the simplest combinatorial auction settings (in particular for the CAP), much work is still needed to extend the results to more complicated (and useful) auction models that take into account high-level bidding languages and market side constraints.

Acknowledgements This project has received financial support from the Bell Canada University Laboratory (LUB), the Natural Sciences and Engineering Research Council of Canada through its Basic Research and Collaborative Industry-University Research grant programs, and the FCAR research development fund of the Province of Québec.

Chapitre 4

A New Bidding Framework for Combinatorial E-Auctions

La référence de cet article est :

J. Abrache, B. Bourbeau, T.G. Crainic, and M. Gendreau, “A New Bidding Framework for Combinatorial E-Auctions”, à paraître dans *Computers and OR*.

A New Bidding Framework for Combinatorial E-Auctions

Jawad Abrache^{1,3}, Benoît Bourbeau⁴, Teodor Gabriel Crainic^{2,3,4}, and Michel Gendreau^{1,3,4}

¹ Département d'informatique et de recherche opérationnelle, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada, H3C 3J7.

² Département management et technologie, Université du Québec à Montréal, CP 6192, Succursale Centre-ville, Montréal, Canada H3C 4R2.

³ Centre de recherche sur les transports, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada, H3C 3J7.

⁴ CIRANO, 2020 University, Montréal, Canada, H3A 2A5.

Abstract. Bidding languages define the means through which participants in an electronic auction define bids and express requirements on their execution. The current state of combinatorial auction market design indicates that no existing bidding language is general enough to support auctions of both divisible and indivisible commodities. In this paper, we propose a novel bidding framework based on a two-level representation of a combined bid. At the inner level, bidding operators impose conditions on the executed proportions of packages of atomic single-item bids. Partial bids defined this way are then recursively combined through logical operators to produce a final combined bid that is submitted to the auctioneer. We present a formal specifi-

cation of the framework, and analyze how it impacts the mathematical programming formulation of the allocation problem. An application in the context of combinatorial auctions of financial assets illustrates the utilization of the proposed bidding framework.

Keywords : Auction design, Combinatorial e-auctions, Bidding models, Financial markets

4.1 Introduction

In many sectors of activity, sellers and buyers need to interact in order to exchange goods and services. Examples of such structured trading occur in many important contexts, including public and private procurement, trade of financial assets, allocation of telecommunication or transportation services, capacity and right-of-way allocation, privatization of public assets, etc. With the emergence of e-commerce platforms and tools, and the recent deregulation wave within several industries, overwhelming evidence starts to accumulate regarding the significant social and economic impact of structured electronic markets.

Multi-lateral multicommodity markets constitute an important model for the trade of heterogeneous items. The aim of multi-lateral markets is to streamline and accelerate the finding of partners willing to buy and sell items, the establishment of many-to-many exchange associations between them, and the determination of quantities traded and associated prices that are acceptable to the traders. In many multilateral markets, called *optimized*, or *smart* markets, a market maker, acting as a mediator, receives information that specify desired quantities and acceptable prices from traders, and must design a *market-clearing mechanism*. The latter consists of an optimization model and method that determine an allocation of items and corresponding payments, such that the traded supply and demand quantities are matched, and a given market objective is achieved.

An optimized market is essentially an “idealized” exchange model that can be hardly implemented directly, even in strongly regulated environments (an example from the province of Québec wood chip industry is described in [22]). In practice, traders are generally *unwilling* to disclose the private information, such as their true valuation of items, needed by market-clearing mechanisms. Many traders would consider true information revelation as a competitive edge given away to concurrents, notwithstanding that, in actual facts, this information will only be disclosed to a virtual e-auctioneer. In some particular contexts, traders could even be *unable* to define and communicate that information, for it is too complex [106], or there is uncertainty in its estimation (see [109], for instance). Market mechanisms organized around multi-round auctions somewhat alleviate the problem by allowing progressive revelation of

information. A particularly important class of iterative mechanisms are *price-driven* multi-round auctions, which proceed as follows. In each negotiation round, provisional allocations and prices are determined by the auctioneer on the basis of *bids* submitted by participants in the negotiation, and information related to these temporary results is posted at the end of the round. The participants may then update their bids according to that account of the auction state and their own objectives. The process eventually ends when an appropriate stopping rule is satisfied, and final allocations and prices are announced.

Combinatorial auctions are special cases of multi-item negotiations in which items are traded in *bundles*. Hence, the bids submitted to the auctioneer are combined offers to sell, buy, or simultaneously sell and buy several different items. Combinatorial auctions have received much attention in the literature, with the early works mainly focusing on solving the allocation problem (e.g., [125, 130]). Recent contributions by Abrache, Crainic, and Gendreau [2] and de Vries and Vohra [39] have emphasized, however, that the design of combinatorial auctions is a multi-faceted problem, involving many challenging issues. In this paper, we focus on one of these issues, the definition of bidding languages that allow participants in a combinatorial auction to formulate their bids, to express their trading needs and requirements, and to communicate them to the auctioneer.

The most valid argument in favor of using expressive languages for bidding in combinatorial auctions is the need to express *concisely* the bidders' preferences. Thus, if n indivisible items are traded, a bidder can of course express any preference by submitting a bid on each one of the $2^n - 1$ non empty subsets of items. However, preference functions often display some special structure that makes enumeration of all bundles of items unnecessary. Consider, for instance, a combinatorial market of freight transportation in which carriers bid on bundles of loads to move. A carrier with a single available truck may service one of five bundles of loads A, B, C, D , and E could avoid to formulate bids on all subsets of $\{A, B, C, D, E\}$ if a bidding language that supports the appropriate exclusive OR (XOR) logic was used.

The first bidding languages based on logical operators have been suggested by researchers as part of their efforts to tackle the winner determination problem in

one-sided (single seller, multiple buyers) combinatorial auctions of indivisible goods. Following the basic formulation of the problem [125], in which an inherent OR logic is present, XOR bids are implicitly introduced in [55] (through “dummy” items that force execution exclusiveness among bundle bids), as well as in [130]. The OR and XOR logics are combined by Sandholm [131] in the OR-of-XORs language, and especially by Hoos and Boutilier [66], who introduce two logical bidding languages : the \mathcal{L}_{CA}^{cnf} language, which consists of bids formulated as logical clauses (obtained by recursive application of the conjunctive and disjunctive operators on single items), together with a price the bidder is willing to pay if the logical clause associated to the bid is satisfied by the auction’s allocation of items ; and the \mathcal{L}_{CA}^{k-of} language, an extension of \mathcal{L}_{CA}^{cnf} that, in addition to the basic conjunctive and disjunctive operators, allows bidders to use the selection operator $k - of - n$ (“select k items among n ”).

The first consistent effort to analyze the strengths and limitations of bidding languages for combinatorial auctions is due to Nisan [105], who suggests to evaluate a bidding language from the perspective of its *expressiveness* (the ability to express concisely various preferences) and *simplicity* (how easily bids formulated in the language are conceptualized and dealt with by the bidders and the auctioneer, notably during the market-clearing process). Nisan then evokes a sample of important preference functions to analyze some of the classical bidding languages of the literature, and reaches the conclusion that the OR/XOR formulae and the OR-with-dummy-items (OR*) languages are the most expressive, with the OR* language achieving the best trade off between expressiveness and simplicity.

Recently, Boutilier and Hoos [23] have suggested a new bidding model that embodies all the previously suggested languages. The proposed language is also based on logical formulae involving single items as their basic components but, additionally, bidders are allowed to specify price valuations at any level of the formulae. Price semantics are governed by an additive logic, that is, price valuations corresponding to components (sub-formulae) of a logical formula are summed up with the price valuation to indicate the price the bidder is ready to pay if the logical condition of the formula is fulfilled. This logic is key to a concise expression of certain preferences

displaying complementarities and substitutabilities. The following example from [23] sheds lights on the expressiveness of the language. A shipper that needs to send a load (valued basically at 50), could bid on two standard containers a and b , or on an oversized container c (with a “convenience” premium of 5). A bid $((a \text{ AND } b, 0) \text{ OR } (c, 5))$ conveniently expresses that containers a and b are complements, container c and bundle $\{a, b\}$ are substitutes, and captures the price premium of using the oversized container.

An interesting alternative to bidding languages that has been recently contemplated is *preference elicitation*. Rather than providing participants with tools and mechanisms to express their bids, elicitation-based approaches try to extract relevant information by asking cleverly formulated *questions* related to the participants’ utility functions. More specifically, Conen and Sandholm [32] present a preference elicitation framework for determining welfare-maximizing allocations in basic combinatorial auctions of indivisible items (extended to combinatorial exchanges in [137]). The framework consists of search algorithms that explore the space of feasible allocations, and relies on data structures in which gathered information about the participants’ preferences is stored. During the search process, participants are asked various questions (e.g., the exact or approximate value of a bundle, the preference “rank” of a bundle, etc.) and the provided answers help constrict the search space. Numerical validation in [67] indicates that elicitation schemes are quite effective in reducing information revelation.

The bidding languages proposed so far in the literature have all been formalized for one-sided combinatorial auctions of indivisible single-unit items. While extensions to multi-unit combinatorial auctions and to combinatorial exchanges can easily be made ([131], for instance), a more potent limitation lies in the fact that these languages support the trade of indivisible goods only (The *eAuctionHouse* auction server prototype presented in [131] arguably allows participants to formulate bids as price-quantity graphs, but only in *multi-unit* auctions). With the emergence of important markets trading intrinsically divisible commodities (telecommunications bandwidth, electricity power, raw materials, etc.), or physically indivisible items that may safely assumed to be divisible due to large trading volumes (e.g., assets in financial

markets), it is legitimate to think that a unified bidding framework encompassing both the divisible and indivisible cases would prove appropriate. As a matter of fact, many suggested applications of combinatorial auctions involving divisible commodities, notably in the finance sector, have implicitly used a form or another of bidding languages [138, 51, 3]. Yet no attempt to conceptualize and analyze a bidding language for combinatorial auctions of divisible commodities has been made to date.

The contribution of this paper consists of a new bidding language framework for combinatorial auctions. The framework has several interesting properties. It is intended to be independent of the physical nature of the items traded within the market, and of the divisibility of the items in particular ; it would allow participants to submit complex bidding definitions, requirements, and conditions ; and finally, it is succinct, and avoids asking participants to do explicit enumerations of their preferences for various bundles of items, even for the most complex queries. More precisely, we show that, in the indivisible case, the framework can produce bidding languages that are at least as expressive as the OR/XOR language of Nisan [105]. In a divisible context, we establish that the framework is *fully expressive* (in the sense that any preference can be supported), and many preferences can be expressed in concise manner.

This paper is organized as follows. In Section 4.2, the purpose of a bidding language for combinatorial auctions is defined, and the broad methodological lines of our new bidding framework are discussed. Sections 4.3 and 4.4 present a formal specification of bidding languages within the framework. Section 4.5 investigates the interpretation of the pricing information that participants include in their bids, and proposes general rules and policies intended to make this information complete and consistent. We analyze the expressive power of the framework in Section 4.6. In Section 4.7, we consider the allocation problem and we analyze the impact of the framework of its mathematical programming formulation. Section 4.8 illustrates the use of the language for a specific application in finance. Concluding remarks and perspectives of further research follow in Section 4.9.

4.2 Purpose and methodology

A bidding language may be defined as the set of means by which participants in an auction define their bids, formulate requirements on their execution, and communicate them to the auctioneer. In this paper, we specifically investigate bid definition and associated formulation issues. Concerning the communication aspects, they are more closely related to issues such as human/machine interactions and communication protocols, and consequently are out of the scope of our study.

The “bundle” is a key concept in combinatorial auctions. Basically, a bundle specifies several items that should be traded simultaneously, as a package. Very often, individual items in a bundle are *interdependent*, i.e., a participant’s actual valuation of a given item depends of whether other items in the bundle are also traded or not. More specifically, item interdependency may take two forms : if A and B are two different items, and $v(\cdot)$ is the participant’s valuation function, then A and B are *complementary* of each other if $v(\{A, B\}) > v(\{A\}) + v(\{B\})$, and *substitutable* if $v(\{A, B\}) < v(\{A\}) + v(\{B\})$. For instance, if the items traded are airport take-off and landing rights, a take-off time slot and a landing time slot that correspond to the departure and the arrival airports of a flight complement each other, while two “identical” pairs of take-off and landing slots (e.g., same take-off and landing airports within the same period of time) are likely to be substitutable.

In order to account for item interdependency, a bidding language should allow a bidder not only to define the basic structure of its bundles (i.e., to specify the items that compose each bundle and the corresponding prices it would be ready to pay or receive if the bundle is traded), but also to translate item interdependency relationships into *trade execution conditions* under which the outcomes of the auction may be acceptable to the bidder. While the formulation of these conditions can rely completely on logical connectives in the indivisible case, their characterization is more complex when items or bids are divisible. Let us consider the bidding example of Fig. 4.1 to clarify our point. Bids B1, B2 and B3 represent trade orders to buy 5000 shares of Oracle, buy 1000 shares of Microsoft, and sell 3000 shares of Cisco, respectively, at the specified unit prices. The orders are infinitely divisible so any *proportions* of the

5000, 1000, and 3000 quantities can a priori be traded. The trader nonetheless specifies three execution constraints. Constraint C1 states that a trade must ensure at least 10% of the sell order is executed to be acceptable to the trader. Constraint C2 asks for sell/buy bundles of Microsoft and Cisco such that there is one Microsoft share bought for three Cisco shares sold. Finally constraint C3 translates the trader's aversion to fragmented portfolios by requiring that only Oracle or Microsoft shares might be bought, but not both. While these bidding constraints indeed display fundamentally different structures, they can all be seen as patterns to express concisely the specific structure of the bidder's preference.

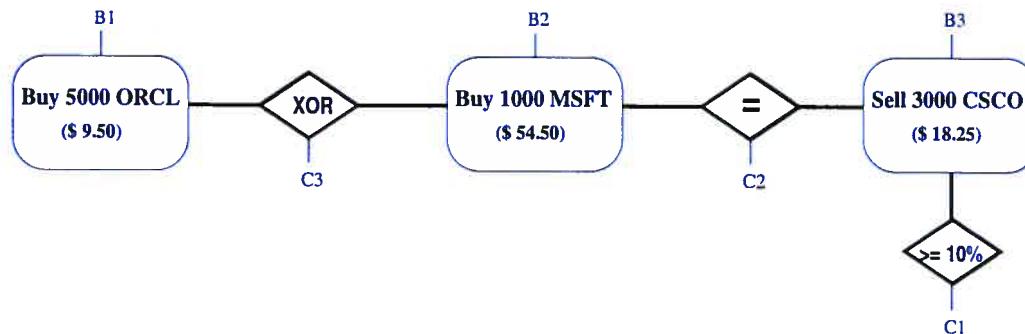


FIG. 4.1 – Trade execution constraints in the divisible case.

We propose a new bidding framework that accounts for bidding requirements in both divisible and indivisible contexts. The framework relies on a two-level representation of a bid, as depicted in Fig. 4.2. The elementary components are the physical items traded in the market. At the lower, *inner* level, we define the *atomic bid* as a sell or buy request of a quantity q of a given item. In the divisible case, the atomic bid can be “subdivided” into arbitrarily small fractions and its *execution* within a trade that is acceptable to the participant means essentially that a positive proportion of the quantity q is traded; otherwise, in the indivisible case, the whole quantity q should be traded for the atomic bid to be executed. Clear distinction should nevertheless be made between *item divisibility*, which is related to the physical nature of items, and *bid divisibility*, which results from modeling choices made by the auctioneer in a given trading context. *Partial bids* are also introduced at the inner level to formalize the combination of atomic bids and, in divisible case, the expression of conditions related to their execution proportions. Therefore, a partial bid refers to a collection

of atomic bids and relies on a *bidding operator* that represents the execution conditions. Fig. 4.2 shows an example of a partial bid, which may be translated into the following request : “I desire to sell up to 40 units of item r_1 at \$100 and to buy up to 20 units of item r_3 at \$90. Moreover, I want *equal proportions* of these orders to be traded”. Partial Bid 1 is then built by combining Atomic Bid 1 and Atomic Bid 2 that represent the single-item sell and buy orders, respectively, and applying the bidding operator **EQUAL** on them. A partial bid is *executed* if all the conditions included in its associated bidding operator are satisfied.

The *outer* level provides means to define and express logical conditions related to the execution of partial bids. At the outer level, we define the important concept of a *combined bid*, which is basically a collection of partial bids that are combined by a logical bidding operator. For instance, Combined Bid 1, in Fig. 4.2, may translate into the statement : “Execute Partial Bid 1 *or* Partial Bid 2, *but not both*”. Recursive application of bidding operators is indeed possible at the outer level and gives rise to logical execution formulae. Each participant formulates a single final combined bid (Combined Bid 2 in Fig. 4.2) that carries all the relevant bidding information and is submitted in the market.

The bidding information formulated by a participant may be represented by a natural graph structure (Fig. 4.3). Nodes of the graph correspond to the atomic bids, partial bids, and combined bids defined by the participant. Arcs indicate bid combinations needed to build partial and combined bids. For instance, arcs $(B1, B)$, $(B2, B), \dots, (Bn, B)$ in Fig. 4.3 mean that bids corresponding to $B1, B2, \dots, Bn$ are combined by the application of a bidding operator to build the bid associated with B . Nodes with no predecessors in the graph, called “leaves”, correspond to atomic bids. The graph also contains a single node with no successors (the “root”) that represents the final combined bid submitted by the participant.

In the next two sections, we present a formal specification of bidding languages in which the principles above are materialized. We give general definitions of atomic bids, partial bids, combined bids, and bidding operators. We also propose specific operators at each level of the language.

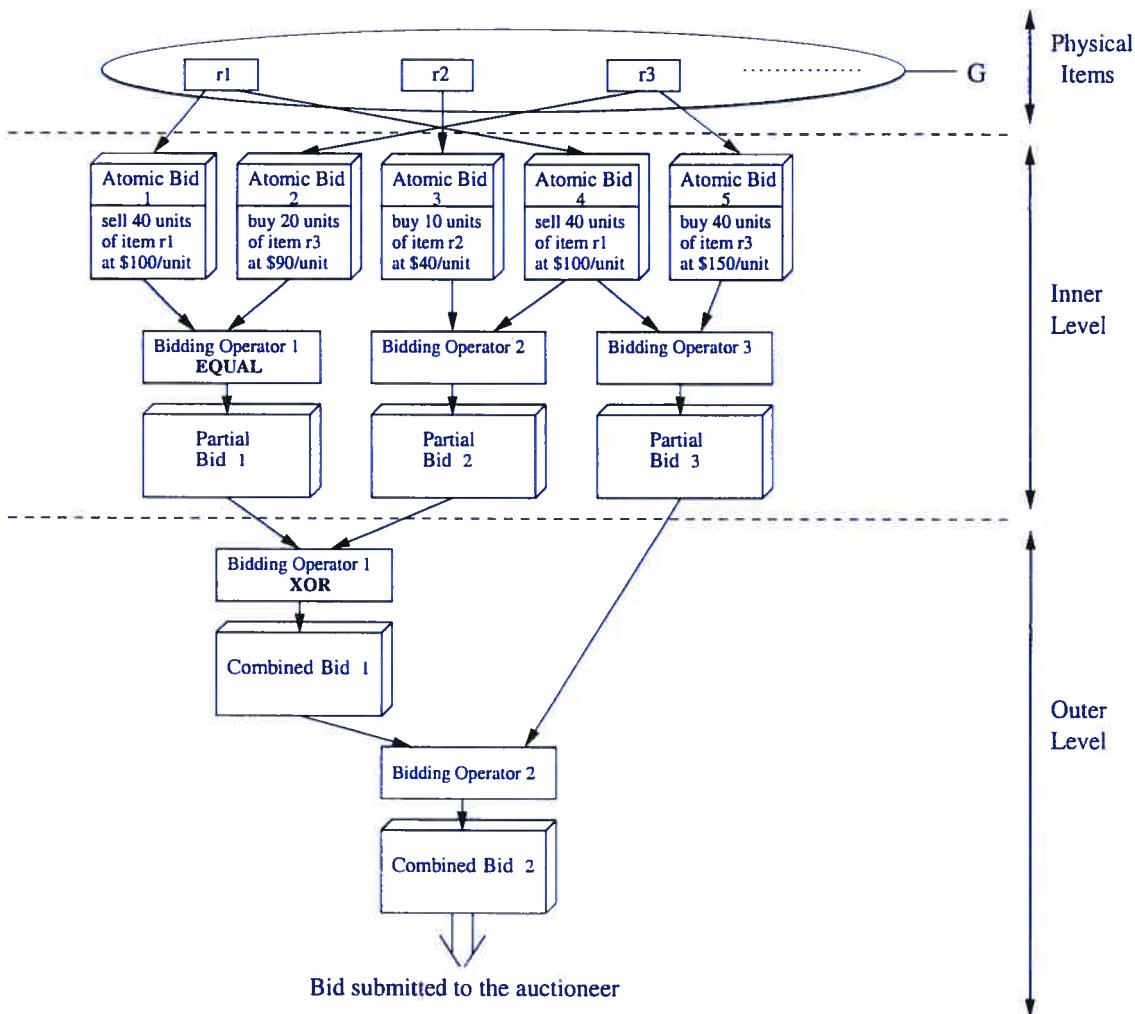


FIG. 4.2 – The two-level bidding framework.

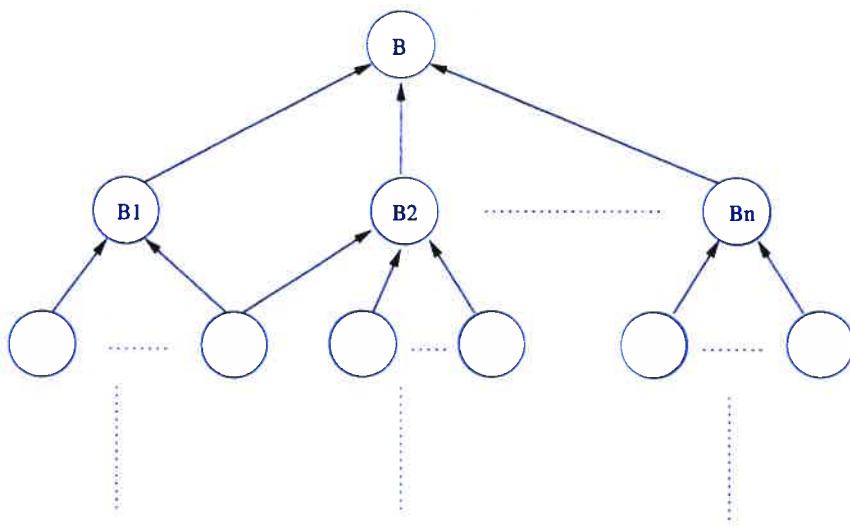


FIG. 4.3 – A graph representation of the bidding structure.

4.3 The inner level

4.3.1 Basic notation and definitions

Let G be the set of items traded in the market and L be the set of participants.

Definition 4.1 (Atomic bid) An atomic bid is a tuple $\delta = (r, \epsilon, q, b, p(.))$ where

- $r \in G$ is a reference to an item;
- $\epsilon = +1$ if δ is a bid to purchase items r , and $= -1$ if δ is a bid to sell items r .
- q is the maximum quantity of item r to be traded in δ ;
- b is a lower bound on the execution proportion x of atomic bid δ . That is, the participant requires at least the proportion b of the maximum quantity q (i.e., $x \geq b$) to be traded; otherwise, nothing should be traded at all ($x = 0$);
- $p(.)$ is a mapping defined on $[b, 1]$ such that $p(x)$ is a price valuation related to atomic bid δ when proportion x of the bid is executed.

The interpretation of the quantity q depends on the divisibility of the atomic bid. If the bid is indivisible, it should be executed entirely and the whole quantity q of item r should be traded, or nothing at all ($b = 1$). In the divisible case, an *execution proportion* $x \in [0, 1]$ may be associated to atomic bid δ to indicate that a quantity xq of item r is traded in δ . We say that atomic bid δ is *executed* if the lower bound condition $x \geq b$ is satisfied by the outcome of the trade. If mapping $p(.)$ defines the

relevant pricing information in δ , then $p(x)$ indicates the price the bidder is ready to pay ($p(x) > 0$) or receive ($p(x) < 0$) if proportion x of atomic bid δ is executed. A special value *NIL* indicates, on the other hand, that no relevant pricing information is specified by the bidder in atomic bid δ .

A partial bid is a concept that allows a participant to combine atomic bids by formulating conditions related to their execution proportions. Partial bids may be defined as follows.

Definition 4.2 (Partial bid) Let \mathcal{A}_l be the set of atomic bids defined by participant l . A partial bid θ_i formulated by participant l may take one of the two following forms :

1. $\theta_i = \delta_h, h \in \mathcal{A}_l$ (θ_i is simply an atomic bid) ;
2. $\theta_i = (\Delta_i, \mathcal{X}_i, p_i(.))$ where
 - $\Delta_i = \{\delta_k\}_{k \in K_i}, K_i \subseteq \mathcal{A}_l$ is a subset of atomic bids defined by participant l ;
 - \mathcal{X}_i is an instance of a **bidding operator** applied to Δ_i ;
 - $p_i(.)$ is a mapping defined on $[0, 1]^{|K_i|}$, such that $p_i(x_1, x_2, \dots, x_{|K_i|})$ is a price valuation related to partial bid θ_i when proportions $x_k, k \in K_i$ of atomic bids $\delta_k, k \in K_i$ are executed.

In general, partial bid θ_i combines atomic bids in Δ_i through bidding operator \mathcal{X}_i , which function is to formulate various constraints on the execution proportions of atomic bids. For example, let $\delta_1 = (r_1, -1, 40, 0, \{x \mapsto -100x\})$ and $\delta_2 = (r_3, +1, 20, 0, \{x \mapsto 90x\})$ represent Atomic Bid 1 and Atomic Bid 2 in Fig. 4.2, respectively, and x_1 and x_2 denote their execution proportions. Bidding Operator 1 (**EQUAL**) applied to δ_1 and δ_2 indicates that the participant would only accept a trade in which proportions x_1 and x_2 are equal. The participant may include a price-mapping $p_i(.)$ in partial bid θ_i . When pricing information is relevant at the partial bid level, $p_i(x_1, x_2, \dots, x_{|K_i|})$ is the price the participant is willing to pay or receive if the constraints corresponding to bidding operator \mathcal{X}_i are satisfied. In relatively simple bidding contexts that do not require the expression of such constraints, the trivial form of partial bids (as atomic bids) allows to circumvent the complexity of the partial bid layer.

Formally, an instance \mathcal{X}_i of an inner-level bidding operator can be associated to a *condition subset* $E_{\mathcal{X}_i} \subseteq [0, 1]^{|K_i|}$, which is an analytical representation (as a set of mathematical constraints) of the execution conditions induced by the bidding operator. The execution of a partial bid may then be defined as follows.

Definition 4.3 Let $\theta_i = (\Delta_i, \mathcal{X}_i, p_i(\cdot))$ be a partial bid, and $x^i = \{x_k\}_{k \in K_i}$ denote the vector of execution proportions of atomic bids in Δ_i . Partial bid θ_i is **executed** if the constraints associated to \mathcal{X}_i are satisfied; that is, if $x^i \in E_{\mathcal{X}_i}$.

In case the bidding operator constraints are not satisfied by the auction outcome, partial bid θ_i is not executed, and no atomic bid in Δ_i should be executed.

4.3.2 Inner-level bidding operators

In practice, a bidding language needs to provide auction participants with a comprehensive set of useful bidding operators. The latter should also be easy to understand and not too burdensome to use by the bidders. An exhaustive enumeration of all these operators is hardly a viable option, however, since each application would probably come with its own lot of operators. In order to remain as context-independent as possible, we rather contemplate a number of bidding operators with high levels of abstraction and practical usability. The operators we propose may be classified from a functional point of view into three categories :

Composition operators

- The proportion ordering operator.

One of the simplest conditions that may reasonably be formulated by a participant are *orderings* on the executed proportions on atomic bids. With these conditions, a participant may formulate a set of atomic bids $\Delta = \{\delta_1, \delta_2, \dots, \delta_n\}$ and indicate to the auctioneer an execution proportion “ranking”, i.e., that the traded proportion x_1 of atomic bid δ_1 should be greater than the traded proportion x_2 of atomic bid δ_2 , which in turn is greater than the traded proportion x_3 of atomic bid δ_3 , and so on. More precisely, let us consider a partial bid θ_i , and

define an ordering \succ_i on the atomic bid index set K_i of θ_i . An instance \mathcal{X}_i of the proportion ordering operator **ORDERING** corresponds to the condition subset

$$E_{\mathcal{X}_i} = \{\{x_k\}_{k \in K_i} \in [0, 1]^{|K_i|} : x_{k_1} \geq x_{k_2}, \forall k_1, k_2 \in K_i \text{ s.t. } k_1 \succ_i k_2\}.$$

- The **EQUAL** operator, which corresponds to the condition subset

$$E_{\mathcal{X}_i} = \{\{x_k\}_{k \in K_i} \in [0, 1]^{|K_i|} : x_{k_1} = x_{k_2}, \forall k_1, k_2 \in K_i\}.$$

An instance of the **EQUAL** operator formulates the requirement that equal proportions of atomic bids in partial bid θ_i should be executed. The **EQUAL** operator may therefore be useful in contexts where items are *perfect complements* of each other. That is, items are worth something to the bidder only if they are traded in precise proportions within bundles.

- The **SIMPLEX** operator, which corresponds to the condition subset

$$E_{\mathcal{X}_i} = \{\{x_k\}_{k \in K_i} \in [0, 1]^{|K_i|} : \sum_{k \in K_i} x_k = 1\}.$$

The **SIMPLEX** operator is motivated by situations in which participants consider atomic bids $\{\delta_i\}_{i \in K_i}$ in partial bid θ_i as “equivalent” to each other. Moreover, they would accept proportions of atomic bids to be executed only if these proportions sum up to one, which means that the combination obtained is equivalent to the entire execution of *any one* of the atomic bids. For example, if $\Delta_i = \{\delta_1, \delta_2\}$, where $\delta_1 = (r_1, \epsilon_1, q_1, b_1, p_1(\cdot))$ and $\delta_2 = (r_2, \epsilon_2, q_2, b_2, p_2(\cdot))$, the application of the **SIMPLEX** operator to Δ_i means the participant would accept one of the following outcomes : a) q_1 units of item r_1 are traded, b) q_2 units of item r_2 are traded, or c) an equivalent “mix” of items r_1 and r_2 , following weights that are equal to q_1 and q_2 , respectively, is traded.

- Quantity operators.

Let us first define the **QTY-EQ** operator. An instance $\mathcal{X}_i = \text{QTY-EQ}(\beta_i)$ of the **QTY-EQ** operator corresponds to the condition subset

$$E_{\mathcal{X}_i} = \{\{x_k\}_{k \in K_i} \in [0, 1]^{|K_i|} : \sum_{k \in K_i} \epsilon_k q_k x_k = \beta_i\},$$

where β_i is a quantity to trade.

The participant indicates its willingness to buy β_i units of *any* items in Δ_i if $\beta_i > 0$, to sell $-\beta_i$ units if $\beta_i < 0$, or to realize a balanced trade if $\beta_i = 0$. The QTY-EQ operator would be useful in circumstances where participants consider *items* as “equivalent”, but are indifferent toward these items actually sold or purchased when the partial bid is executed as long as the desired quantity β_i is traded.

Inequality variants of the QTY-EQ operator are defined similarly. Hence, instances of the **QTY-MORE** and **QTY-LESS** operators can be associated respectively to condition subsets

$$E_{\mathcal{X}_i} = \{\{x_k\}_{k \in K_i} \in [0, 1]^{|K_i|} : \sum_{k \in K_i} \epsilon_k q_k x_k \geq \beta_i\}$$

and

$$E_{\mathcal{X}_i} = \{\{x_k\}_{k \in K_i} \in [0, 1]^{|K_i|} : \sum_{k \in K_i} \epsilon_k q_k x_k \leq \beta_i\}$$

– Price operators.

Price operators express bid execution conditions in terms of constraints on the corresponding prices specified in atomic bids. In formal terms, a price operator \mathcal{X}_i can generally be associated to the condition subset

$$E_{\mathcal{X}_i} = \{\{x_k\}_{k \in K_i} \in [0, 1]^{|K_i|} : \{p_1(x_1), p_2(x_2), \dots, p_{|K_i|}(x_{|K_i|})\} \in \mathcal{P}_i\},$$

where $\mathcal{P}_i \subseteq \mathbb{R}^{|K_i|}$ is a set of admissible price vectors.

Although many different price operators can be derived, we only consider one of the most useful, the **BUDGET** operator. An instance $\mathcal{X}_i = \text{BUDGET}(\mathcal{B}_i)$ corresponds to

$$E_{\mathcal{X}_i} = \{\{x_k\}_{k \in K_i} \in [0, 1]^{|K_i|} : \sum_{k \in K_i} p_k(x_k) \leq \mathcal{B}_i\},$$

where $\mathcal{B}_i > 0$, and indicates that the bidder does not want to exceed the budget limit \mathcal{B}_i in its overall spending on executed proportions of atomic bids $\{\delta_k\}_{k \in K_i}$.

The selection operator

The inner-level selection operator, as the name suggests, allows participants to specify the *number* of atomic bids in a partial bid that should be executed when the partial bid is executed.

More precisely, let $\theta_i = (\Delta_i, \mathcal{X}_i, p_i(\cdot))$ be a partial bid, where $\Delta_i = \{\delta_k\}_{k \in K_i}$. Denote by $\Pi_i = \{k \in K_i : \delta_k \text{ is executed}\}$ the set of atomic bids in Δ_i that are executed in the trade, and consider the logical operator

$$S_{k^{\min}, k^{\max}}(\Delta_i) = \begin{cases} 1 & \text{if } k^{\min} \leq |\Pi_i| \leq k^{\max}, \\ 0 & \text{otherwise,} \end{cases}$$

where k^{\min} and k^{\max} are integer parameters such that $0 \leq k^{\min} \leq k^{\max} \leq |K_i|$. An instance $\mathcal{X}_i = \text{SELECT-INNER}(k^{\min}, k^{\max})$ of the selection operator corresponds to the condition subset

$$E_{\mathcal{X}_i} = \{\{x_k\}_{k \in K_i} \in [0, 1]^{|K_i|} : S_{k^{\min}, k^{\max}}(\Delta_i) = 1\}.$$

Simply put, the selection operator \mathcal{X}_i expresses the condition that no less than k^{\min} atomic bids and no more than k^{\max} atomic bids should be executed when θ_i is executed.

Hybrid operators

Hybrid operators combine functions of the selection operator and composition operators. Let \mathcal{X}_1 be an instance of the selection operator, and \mathcal{X}_2 an instance of a composition operator. Consider a partial bid $\theta_i = (\Delta_i, \mathcal{X}_i, p_i(\cdot))$, where $\Delta_i = \{\delta_k\}_{k \in K_i}$, and $\mathcal{X}_i = \mathcal{X}_1 + \mathcal{X}_2$ denotes an hybrid operator. The action of \mathcal{X}_i consists in : 1) selecting atomic bids in Δ_i for execution, according to operator \mathcal{X}_1 ; and 2) applying execution constraints of operator \mathcal{X}_2 on the *selected atomic bids*. In more formal terms, the condition subset $E_{\mathcal{X}_i}$ is the *projection* of the composition operator condition subset $E_{\mathcal{X}_2}$ on the selection operator condition subset $E_{\mathcal{X}_1}$. For example, the condition subset corresponding to a **SELECT-INNER**(k^{\min}, k^{\max}) + **EQUAL** hybrid operator (Fig. 4.4) corresponds to

$$E_{\mathcal{X}_i} = \{\{x_k\}_{k \in K_i} \in [0, 1]^{|K_i|} : S_{k^{\min}, k^{\max}}(\Delta_i) = 1; x_{k_1} = x_{k_2}, \forall k_1, k_2 \in \Pi_i\}.$$

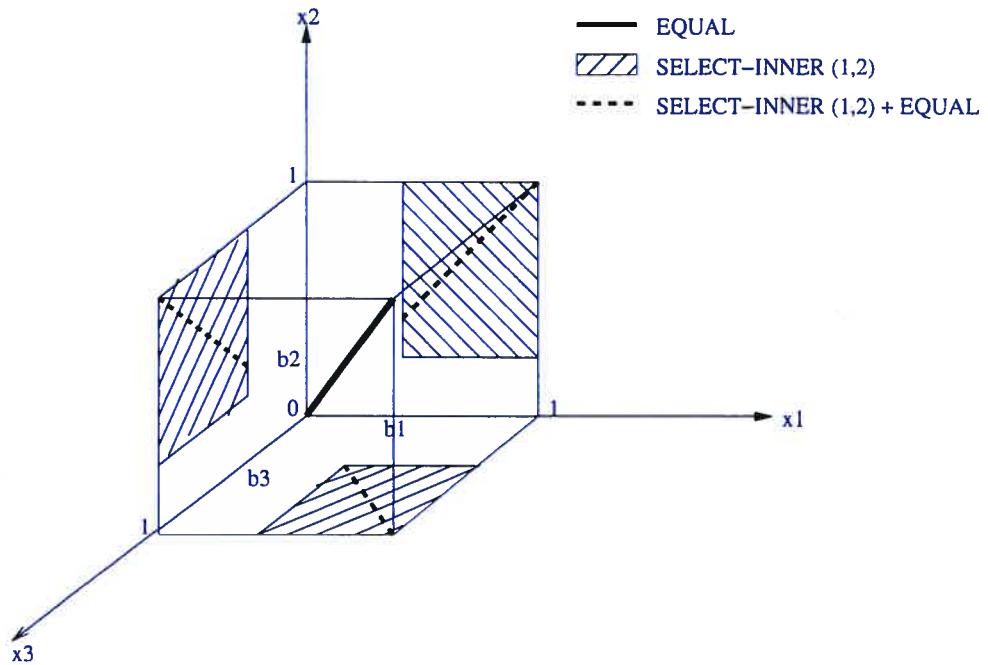


FIG. 4.4 – Condition subsets corresponding to instances of the EQUAL, SELECT-INNER, and SELECT-INNER + EQUAL operators.

When the projection of $E_{\mathcal{X}_2}$ on $E_{\mathcal{X}_1}$ coincides with the intersection of the two subspaces, the definition of an hybrid operator is obviously superfluous. This is the case when \mathcal{X}_2 is the SIMPLEX operator or one of the three quantity operators.

Let us summarize the main elements of the inner level. Atomic bids define simple requests to sell or buy items. In the indivisible case, they cannot be subdivided and may therefore be considered as trivial partial bids. In the divisible case, the participants may use bidding operators to combine atomic bids and set requirements on their execution proportions.

4.4 The outer level

We proceed now to specify the outer level. The main concern will be the formulation of bidding constraints involving partial bids defined in terms of logical expressions. By way of illustration, consider two partial bids θ_1 and θ_2 and the usual logical operators **AND**, **OR**, and **XOR**. Associated execution conditions are

- θ_1 AND θ_2 : both θ_1 and θ_2 should be executed ;
- θ_1 OR θ_2 : one of the two partial bids θ_1 and θ_2 should be executed ;

- $\theta_1 \text{ XOR } \theta_2 : \theta_1$ or θ_2 should be executed, but not both.

It is indeed possible to define more complex logical constructions involving recursive application of logical operators. For instance, let θ_1 , θ_2 and θ_3 be partial bids and consider the logical expression $((\theta_1 \text{ AND } \theta_2) \text{ XOR } \theta_3)$. $\Theta = \theta_1 \text{ AND } \theta_2$ may be seen as an intermediary bid formulated by the participant that is executed if both θ_1 and θ_2 are executed in the trade. Θ is then combined with θ_3 in an exclusive execution association through the XOR operator. The concept of the *combined bid* formalizes the recursive construction of logical expressions.

Definition 4.4 (Combined bids) Let I_l be the set of partial bids defined by participant l , $l \in L$.

A combined bid Θ_j that participant l formulates may take one of the two following forms :

1. $\Theta_j = \theta_i$, $i \in I_l$ (Θ_j is a partial bid);
2. $\Theta_j = (\Omega_j, \mathcal{X}_j, p_j)$ where
 - $\Omega_j = \{\Theta_j\}_{j \in J_j}$ is a subset of other previously defined **combined bids** formulated by participant l , J_j being the index set of these combined bids;
 - \mathcal{X}_j is an instance of a **logical bidding operator** applied to Θ_j ;
 - p_j is a price valuation related to Θ_j .

At its most simple, a combined bid is a partial bid. Otherwise, it arises from the application of an outer-level bidding operator to a set of other combined bids. We allow a price valuation p_j to be specified in combined bid Θ_j . When this valuation is relevant, it indicates the price that the bidder is willing to pay or receive if the combined bid is executed. The final combined bid formulated by the participant is the bid that is submitted to the auctioneer. Hence, let us define

- J_l : the set of all combined bids formulated by participant l , $l \in L$; and
- $\Theta_l \in J_l$: the combined bid submitted by participant l , $l \in L$.

At the outer level, the logical structure of bidding constraints naturally suggests the use of a selection operator. Hence, consider combined bid $\Theta_j = (\Omega_j, \mathcal{X}_j, p_j)$, where $\Omega_j = \{\Theta_j\}_{j \in J_j}$, and denote by $\Psi_j = \{\bar{j} \in J_j : \Theta_{\bar{j}} \text{ is executed}\}$ the set of combined

bids in the expression of Θ_j that are executed when Θ_j is executed in the trade. An instance $\mathcal{X}_j = \text{SELECT-OUTER}(N^{\min}, N^{\max})$ of the outer-level selection operator corresponds to the following logical operator

$$S_{N^{\min}, N^{\max}}(\Omega_j) = \begin{cases} 1 & \text{if } N^{\min} \leq |\Psi_j| \leq N^{\max}, \\ 0 & \text{otherwise.} \end{cases}$$

Here N^{\min} and N^{\max} are integer parameters such that $0 \leq N^{\min} \leq N^{\max} \leq |J_j|$. The \mathcal{X}_j operator indicates that no less than N^{\min} combined bids in Ω_j and no more than N^{\max} combined bids have to be executed, should combined bid Θ_j be executed. Otherwise, if the selection condition is not satisfied, then no combined bid in Ω_j should be executed.

Since the usual logical operators are indeed special cases of the selection operator ($\text{AND} \equiv \text{SELECT-OUTER}(|J_j|, |J_j|)$, $\text{OR} \equiv \text{SELECT-OUTER}(1, |J_j|)$, and $\text{XOR} \equiv \text{SELECT-OUTER}(1, 1)$), the SELECT-OUTER operator is therefore, strictly speaking, sufficient to express any logical formula. For the sake of a simpler notation, we will nevertheless continue to make use of operators AND, OR, and XOR in the remainder of the paper.

4.5 Price consistency

The pricing information that participants include in their bids is a fundamental issue of any bidding framework. It is the medium through which bidders specify their willingness to pay or receive certain amounts of money when the bids they submit are executed in a trade. They may also use it to simply convey indications that “signal” to the auctioneer the worth of items and bundles to them. In a complex bidding language, pricing information raises, however, fundamental difficulties related to its interpretation by the bidders and the auctioneer.

To illustrate, consider the following two bidding situations. In the first example, the participant bids to buy a bundle made of 1 unit of item r_3 and 1 unit of either item r_1 or item r_2 . The bid is represented in Fig. 4.5. The participant also indicates, in the expression of Combined Bid 2, that it is willing to pay \$100 if its bid is executed, that is, if it gets the bundle. In that case, that bundle price is consistent with what

the participant really wants - to get the entire bundle - and hence is the only pricing information the participant needs to submit. Any atomic bid prices, for instance, would have little significance for the participant here.

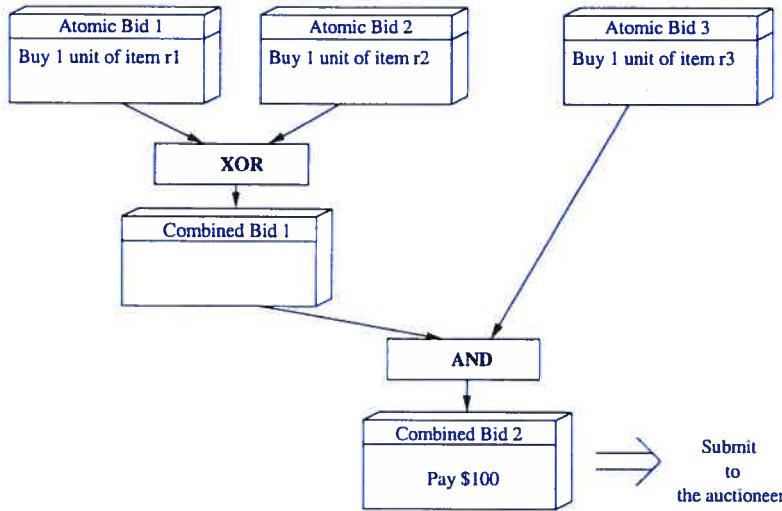


FIG. 4.5 – Bidding Example 1.

Consider now the second example presented in Fig. 4.6. The participant formulates *divisible* atomic bids to buy up to 10 units of r_1 , buy up to 20 units of r_2 , and sell up to 10 units of r_3 . It also requires the proportions of Atomic Bid 1 and Atomic Bid 2 that are actually executed to sum up to 1, and at least 40% of Atomic Bid 3 to be executed. Concerning prices, the participant indicates only single-item prices in the expression of each atomic bid. Indeed, it does not make much sense here for the participant to specify a bundle price in Combined Bid 1, since the OR operator indicates that it would, for instance, accept a trade in which all the atomic bids are executed, as well as a trade in which Atomic Bid 3 only is executed, and these outcomes are very unlikely to have the same “worth” to the participant. Moreover, the divisibility of the atomic bids is another issue that makes a single bundle price of little significance. Thus, the price the bidder would pay or receive is dependent on the prices specified at the atomic bid level and on the logic defined by the bidding operators.

These bidding situations correspond to extremal cases, in the sense that prices are indicated either at the atomic bid level, or in the expression of the final combined bid submitted by the participant. It is not difficult, however, to think of other situations

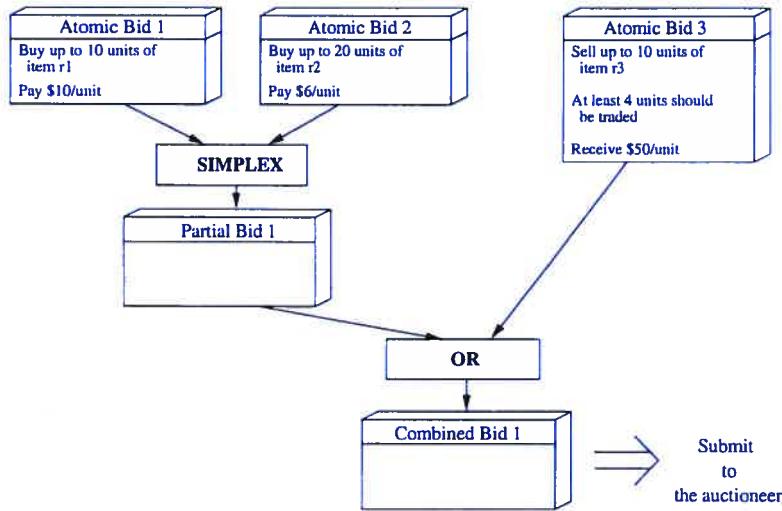


FIG. 4.6 – Bidding Example 2.

where a participant would rather specify relevant prices at intermediary levels of their bidding structures. For that reason, it is important to set rules that guarantee the *consistency* of the pricing information formulated by a bidder. Consistency embodies the concept of pricing *exhaustiveness*, which means that the auctioneer will always know precisely how much the bidder is ready to pay or receive, whatever the auction outcome. It also ensures that the pricing information is *non-conflicting*, i.e., that the auctioneer would never face, for example, overlapping and contradictory price valuations. It is not concerned, however, with the congruence of the submitted prices and the participant's business activities - the bidding language is, after all, a simple medium to formulate and communicate bids, not an advisor to participants.

In order to formalize price consistency, we recall in Fig. 4.7 the graph representation T of a participant's bidding structure. At a given node of T , for example node B in Fig. 4.7, denote by $S(B)$ the sub-graph of T made of B , all its predecessors in T , and the arcs linking them. We obviously have

- $S(R) = T$ if R is the root of the graph ; and
- $S(L) = (L, \emptyset)$ if L is a leaf.

Definition 4.5 *Pricing information is completely defined on sub-graph $S(B)$ if one of following conditions holds :*

- (C-1) *The participant indicates a relevant price valuation in the expression bid*

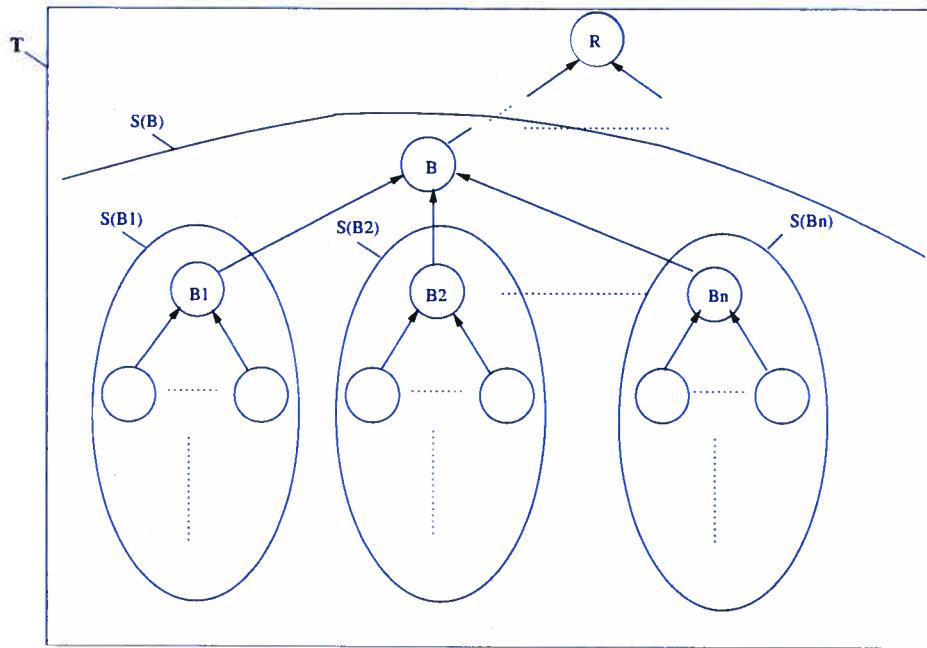


FIG. 4.7 – Price consistency in the graph representation of the bidding structure.

corresponding to node B .

(C-2) The pricing information is **completely defined** on $S(B_1), S(B_2), \dots, S(B_n)$, where B_1, B_2, \dots, B_n are immediate predecessors of B in T .

The definition above needs some words of explanation. When condition (C-1) holds at the root R of the bidding structure (this is the case of the bidding example of Fig. 4.5), a relevant bundle price is indicated by the bidder at the deepest level of the graph, starting from the leaves (atomic bids). If it is condition (C-2) which is verified, relevant pricing information is specified somewhere else in the bidding structure.

We may then state two price consistency rules, which 1) ensure that a complete pricing information is defined by the participant within the bidding structure sent to the auctioneer ; and 2) prevent the participant from defining relevant prices that overlap at two different layers of its bidding structure, which may be confusing to the auctioneer :

Rule 1 - Pricing information is exhaustive. Pricing information is completely defined on the bidding graph T .

Rule 2 - Pricing information is non-conflicting. At any given node B of the bidding

graph T , if a relevant price valuation is indicated in the bid represented by B , then no relevant price valuations should be indicated in any node of $S(B1), S(B2), \dots, S(Bn)$, where $B1, B2, \dots, Bn$ are immediate predecessors of B in T .

We need to emphasize that the rules presented in this section should be considered as minimal conditions for price consistency. In practice, each application has its own pricing context, and the auctioneer may often need to agree with the participants on additional and more specific policies to define and interpret prices.

4.6 Analysis of the framework

We devote this section to an empirical analysis of the expressiveness of our bidding framework in indivisible and divisible settings. In the indivisible case, it is shown that a bidding language that generalizes the OR/XOR formulae language of Nisan [105] can be defined within the framework. More interestingly, when divisible bids are considered, we establish that the framework supports bidding according to general (continuous) valuations functions. Furthermore, we suggest several valuation functions that can be expressed succinctly in specific bidding languages of the framework.

4.6.1 The indivisible case

Let us first give a brief overview of Nisan's OR/XOR formulae language. The auctioneer has a set G of indivisible items available for sale. Basic bundle bids (called *atomic bids* in [105]) are formulated as pairs (S, p) , where $S \subseteq G$ and p is the price the bidder is willing to pay for items in S . OR and XOR bids are collections of bundle bids such that the bidder may accept the execution of any number of bundle bids in an OR bid, and only one bundle bid in a XOR bid. OR/XOR-formulae bids are formulated through recursive combination of the OR and XOR logics in the usual way.

Proposition 4.1 *The bidding framework presented in Sections 4.3 to 4.5 is at least as expressive as Nisan's OR/XOR formulae language.*

Proof. In Section 4.4, we have seen that OR and XOR operators are special cases of the selection SELECT-OUTER bidding operator. Since combined bids are

recursively combined at the outer level, all we need to prove is that basic bundle bids can be expressed as combined bids. Let (S, p) be a basic bundle bid. The following language can be suggested. Define $|S|$ atomic bids $\delta_k = (k, +1, 1, 1, NIL)$, $k \in S$, which correspond to single-item bids to buy items in S , and let each atomic bid δ_k correspond to a trivial partial bid θ_k , $k \in S$. Then define combined bid $\Theta = (\{\theta_k\}_{k \in S}, \text{AND}, p)$ asking for all partial bids θ_k , $k \in S$ to be executed at a bundle price p . ■

4.6.2 The divisible case

Let G be a set of divisible items. A bundle Q is a $|G|$ -vector $Q = \{Q_1, Q_2, \dots, Q_{|G|}\}$, where Q_k , $k \in G$ is a quantity of item k traded in the bundle. By convention, $Q_k > 0$ if item k is bought in bundle Q , and $Q_k < 0$ if it is sold. We define the valuation $v(Q)$ as the bidder's preference for trading bundle Q . We make the following assumptions :

- (A1) G can be partitioned into two subsets : G_B , the subset of items that are bought, and G_S , the subset of items that are sold.
- (A2) Quantity $|Q_k|$ of item k that is to be traded in the bundle is bounded by Q_k^{\max} .

Thus valuation function $v(\cdot)$ is defined on $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_{|G|}$, where $\mathcal{D}_k = [0, Q_k^{\max}]$ if $k \in G_B$ and $\mathcal{D}_k = [-Q_k^{\max}, 0]$ if $k \in G_S$.

Proposition 4.2 *Any valuation $v(\cdot)$ that complies to assumptions (A1) and (A2) can be expressed in a bidding language of the framework presented in Sections 4.3 to 4.5.*

Proof. The following bidding language can be suggested. Consider atomic bids $\{\delta_k\}_{k \in G}$ such that $\delta_k = (k, \epsilon_k, Q_k^{\max}, 0, NIL)$ where $\epsilon_k = +1$ if $k \in G_B$ and $\epsilon_k = -1$ if $k \in G_S$. Let $\mathcal{X} = \text{SELECT-INNER}(1, |G|)$ be an instance of the inner level selection operator (OR), and define partial bid $\theta = (\{\delta_k\}_{k \in G}, \mathcal{X}, p(\cdot))$ where $p(\cdot)$ is a price-mapping defined on $[0, 1]^{|G|}$ s.t. $p(x_1, \dots, x_{|G|}) = v(\epsilon_1 Q_1^{\max} x_1, \dots, \epsilon_{|G|} Q_{|G|}^{\max} x_{|G|})$. Finally, formulate combined bid Θ trivially as partial

bid θ . ■

In other words, a bidder can bid its preference according to a general valuation function by conveying directly that function in the price-mapping of a partial bid. However, the virtue of generality of the language conflicts with conciseness of representation, since the bidder needs to specify explicitly its valuation for each possible bundle Q . In the following, we suggest some specific valuation functions and show how our bidding framework provides the necessary tools for concise bid formulation according to the proposed valuations.

Additive valuations

Additive valuations refer to preferences with no interdependencies between the different items. That is, $v(Q) = \sum_{k \in G} v_k(Q_k)$, where $v_k(Q_k)$ denotes the bidder's preference for trading quantity Q_k of item k , $k \in G$. The following bidding language, in which relevant price information is indicated by the bidder at the atomic bid level, may be suggested in this case.

(a) The inner level :

- Formulate $|G|$ atomic bids $\{\delta_k\}_{k \in G}$, s.t. $\delta_k = (k, \epsilon_k, Q_k^{\max}, 0, p_k(\cdot))$ with $\epsilon_k = +1$ if $k \in G_B$ and $\epsilon_k = -1$ if $k \in G_S$, and $p_k(x_k) = v_k(\epsilon_k Q_k^{\max} x_k)$.
- Define partial bid $\theta = (\{\delta_k\}_{k \in G}, \text{OR}, \text{NIL})$.

(b) The outer level : submit trivially-defined combined bid $\Theta = \theta$.

Price-mappings $p_k(\cdot)$ may be simplified further when additive valuations exhibit certain logical structures, notably disjunctive relations. This simplification may come, however, at the expense of more atomic bids and additional complexity at the partial and combined bid levels. Piece-wise linear additive valuations provide a classical example. The principle is as follows. An atomic bid (with a linear price-mapping and a properly defined lower bound) is defined for each segment of the piece-wise linear functions $v_k(\cdot)$, $k \in G$. A XOR bid is formulated for each item k to ensure that one segment is selected. A final OR bid that combines the XOR bids is submitted by the bidder.

Production-recipe valuations

Many important industries (oil refining, food processing , etc.) involve transformation of basic (raw) commodities into processed (end) products. Consequently, producers often need to acquire the raw commodities in precise proportions, according to certain production recipes. By way of illustration, we only consider a single end product and we limit ourselves to elementary recipes. Let G be the set of raw commodities and P the end product. To process one unit of P , the producer needs to mix quantities s_k , $k \in G$ of the different raw commodities. Moreover, technological constraints faced by the producer are such that no more than Q^{max} units of product P , and no less than Q^{min} units can be processed. The producer's valuation $v(Q)$ of a bundle Q of raw commodities depends of whether the bundle has the right composition specified in the recipe, and it allows to process an admissible volume of the end product. That is,

$$v(Q_1, \dots, Q_{|G|}) = \begin{cases} \bar{v}(Q^P) & \text{if } Q^P = \frac{Q_k}{s_k}, \forall k \in G, \text{ and } Q^{min} \leq Q^P \leq Q^{max}; \\ 0 & \text{otherwise.} \end{cases}$$

Where $\bar{v}(Q^P)$ denotes the producer's preferences for processing Q^P units of product P .

The following language supports bidding according to valuation function $v(\cdot)$:

(a) The inner level :

- Formulate $|G|$ atomic bids $\{\delta_k\}_{k \in G}$, s.t. $\delta_k = (k, +1, s_k Q^{max}, \frac{Q^{min}}{Q^{max}}, NIL)$.
- Formulate partial bid $\theta = (\{\delta_k\}_{k \in G}, EQUAL, p(\cdot))$, where price-mapping $p(\cdot)$ is defined such that $p(x_1, \dots, x_{|G|}) = \bar{v}(x_1 s_1 Q^{max})$

(b) The outer level : submit combined bid $\Theta = \theta$.

It should be underlined that more elaborate bid formulation at the outer level of the language above would allow the support of many variants of the basic production-recipe valuation. These include complex production plans featuring logically interdependent production recipes, the availability of several concurrent production plans to choose from, and so on.

Categories-of-items valuations

There are trading contexts in which bidders are indifferent towards the trade of several items they consider as “equivalent”. In such cases, it may be convenient to classify items into categories, and formulate category-based bidding requirements. Fixed-income financial markets are a case in point. Typically, these markets trade, among other debt instruments, a huge variety of bonds that differ in type (treasury, municipal, and corporate) and in safety levels, and have different issue and maturity dates. With the yield to maturity and the price of a bond indicating its intrinsic value, it is very common to encounter trading requests that group desired bonds by safety levels and by maturity date intervals (e.g., “invest \$30K in high-quality (A) corporate bonds maturing between January 2004 and September 2005”).

More generally, suppose that the set of items G can be subdivided into S subsets (categories) G_s , $s = 1, \dots, S$. To simplify the presentation, we assume items can only be purchased. The bidder requires that exactly M_s units of *any* items in category s should be acquired. Furthermore, it only values positively outcomes that provide items in no more than S^{max} categories, and no less than S^{min} categories. This valuation may be associated to the following language :

(a) The inner level :

- Formulate $|G|$ atomic bids $\{\delta_k\}_{k \in G}$, s.t. $\delta_k = (k, +1, Q_k^{max}, 0, p_k(.))$, where $p_k(x_k)$ denotes the bidder’s preference for acquiring $Q_k^{max}x_k$ units of item k .
- Let $\mathcal{X}_s = \text{QTY-EQ}(M_s)$, $s \in S$ be an instance of the quantity operator. Define partial bids $\theta_s = (\{\delta_k\}_{k \in G_s}, \mathcal{X}_s, NIL)$, $s \in S$.

(b) The outer level : submit combined bid $\Theta = (\{\theta_s\}_{s \in S}, \mathcal{X}, NIL)$, where $\mathcal{X} = \text{SELECT-OUTER}(S^{max}, S^{min})$ is an instance of the outer-level selection operator.

4.7 Impact on the allocation problem

The allocation problem in combinatorial auctions consists in determining the winning bids among those submitted to the auctioneer and, for each winning bid, the corresponding traded proportions of its atomic bids, such that a market objective is

optimized. In this section, we examine the impact of the bidding framework on the mathematical programming formulation of the allocation problem.

The winner determination problem in the basic formulation of combinatorial auctions has received much attention in the literature (e.g., [39]). The model deals with the one-sided, single-unit, indivisible case. One seller (the auctioneer) auctions off a set of m different items, one unit of each available. Several buyers bid on bundles S , which are non-empty collections of items. On each possible bundle $S \subseteq G, S \neq \emptyset$, a buyer $l \in L$ has to submit a price $p_l(S) \geq 0$, which is the amount of money the buyer is ready to pay if it gets S . The auctioneer then determines the winning bids that maximize the total revenue of the auction. So, let $x_{l,S} = 1$ if bundle S is allocated to buyer l , and 0 otherwise. The allocation model can be formulated as follows :

$$\max_{l \in L} \sum_{S \subseteq G; S \neq \emptyset} p_l(S)x_{l,S} \quad (4.1)$$

$$s.t. \quad \sum_{S \subseteq G; S \neq \emptyset} x_{l,S} \leq 1, \quad l \in L \quad (4.2)$$

$$\sum_{l \in L} \sum_{S \subseteq G; S \neq \emptyset} \delta_{r,S} x_{l,S} \leq 1, \quad r \in G \quad (4.3)$$

$$x_{l,S} \in \{0, 1\}, \quad S \subseteq G, S \neq \emptyset, \quad l \in L \quad (4.4)$$

where $\delta_{r,S} = 1$ if item r belongs to S and 0 otherwise. Constraints (4.2) enforce the requirement that at most one bundle is allocated to any buyer. Relations (4.3) ensure that no single item is allocated to more than one buyer.

The “bidding language” used in this basic model is elementary : participants make indivisible bids on all the possible bundles that are composed of single units of items. By comparison, the settings implied by the general bidding language are far more complex. Both the divisible and indivisible cases need to be considered. Participants need not bid on every possible combination but rather decide on the composition of the combined bids they submit to the market, and these bids can be composite bids to sell and buy several units of each item. More importantly, they use the operators of the language to specify execution conditions on their bids. All these aspects have to be reflected on the formulation of the allocation problem.

Following our bidding framework, the allocation problem decisions are : 1) which final combined bids among those submitted by participants will be executed ? and 2) if a combined bid is executed, what does it mean for atomic bids ? In other terms, what are the corresponding execution proportions of atomic bids ? Therefore, we need to define the following *primary* decision variables :

- y_l : a binary variable characterizing the execution of the combined bid Θ_l submitted by participant l , i.e.,

$$y_l = \begin{cases} 1 & \text{if } \Theta_l \text{ is executed, } l \in L; \\ 0 & \text{otherwise,} \end{cases}$$

- x_k : the execution proportion of atomic bid δ_k , $k \in \mathcal{A}_l$, $l \in L$.

Auxiliary decision variables are required to express the constraints of the model :

- \tilde{y}_k : a binary variable characterizing the execution of atomic bid δ_k , i.e.,

$$\tilde{y}_k = \begin{cases} 1 & \text{if } \delta_k \text{ is executed, } k \in \mathcal{A}_l, l \in L; \\ 0 & \text{otherwise,} \end{cases}$$

- \bar{y}_i : a binary variable characterizing the execution of partial bid θ_i , i.e.,

$$\bar{y}_i = \begin{cases} 1 & \text{if } \theta_i \text{ is executed, } i \in I_l, l \in L; \\ 0 & \text{otherwise,} \end{cases}$$

- y_j : a binary variable characterizing the execution of combined bid Θ_j , i.e.,

$$y_j = \begin{cases} 1 & \text{if } \Theta_j \text{ is executed, } j \in J_l, l \in L. \\ 0 & \text{otherwise,} \end{cases}$$

Bid composition and the application of bidding operators add several categories of constraints to the market's own constraints (e.g., constraints (4.3) in the basic formulation of the allocation problem, which express the physical "conservation" of single items). The new constraints are as follows.

4.7.1 Lower bound constraints

The following constraints must hold :

$$\tilde{y}_k l_k \leq x_k \leq \tilde{y}_k, \quad \forall k \in \mathcal{A}_l, \forall l \in L \tag{4.5}$$

If atomic bid δ_k is executed, i.e., $y_k = 1$, then the lower bound condition $l_k \leq x_k \leq 1$ must be satisfied; otherwise, $x_k = 0$ and nothing is traded at all. In the indivisible case, we obviously have $x_k = \tilde{y}_k, k \in \mathcal{A}_l, l \in L$

4.7.2 Partial bid execution

If $\theta_i, i \in I_l, l \in L$ is a trivial partial bid corresponding to atomic bid $\delta_k, k \in \mathcal{A}_l$, then

$$\bar{y}_i = \tilde{y}_k \quad (4.6)$$

Otherwise, $\theta_i = (\{\delta_k\}_{k \in K_i}, \mathcal{X}_i, p_i(\cdot))$, and we have

$$0 \leq \sum_{k \in K_i} \tilde{y}_k \leq \bar{y}_i |K_i|, \quad (4.7)$$

which states that no atomic bid in $\Delta_i = \{\delta_k\}_{k \in K_i}$ should be executed if θ_i is not executed.

4.7.3 The ORDERING operator

An instance \mathcal{X}_i of the ORDERING operator associated to partial bid $\theta_i, i \in I_l, l \in L$ gives rise to the constraints

$$x_{k_1} \geq x_{k_2}, \quad \forall k_1, k_2 \in K_i, \text{s.t. } k_1 \succ_i k_2 \quad (4.8)$$

4.7.4 The EQUAL operator

An instance \mathcal{X}_i of the EQUAL operator associated to a partial bid $\theta_i, i \in I_l, l \in L$ induces

$$x_{k_1} = x_{k_2}, \quad k_1, k_2 \in K_i \quad (4.9)$$

4.7.5 The SIMPLEX operator

For an instance of the SIMPLEX operator associated to partial bid $\theta_i, i \in I_l, l \in L$, we have

$$\sum_{k \in K_i} x_k = \bar{y}_i \quad (4.10)$$

By virtue of constraint (4.10), if partial bid θ_i is not executed, i.e., $\bar{y}_i = 0$, then $x_k = 0, \forall k \in K_i$, and no atomic bid in Δ_i should be executed. Otherwise, the simplex condition $\sum_{k \in K_i} x_k = 1$ is satisfied.

4.7.6 Quantity operators

An instance $\mathcal{X}_i = \text{QTY-EQ}(\beta_i)$ of the QTY-EQ operator, associated to partial bid $\theta_i, i \in I_l, l \in L$, corresponds to the following constraint

$$\bar{y}_i \beta_i \leq \sum_{k \in K_i} \epsilon_k q_k x_k \leq \beta_i + (1 - \bar{y}_i) M_i \quad (4.11)$$

where $M_i = -\beta_i + \sum_{k \in K_i} |q_k|$.

If partial bid θ_i is executed ($\bar{y}_i = 1$), then constraint (4.11) reduces to $\sum_{k \in K_i} \epsilon_k q_k x_k = \beta_i$. Otherwise, constraint (4.7) implies that no atomic bid in Δ_i is executed, while (4.11) holds by definition of M_i .

Similarly, an instance $\mathcal{X}_i = \text{QTY-MORE}(\beta_i)$ corresponds to

$$\sum_{k \in K_i} \epsilon_k q_k x_k \geq \bar{y}_i \beta_i, \quad (4.12)$$

while an instance $\mathcal{X}_i = \text{QTY-LESS}(\beta_i)$ adds the constraint

$$\sum_{k \in K_i} \epsilon_k q_k x_k \leq \beta_i \quad (4.13)$$

4.7.7 The BUDGET operator

For an instance $\mathcal{X}_i = \text{BUDGET}(\mathcal{B}_i)$ associated to partial bid $\theta_i, i \in I_l, l \in L$, we have

$$\sum_{k \in K_i} p_k(x_k) \leq \mathcal{B}_i \quad (4.14)$$

4.7.8 The SELECT-INNER operator

An instance $\mathcal{X}_i = \text{SELECT-INNER}(k^{\min}, k^{\max})$ associated to a partial bid θ_i , $i \in I_l$, $l \in L$ induces

$$\bar{y}_i k^{\min} \leq \sum_{k \in K_i} \tilde{y}_k \leq \bar{y}_i k^{\max} \quad (4.15)$$

Constraint (4.15) imposes the selection condition $k^{\min} \leq \sum_{k \in K_i} \tilde{y}_k \leq k^{\max}$ when partial bid θ_i is executed ($\bar{y}_i = 1$). Otherwise, $\tilde{y}_k = 0$, $\forall k \in K_i$, and no atomic bid in Δ_i is executed.

4.7.9 Hybrid operators

In addition to constraints (4.15), an hybrid operator $\mathcal{X}_i = \text{SELECT-INNER}(k^{\min}, k^{\max}) + \text{EQUAL}$ associated to a partial bid θ_i , $i \in I_l$, $l \in L$ gives rise to

$$\tilde{y}_k - 1 \leq x_k - z \leq 0, \quad \forall k \in K_i \quad (4.16)$$

$$0 \leq z \leq 1 \quad (4.17)$$

where z is an auxiliary variable.

Let us analyze constraints (4.16) and (4.17). For all the atomic bids selected by virtue of (4.15), (4.16) implies that $x_k = z$, $\forall k \in K_i$, which corresponds to the equal-proportions conditions. For those atomic bids that are not selected, i.e., $\tilde{y}_k = 0$, the atomic bid lower bound constraints (4.5) state that these bids are not executed at all ($x_k = 0$), while constraint (4.16) reduces to $0 \leq z \leq 1$ which is obviously true, thanks to constraint (4.17).

4.7.10 Trivial combined bids

If Θ_j , $j \in J_l$, $l \in L$ is a trivial combined bid corresponding to partial bid θ_i , $i \in I_l$, then

$$y_j = \bar{y}_i \quad (4.18)$$

4.7.11 The SELECT-OUTER operator

$$y_j N^{\min} \leq \sum_{\bar{j} \in J_j} y_{\bar{j}} \leq y_j N^{\max}, \quad \forall j \in J_l, \forall l \in L \quad (4.19)$$

Constraints (4.19) impose the selection condition $N^{\min} \leq \sum_{\bar{j} \in J_j} y_{\bar{j}} \leq N^{\max}$ when combined bid Θ_j is executed. If it is not, $y_{\bar{j}} = 0, \forall \bar{j} \in J_j$, and no combined bid that is part of the bundle Ω_j is executed.

The impact of the language on the objective function of the allocation model is not considered in this section since it is strongly linked to the specific price interpretation policy that the auctioneer and the participants set up in each application, as well as to the general objectives of the auction.

4.8 An application to portfolio bundle trading

Traders in financial markets often need to re-balance their portfolios of assets at “the end of the day” to reflect customer or company guidelines regarding the sectoral compositions of the portfolios, or to match a certain performance index. Most financial marketplaces trade assets on an individual basis, i.e., asset by asset. So a trader has to be involved in “combined negotiations”, possibly across several different marketplaces, in order to re-balance its portfolios. This practice typically induces important transaction costs and a significant risk to end up with unbalanced portfolios. In this context, many-to-many *bundle-based* e-markets dedicated to re-balancing portfolios offer an interesting alternative. In these markets, traders submit to the market maker consolidated orders to simultaneously sell and buy various assets. Since a given asset in a bundle may be traded only if some other assets are traded as well, traders at least guarantee that the executed proportions of their orders correspond to their objectives. After it receives all the bundle orders, along with corresponding price offers, the market maker determines the executed proportions of each order and payments the traders should make or receive. The objective of the market maker is to maximize the total trade surplus of the market.

Among the many additional bidding requirements traders may formulate in this

context, two features stand out. First, due to fixed-cost transaction fees, a trader may indicate a lower bound on the executed proportion of the order such as it prefers the order not to be executed at all below the lower bound. Also, a trader may identify certain bundle orders as being “equivalent”, and indicate its willingness to accept execution of at most one of them (this may translate the trader’s aversion to excessive fragmentation of its portfolio).

We now introduce the notation of the model. A bundle order j defined by trader l is a vector $O_{lj} = (\{(\epsilon_{ljr}, q_{ljr})\}_{r \in G}, b_{lj}, c_{lj})$ where

- $\epsilon_{ljr} = +1$ if asset r is purchased in order j , and $\epsilon_{ljr} = -1$ if asset r is sold in order j ;
- q_{ljr} is the maximum number of units of asset r that may be traded in order j ($q_{ljr} = 0$ if item r is not traded in order j) ;
- b_{lj} is the lower bound on the execution proportion of order j ; $b_{lj} = 0$ indicates that no lower bound is specified ;
- c_{lj} is the bundle price the trader is willing to pay or receive if order j is entirely executed.

Let J_l be the set of bundle orders formulated by trader l , $l \in L$. A XOR association Υ is a subset of J_l such that at most one order in Υ should be executed. Each trader $l \in L$ may thus formulate a (possibly empty) set Υ_l of XOR associations.

Bidding in the financial market specified above can be achieved through the following language :

(a) The inner level :

- Atomic bids correspond to single-asset sell or buy orders of each bundle. Hence formulate atomic bids $\delta_{ljr} = (r, \epsilon_{ljr}, q_{ljr}, b_{lj}, NIL)$, $r \in G_{lj}$, $j \in J_l$, where $G_{lj} \subseteq G$ of items that are traded in order j (such that $q_{ljr} \neq 0$).
- Formulate a partial bid $\theta_{lj} = (\{\delta_{ljr}\}_{r \in G_{lj}}, EQUAL, p_{lj}(\cdot))$ for each bundle order $j \in J_l$. Here p_{lj} is a linear price-mapping defined such that $p_{lj}(x_1, \dots, x_{|G_{lj}|}) = c_{lj}x_1$.

(b) The outer level :

- Formulate a combined bid $\Theta_{l,\Upsilon} = (\{\theta_{lj}\}_{j \in \Upsilon}, XOR, NIL)$ for each XOR asso-

ciation $\Upsilon \in \Upsilon_l$.

- Let $\bar{J}_l = J_l - \cup_{\Upsilon \in \Upsilon_l} \{\Upsilon\}$ be the set of bundle orders that appear in no XOR association. Define a final combined bid $\Theta = (\{\theta_{lj}\}_{j \in \bar{J}_l} \cup \{\Theta_{l,r}\}_{r \in \Upsilon_l}, \text{OR}, NIL)$ that encapsulates all the bundles and the XOR associations formulated by the trader. This combined bid is submitted by trader l to the market maker.

Decision variables of the allocation problem are

- x_{ljr} : the traded proportion of item r in bundle order j formulated by trader l ;
- \bar{x}_{lj} : the traded proportion of bundle order j formulated by trader l ;
- \bar{y}_{lj} : a binary variable characterizing the execution of bundle order j formulated by trader l ; i.e., $\bar{y}_{lj} = 1$ if it is executed, $\bar{y}_{lj} = 0$ otherwise;
- y_l : a binary variable characterizing the execution of the OR bid submitted by trader l to the market maker; $y_l = 1$ if it is executed, $y_l = 0$ otherwise.

The allocation problem can then be formulated as the following optimization model :

$$\max_{l \in L} \sum_{j \in J_l} p_{lj} \bar{x}_{lj} \quad (4.20)$$

$$\text{s.t.} \quad \sum_{l \in L} \sum_{j \in J_l} \epsilon_{ljr} q_{ljr} x_{ljr} = 0, \quad r \in G \quad (4.21)$$

$$\bar{x}_{lj} = x_{ljr}, \quad r \in G, l \in L, j \in J_l \quad (4.22)$$

$$\bar{y}_{lj} b_{lj} \leq \bar{x}_{lj} \leq \bar{y}_{lj}, \quad l \in L, j \in J_l \quad (4.23)$$

$$\sum_{j \in \mathcal{X}} \bar{y}_{lj} \leq 1, \quad \mathcal{X} \in \mathcal{X}_l, l \in L \quad (4.24)$$

$$y_l \leq \sum_{l \in J_l} \bar{y}_{lj} \leq |J_l| y_l, \quad l \in L \quad (4.25)$$

The objective reflects the desire of the market maker to seek the maximum market surplus, so highly priced buy orders and lowly priced sell order are given high execution priority. Constraints (4.21) express the balance of the market : the quantity of asset r bought equals the quantity sold. Relations (4.22) arise from the application of the EQUAL operator at the inner level. Constraints (4.23) correspond to the lower

bound on the execution proportion of each bundle order, and constraints (4.24) to the application of the XOR operator at the outer level. Finally, constraints (4.25) express the OR condition which states that at least one bundle order should be executed if the submitted OR bid is executed in a trade.

4.9 Conclusion and perspectives for future work

In this paper, we have presented a novel framework for bidding in combinatorial auctions. The framework relies on a two-level representation of the combined bid, in which the definition of the atomic bids and the conditions related to their execution proportions are separated from the logical execution constraints. We have suggested interesting classes of bidding operators that allow the definition of very general bidding languages. The framework is generic and flexible in the sense that auction designers can select and adapt the operators, suggest new ones, define their price interpretation rules and policies, and choose the appropriate levels of recursivity of the language, depending of the specific context of each application. After an analysis of the expressiveness of the framework and an examination of its impact on the allocation problem, we have used an application in the finance sector to illustrate the usefulness of the bidding framework.

Many challenges remain to be addressed, though. In the following, we briefly discuss three that clearly stand out. First, once the auctioneer has received the submitted bids, it must formulate and solve the allocation problem. The latter generally takes the form of a mixed integer programming model, with potentially large numbers of variables and constraints that depend on the number of bids submitted to the auctioneer, the number of items traded, the particular characteristics and constraints of the application, and the complexity of the bids. To gain further insights into the complexity of the allocation problem, we have undertaken in [3] an experimental study on the basis of the portfolio bundle trading application of Section 4.8, in which sample allocation problems are solved by a commercial MIP solver (ILOG CPLEX 7.1). We have found that, on some of our largest problems (1000 different assets, 100 traders, up to 20 bundle orders per trader) it takes more than 4 days to achieve optimality on a reasonably fast machine. Significant work is thus needed to design more effi-

cient exact and heuristic methods than simple Branch and Bound. This is of utmost importance for iterative multi-round combinatorial auctions, in which the allocation problem has to be solved repeatedly.

Determining the payments that participants need to make or receive when the trade is realized is another challenging problem that faces the auctioneer. These prices might be simply the prices indicated by participants in their bids, but could be (and often are) different. In fact, the way the auctioneer determines prices depends strongly on the auction objectives (maximize the revenue of the auctioneer ? the social welfare of the bidders ? etc.) and its economic properties (e.g., pricing equilibria, incentive-compatibility, etc.). Although some interesting progress in the study of pricing schemes has been made recently [149, 41, 18, 114], the understanding of pricing is still at its very beginnings and much work is needed in this area.

On the participant side, there is, of course, the need to construct and price bids. Given the relative sophistication of the bidding language and the fact that participants have to make bids that are consistent with their business processes (cost policies, current and forecast operations, knowledge of the economical sector and the competitors), bid construction is a hard task. It becomes even harder when participants are involved in combined negotiations on several markets, each running dynamic auctions (multi-round discrete or continuous), but with some not allowing bundle trading. Participants then face the additional burden of designing complex bidding strategies in terms of both bid composition and fallback scenarios [29]. Therefore, there is a real need to develop optimization-based decision support tools, so-called *advisors*, to assist participants in tackling these decisions.

Acknowledgements This project has received financial support from the Bell Canada University Laboratories (LUB), the Natural Sciences and Engineering Research Council of Canada through their Basic Research and Collaborative Industry-University Research grant programs, the FCAR research development fund of the Province of Québec, and the Valorisation-Recherche Québec fund.

We would like to thank our colleagues Jacques Robert, Yves Richelle, Martin Boyer, and Robert Gérin-Lajoie for valuable discussions on the matter of this paper.

Chapitre 5

Models for Bundle Trading in Financial Markets

La référence de cet article est :

J. Abrache, T.G. Crainic, and M. Gendreau, “Models for Bundle Trading in Financial Markets”, à paraître dans *EJOR*.

Models for Bundle Trading in Financial Markets

Jawad Abrache^{1,3}, Teodor Gabriel Crainic^{2,3}, and Michel Gendreau^{1,3}

¹ Département d'informatique et de recherche opérationnelle, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada, H3C 3J7.

² Département management et technologie, Université du Québec à Montréal, C.P. 6192, succursale Centre-ville, Montréal, Canada H3C 4R2.

³ Centre de recherche sur les transports, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada, H3C 3J7.

Abstract. Bundle trading is a new trend in financial markets that allows traders to submit consolidated orders to sell and buy packages of assets. We propose a new bundle-based market-clearing formulation for portfolio balancing that extends the previous models in the literature through a more detailed representation of portfolios and the formulation of new bidding requirements. We also present post-optimality tie-breaking procedures intended to discriminate between equivalent orders on the basis of submission times. Numerical results evaluate the “bundle” effect as well as the bidding flexibility and the computational complexity of the formulation.

Keywords : Auction design, Financial markets, Bundle trading, Discrimination procedures

5.1 Introduction

The emergence of the Internet as a medium for trading goods and services has changed in an unpreceded manner the way financial services are offered. The most obvious manifestation of these changes is the phenomenal growth in popularity of online asset trading. More and more individuals and corporate investors have access to a wide selection of cyber-brokerage services that set themselves up as an alternative to full-service brokers. This has greatly enhanced competitiveness in the financial services industry and contributed to increase the quality of service and to lower transaction costs.

The most significant changes, however, probably are those of the financial marketplaces themselves. Traditional stock market models, which are generally specialist-run, on-the-floor exchanges (the NYSE model), or market-maker, over-the-counter auctions (the NASDAQ model), increasingly face the challenge of Electronic Communications Networks (ECNs). The ECNs are fully-automated, computerized networks that can efficiently match sell and buy orders of financial assets, while offering customers several advantages such as anonymous access and after-hours trading. For traditional financial marketplaces, structural changes (mergers, strategic alliances, etc.) and, most importantly, re-organization of internal policies and procedures were necessary actions for survival. Meanwhile, we witnessed an increased interest in auction-based mechanism design for financial marketplaces (e.g., Domowitz [42], Madhavan [90]), which reflects a growing awareness that innovative and efficient market mechanisms are key to successful financial marketplaces.

One of the most critical design issues financial marketplaces need to consider is how to make their procedures reflect as much as possible the trading needs and requirements of their users. Unfortunately, there is still a large gap between market procedures and what traders may actually want to do. For instance, while financial portfolios tend to be increasingly more diversified, comprising notably stocks, futures, bonds, and foreign currencies of different kinds, the current organization of financial marketplaces remains heavily sectoral. This, together with the lack of institutionalized links between different marketplaces, increases the dependency of traders with large portfolios on brokerage institutions. Moreover, most financial marketplaces trade as-

sets on an individual basis, making it difficult for investors to maintain a precise and timely control of the composition of their portfolios. In that regard, a market mechanism based on *bundle trading*, which would allow traders to submit *consolidated* bids to sell and buy different quantities of various assets, such that the single-asset orders in a consolidated bid are executed in the same proportions, would be an extremely interesting feature.

Bundle trading is not, strictly speaking, a new concept in financial markets, as its origins could be traced back to Markowitz's seminal paper (Markowitz [93]) setting the foundations of modern portfolio selection theory. It is not specific to financial markets, either, since it may be encountered in many other contexts where items of different physical nature are traded and the traders' valuations of a given item depend on whether or not other items are traded as well (e.g., Graves, Schrage, and Sankaran [60], Brewer and Plott [24], Eso [47], Jones and Koehler [68]). In all generality, item interdependency takes two basic forms : two items A and B are *complementary* if the trader's valuation of the bundle $\{A, B\}$ is greater than the valuations of A and B taken separately ; they are *substitutable* if the valuation of $\{A, B\}$ is lower than the valuations of A and B taken separately. It is well known in the literature on combinatorial auctions (e.g., Rothkopf, Pekeč, and Harstad [125]) that an “exposure problem” may arise when items are traded using simultaneous ascending auctions, especially when complementarities prevail : hoping to obtain a desired “package”, a trader may continue submitting bids on the single items in the package even if the price on the market exceeds the trader's valuation for a given item. The trader then faces the risk of obtaining only a few of the desired items which it would have paid more than their worth. When bidders decide to avoid such risk by restraining themselves from bidding aggressively, they induce economically inefficient outcomes. *Combinatorial bidding*, where bids and allocations are based on *bundles* of items, alleviates the problem by allowing traders to reflect their preferences in the bids they submit to the market.

Bundle trading offers many other potential benefits in the specific context of financial markets including :

1. Opportunities for cumulative aggregation of value. By submitting consolidated

orders to sell and buy assets, a trader could combine trade orders with very competitive prices (for highly sought-after assets it desires to sell, for instance) and orders with less competitive prices. Even without taking complementarity effects into consideration, the trader would obviously increase its chances of executing all its orders. Fan *et al.* [49] illustrate cumulative value aggregation in Table 5.1. Here, the trader builds its ask-offer prices on the basis of the last trading day closure prices. If trading were to be done asset by asset, the trader's portfolio would remain over-exposed in the car industry sector and under-exposed in the technology sector. Bundle trading, on the other hand, permits to completely balance the portfolio, even with respect to the worst prices of the day.

Stock	Quantity	Yesterday's close price	Today's price range	Worst price	Trade executed ?
IBM	+100	74 3/4	75 1/8 - 75 5/8	75 5/8	no
Microsoft	+200	148	144 1/4 - 146 3/4	146 3/4	yes
Cisco	+50	76 1/8	75 - 76 1/4	76 1/4	yes
GM	-200	84 1/4	84 1/2 - 85 3/4	84 1/2	yes
Ford	-100	122 1/2	121 3/8 - 122 3/8	121 3/8	no
Chrysler	-50	99	98 3/4 - 102 1/2	98 3/4	yes
Bundle		6831 1/4		6750	yes

TAB. 5.1 – Example of portfolio bundle trading

2. Bundle trading often involves large packages of assets, bringing more liquidity to the marketplace. Popper [118] reports that, according to brokers' estimates in the UK, the majority of bundle trades customers ask them to perform are worth between \$15M and \$80M, and huge transactions involving packages of \$1B and more are encountered from time to time.
3. Bundle trading should lower commission and transaction fees. Since fund managers and private traders would be able to submit directly their orders to the market, they would be less dependent on intermediaries. Moreover, end users would retain the largest part of execution risks, and commissions should be lower. Finally, bundle trading generates fewer orders, which means lower operation and administrative costs.

Bundle-based market mechanisms for financial markets generally consist of order matching algorithms to compute allocations and determine payments traders make or receive, so that the traders' requirements are satisfied and the economic surplus of the market is maximized. Recent years have seen growing interest in developing and implementing such mechanisms. Among these efforts, Fan *et al.* [49] propose a bundle trading market mechanism based on a simple market-clearing linear programming formulation. They also discuss the qualitative advantages of bundle trading mechanisms over elaborate trading mechanisms based on single-asset orders, giving the example of the OptiMark Trading System (<http://www.optimark.com>). Fan, Stallaert, and Winston ([51] [52]) present FBTS, which is an experimental web-based bundle trading system employing a real-time order matching and execution mechanism. Bossaerts, Fine, and Ledyard [21] exhibit another advantage of portfolio bundle trading mechanisms. Starting from the observation that thin financial markets often fail to fully equilibrate due to a lack of liquidity (according to the maximum reward/risk ratio criterion of the well-known Capital Asset Pricing Model, Sharpe [136]), the authors experimentally show that implementing a portfolio bundle trading mechanism can "induce" extra liquidity in the market and boost equilibration. Finally, Polk and Schulman [117] analyze specifically the bond market context and conclude that a proper use of the combined-value logic inherent to bundle trading mechanisms enhances liquidity in that kind of markets.

These models, however, are arguably not sophisticated enough to allow traders to control with flexibility the composition of their portfolios after the trade. For example, no known financial e-market model would permit a trader to indicate its willingness to trade a bundle *A* or a "substitutable" bundle *B*, but *not both of them*. Another aspect on which the literature has been quite elusive is the post-optimality discrimination of solutions, when multiple optimal allocations and prices exist. These two issues constitute the core contribution of the present paper.

We propose a bundle-based market-clearing model in which we introduce new categories of constraints representing various trader order execution requirements. While it is possible to envision the basic bundle model and the suggested extensions as special instances of a general combinatorial bidding framework for *divisible* com-

modities and extrapolate their usage to other market contexts, we take a different approach in this paper. Thus, we will not try to address the full complexity of a complete bidding language, but rather focus on a number of bidding requirements that have special interest for financial markets such as setting limitations on volumes of assets traded in a portfolio, requiring minimal proportions of bundles to be traded, and defining exclusive OR relations between traded bundles. We also present post-optimality tie-breaking procedures intended to discriminate optimal allocations and prices with respect to an “ethical” criterion. Experimental results analyze empirically bundle trading effects from the perspective of the economic surplus. They also verify the impact of bidding requirements introduced in our market-clearing formulation on allocation complexity and solution times.

The article is organized as follows. In Section 5.2, we present our market-clearing formulations and use dual information to compute acceptable market prices in the continuous case. In Section 5.3, we introduce allocation and price discrimination procedures and prove their validity. Section 5.4 is devoted to the experimental study. Finally, Section 5.5 sums up our contributions and discusses directions for future research.

5.2 Portfolio bundle trading market mechanisms

Bundle trading is best motivated in the context of “end-of-day” balancing of financial portfolios. Traders, which are private investors or professional managers of portfolios, need to *simultaneously* sell and buy various financial assets (stocks, futures, bonds, foreign currencies, etc.) to reflect customer or company guidelines concerning the composition of their portfolios. The structure of current financial marketplaces is not well suited to portfolio balancing, however, having been designed with other purposes in mind. Hence, traders that want to balance their portfolios must be involved in several combined negotiations, possibly across different marketplaces, and submit bids that are good enough to ensure *all* the corresponding single-asset trade orders are executed. This practice typically induces important transaction costs, a burdensome and complex strategic analysis, and most importantly, a significant risk to end up with unbalanced portfolios.

Bundle-based *exchanges* offer an interesting alternative to combined negotiations. In this market model, traders willing to balance their portfolios register in an electronic marketplace. A market maker, which may be a human or a virtual software agent, organizes a single-round sealed-bid auction between traders in which it acts as a mediator. The traders submit to the market bundle orders to simultaneously sell and buy different assets, along with limit prices they are willing to pay or receive if these orders are executed. When they submit their trade orders, traders understand and accept that the market maker may only execute *proportions* of these orders. After it receives all the orders, the market maker invokes a market-clearing mechanism, which consists in an optimization model and an algorithm that solve two problems : the *allocation* problem and the *pricing* problem. The allocation problem consists in determining many-to-many sell and buy associations between traders (matching the bundle orders) and the executed proportion of each order, such that the total purchase and sale volumes are equal and the market surplus is maximized. The pricing problem consists in determining acceptable payments to be made or received by the traders once the trade is completed.

5.2.1 The allocation problem

Prior to presenting formulations of the market-clearing allocation problem, we introduce some basic notation and definitions.

Let

- I = the set of assets traded in the market ; and
- K = the set of traders.

Definition 5.1 (Bundle Order) A bundle order j defined by trader $k \in K$ is submitted to the market as a vector $O_j = (\{q_{ji}\}_{i \in I}, P_j)$ where :

- q_{ji} is the maximum volume of asset $i \in I$ that may be traded in order j ; $q_{ji} > 0$ corresponds to a buy order, $q_{ji} < 0$ to a sell order, and $q_{ji} = 0$ if asset i is not traded in order j ;
- $|P_j|$ is the maximum (minimum) price trader k is willing to pay (receive) if order j is entirely executed ; $P_j > 0$ if the trader is willing to pay P_j , $P_j < 0$ if the trader is willing to receive $-P_j$.

A bundle order j is said to be *executed* in a trade if a positive proportion of the maximum volumes q_{ji} , $i \in I$, requested in order j is traded.

Let us also define

- $J_k =$ the set of bundle orders formulated by trader $k \in K$;
- $J = \bigcup_{k \in K} J_k =$ the set of all bundle orders formulated by traders.

The basic formulation

The basic formulation of the allocation problem considers only a minimal set of constraints that express the physical conservation of assets in the trade. Decision variables are :

$x_j =$ the traded proportion of bundle order j , $j \in J_k, k \in K$.

The allocation problem corresponds in this case to the optimization model (**MC-B-1**) :

$$\max \quad \sum_{k \in K} \sum_{j \in J_k} P_j x_j \quad (5.1)$$

$$s.t. \quad \sum_{k \in K} \sum_{j \in J_k} q_{ji} x_j = 0, \quad i \in I \quad (5.2)$$

$$x_j \leq 1, \quad j \in J_k, k \in K \quad (5.3)$$

$$x_j \geq 0, \quad j \in J_k, k \in K \quad (5.4)$$

Model (**MC-B-1**) is very similar to the bundle order matching formulations already proposed by Fan, Stallaert, and Whinston [52] and Fan *et al.* [49]. Constraints (5.2) express the balance of the market : the volume of an asset $i \in I$ that is sold in bundle orders equals the volume that is bought. Constraints (5.3) define valid traded proportions. The objective reflects the market maker's desire to seek the maximum market surplus, so that highly-priced buy orders (bundle orders with positive limit prices) and lowly-priced sell orders (bundle orders with negative limit prices) are given high execution priority.

Extensions

The only element from the trader side considered thus far in the formulation of the allocation problem is the definition of the basic bundle order structure, that is, the specification of the maximum volumes of assets to sell and buy, and the limit prices the trader is ready to pay or receive if the order is executed. It can safely be assumed, however, that the traders would need to formulate more complex trading conditions and requirements. In their simplest form, these requirements directly reflect traders' valuations of single assets and bundles translating, in particular, various complementarity and substitutability relationships. They could also express, however, constraints derived from more elaborate business policies and practices.

Bidding languages (Nisan [105], Abrache *et al.* [1]) address the issue of bidding requirements by providing participants in general combinatorial auctions with the means to define bids, formulate complex requirements on their execution, and communicate them to the auctioneer. In Abrache *et al.* [1], we have proposed a new bidding framework that relies on a two-level representation of a combined bid. At the inner level, *atomic bids*, which are single-item sell or buy orders, are defined and combined with the help of bidding *operators* that represent continuous constraints on the traded proportions of atomic bids. So called *partial bids* created this way are then recursively combined at the outer level with the help of logical bidding operators. In this section, we specifically consider three classes of operators that can be particularly useful in the context of portfolio bundle trading, and we analyze how the corresponding bidding requirements impact the formulation of the allocation problem.

Global upper bounds on the traded volume of an asset. These bounds correspond to limitations (due to internal trading policies, liquidity issues, etc.) traders may have on the total volumes of some assets to be bought or sold as part of a trade. More precisely, let us define :

M_{ki} = the maximum volume of asset i trader k is ready to trade.

A bound M_{ki} adds the following constraint to the formulation of the allocation problem :

$$\epsilon_{ki} \sum_{j \in J_k} q_{ji} x_j \leq M_{ki} \quad (5.5)$$

where $\epsilon_{ki} = +1$ if the bound corresponds to a buy limitation, and $\epsilon_{ki} = -1$ if the bound corresponds to a sell limitation.

Lower bounds on the traded proportions of an order. There exist circumstances under which traders may consider that a bundle should be executed only if a minimal proportion of the bundle is traded ; otherwise, they prefer not to trade the bundle. This is notably the case when fixed-charge execution commissions and fees make execution of marginally small proportions of some orders non profitable. Hence, consider

l_j = the minimum proportion of bundle order j to be executed.

In order to formulate the bidding conditions corresponding to these bounds, we need to define the following auxiliary binary variables :

$y_j = 1$ if bundle order j is executed, and $y_j = 0$ otherwise.

Constraints of the allocation problem corresponding to a lower bound l_j are then :

$$l_j y_j \leq x_j \leq y_j \quad (5.6)$$

Constraints (5.6) impose that the traded proportion x_j of order j be greater than the lower bound l_j when the order is executed ; otherwise, $x_j = 0$ and nothing at all is traded.

XOR relations. XOR (exclusive OR) relations are best explained with the help of an example. Consider the following trade situation :

Asset	Bundle Order 1	Bundle Order 2
GM	1000 (Sell)	
Toyota	200 (Sell)	
Ford		1500 (Sell)
IBM	1000 (Buy)	
AMD	1000 (Buy)	
Cisco		2500 (Buy)

The trader in this example formulates two “equivalent” trade orders in the sense that both of them sell assets in the automobile sector and buy assets in the technology sector. In order to preserve its portfolios from unnecessary fragmentation, the trader may ask that at most one of the two bundle orders be executed as part of the trade, but without specifying which one.

More generally, we define an XOR relation \mathcal{X} , formulated by trader k , as a subset of bundle orders in J_k which indicates that at most one order in \mathcal{X} should be executed. Let \mathcal{X}_k be the set of all XOR relations defined by trader k , $k \in K$.

The corresponding XOR constraints in the formulation of the allocation problem are

$$\sum_{j \in \mathcal{X}} y_j \leq 1, \quad \mathcal{X} \in \mathcal{X}^k, k \in K \quad (5.7)$$

In summary, the market-clearing allocation problem can be formulated as model **(MC-B-2)** :

$$\max \quad \sum_{k \in K} \sum_{j \in J_k} P_j x_j \quad (5.8)$$

$$s.t. \quad \sum_{k \in K} \sum_{j \in J_k} q_{ji} x_j = 0, \quad i \in I \quad (5.9)$$

$$\epsilon_{ki} \sum_{j \in J_k} q_{ji} x_j \leq M_{ki}, \quad k \in K, i \in I \quad (5.10)$$

$$l_j y_j \leq x_j \leq y_j \quad j \in J_k, k \in K \quad (5.11)$$

$$\sum_{j \in \mathcal{X}} y_j \leq 1, \quad \mathcal{X} \in \mathcal{X}^k, k \in K \quad (5.12)$$

$$x_j \geq 0, \quad j \in J_k, k \in K \quad (5.13)$$

$$y_j = \{0, 1\}, \quad j \in J_k, k \in K \quad (5.14)$$

5.2.2 The pricing problem

The pricing problem answers the following question : are there acceptable payments that traders can make or receive such that the market is *budget-balanced*, that is, the economic surplus is redistributed to the traders? To answer this question, let us consider first the dual **(D)** of the continuous version of the allocation model **(MC-B-1)** :

$$\min \quad \sum_{k \in K} \sum_{j \in J_k} \mu_j \quad (5.15)$$

$$s.t. \quad \sum_{i \in I} q_{ji} \Pi_i + \mu_j \geq P_j, \quad j \in J_k, k \in K \quad (5.16)$$

$$\mu_j \geq 0, \quad j \in J_k, k \in K \quad (5.17)$$

where $\{\Pi_i\}_{i \in I}$ and $\{\mu_j\}_{j \in J_k, k \in K}$ are dual variables corresponding to constraints (5.2) and (5.3), respectively. Variables $\{\Pi_i\}_{i \in I}$ and $\{\mu_j\}_{j \in J_k, k \in K}$ are interesting because of the following result.

Proposition 5.1 *Let $\{\Pi_i^*\}_{i \in I}$ and $\{\mu_j^*\}_{j \in J_k, k \in K}$ be optimal solutions of (D). Bundle prices $\bar{P}_j = \sum_{i \in I} q_{ji} \Pi_i^*$, $j \in J_k, k \in K$ have the following properties :*

- (a) *When a bundle order j formulated by trader k is executed, i.e., $x_j > 0$, the payment trader k makes or receives, computed on the basis of the bundle price \bar{P}_j , is always at least as good as the trader's limit price P_j .*
- (b) *Payments determined on the basis of bundle prices $\bar{P}_j, j \in J_k, k \in K$, make the market budget-balanced.*

Proof. Let $\{x_j^*\}_{j \in J_k, k \in K}$ be an optimal solution of the allocation model (**MC-B-1**). Statement (a) of the proposition is a result of complementarity slackness conditions $x_j^*(P_j - \sum_{i \in I} q_{ji} \Pi_i^*) = x_j^* \mu_j^* \geq 0, j \in J_k, k \in K$, which is equivalent to $x_j^* \bar{P}_j \leq x_j^* P_j$. Statement (b) follows immediately from constraints (5.2). ■

Solving the pricing problem is far more complicated when the general market-clearing model (**MC-B-2**) is considered. Hence, due to non-convexities introduced by constraints corresponding to lower bounds and XOR relations, the existence of *single-asset* market-clearing pricing is not assured. While determining payments that are acceptable to the traders is not an issue ("the traders could pay or receive their limit bundle prices"), the market maker needs to redistribute - as "fairly" as possible - the resulting surplus to the traders. In that regard, the Vickrey-Clarke-Groves mechanism (VCG) (Vickrey [140], Clarke [31], Groves [61]) is capable of (partially) redistributing

the market surplus by returning to each trader a “discount” that corresponds to the impact of the trader’s own bid on the market surplus. This mechanism has the interesting property of being *incentive-compatible* (i.e., it is a dominant strategy for a trader to bid its true values for the desired bundles). More precisely, given bundle orders $O_j = (\{q_{ji}\}_{i \in I}, P_j)$, lower bounds l_j , and XOR relations $\mathcal{X}_k, j \in J_k, k \in K$, the VCG mechanism consists in : (a) determining a surplus-maximizing allocation x^* by solving **(MC-B-2)**; and (b) asking trader $k, k \in K$, to pay (or receive) $\sum_{j \in J_k} P_j x_j^* - (z^* - z_{-k}^*)$, where z^* is the surplus achieved by the optimal allocation x^* and z_{-k}^* is the maximum surplus achieved *without* trader k ’s orders.

Yet, it is well known that the payment rule of the VCG mechanism cannot guarantee the *budget-balance* of the market in combinatorial exchanges, in the sense that the market maker may ultimately have to “feed” the market in order to cover the monetary discounts to the traders. This critical issue has been notably addressed by Parkes, Kalagnanam, and Eso [114] who suggest to impose budget-balance as a hard constraint and to compute discounts that minimize the distance to VCG payments.

5.3 Discrimination procedures

A practical problem may arise when the market mechanism determines an optimal market-clearing allocation and a set of acceptable prices : what happens if the allocation, or the prices, or both of them, are not unique ? This uniqueness issue is particularly critical in the context of financial markets where the market maker should provide traders with a satisfactory justification of the auction outcomes. Therefore, since an arbitrary choice between possibly multiple optimal solutions is clearly unacceptable, what is required is a discrimination procedure based on “ethical” criteria.

Submission time of bundle orders is such a criterion that can reasonably be used to separate equivalent orders. Suppose, for instance, that two traders A and B submit one bundle order each : trader A sends bundle order O_A , which is received by the market maker at time t , while trader B sends bundle order O_B that the market maker receives later, at time $t' > t$. On the basis of the submission time criterion, trader A has got an advantage. In this case, discriminating between traders A and

B means that, when there are multiple optimal solutions, the market maker will use the following choice strategy :

1. Ensure that the selected optimal allocation gives the largest volume possible to trader A that submitted the earliest bid.
2. Guarantee trader A of getting as much pricing “reward” as possible from the trade, that is, paying the less if it buys and receiving the most if it sells.

Lexicographical orderings of optimal solutions conceptualize best allocations and price preferences on the basis of submission time. Denote by t_j the submission time of order j , $j \in J$. Suppose that the submission times of any two orders can be compared in a strict way ($t_{j_1} > t_{j_2}$ or $t_{j_2} > t_{j_1}$, $\forall j_1, j_2 \in J, j_1 \neq j_2$). We may also assume with no loss of generality that the ordering of the index set of orders J is the same as that of submission times, that is, $j_1 > j_2 \equiv t_{j_1} > t_{j_2}, \forall j_1, j_2 \in J$.

Definition 5.2 (*Primal lexicographical ordering*) Let $X^{(1)*} = \{x_{j,k}^{(1)*}\}_{j \in J_k, k \in K}$ and $X^{(2)*} = \{x_{j,k}^{(2)*}\}_{j \in J_k, k \in K}$ be two different optimal allocations (i.e., optimal solutions of model (MC-B-1)). We say that $X^{(1)*}$ is lexicographically better than $X^{(2)*}$ with respect to submission times, and denote $X^{(1)*} \succ_P X^{(2)*}$, if there exists $j \in J$ such that

1. $x_{j',j'}^{(1)*} = x_{j',j'}^{(2)*}, \forall j' \leq j - 1$; and
2. $x_{j,j}^{(1)*} > x_{j,j}^{(2)*}$.

In other terms, the optimal allocation $X^{(1)*}$ is lexicographically better than the optimal allocation $X^{(2)*}$ with respect to submission times if there exists an index j' such that the j'^{th} order has a larger execution proportion in $X^{(1)*}$ than in $X^{(2)*}$, while the first $j' - 1$ orders have equal execution proportions in $X^{(1)*}$ and $X^{(2)*}$. A similar definition can be proposed for a dual lexicographical ordering, where bundle prices are compared instead of execution proportions.

Definition 5.3 (*Dual lexicographical ordering*) Let $Y^{(1)*} = [\{\Pi_i^{(1)*}\}_{i \in I}, \{\mu_j^{(1)*}\}_{j \in J_k, k \in K}]$ and $Y^{(2)*} = [\{\Pi_i^{(2)*}\}_{i \in I}, \{\mu_j^{(2)*}\}_{j \in J_k, k \in K}]$ be two different optimal set of prices (i.e., dual solutions of model (MC-B-1)), and consider bundle prices computed on the basis

of $Y^{(1)*}$ and $Y^{(2)*}$, that is, $\bar{P}_j^{(1)} = \sum_{i \in I} q_{ji} \Pi_i^{(1)*}$, $\bar{P}_j^{(2)} = \sum_{i \in I} q_{ji} \Pi_i^{(2)*}$, $j \in J_k, k \in K$. We say that $Y^{(1)*}$ is lexicographically better than $Y^{(2)*}$ with respect to submission times, and denote $Y^{(1)*} \succ_D Y^{(2)*}$, if there is $j \in J$ such that

1. $\bar{P}_{j'}^{(1)} = \bar{P}_{j'}^{(2)}, \forall j' \leq j - 1$; and
2. $\bar{P}_j^{(1)} > \bar{P}_j^{(2)}$.

We next propose a post-optimality procedure (Algorithm 1) that discriminates between optimal allocations of model **(MC-B-1)** using lexicographical ordering \succ_P . Let \mathcal{P} denote the polyhedron that represents the set of optimal allocations of model **(MC-B-1)**. The procedure consists in a search algorithm that works on the vertices of \mathcal{P} . Starting at an arbitrary vertex of \mathcal{P} , which corresponds to a basic optimal solution of **(MC-B-1)**, the algorithm constructs a sequence of vertices of \mathcal{P} in which a move from a vertex to the next one is performed similarly to the simplex algorithm, and in such a way that ordering \succ_P is improved. In practice, these moves correspond to (degenerate) pivots, driven by \succ_P , of the simplex method. Once the algorithm is unable to locally improve the current optimal solution with respect to \succ_P , it returns it as the best allocation found.

Algorithm 1 Primal Discrimination

Require: Solve **(MC-B-1)**; Let X^* be an arbitrary basic optimal solution of **(MC-B-1)**.
if X^* is unique **then**
 Stop; **return** X^*
else {There are multiple optimal solutions}
 $SOL \Leftarrow X^*$; BestAllocFound \Leftarrow FALSE
 while NOT (BestAllocFound) **do**
 $\mathcal{A} \Leftarrow$ set of basic optimal solutions **(MC-B-1)** that are adjacent to SOL
 if $SOL \succ_P S^*, \forall S^* \in \mathcal{A}$ **then**
 BestAllocFound \Leftarrow TRUE
 else
 Choose $S^* \in \mathcal{A}$ such that $S^* \succ_P SOL$
 $SOL \Leftarrow S^*$
 end if
 end while
 return SOL
end if
Ensure: SOL is lexicographically the best optimal allocation according to submission times

To prove the validity of the primal discrimination procedure, let us first consider the following result.

Lemma 5.1 Let \mathcal{P} be a polytope (bounded polyhedron) and x^0 a vertex of \mathcal{P} . Suppose there are l other vertices of \mathcal{P} , x^1, \dots, x^l that are adjacent to x^0 . Now consider x^{l+1} , a vertex of \mathcal{P} that is not adjacent to x^0 . Then, there exists $z = \lambda x^{l+1} + (1 - \lambda)x^0, 0 < \lambda < 1$ such that $z \in \mathcal{P}' = \text{CONV}(x^0, \dots, x^l)$.

Proof. Consider polyhedron \mathcal{R} defined by vertex x^0 and extremal rays $r^i = x^i - x^0, 1 \leq i \leq l$, that is $\mathcal{R} = \{x : x = x^0 + \sum_{1 \leq i \leq l} \alpha_i r^i, \alpha_i \geq 0, 1 \leq i \leq l\}$. We have indeed that $\mathcal{P} \subset \mathcal{R}$: if $\{x : Ax \leq b\}$ is a representation of \mathcal{P} , then \mathcal{R} is the intersection of all half-spaces defined by facets of \mathcal{P} that are binding at x^0 . Hence, there exist $\alpha_1, \dots, \alpha_l \geq 0$, such that $x^{r+1} = x^0 + \sum_{1 \leq i \leq l} \alpha_i r^i$, and $\sum_{1 \leq i \leq l} \alpha_i > 0$ since $x^{r+1} \neq x^0$. Consider now $z = \lambda x^{r+1} + (1 - \lambda)x^0$. It is easy to verify that if $0 < \lambda \leq \frac{1}{\sum_{1 \leq i \leq l} \alpha_i r^i}$, then $z \in \text{CONV}(x^0, \dots, x^l)$. ■

Proposition 5.2 The discrimination procedure of Algorithm 1 terminates after a finite number of iterations, and provides an optimal solution that is lexicographically the best with respect to submission times.

Proof. Consider the polytope \mathcal{P} of optimal solutions of model (MC-B-1). It is easy to verify that the discrimination procedure terminates in a finite number of iterations: at an optimal solution SOL , either the algorithm moves to an adjacent basic solution that is lexicographically better than SOL , or it returns SOL as the best solution found. Since there is a finite number of basic solutions of (MC-B-1), the algorithm cannot improve indefinitely optimal solutions with respect to lexicographical ordering \succ_P .

We need to prove that SOL , the optimal solution the algorithm returns, is really the best according to lexicographical ordering \succ_P . First, we note that no basic optimal solution of (MC-B-1) that is adjacent to SOL is better than SOL with respect to \succ_P . Suppose now that x^1, \dots, x^l are basic optimal solutions adjacent to SOL , and there exists x^{l+1} , an optimal basic solution that is not adjacent to SOL , which verifies $x^{l+1} \succ_P SOL$. According to Lemma 5.1, there exists $z = \lambda x^{l+1} + (1 - \lambda)SOL, 0 < \lambda < 1$, such that $z \in \mathcal{P}' = \text{CONV}(SOL, x^1, \dots, x^l)$. It is

easy to establish that $z \succ_P SOL$. On the other hand, since $SOL \succ_P x^i, \forall i, 1 \leq i \leq l$, then $\forall i \in \{1, \dots, l\}, \exists j'_i \in J$, such that $SOL_j = x_j^i, \forall j \leq j'_i - 1$, and $SOL_{j'_i} > x_{j'_i}^i$. Now, if $j_{min} = \min\{j'_i\}_{1 \leq i \leq l}$, then $SOL_j = z_j, \forall j \leq j_{min} - 1$, and $SOL_{j_{min}} > z_{j_{min}}$, which contradicts the previous statement that $z \succ_P SOL$. Therefore, SOL is the best basic optimal solution according to lexicographical ordering \succ_P , and overall the best optimal solution since all optimal solutions in \mathcal{P} can be represented as convex combinations of basic optimal solutions of \mathcal{P} . ■

The dual discrimination procedure presented in Algorithm 2 is quite similar to the primal procedure of Algorithm 1. The only differences are that the dual procedure works on the space of the dual optimal solutions of model **(MC-B-1)** and relies on lexicographical ordering \succ_D to separate optimal set of prices.

Algorithm 2 Dual Discrimination

Require: Solve **(D)**; Let Y^* be an arbitrary basic optimal solution of **(D)**.

```

if  $Y^*$  is unique then
    Stop; return  $Y^*$ 
else {There are multiple optimal solutions}
     $SOL \leftarrow Y^*$ ; BestPricesFound  $\leftarrow$  FALSE
    while NOT (BestPricesFound) do
         $\mathcal{D} \leftarrow$  set of basic optimal solutions of (D) that are adjacent to  $SOL$ 
        if  $SOL \succ_D S^*, \forall S^* \in \mathcal{D}$  then
            BestPricesFound  $\leftarrow$  TRUE
        else
            Choose  $S^* \in \mathcal{D}$  such that  $S^* \succ_D SOL$ 
             $SOL \leftarrow S^*$ 
            end if
        end while
    return  $SOL$ 
end if

```

Ensure: SOL is lexicographically the best optimal set of prices according to submission times

5.4 Experimental analysis

In this section, we present the main numerical results and conclusions of our computational study. The experiments involved several data sets corresponding to instances of bundle trading allocation models **(MC-B-1)** and **(MC-B-2)**. Each data set comprised several series of randomly generated test problems, with the following general characteristics : 200 to 2000 assets in 27 different sectors, 100 traders, and up to 40 bundle orders per trader. All computational testing was carried out on a SUN

Enterprise 10000, with SunOS 5.8 as the operating system. Test problems generation procedures and the solution algorithms were coded in C++, and used the ILOG CPLEX 7.1 MIP solver callable library, with no particular parameter tuning.

Bundle order generation is a key aspect of our experimental study and needs some explanation. Algorithm 3 describes the implemented procedure for the generation of trader k 's bundles, $k \in K$. The first part of the procedure, intended to define the composition of the bundles, proceeds as follows. First, the bundle size (the number of assets traded in the bundle) s is drawn according to a discrete uniform distribution on $\{s_{\min}, \dots, s_{\max}\}$; s different assets are then uniformly chosen at random in I and the corresponding maximum sale or buy volumes are generated with respect to a uniform distribution over the interval $[Q_{\min}, Q_{\max}]$. The formation of the bundle price follows the model depicted in Figure 5.1. The model assumes the availability to all the traders of a common source of information providing them with freely observable prices (“quotations”) $\{p_i\}_{i \in I}$ of the assets. For instance, in the “secondary” market context particularly relevant for the bundle trading model, that is, a private exchange set up by a group of traders in order to balance their portfolios, these quotations could be the last day's close prices in the “main” markets in which the assets are publicly exchanged. Based on the observed prices $\{p_i\}_{i \in I}$ and its own analysis, trader k derives through Module 1 prices $\{\hat{p}_i^k\}_{i \in I}$, which represent what the trader believes are fair unit prices to pay or receive when the assets in question are traded on an *individual basis*. Finally, a bundle-specific layer of analysis (Module 2) evaluates the price that should be associated to a particular bundle, given single-asset prices $\{\hat{p}_i^k\}_{i \in I}$ and the composition of the bundle.

In practice, the decision process according to which traders decide on the composition of their bundle orders and the corresponding bid prices needs to incorporate many considerations such as the traders' speculations about the future payoffs of the assets, their willingness to bear risk, and the complex interdependencies between different assets. Consequently, the bundle formation procedures could actually be far more sophisticated than what the above-mentioned model suggests. Nevertheless, whether the bidding strategies of the traders are elaborated or not is only peripheral to the scope of the paper and does not impact the conclusions of our numerical study.

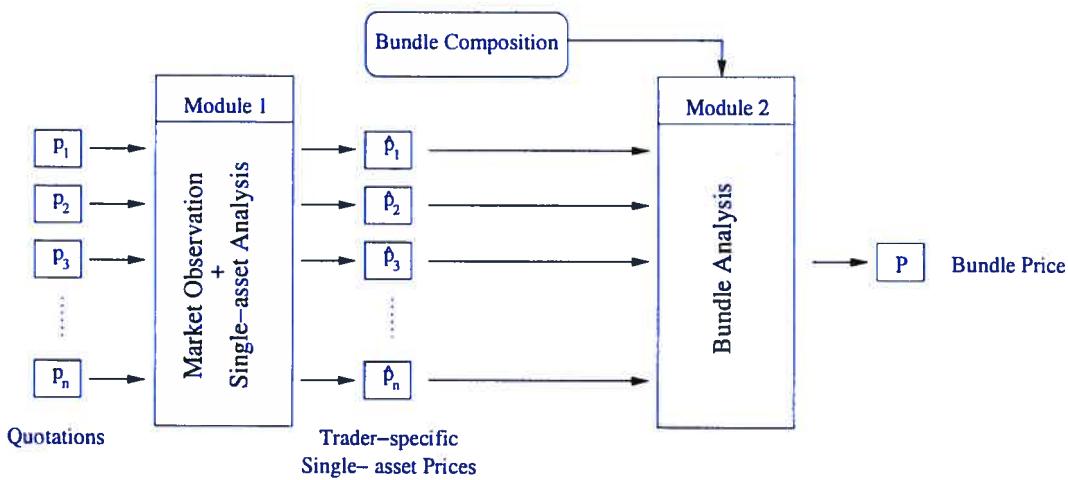


FIG. 5.1 – A model for bundle price formation

The following simple implementation is therefore carried out in Algorithm 3 :

- The single-asset analysis module is “simulated” through trader-specific perturbations of the quotations. That is, $\hat{p}_i^k = p_i(1+\tilde{\delta}_i^k)$, $\forall i \in I$, where $\tilde{\delta}_i^k$ is a uniformly distributed random variable over the interval $[-\delta_i^k, \delta_i^k]$, and δ_i^k is the maximal amplitude of the perturbation by trader k of the quotation p_i , $i \in I$.
- *Linear* bundle prices. $P_j = (\sum_{i \in I} q_{ji} \hat{p}_i^k)$, $\forall j \in J_k, \forall k \in K$. In particular, we do not consider the possibility that an investor reflect any “added-value” that may result from the opportunity of executing the single-asset trade orders in a bundle in equal proportions. Our rationale for making this assumption is the desire to realize a fair and unbiased evaluation of the gross effect of trading assets in bundles.

The experiments were conducted with two objectives in mind. We first intended to appraise the “bundle effect” in the basic allocation model (**MC-B-1**). A formulation of traditional single-asset market-clearing mechanisms is thus required as a benchmark. In that regard, we adopted a simple *disaggregation* approach that consists in decomposing each bundle order into its elementary components, that is, the single-asset sell and buy orders from which the bundle is made. More precisely, consider a bundle order $j : O_j = (\{q_{ji}\}_{i \in I}, P_j)$ formulated by trader $k \in K$, and let I_j denote the subset of the assets that are traded in order j , i.e. $I_j = \{i \in I : q_{ji} \neq 0\}$. Bundle

Algorithm 3 Bundle order generation

Require: I ; k ; J_k ; $\{p_i\}_{i \in I}$ = asset quotations; Q_{min}, Q_{max} = limits on the max. volumes q_{ji} ; s_{min}, s_{max} = limits on the bundle size; δ_i^k = max. perturbation of quotation p_i for trader k

{Bundle Composition}

```
for  $j \in J_k$  do
     $s \leftarrow Uniform(s_{min}, \dots, s_{max})$ ; choose  $I_j \subseteq I$  s.t.  $|I_j| = s$ 
    for  $i \in I$  do
        if  $i \in I_j$  then
             $q_{ji} \leftarrow \pm Uniform(Q_{min}, Q_{max})$ 
        else
             $q_{ji} \leftarrow 0$ 
        end if
    end for
end for
```

{Simulate quotations observation and single-asset analysis}

```
for  $i \in I$  do
     $\hat{p}_i^k \leftarrow p_i * (1 + Uniform(-\delta_i^k, \delta_i^k))$ 
end for
```

{Bundle prices}

```
for  $j \in J_k$  do
     $P_j \leftarrow 0$ 
    for  $i \in I$  do
         $P_j \leftarrow q_{ji} \hat{p}_i^k$ 
    end for
end for
```

Ensure: Vectors $O_j = (\{q_{ji}\}_{i \in I}, P_j)$, $j \in J_k$

order j gives rise to the collection $\{o_{ji} = (q_{ji}, q_{ji}\hat{p}_i^k)\}_{i \in I_j}$ of single-asset trade orders. Let x_{ji} be the execution proportion of order o_{ji} , $j \in J, i \in I_j$. The allocation problem of the single-asset market corresponds to

$$(\text{MC-SA}) \quad \max_{k \in K} \sum_{j \in J_k} \sum_{i \in I_j} q_{ji} \hat{p}_i^k x_{ji} \quad (5.18)$$

$$\text{s.t.} \quad \sum_{k \in K} \sum_{j \in J_k : q_{ji} \neq 0} q_{ji} x_{ji} = 0, \quad i \in I \quad (5.19)$$

$$0 \leq x_{ji} \leq 1, \quad k \in K, j \in J_k, i \in I_j \quad (5.20)$$

The market-clearing formulation **(MC-SA)** thus preserves the main trading objectives of each trader, but relaxes the requirement that sell and buy orders in a bundle must be executed in equal proportions. Comparisons between models **(MC-B-1)** and **(MC-SA)** relied on two evaluation metrics : (i) the total economic surplus achieved by the market ; and (ii) the number of cumulative value aggregation occurrences, which are the cases where a trader obtains better execution proportions on some of the assets it desires to sell or buy when bundle-based market clearing is used instead of single-asset allocation. Explicitly, the latter are the number of bundle orders $j \in J$ such that $x_j^* > \min_{i \in I_j} x_{ji}^*$, where $\{x_j^*\}_{j \in J}$ and $\{x_{ji}^*\}_{j \in J, i \in I_j}$ denote optimal allocations achieved by **(MC-B-1)** and **(MC-SA)**, respectively. We also investigated how various parameters of the allocation problem (number of assets and orders, bundle size, etc.) influence the efficiency of the market mechanisms with regard to economic surplus and value aggregation.

Our second objective was to estimate, from economic and computational perspectives, the impact of the bidding requirements added to the combinatorial market-clearing formulation **(MC-B-2)**. We focused on lower bounds and XOR relations and considered instances of model **(MC-B-2)** in which constraints corresponding to one or the other of these two classes of requirements are generated. We then measured relative gains or losses in economic surplus, using the basic allocation model **(MC-B-1)** as the formulation of reference. We also report integrality gaps and CPU solution times of the corresponding MIP problems, which helps gain insights about

the complexity of the allocation model (**MC-B-2**).

5.4.1 Basic bundle-based problems

DATASET-1 consists of several series of test problems that correspond to the basic bundle-based market-clearing formulation (**MC-B-1**). The structure of these problems is shown in Table 5.2. Each series is made of 150 test problems, equally divided into three groups : problems with small bundle orders (3 to 5 assets), medium bundles (10 to 20 assets), and large bundles (30 to 50 assets). For all the experiments, we use the following values : $Q_{min} = 1000$, $Q_{max} = 10000$, and $\delta_i^k = 10\%, \forall i \in I, \forall k \in K$.

Problem series	Problem description			
	#assets	#traders	#orders per trader	(#orders/#assets)
$S_B - 01$	200	100	3	3/2
$S_B - 02$	300	100	4	4/3
$S_B - 03$	400	100	5	5/4
$S_B - 04$	400	100	10	5/2
$S_B - 05$	500	100	10	2
$S_B - 06$	1000	100	15	3/2
$S_B - 07$	1000	100	20	2
$S_B - 08$	2000	100	30	3/2
$S_B - 09$	2000	100	40	2

TAB. 5.2 – DATASET-1 - Basic bundle trading allocation problems

The first series of results obtained on DATASET-1 problems, summarized in Figure 5.2, uses the economic surplus achieved by the single-asset and the bundle trading mechanisms as the criterion to characterize the *liquidity* of the corresponding markets. Hence, we measured the ratio of the surplus of the basic bundle-trading model (**MC-B-1**) to that of the corresponding disaggregated formulation (**MC-SA**), and took the mean of that ratio over the 50 problems of each series. In light of the results, the following observations can be made :

- The market liquidity of the bundle-based mechanism is relatively poor in general. More to the point, the average surplus ratios are always under 30%, and the mean of these ratios over all DATASET-1 problems does not exceed 7%. Given the structure of the allocation models (**MC-B-1**) and (**MC-SA**), this

decrease in market liquidity may be logically attributed to the additional requirement in the bundle-based formulation that the trade orders in a same bundle be executed according to the same proportions.

- Two factors that significantly influence market liquidity are the bundle size and the (#orders/#assets) ratio of the number of orders submitted to the market, to the number of assets traded. Concerning the bundle size, we report average surplus ratios of 16.26%, 3.19%, and 1.01% over all the small, medium, and large bundle problems, respectively. This clearly indicates that the liquidity of the bundle-based market decreases as the traders are allowed to submit larger bundles. The effect of the (#orders/#assets) ratio is similar. These findings are not surprising, given that (a) the bundle size is directly related to the number of equal execution proportion constraints inherent to the formulation (**MC-B-1**), and (b) increasing the number of bundle orders relative to the number of assets improves the chances of matching these orders.

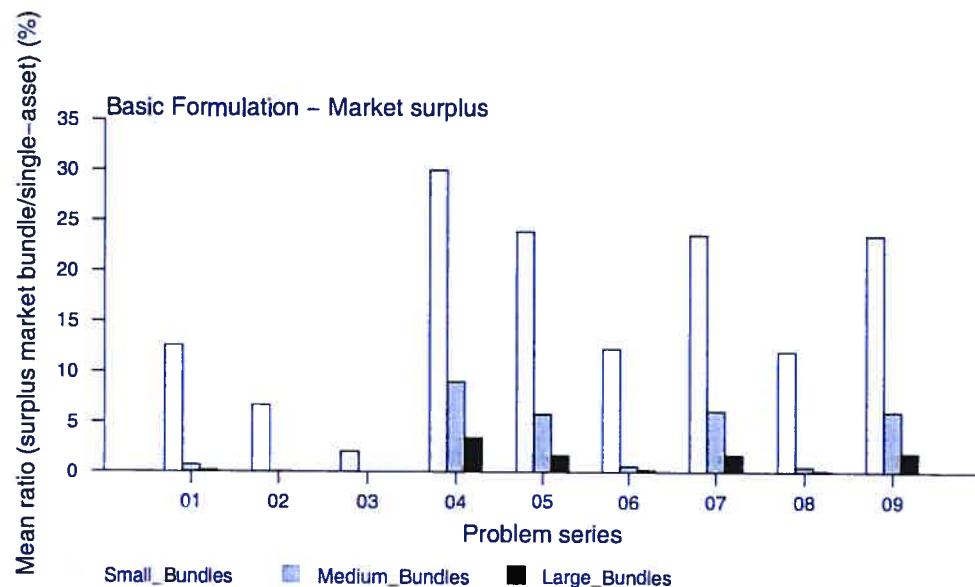


FIG. 5.2 – DATASET-1 : Market surplus

The results shown in Figure 5.3 also indicate that cumulative value aggregation occurs quite frequently. For instance, out of 4000 bundle orders, 2627, 2602, and 2580

occurrences of value aggregation could be enumerated on average for small, medium, and large bundle problems of the $S_B - 09$ series, respectively. Overall, 64% of all the submitted orders display value aggregation on average. Furthermore, it is interesting to note that value aggregation, unlike market liquidity, seems not to depend on the bundle size or the (#orders/#assets) ratio. This observation tends to confirm the intuitive idea that, an investor with a primary objective is to balance (even partially) its portfolios will generally achieve a better satisfaction of this objective with a bundle-based mechanism, irrespective of the number of bundle orders it submits to the market and their size. The corresponding advantage of bundle-based markets may thus be solely accredited to the fact that traders are allowed to “mix in a same bag” highly-priced single-asset orders with lowly-priced ones.

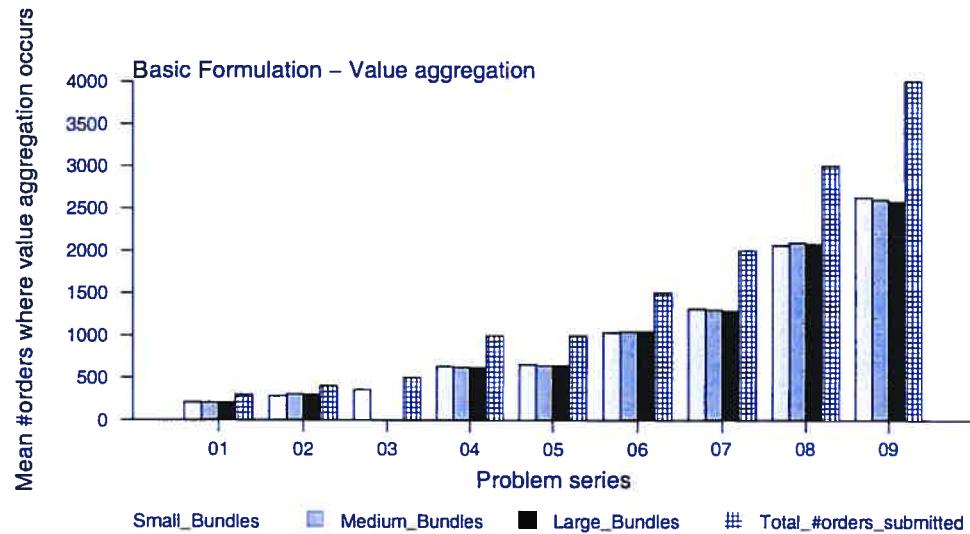


FIG. 5.3 – DATASET-1 : Cumulative value aggregation

5.4.2 Lower bound problems

DATASET-2 is a set of problems representing instances of the combinatorial market-clearing model (**MC-B-2**) with lower bound constraints. The attributes of DATASET-2 problems are shown in Table 5.3. Two additional parameters characterize the generation of lower bounds : F_{LB} , the lower bound *frequency*, indicates the

probability that a lower bound constraint is attached to a bundle order; and the lower bound *amplitude* LB_{max} , which is the maximum value (comprised between 0 and 1) that the corresponding bound l_j in formulation (MC-B-2) can take (in our implementation, l_j is uniformly distributed over $[0, LB_{max}]$).

Problem series	Problem description					
	#assets	#traders	#orders per trader	(#orders/#assets)	F_{LB}	LB_{max}
$S_{LB} - 01$	200	100	3	3/2	1/10	0.2
$S_{LB} - 02$	200	100	4	2	1/10	0.2
$S_{LB} - 03$	500	100	8	8/5	1/10	0.2
$S_{LB} - 04$	500	100	10	2	1/10	0.2
$S_{LB} - 05$	1000	100	20	2	1/10	0.2
$S_{LB} - 06$	200	100	3	3/2	1/10	0.9
$S_{LB} - 07$	500	100	10	2	1/10	0.9
$S_{LB} - 08$	1000	100	20	2	1/10	0.9
$S_{LB} - 09$	200	100	3	3/2	1/3	0.2
$S_{LB} - 10$	500	100	10	2	1/3	0.2
$S_{LB} - 11$	1000	100	20	2	1/3	0.2
$S_{LB} - 12$	200	100	3	3/2	1/3	0.9
$S_{LB} - 13$	500	100	10	2	1/3	0.9
$S_{LB} - 14$	1000	100	20	2	1/3	0.9

TAB. 5.3 – DATASET-2 - Lower bound allocation problems

Before proceeding with the analysis of the results obtained on DATASET-2 problems, it is important to note that integrality gaps of lower bound problems can actually be related to the economic surplus of the market. Let (P_{LB}) be the MIP allocation problem corresponding to a DATASET-2 instance, and (P) the allocation problem that results from the relaxation of the lower bound constraints in (P_{LB}) . It is easy to verify that an optimal solution $(\tilde{x}^*, \tilde{y}^*)$ of the *linear programming relaxation* of (P_{LB}) can be straightforwardly derived from an optimal solution x^* of (P) : let $\tilde{x}^* = x^*$, and \tilde{y}^* be such that $x_j^* \leq \tilde{y}_j^* \leq \min\{\frac{x_j^*}{l_j}, 1\}$, $\forall k \in K, \forall j \in J_k$, s.t., $l_j \neq 0$. It can therefore be argued that the integrality gap of a lower bound problem provides an indication on the relative “loss” in the economic surplus of the market when lower bound constraints are taken into account.

Integrality gaps and CPU solution times, averaged over the DATASET-2 instances, are visualized in Figure 5.4. As one may expect, integrality gaps are clearly influen-

ced by the values taken by parameters F_{LB} and LB_{max} . In that regard, DATASET-2 problems are separated in 4 groups, each group being characterized by a pair (F_{LB}, LB_{max}) for $F_{LB} \in \{1/10, 1/3\}$ and $LB_{max} \in \{0.2, 0.9\}$. The average integrality gaps are equal to 20.53%, 38.59%, 27.52%, and 75.13% for the four groups, respectively. These average gaps confirm numerically that lower bound problems become “tighter” as lower bounds constraints are more frequently associated by bidders to their bundle orders, and as the corresponding bounds become larger. Bundle size has also a noticeable impact on the magnitude of integrality gaps, as problems with larger bundles consistently display larger gaps than similar problems (using the same values of F_{LB} and LB_{max}) with smaller bundles. A closer examination of the results enlightens the influence of a fourth parameter, the (#orders/#assets) ratio : for example, the average integrality gaps over instances in series $S_{LB} - 01$ to $S_{LB} - 05$ (58.50%, 1.86%, 13.33%, 1.68%, 1.40%) seem to vary in accordance with the corresponding ratios (3/2, 2, 8/5, 2, 2). This suggests that “sparse” instances - from the perspective of the ratio of the number of bundles to the number of assets - are likely to be more sensitive to the presence of lower bounds than “dense” ones.

The CPU solution times grow exponentially with respect to the number of assets traded in the market. One may also notice that problems with larger bundles are in general significantly more costly to solve, which is consistent with our previous observations on integrality gaps, with the notable exception of problem series $S_{LB} - 13$ and $S_{LB} - 14$. The rationale for this behavior : due to the combined effects of the bundle size and the large bounds frequently associated to the bundles, the Branch & Bound procedure generates a relatively large number of nodes at which the corresponding linear programming relaxations are infeasible, which favors the pruning of the enumeration tree and thus limits its size.

5.4.3 XOR problems

XOR relations formulate traders’ bidding requirements on the execution of “equivalent” bundle orders. A complete characterization of the equivalence relationship is a difficult task that relies on the good understanding of the traders’ profiles, the objectives that drive their trade operations, etc. Notwithstanding its importance, this issue

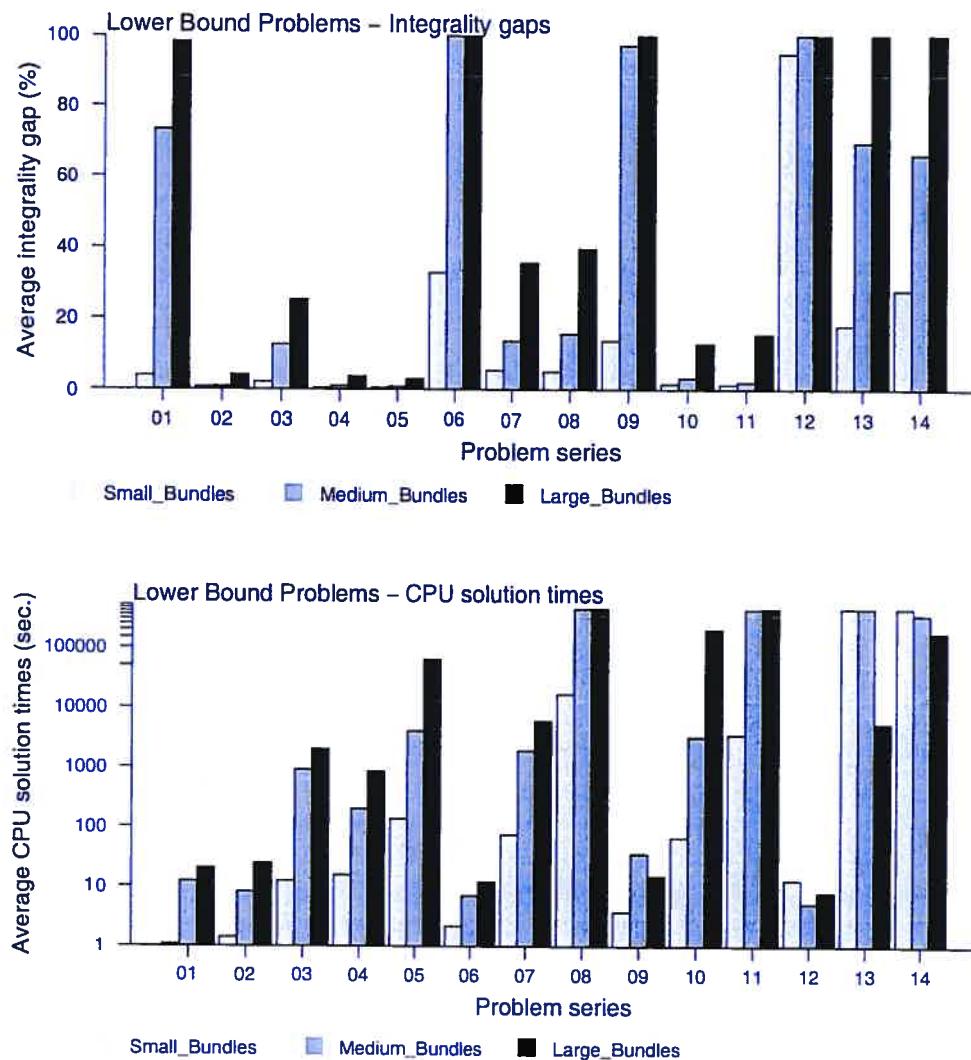


FIG. 5.4 – Integrality gaps and CPU times for DATASET-2 test problems

is well beyond the scope of our study, and we have thus used simple empirical techniques to generate XOR relations. Such a technique, which we call *asset-switching*, can be described as follows. Given a bundle order $O_j = (\{q_{ji}\}_{i \in I}, P_j)$, an asset $i_0 \in I$ belonging to activity sector A and traded in O_j ($q_{j i_0} > 0$, for instance) is arbitrarily selected. Asset-switching consists in choosing an asset $i_1 \in I$ that also belongs to activity sector A but is not traded in O_j , and building a new bundle $O_{j'} = (\{q_{j'i}\}_{i \in I}, P_{j'})$ such that $q_{j'i_0} = 0$, $q_{j'i_1} > 0$, and $q_{j'i} = q_{j i_0}$, $\forall i \in I$, s.t. $i \neq i_0, i \neq i_1$. The choice of asset i_1 and the computation of $q_{j'i_1}$ are carried out such that orders O_j and $O_{j'}$ have approximately equivalent monetary values and some predefined conditions on demand and supply of assets in the portfolio are satisfied. Finally, a XOR relation \mathcal{X} , involving bundle orders O_j and $O_{j'}$, is generated. The process may be easily extended to generate XOR relations between more than two bundle orders.

DATASET-3, whose structure is shown in Table 5.4, is a set of problem series that represent instances of model (MC-B-2) with XOR relations. Parameter F_{XOR} , denoting the generation frequency of XOR relations, varies in the range $\{1/10, 1/3, 3/4\}$.

Problem series	Problem description				
	#assets	#traders	#orders per trader	(#orders/#assets)	F_{XOR}
$S_{XOR} - 01$	200	100	3	3/2	1/10
$S_{XOR} - 02$	500	100	10	2	1/10
$S_{XOR} - 03$	1000	100	20	2	1/10
$S_{XOR} - 04$	200	100	3	3/2	1/3
$S_{XOR} - 05$	500	100	10	2	1/3
$S_{XOR} - 06$	1000	100	20	2	1/3
$S_{XOR} - 07$	200	100	3	3/2	3/4
$S_{XOR} - 08$	500	100	10	2	3/4
$S_{XOR} - 09$	1000	100	20	2	3/4

TAB. 5.4 – DATASET-3 - XOR allocation problems

XOR relations clearly enhance the bidding flexibility of the market from the perspective of the traders and result in additional liquidity on the marketplace. Hence, we define the *economic gap* as a measure of the extra market liquidity induced by the XOR relations in a DATASET-3 instance. More precisely, let z_{XOR}^* and z^* denote the economic surpluses respectively achieved by the allocation problem (P_{XOR}) corresponding to a XOR instance and its associated basic bundle formulation (P) (no

XOR relations allowed). The economic gap of the XOR instance designates the ratio $(z_{XOR}^* - z^*)/z^*$. It is noteworthy that, contrary to the lower bound case, economic gaps of XOR problems need to be differentiated from integrality gaps.

Average economic gaps over DATASET-3 problems are represented in Figure 5.5. Basically, the results show that fairly high levels of additional liquidity, attributed to XOR relations, have been induced in the market : none of the average economic gaps we have obtained is below 12%, and the average gap is around 50% over all DATASET-3 instances. Expectedly, economic gaps are closely related to the generation frequency of XOR relations, as indicated by the average economic gaps (28.6%, 52.5%, and 69.4%) over the three classes of problems with $F_{XOR} = 1/10, 1/3, 3/4$, respectively. Furthermore, they seem to be more pronounced for instances with larger bundles.

Integrality gaps and CPU solution times, depicted in Figure 5.6, broadly follow the same tendency reported for lower bound problems, with respect to the impact of the bundle size, the (#orders/#assets) ratio, and the F_{XOR} frequency. It is however interesting to note that, while integrality gaps are comparatively small (below 15%), DATASET-3 problems are nevertheless as costly to solve to optimality as lower bound instances.

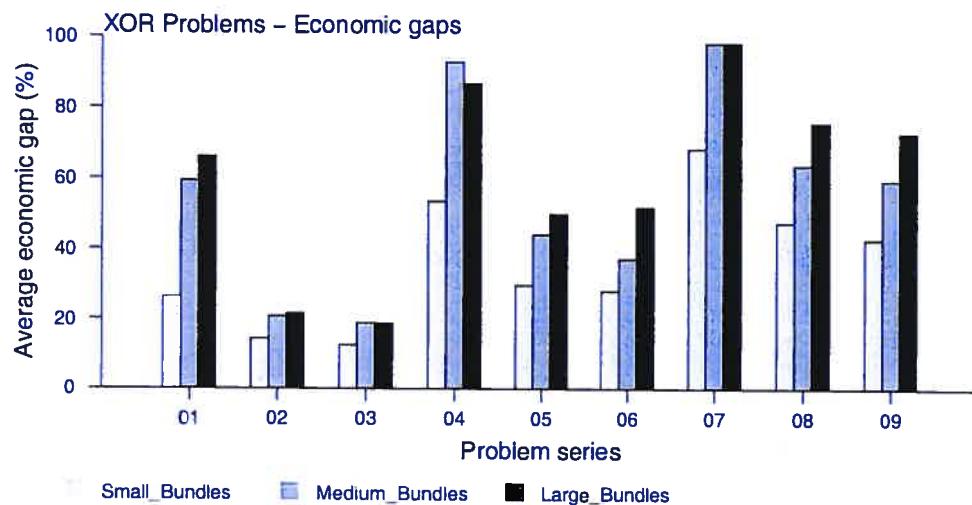


FIG. 5.5 – Economic gaps for DATASET-3 test problems

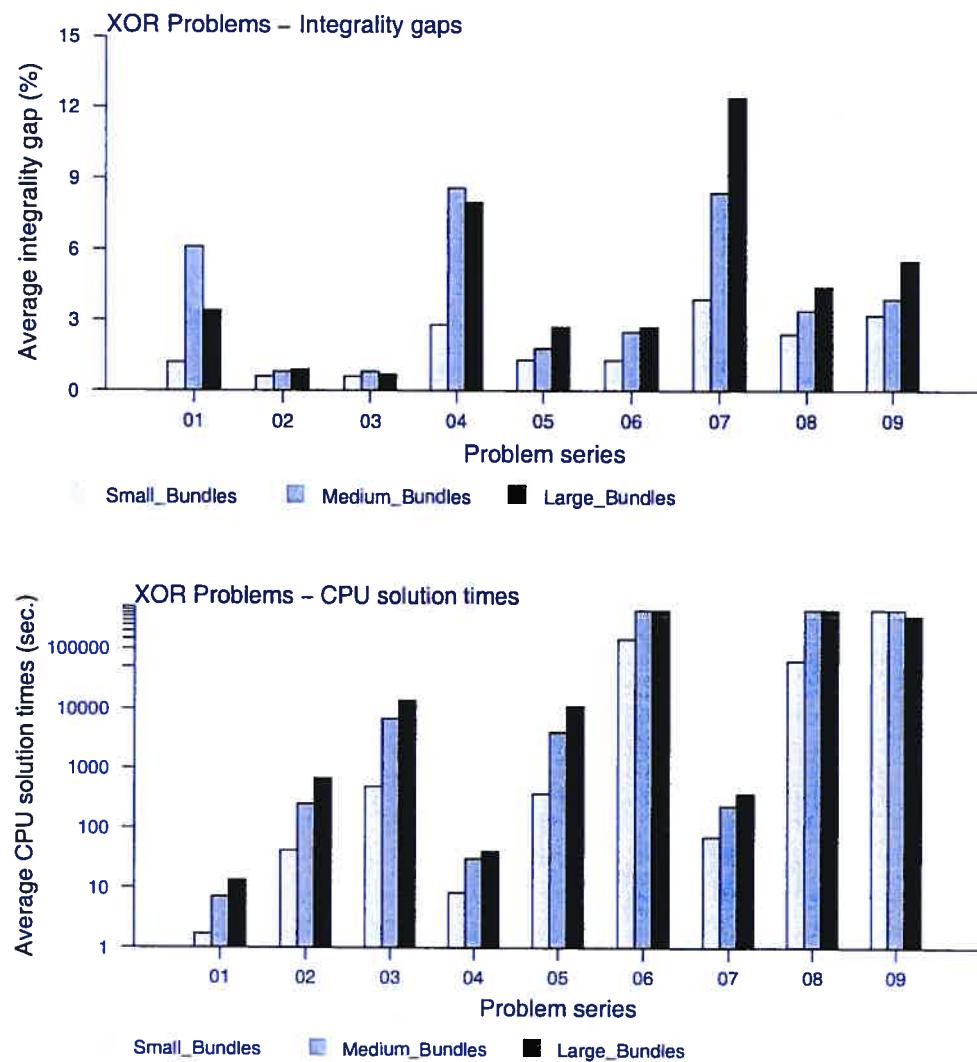


FIG. 5.6 – Integrality gaps and CPU times for DATASET-3 problems

5.5 Concluding remarks

In this paper, we have presented a new bundle-based market-clearing allocation model for portfolio balancing in financial marketplaces. Our formulation introduced classes of constraints that correspond to various bidder trading requirements. We have focused on three bidding operators that allow traders to set limitations on volumes of assets to be traded and on bundle execution proportions, as well as to define exclusive OR execution relations between “equivalent” bundles. We believe that, despite being very simple, these requirements are particularly important in the context of portfolio bundle trading, as they make it possible for traders to design more elaborate bidding strategies and enhance considerably the accuracy of their control over the composition of their portfolios. Interestingly enough, there is space in the proposed bidding operators for further generalization. In particular, it is possible to use a selection operator “ k -out-of- n ” as a natural extension of the XOR operator, and to go beyond one level of recursivity in the definition of XOR relations (e.g., accept XOR of XORs). With these extensions, traders should ultimately be able to express more detailed bidding requirements (e.g., devise complex “programs” involving several concurrent investment plans to choose from). Empirical investigation of the impact of recursively defined XOR relations on market-clearing formulations (Abrache *et al.* [1]) seems to indicate such constructions are unlikely to significantly raise the complexity of solving the allocation problem, as the corresponding new constraints are essentially the same as those representing XOR constraints (5.12) in model **(MC-B-2)**. This conjecture needs to be further investigated by numerical experimentation.

The experimental study examined continuous and combinatorial variants of the market-clearing model and verified the impact of several parameters of the allocation problem on allocation efficiency and computational complexity. Among the important findings of this study is the fact that bundle trading mechanisms, by their very nature, improve bidders’ chances of balancing portfolios but achieve relatively poor market surplus, which makes them appealing for private, inter-institution secondary markets specifically intended to realize portfolio balancing operations. We have furthermore confirmed that additional bidding requirements, especially XOR relations, can have

a significant economic impact on market surplus.

The combinatorial bundle-based market-clearing formulation (**MC-B-2**) raises theoretical and practical challenging issues. First of all, solutions times for large problems (1000 assets and more) have been high : for instance, most of the instances in series $S_{LB} - 14$ and $S_{XOR} - 9$ needed more than 12 hours of CPU time. This is a definitive bottleneck for a mechanism supposed to run at least once a day, especially if we account for the fact that the determination of payments may require *repeated* solution of the allocation problem (this is the case, for instance, of the VCG mechanism suggested in Section 5.2, which commands the determination of a surplus-maximizing allocation once without each trader). Thus, the necessity to put more efforts in developing efficient solution methods needs to be emphasized. Several avenues seem quite promising. An obvious one is the fine-tuning of the mixed-integer programming engine of CPLEX. More interestingly, we believe that the adaptation of the Branch & Bound algorithm to the special structure of the problem and the integration of strong valid inequalities should prove of prime importance. To give an idea of the last point, let us consider formulation (**MC-B-2**) and note that, if there exists a bundle order $j_0 \in J$ in which a given asset i is purchased ($q_{j_0i} > 0$) and such that *no* subset $J^{(s)}$ of s sell orders ($q_{ji} < 0$) could supply enough volume of asset i to execute order j_0 (i.e., $l_{j_0} q_{j_0i} + \sum_{j \in J^{(s)}} q_{ji} < 0$), then the cut $(s+1)y_{j_0} \leq \sum_{j \in J} y_j$ may be added to (**MC-B-2**).

We are also much interested in evaluating bundle trading market mechanisms over longer periods of time. Simulation techniques can then be extremely valuable in creating controlled environments for manipulating market conditions and trader profiles regarding asset preference, risk aversion, bid sophistication, and so on.

Acknowledgements This project has received financial support from the Bell Canada University Laboratories (LUB), the Natural Sciences and Engineering Research Council of Canada through its Basic Research and Collaborative Industry-University Research grant programs, the FCAR research development fund of the Province of Québec, and the Valorisation-Recherche Québec fund.

We would like to thank our colleagues Jacques Robert, Yves Richelle, Martin

Boyer, and Robert Gérin-Lajoie for valuable discussions on the matter of this paper. We also express our gratitude to Zine Khelil who helped in conducting the numerical tests.

Chapitre 6

Decomposition Methods and Iterative Combinatorial Auctions

La référence de cet article est :

J. Abrache, T.G. Crainic, and M. Gendreau, “Decomposition Methods and Iterative Combinatorial Auctions”, en préparation.

Decomposition Methods and Iterative Combinatorial Auctions

Jawad Abrache^{1,3}, Teodor Gabriel Crainic^{2,3}, and Michel Gendreau^{1,3}

¹ Département d'informatique et de recherche opérationnelle, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada, H3C 3J7.

² Département management et technologie, Université du Québec à Montréal, C.P. 6192, Succursale Centre-ville, Montréal, Canada H3C 4R2.

³ Centre de recherche sur les transports, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada, H3C 3J7.

Abstract. Iterative auctions are motivated by the fact that the market-maker often lacks 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 decomposition methods. We focus specifically on two well-known approaches : Lagrangian relaxation, and Dantzig-Wolfe decomposition.

Keywords : Combinatorial auctions ; Iterative auction mechanisms ; Primal-dual approaches ; Mathematical programming decomposition methods.

6.1 Introduction

The spectacular growth of electronic commerce and information technology has made the Internet a place of choice for conducting business. An immediate consequence of this trend is the creation of many public and private electronic marketplaces where huge volumes of goods and services are exchanged every day. The apparition of these new entities - the electronic marketplaces - has also brought forth many fundamental issues related to their structure and organization. Among these, the question of how the economic rules that govern the market should be designed is of prime importance, and contributed to the renewed interest in the economic, game-theoretic, and computational aspects of market mechanism design.

A market mechanism can be defined, in informal terms, as a set of deterministic rules that specify an allocation of the items traded in the market to the participants, and the corresponding payments the latter should make or receive. According to the economic theory of mechanism design (Mas-Colell, Whinston, and Green [95]), the information exchange between the market-maker and the participants distinguishes basically between two classes of market mechanisms, *direct-revelation* and *indirect-revelation* mechanisms. Hence, direct-revelation mechanisms require the participants to completely and truthfully report their *types* to the market-maker ; that is, in the context of a production-consumption economy, for instance, sellers should disclose their production technologies and cost functions, and the buyers their preferences for the goods and the consumption constraints. Assuming such information is available, the market-maker could formulate an optimization model that clears the market such that an objective of “effectiveness” (that could be the maximization of the total social welfare of the participants, the revenue of a seller, etc.) is achieved.

One could argue, however, that complete and truthful information revelation is too strong an assumption in most market contexts. First and foremost, participants are generally self-interested agents, which pursue their own goals independently of the market objectives. They may thus be unwilling to disclose, without careful screening, all their private data. In many cases, this data may simply be too complex to assess and to communicate in its entirety (Nisan [106]). Finally, a participant may possess only uncertain valuations of the goods, which it hopes to adjust by ob-

serving how other participants behave (e.g., auctions of art objects or rare items). Indirect-revelation mechanisms help circumvent some of these problems by departing from the paradigm of an omniscient market-maker, with systematic access to all the information. Iterative auctions constitute an important class of indirect-revelation mechanisms that allow for progressive and considerably less forceful disclosure of information. In each round of an iterative auction, participants 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 that give participants impetus to commit themselves further. It is noteworthy that bids do not have to represent complete preferences (they only reflect participants’ needs *given the observed signals*), nor to convey truthful information (unless the auction incentives prevent preference misrepresentation).

The design of iterative auction mechanisms for the optimization of distributed systems has attracted much attention recently. In particular, the Combinatorial Allocation Problem (CAP), which consists in determining a socially-efficient allocation of multiple indivisible goods to several potential buyers, has been thoroughly investigated. For the CAP, Parkes’ doctoral thesis (Parkes [113]) surveys the research path toward incentive-compatible, multi-round auctions that reconcile the market’s welfare maximization objective with self-interested buyers aiming at maximizing their profits. Most of these auctions can be interpreted as primal-dual algorithms that capitalize on linear programming formulations of the problem. Mathematical programming decomposition approaches are a second interesting avenue (Geoffrion [57]). These methods have been used for decades to address large-scale structured optimization problems. Nevertheless, 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 market properties of two important decomposition approaches, that is, Lagrangian relaxation and Dantzig-Wolfe (DW) decomposition.

We consider a general combinatorial exchange economy in which self-interested agents trade several different divisible goods. For that economy, we initially formulate the allocation and payment rules of an ideal direct-revelation mechanism aiming at

maximizing the overall social efficiency of the market. We then show that, under appropriate assumptions, the application of Lagrangian relaxation and Dantzig-Wolfe decomposition to the allocation problem in the centralized market leads to indirect mechanisms in which the agents' own interests are ultimately reconciled with the "global" objective of the market. We furthermore establish that subtle but important differences exist between Lagrangian relaxation and DW-based auction processes, in their use of the information disclosed by the agents through their bids and the assumptions these mechanisms make on the agents' strategic behavior.

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, De Vries and Vohra [39]), 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 combinatorial exchange context, in which many sellers and buyers interact. Moreover, the paper presents an original comparison of these methods from informational and strategic perspectives. Finally, the numerical results provide interesting insights into the potential and limitations of auction mechanisms based on the different methods.

The remainder of the paper is organized as follows. In Section 6.2, we survey the existing literature on iterative auction mechanisms for combinatorial markets. We put the emphasis on classical tâtonnement processes and auction mechanisms based on the primal/dual approach. We formulate in Section 6.3 the problems of determining a socially efficient allocation and corresponding equilibrium prices in a centralized many-to-many direct-revelation market. In the following two sections, we present two relaxation-based methods (using the subgradient algorithm and a bundle method, respectively) and the Dantzig-Wolfe decomposition scheme, and interpret them as iterative auction processes. Finally, we devote Section 6.6 to a preliminary experimental study aiming at the evaluation of the corresponding auctions.

6.2 Prior work

The theory of general competitive equilibrium (Mas-Colell, Whinston, and Green [95]) institutes the framework for the econometric analysis of markets involving the trade of several interdependent divisible goods. General equilibrium is specifically concerned with the study of the social efficiency of the allocation of the goods in the market, as well as the determination of states of equilibrium from which the participants would not depart. In particular, *Walrasian* equilibria, defined in terms of payments based on single-item prices, has been extensively studied in the economic literature and a focal result (Arrow and Debreu [6]) establishes their existence under certain conditions.

The price-tâtonnement process of Walras [141] is probably the first iterative scheme aiming at the determination of competitive equilibria. Walrasian tâtonnement proceeds by adjusting prices of the various goods one by one until equilibration of supply and demand of all the goods is realized, which is only guaranteed if additional conditions that insure the equilibrium *stability* are satisfied (for instance, the *gross-substitutes* condition, which states that the demand for a given item does not decrease if prices of other items increase, is sufficient for equilibrium stability (Arrow and Hahn [7])). Price-tâtonnement has been recently implemented in the WALRAS algorithm (Cheng and Wellman [30]). More generally, it has given rise to computational paradigms such as *market-oriented programming* (Wellman [143]), which model and implement resource allocation problems as distributed systems of autonomous agents reacting to a pricing system. Market-oriented programming approaches have been applied to many areas, among which the allocation of transportation services (Wellman [144]), quality-of-service allocation in multimedia applications (Yamaki, Wellman, and Ishida [152]), vehicle routing (Sandholm [127]), power load management (Ygge [153]), and decentralized scheduling (Wellman *et al.* [145]).

The other important class of market mechanisms based on tâtonnement are resource-directed quantity-tâtonnement processes. In an iteration of a quantity-tâtonnement process, the market-maker determines provisional allocations of goods and participants submit their *marginal* utilities, that is, the utilities for pro-

ducing or consuming one additional unit of each good. The allocations are adjusted by the market-maker such that participants with higher (lesser) marginal utilities are given more (less) of the goods, until a state in which everybody is ready to pay (buyers) or receive (sellers) the same marginal prices. Simple fixed and variable-stepsize adjustment methods have been suggested by Kurose and Simha [79], and a more efficient Newton-Raphson search method by Ygge and Akkermans [154].

Bertsekas [15] auction algorithm for the assignment problem is one of the first and most successful attempts to use an iterative distributed process for solving optimization problems. The algorithm develops basically a multi-round auction and capitalizes on the economic interpretation of complementary slackness conditions in linear programming as equilibrium conditions for self-interested agents. Hence, in a given iteration of the algorithm, each agent chooses the most “interesting” object at the current prices and bids a new value that beats the current price by a certain increment (computed to maintain complementary slackness throughout the auction). The market-maker then simply allocates the objects to their respective highest bids. Extensions of the auction algorithms to other network flow problems (notably the shortest path and the minimum cost flow problems) have been presented (Bertsekas [16], [17]). Another important iterative process for the assignment problem is due to Demange, Gale, and Sotomayor [40]. The process is a price-ascending auction in which prices on over-demanded goods are raised across the auction rounds, and can be seen as a particular implementation of the primal-dual algorithm (Bikhchandani *et al.* 2001). More importantly, the authors show that by choosing *minimal* over-demanded sets of items, the auctioneer produces “second-price” Vickrey payments. Demange, Gale, and Sotomayor mechanism is thus an incentive-compatible iterative auction for the assignment problem.

The study of iterative auctions for the Combinatorial Auction Problem (CAP) has been at the heart of recent research. The CAP is the problem of determining welfare-maximizing allocations of several indivisible items to buyers that have preferences for bundles of items. Bikhchandani and Mamer [19] investigate the existence of Walrasian equilibrium prices for the CAP and establish that it relies on the integra-

lity of the LP relaxation of the CAP. For their part, Bikhchandani and Ostroy (2000) suggest extended formulations of the CAP which imply that nonlinear (and possibly discriminatory) bundle equilibrium prices always exist and correspond to optimal dual solutions of the associated formulations. This allowed the development of primal-dual implementations (e.g., the iBundle family of iterative auctions, Parkes [111]). The existence of LP formulations of the CAP leads to iterative auctions implementing Vickrey payments in particular cases. Thus, when an “agents-are-substitutes” condition holds, which states that the impact on the optimal allocation of a group of participants exceeds the sum of the individual impacts of the participants, Bikhchandani and Ostroy [20] show that minimal competitive equilibrium prices correspond to Vickrey payments. Under the additional assumption of *gross substitutes*, primal-dual implementations have been suggested by Gul and Stacchetti [63] and Ausubel [9].

Decentralized market-based mechanisms have also been suggested for coordinating the flow of goods, information, or services within organizations. Important methodological developments in the analysis of information and decision distribution within organization include Malone, Yates, and Benjamin [91], which compares empirically decentralized decision mechanisms (more specifically markets) and centrally-coordinated one (hierarchies), and predicts that advances in IT will induce firms to shift from hierarchies to market-centric processes, by making information more accessible and less costly. Tan and Harker [139] focus on operation scheduling, for which they quantify centralized and decentralized workflow structures with respect to production, coordination and disruption costs. The results of their study indicate, notably, that a distributed design for operation scheduling is more attractive than scheduling by a central authority when job processing times are relatively long, communication costs are low, and agents are reliable. A distributed scheduling framework, based on the Lagrangian relaxation technique, appears in Kutanoglu and Wu [80]. The authors develop an iterative combinatorial auction scheme (*adaptative tâtonnement*) in which jobs requiring processing time submit utility-maximizing bids on time slots on the corresponding machines, and a coordinating agent allocates time slots and updates prices in proportion to the excess demand arising for time slots. The authors show also that Lagrangian relaxation using the subgradient algorithm

corresponds to a special case of adaptative tâtonnement. Lagrangian relaxation has been similarly suggested by Kim *et al.* [74] in the design of an electronic brokering process for truckload freight.

6.3 Centralized market-clearing for a combinatorial exchange economy

Consider an economy with a set of divisible goods on sale, and two categories of agents : sellers and buyers. Sellers have the capacity to produce the goods according to their own technology and production cost functions. Buyers, for their part, consume goods either directly or as inputs to a transformation process. Hence, they have preferences for bundles of goods on sale and may also face technological requirements that constrain their consumption.

Let us introduce the following notation :

- \mathcal{L} : set of goods ;
- \mathcal{S} : set of sellers ;
- \mathcal{J} : set of buyers ;
- $q_{s,l}$: quantity of good $l, l \in \mathcal{L}$ produced by seller $s, s \in \mathcal{S}$;
- $q_{j,l}$: quantity of good $l, l \in \mathcal{L}$ consumed by buyer $j, j \in \mathcal{J}$;
- \mathcal{D}_s : production feasibility set of seller s , containing all admissible quantity vectors $q_s = \{q_{s,l}\}_{l \in \mathcal{L}}$ that seller s may produce ;
- \mathcal{D}_j : consumption feasibility set of buyer j , containing all admissible quantity vectors $q_j = \{q_{j,l}\}_{l \in \mathcal{L}}$ that buyer j may consume ;
- $C_s(\cdot)$: 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 ;
- $V_j(\cdot)$: valuation function of buyer $j, j \in \mathcal{J}$; similarly, $V_j(q_j)$ is buyer j 's preference for consuming q_j .

We make the following assumptions on the feasibility sets and the cost and valuation functions :

(A1) Production cost functions are *continuous, convex, monotone increasing*

(i.e., $C_s(q_s^{(1)}) \geq C_s(q_s^{(2)}), \forall q_s^{(1)}, q_s^{(2)} \in \mathcal{D}_s$ s.t. $q_s^{(1)} \geq q_s^{(2)}$). Similarly, buyer valua-

tion functions are *continuous, concave, monotone increasing*.

(A2) Production and consumption feasibility sets are *convex* and *bounded*.

A direct-revelation market mechanism (see Figure 6.1) would be in this case a one-shot market-clearing process aimed at matching the supply and demand for the various goods. The 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 the goods and payments the sellers need to make to the buyers.

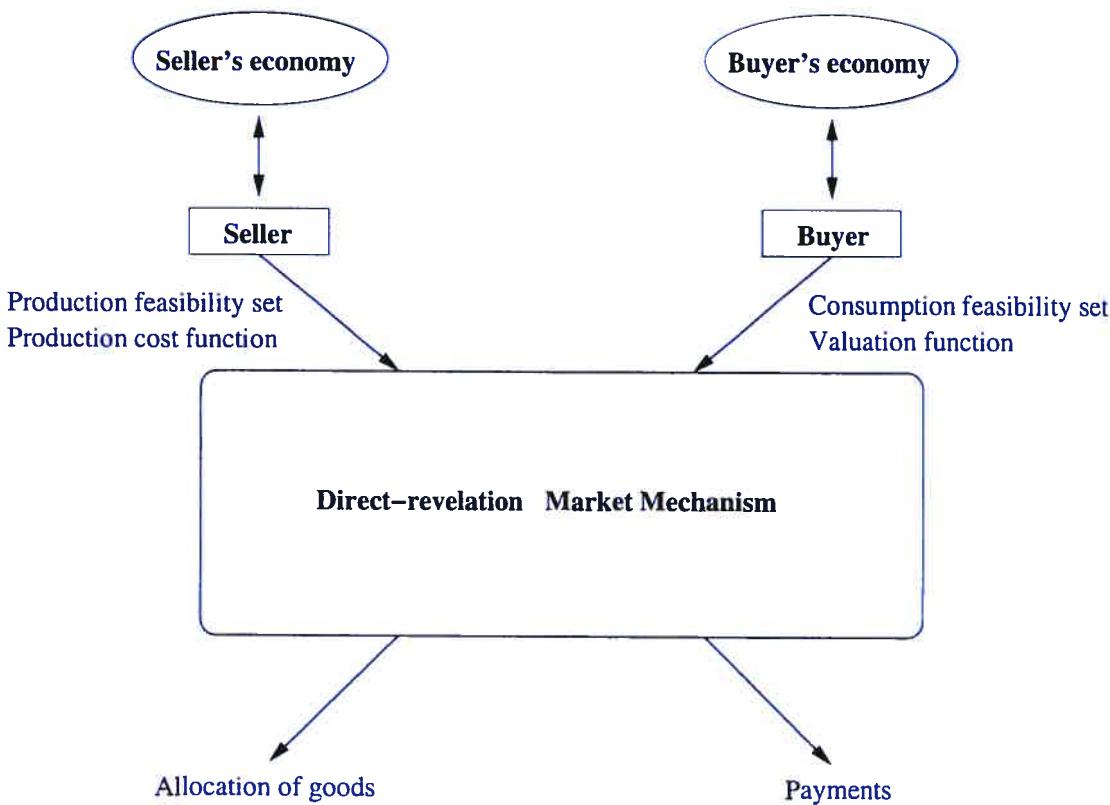


FIG. 6.1 – A direct-revelation market mechanism.

The primary goal of the market-maker is to determine a *socially-efficient* allocation, that is, a feasible allocation of the goods that maximizes the global welfare of the sellers and buyers. With the notation above, a socially-efficient allocation is a solution of :

$$\max \quad \sum_{j \in \mathcal{J}} V_j(q_j) - \sum_{s \in \mathcal{S}} C_s(q_s) \quad (6.1)$$

$$(\text{MC}) \quad s.t. \quad \sum_{j \in \mathcal{J}} q_{j,l} - \sum_{s \in \mathcal{S}} q_{s,l} = 0, \quad l \in \mathcal{L} \quad (6.2)$$

$$q_j \in \mathcal{D}_j, \quad j \in \mathcal{J} \quad (6.3)$$

$$q_s \in \mathcal{D}_s, \quad s \in \mathcal{S} \quad (6.4)$$

The objective of model (MC) is to maximize the market surplus, computed as the difference between the buyers' valuations and the buyers' production costs. Equations (6.2) matches the demand with the supply, while relations (6.3) and (6.4) constrain the quantities traded to be within admissible production and consumption limits.

The concept of pricing equilibria serves the market-maker's second important purpose, which is to reconcile the global objective of the market (social efficiency) with the sellers and buyers own interests. Prior to the definition of a Walrasian equilibrium, it is convenient to make the following classical assumptions :

(A3) Sellers and buyers have *quasi-linear* utilities, in which money is considered as an exogenous resource. That is, the utility of a buyer j with a wealth level m for consuming q_j is $u_j(q_j, m) = V_j(q_j) + m$, and that of a seller s producing q_s is $u_s(q_s, m) = -C_s(q_s) + m$.

(A4) Sellers and buyers are *price-takers*, i.e., they take market prices as given and ignore the effect of their own actions on prices.

Definition 6.1 (Walrasian Equilibrium) *The price vector $\{p_l\}_{l \in \mathcal{L}}$ constitutes a set of Walrasian prices if there is a feasible allocation $\tilde{q} = [\{\tilde{q}_j\}_{j \in \mathcal{J}}; \{\tilde{q}_s\}_{s \in \mathcal{S}}]$ such that*

$$1. \quad V_j(\tilde{q}_j) - p_l \cdot \tilde{q}_j = \max_{q_j \in \mathcal{D}_j} (V_j(q_j) - p_l \cdot q_j), \quad \forall j \in \mathcal{J}; \text{ and}$$

$$2. \quad p_l \cdot \tilde{q}_s - C_s(\tilde{q}_s) = \max_{q_s \in \mathcal{D}_s} (p_l \cdot q_s - C_s(q_s)), \quad \forall s \in \mathcal{S}.$$

6.4 Market-clearing based on Lagrangian relaxation

The particular structure of the centralized market-clearing formulation (6.1) naturally suggests Lagrangian relaxation as an approach to decompose the problem. Hence, consider (MC) as the primal problem and dualize it with respect to the matching constraints (6.2). Let $\lambda = \{\lambda_l\}_{l \in \mathcal{L}}$ be the vector of Lagrangian multipliers associated with (6.2). The corresponding Lagrangian can be defined as :

$$\begin{aligned} L(q; \lambda) &= \sum_{j \in \mathcal{J}} V_j(q_j) - \sum_{s \in \mathcal{S}} C_s(q_s) + \sum_{l \in \mathcal{L}} \lambda_l (\sum_{s \in \mathcal{S}} q_{s,l} - \sum_{j \in \mathcal{J}} q_{j,l}); \\ \forall q : q_j &\in \mathcal{D}_j, j \in \mathcal{J}; q_s \in \mathcal{D}_s, s \in \mathcal{S}; \forall \lambda \in \mathbb{R}^{|\mathcal{L}|} \end{aligned} \quad (6.5)$$

Consider the 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 :

$$\min_{\lambda} \Theta(\lambda) \quad (6.6)$$

It is noteworthy that the Lagrangian dual function can be written as the following problem, $(LR_{(\lambda)})$:

$$\max_q \sum_{j \in \mathcal{J}} \left(V_j(q_j) - \sum_{l \in \mathcal{L}} \lambda_l q_{j,l} \right) + \sum_{s \in \mathcal{S}} \left(\sum_{l \in \mathcal{L}} \lambda_l q_{s,l} - C_s(q_s) \right) \quad (6.7)$$

$$s.t. \quad q_j \in \mathcal{D}_j, \quad j \in \mathcal{J} \quad (6.8)$$

$$q_s \in \mathcal{D}_s, \quad s \in \mathcal{S} \quad (6.9)$$

which decomposes into $|\mathcal{J}| + |\mathcal{S}|$ independent sub-problems, one for each seller and buyer.

The following result states conditions under which an optimal solution of the Lagrangian dual problem and the corresponding primal solutions correspond to an efficient allocation and form a Walrasian equilibrium.

Proposition 6.1 Let λ^* be an optimal solution of the Lagrangian dual problem, and $q^* = \arg \max_q \{L(q; \lambda) : q_j \in \mathcal{D}_j, j \in \mathcal{J}; q_s \in \mathcal{D}_s, s \in \mathcal{S}\}$ the optimal solution of the corresponding dual relaxation. Under conditions of :

1. convexity of production and consumption functions and feasibility sets (Assumptions (A1) and (A2)),
2. stability of problem (MC),
3. feasibility of q^* for (MC),

q^* is a socially-efficient allocation and λ^* is a corresponding vector of Walrasian prices.

Proof. The social-efficiency of the allocation q^* follows immediately from the strong Lagrangian duality theorem (Geoffrion [58]), which precludes the existence of a duality gap between the primal problem and the Lagrangian dual problem under convexity and stability conditions. Concerning the fact that q^* and λ^* constitute a Walrasian equilibrium, one might simply notice that each one of the $|\mathcal{J}| + |\mathcal{S}|$ subproblems of the Lagrangian relaxation $(LR_{(\lambda)})$ corresponds to the maximization of the surplus of a seller or a buyer. ■

The *stability* of problem (MC) (Geoffrion [58]) is a key condition for the “duality gap” between the primal problem (MC) and its Lagrangian dual to vanish. Stability is guaranteed, in particular, when the objective and the constraints of the primal problem are affine or quadratic. A more general condition is the *filling property* (Hiriart-Urruty and Lemaréchal [65]), which only requires compactness of the subsets $\mathcal{D}_j, j \in \mathcal{J}$ and $\mathcal{D}_s, s \in \mathcal{S}$, and continuity of functions $V_j, j \in \mathcal{J}$ and $C_s, s \in \mathcal{S}$, and which is sufficient to ensure the absence of duality gap when combined with convexity of problem (MC).

6.4.1 The subgradient approach

The subgradient algorithm, an approach that finds its inspiration in the gradient method for convex differentiable optimization, has been traditionally used to solve Lagrangian dual problems. At each iteration of the (basic version) of the algorithm, a subgradient of the dual function Θ is computed at the current vector of Lagrangian

multipliers and the multipliers are updated along the direction of the subgradient. For the problem in hand, the subgradient algorithm can be stated as follows :

STEP 0 : Set $k = 0$, $\Theta_*^{(0)} = +\infty$, $k_* = 0$. Initialize the vector of Lagrangian multipliers, e.g., $\lambda^{(0)} = 0$.

STEP 1 : Evaluate $\Theta(\lambda^{(k)})$. Compute a subgradient $g^{(k)} \in \partial\Theta(\lambda^{(k)})$ of Θ at $\lambda^{(k)}$:

$$\text{let } q^{(k)} \in \arg \max_q \{L(q; \lambda^{(k)}) : q_j \in \mathcal{D}_j, j \in \mathcal{J}; q_s \in \mathcal{D}_s, s \in \mathcal{S}\} ;$$

$$g^{(k)} = \sum_{s \in \mathcal{S}} q_s^{(k)} - \sum_{j \in \mathcal{J}} q_j^{(k)}.$$

Update $\Theta_*^{(k)}$: if $\Theta(\lambda^{(k)}) < \Theta_*^{(k)}$ then set $\Theta_*^{(k)} = \Theta(\lambda^{(k)})$, and $k_* = k$.

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 2.

The choice of the stepsize $\xi^{(k)}$ is critical to the convergence of the algorithm. In practice, the following two schemes (Lemaréchal [87]) are the most commonly used.

1. The series $\{\xi^{(k)}\}_{k \in \mathbb{N}}$ is such that $\lim_{k \rightarrow \infty} \xi^{(k)} = 0^+$ and $\sum_{k=0}^{\infty} \xi^{(k)} = +\infty$.
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 $\rho^{(0)} \in]0, 2[$ and $\rho^{(k)}$ is halved every time $\Theta_*^{(k)}$ has not been updated for n (generally equal to 10 or 20) consecutive iterations.

The determination of appropriate stopping criteria is among the weakest points of the subgradient method. A reasonable condition to verify for interrupting the algorithm would be $\|g^{(k)}\| \leq \epsilon$ (the rationale of that criterion is Everett’s theorem (Everett [48])), which states that, as far as the supply/demand constraints 6.2 are concerned, the corresponding primal solution $q^{(k)}$ would be ϵ -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 $\Theta_*^{(k)}$ has not been improved in the last N 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, illustrated in Figure (6.2), as an iterative auction process. 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 bid $B_j^{(k)} = \{q_j^{(k)}\}$, where $q_j^{(k)}$ is a surplus-maximizing consumption level for buyer j at the given prices. Given the bids it receives, the market-maker revises the prices of the goods along a steepest descent direction given by the excess vector $\sum_{s \in \mathcal{S}} q_s^{(k)} - \sum_{j \in \mathcal{J}} q_j^{(k)}$.

Considering that the subgradient-based auction may stop before an “implementable” outcome is reached (one which 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 in 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 [81], 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 :

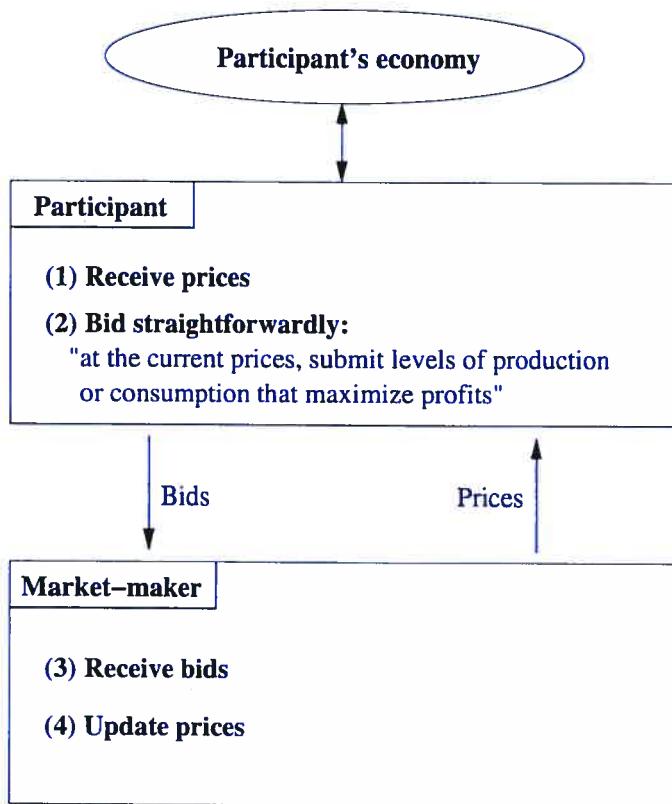


FIG. 6.2 – The subgradient algorithm as an iterative auction.

$$\min_q \quad \|q - \bar{q}^{(K)}\|^2 \quad (6.10)$$

$$s.t. \quad \sum_{j \in \mathcal{J}} q_{j,l} - \sum_{s \in \mathcal{S}} q_{s,l} = 0, \quad l \in \mathcal{L} \quad (6.11)$$

$$q_j \in \mathcal{D}_j, \quad j \in \mathcal{J} \quad (6.12)$$

$$q_s \in \mathcal{D}_s, \quad s \in \mathcal{S} \quad (6.13)$$

The solution to this problem is nevertheless conditioned by the availability to the market-maker of complete knowledge of feasibility sets $\mathcal{D}_j, j \in \mathcal{J}$ and $\mathcal{D}_s, s \in \mathcal{S}$. In our case, 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 :

$$\hat{\mathcal{D}}_j = \{q_j = \sum_{k=0}^K q_j^{(k)} \alpha_j^k : \sum_{k=0}^K \alpha_j^k = 1; \alpha_j^k \geq 0, \forall k = 0, \dots, K\}, j \in \mathcal{J}$$

and

$$\hat{\mathcal{D}}_s = \{q_s = \sum_{k=0}^K q_s^{(k)} \beta_s^k : \sum_{k=0}^K \beta_s^k = 1; \beta_s^k \geq 0, \forall k = 0, \dots, K\}, s \in \mathcal{S}$$

By virtue of 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 set yields the quadratic problem :

$$\min_q \quad \|q - \bar{q}^{(K)}\|^2 \quad (6.14)$$

$$s.t. \quad \sum_{j \in \mathcal{J}} q_{j,l} - \sum_{s \in \mathcal{S}} q_{s,l} = 0, \quad l \in \mathcal{L} \quad (6.15)$$

$$q_j \in \hat{\mathcal{D}}_j, \quad j \in \mathcal{J} \quad (6.16)$$

$$q_s \in \hat{\mathcal{D}}_s, \quad s \in \mathcal{S} \quad (6.17)$$

One should be warned against the “inexact” nature of this approach. Problem (6.14-6.17) is a restriction of (6.10-6.13), and thus is not necessarily feasible and may

not be successful in providing a feasible outcome when only “few” bids are used to approximate the feasibility sets, or when the feasible domain of the allocation problem is tight.

6.4.2 The bundle approach

Bundle methods, originally developed for nonsmooth optimization (Wolfe [147], Lemaréchal [86]), may equally be suggested for solving the Lagrangian dual problem. These methods rely basically on the *bundle of information* concept, 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 $\mathcal{B} = \{(\lambda^{(k)}; \Theta(\lambda^{(k)}); g^{(k)})\}_{k=1,\dots,K}$ represent the information gathered at a given time, where $g^{(k)} \in \partial\Theta(\lambda^{(k)})$, $\forall k = 1, \dots, K$ 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 earlier cutting-plane algorithm (Kelley [71]) 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 stepsize 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_B) :

$$\min_{\nu, \lambda} \quad \nu + \frac{1}{2t^K} \|\lambda - \bar{\lambda}\|^2 \quad (6.18)$$

$$s.t. \quad \nu \geq \Theta(\lambda^{(k)}) + g^{(k)T}(\lambda - \lambda^{(k)}), \quad k = 1, \dots, K \quad (6.19)$$

Let $(\nu^{(K+1)}, \lambda^{(K+1)})$ be the optimal solution of this problem, and let $\Delta^K = \Theta(\bar{\lambda}) - \Theta(\lambda^{(K+1)})$ and $\tilde{\Delta}^K = \Theta(\bar{\lambda}) - \Theta(\lambda^{(K+1)})$ denote the *actual* and the *predicted* (by the model Θ_B) decreases 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 methods perform 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 finer approximation of Θ , is made.

While it does not affect the validity of the approach, the setting of the parameter sequence $\{t_k\}_{k \in \mathbb{N}}$ plays a decisive role in the behavior of a bundle algorithm and its numerical efficiency. Theoretical and practical evidence shows the difficulties that arise from setting a *fixed* parameter $t_k = t, \forall k \in \mathbb{N}$ (Hiriart-Urruty and Lemaréchal [65]). Hence, a small value of t will drive the bundle algorithm to make relatively few null-steps, but also “small” serious-steps resulting in marginal improvement of the dual functions (when $t \rightarrow 0$, the bundle method is nothing else than the subgradient algorithm). On the other hand, the bundle algorithm performs few serious-steps when large values of t are considered, moving toward the cutting-plane algorithm as $t \rightarrow \infty$. 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 [75], Schramm and Zowe [135]) 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. The Kuhn-Tucker optimality conditions for Q_B imply that, for an optimal solution (ν^*, λ^*) of Q_B , there exists a vector of multipliers $\{\delta_k^*\}_{k=1,\dots,K}$ such that :

- (i) $d^* = \lambda^* - \bar{\lambda} = -t^K \sum_{k=1}^K \delta_k^* g^{(k)}$;
- (ii) $\nu^* = -t^K \left\| \sum_{k=1}^K \delta_k^* g^{(k)} \right\|^2 - \sum_{k=1}^K \delta_k^* e_k + \Theta(\bar{\lambda})$, where $e_k = \Theta(\bar{\lambda}) - (\Theta(\lambda^{(k)}) + g^{(k)T}(\bar{\lambda} - \lambda^{(k)}))$, $k = 1, \dots, K$;
- (iii) $\delta_k^*(-\nu^* + g^{(k)T}(\lambda^* - \lambda^{(k)}) + \Theta(\lambda^{(k)})) = 0$, $k = 1, \dots, K$;
- (iv) $\sum_{k=1}^K \delta_k^* = 1$ and $\delta_k^* \geq 0$, $k = 1, \dots, K$.

Conditions (i) and (iv) are particularly instructive. Together, they indicate that bundle methods actually construct *aggregated subgradients* $\tilde{g}^{(K)} = \sum_{k=1}^K \delta_k^* g^{(k)}$ as convex combinations of the subgradients available in the bundle, and move (in the case of a serious-step) in the opposite direction of the aggregated subgradient, to an extent given by stepsize t^K .

6.5 An auction scheme based on column-generation

Another decomposition technique that can be given a market interpretation as an iterative auction is the Dantzig-Wolfe decomposition principle. Let us suppose sets of feasible consumption levels $\{q_j^r\}_{r \in k_j}$, $j \in \mathcal{J}$, and production levels $\{q_s^r\}_{r \in k_s}$, $s \in \mathcal{S}$ are available to the market-maker. Consider the corresponding convex hulls :

$$\tilde{\mathcal{D}}_j^{k_j} = \{q_j = \sum_{r=1}^{k_j} q_j^r \alpha_j^r : \sum_{r=1}^{k_j} \alpha_j^r = 1; \alpha_j^r \geq 0, \forall r = 1, \dots, k_j\}, j \in \mathcal{J}$$

and

$$\tilde{\mathcal{D}}_s^{k_s} = \{q_s = \sum_{r=1}^{k_s} q_s^r \beta_s^r : \sum_{r=1}^{k_s} \beta_s^r = 1; \beta_s^r \geq 0, \forall r = 1, \dots, k_s\}, s \in \mathcal{S}$$

Inner-linearization (Geoffrion [57]) of the production (consumption) feasibility sets and the cost (valuation) functions suggests to solve linear programming approximations of the nonlinear market-clearing problem (MC). More precisely, it proceeds in the following two steps :

- Approximate $\mathcal{D}_j, j \in \mathcal{J}$ and $\mathcal{D}_s, s \in \mathcal{S}$ with $\tilde{\mathcal{D}}_j^{k_j}$ and $\tilde{\mathcal{D}}_s^{k_s}$, respectively.
- For $q_j \in \mathcal{D}_j, j \in \mathcal{J}$, replace $V_j(q_j) = V_j(\sum_{r=1}^{k_j} q_j^r \alpha_j^r)$ with $\sum_{r=1}^{k_j} \alpha_j^r V_j(q_j^r)$. Similarly, replace $C_s(q_s) = C_s(\sum_{r=1}^{k_s} q_s^r \beta_s^r)$ with $\sum_{r=1}^{k_s} \beta_s^r C_s(q_s^r)$.

These approximations yield the restricted master problem (MC-R) :

$$\max_{j \in \mathcal{J}} \quad \sum_{r=1}^{k_j} \alpha_j^r V_j(q_j^r) - \sum_{s \in \mathcal{S}} \sum_{r=1}^{k_s} \beta_s^r C_s(q_s^r) \quad (6.20)$$

$$s.t. \quad \sum_{j \in \mathcal{J}} \sum_{r=1}^{k_j} \alpha_j^r q_{j,l}^r - \sum_{s \in \mathcal{S}} \sum_{r=1}^{k_s} \beta_s^r q_{s,l}^r = 0, \quad l \in \mathcal{L} \quad (6.21)$$

$$\sum_{r=1}^{k_j} \alpha_j^r = 1, \quad j \in \mathcal{J} \quad (6.22)$$

$$\sum_{r=1}^{k_s} \beta_s^r = 1, \quad s \in \mathcal{S} \quad (6.23)$$

$$\alpha_j^r \geq 0, \quad r = 1, \dots, k_j, \quad j \in \mathcal{J} \quad (6.24)$$

$$\beta_s^r \geq 0, \quad r = 1, \dots, k_s, \quad s \in \mathcal{S} \quad (6.25)$$

Let $[\{\bar{\alpha}_j^r\}_{r=1, \dots, k_j, j \in \mathcal{J}}, \{\bar{\beta}_s^r\}_{r=1, \dots, k_s, s \in \mathcal{S}}]$ be an optimal basic solution of (MC-R) and $\{\tilde{\mu}_l\}_{l \in \mathcal{L}}, \{\tilde{\tau}_j\}_{j \in \mathcal{J}}$, and $\{\tilde{\tau}_s\}_{s \in \mathcal{S}}$ the corresponding simplex multipliers. The relative cost optimality conditions for the restricted master problem are :

$$V_j(q_j^r) - \sum_{l \in \mathcal{L}} \tilde{\mu}_l q_{j,l}^r - \tilde{\tau}_j \leq 0, \quad \forall r = 1, \dots, k_j, \forall j \in \mathcal{J}$$

and

$$-C_s(q_s^r) + \sum_{l \in \mathcal{L}} \tilde{\mu}_l q_{s,l}^r - \tilde{\tau}_s \leq 0, \quad \forall r = 1, \dots, k_s, \forall s \in \mathcal{S}$$

Thus, the generation of new feasible consumption levels $q_j^{k_j+1}$, $j \in \mathcal{J}$, or production levels $q_s^{k_s+1}$, $s \in \mathcal{S}$ that eventually improve the approximation can be done by pricing out sets \mathcal{D}_j , $j \in \mathcal{J}$ and \mathcal{D}_s , $s \in \mathcal{S}$; that is, by solving the subproblems

$$(SP)_j : \max_{q_j \in \mathcal{D}_j} \{V_j(q_j) - \sum_{l \in \mathcal{L}} \tilde{\mu}_l q_{j,l} - \tilde{\tau}_j\}, j \in \mathcal{J};$$

and

$$(SP)_s : \max_{q_s \in \mathcal{D}_s} \{-C_s(q_s) + \sum_{l \in \mathcal{L}} \tilde{\mu}_l q_{s,l} - \tilde{\tau}_s\}, s \in \mathcal{S}.$$

The Dantzig-Wolfe decomposition principle has a classical economic interpretation as a decentralized planning procedure (Dantzig [36], chapter 23). A central authority (the headquarters) has to devise an optimal operation plan for an enterprise composed of several subsidiaries. Each subsidiary has private information concerning its technology and how it constrains its contribution to the overall plan. The headquarters deals with the constraints concerning the resource exchange between the subsidiaries. The Dantzig-Wolfe decomposition principle can be viewed as an iterative decision process in which the role of the central authority is to determine an optimal operation plan given a set of partial operation plans suggested by the subsidiaries and to announce corresponding simplex multipliers (interpreted as prices), while the subsidiaries react to the announced prices by proposing new promising partial plans.

In the specific context of the exchange economy considered in this paper, it turns out that this principle translates easily into a two-phase auction :

Phase 1.

In this initialization phase, the market-maker asks the sellers and the buyers to submit the $|\mathcal{L}| + |\mathcal{J}| + |\mathcal{S}|$ bids required to build a first restricted master problem.

Denote by $k_j^{(0)}$ the number of bids submitted by buyer j during the initialization phase, and by $k_s^{(0)}$ the number of bids submitted by seller s .

Set $n = 0$.

Phase 2.

At iteration n :

- Suppose $k_j^{(n)}$ and $k_s^{(n)}$ bids have been submitted, up to iteration n , by buyer j and seller s , respectively. The market-maker solves the restricted master problem (MC-R) and announces prices $\{\tilde{\mu}_l\}_{L \in \mathcal{L}}$, as well as multipliers $\{\tilde{\tau}_j\}_{j \in \mathcal{J}}$ and $\{\tilde{\tau}_s\}_{s \in \mathcal{S}}$ to the participants.
- Each buyer $j, j \in \mathcal{J}$, determines a surplus-maximizing solution $q_j^{k_j^{(n)}}$ of $(SP)_j$.
If $q_j^{k_j^{(n)}}$ improves on the approximation, buyer j submits bid
 $B_j^{(n)} = \{q_j^{k_j^{(n)}}; V(q_j^{k_j^{(n)}})\}.$
- Similarly, each seller $s, s \in \mathcal{S}$, determines a solution $q_s^{k_s^{(n)}}$ of $(SP)_s$. If $q_s^{k_s^{(n)}}$ improves the approximation, seller s submits bid $B_s^{(n)} = \{q_s^{k_s^{(n)}}; C(q_s^{k_s^{(n)}})\}.$

Although Lagrangian relaxation techniques and the Dantzig-Wolfe decomposition method can be seen as duals of each other (see, for example, Lasdon [82], chapter 8), there are important differences between the corresponding auction schemes. First, from an *informational* point of view, auctions based on the subgradient approach are *punctual* mechanisms regarding the determination of allocations, in the sense that only the last bids made by the participants are relevant to the market-maker. By opposition, the Dantzig-Wolfe auction aggregates all the bids it receives into a feasible allocation and may thus be considered as *cumulative* process. A straightforward consequence of the latter point is that the prices announced by the market-maker often do not form an equilibrium with any allocation corresponding to optimal solutions of sellers' and buyers' surplus-maximization subproblems. As previously noticed, bundle methods use more sophisticated pricing schemes in comparison with the subgradient algorithm, exploiting a history of the market made of all the prices computed during a certain time window and the corresponding bidder reactions to these prices.

The *strategic* viewpoint is an equally interesting one. One could argue that auction schemes based on Lagrangian relaxation are iterative mechanisms that determine socially efficient allocations under the assumption that the sellers and buyers follow *myopic* best-response bidding strategies. That is, they bid for production and consumption levels that maximize their surplus given the current prices announced

by the market-maker. Thus, a participant does not consider future rounds of the auction, nor does it exploit “what it knows” about other participants’ private data. The Dantzig-Wolfe auction process, on the other hand, requires the participants to formulate bids that contain, in addition to the desired production and consumption levels at the prices, cost and valuation information. The fact that the participants may misrepresent their true costs and valuations, while continuing to bid on surplus-maximizing quantities, implies that even the myopic best-response strategy is too weak to ensure that a Dantzig-Wolfe auction will produce socially efficient final outcomes.

6.6 Computational study

6.6.1 The experimental setting

In order to illustrate the above-mentioned auction mechanisms, we call upon a more detailed model of multi-lateral multi-commodity markets suggested in Bourbeau *et al.* [22], which has the advantage of being closer to actual applications, 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.

Participants in the market seek essentially to trade a set of *products*. A product is a basic commodity with a specific physical denomination (e.g., a wood species). Products are not available in a “pure” state, however, but come rather as part of different *lots* that are “mixtures” of several products. Hence, let

- \mathcal{K} : set of basic products.
- \mathcal{L} : set of lots.
- b_l^k : the proportion of product k in lot l , $k \in \mathcal{K}$, $l \in \mathcal{L}$.

It is assumed for simplicity that each seller may only procure 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 $C_l(\cdot)$ of lot l is assumed to be such that its corresponding marginal cost function $C'_l(\cdot)$ is continuous, piecewise-linear, and strictly increasing (Bourbeau *et al.* [22]).

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 $U_j(\cdot)$ is a utility function such that $U'_j(\cdot)$ is continuous, piecewise-linear, and strictly decreasing. Buyers need also to express requirements regarding the composition (in terms of the different products) 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 :

$$\max_{q_j} \quad \sum_{j \in \mathcal{J}} U_j \left(\sum_{l \in \mathcal{L}} r_j^l q_{j,l} \right) + \sum_{l \in \mathcal{L}} (s_j^l - t_j^l) q_{j,l} - \sum_{l \in \mathcal{L}} C_l(q_l) \quad (6.26)$$

$$s.t. \quad \sum_{j \in \mathcal{J}} q_{j,l} - q_l = 0, \quad l \in \mathcal{L} \quad (6.27)$$

$$q_l \leq Q^l, \quad l \in \mathcal{L} \quad (6.28)$$

$$\sum_{l \in \mathcal{L}} r_j^l q_{j,l} \leq Q^j, \quad j \in \mathcal{J} \quad (6.29)$$

$$m_j^k \sum_{l \in \mathcal{L}} r_j^l q_{j,l} \leq \sum_{l \in \mathcal{L}} b_l^k r_j^l q_{j,l} \leq M_j^k \sum_{l \in \mathcal{L}} r_j^l q_{j,l}, \quad j \in \mathcal{J}, k \in \mathcal{K} \quad (6.30)$$

$$q_{j,l} \geq 0, \quad j \in \mathcal{J}, l \in \mathcal{L} \quad (6.31)$$

$$q_l \geq 0, \quad l \in \mathcal{L} \quad (6.32)$$

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 we realized were guided by the general objective of comparing auction processes based on various decomposition methods from the

perspective of the economic efficiency achieved. The “benchmark” we used was thus a mechanism based on the centralized market-clearing formulation (6.26-6.32), 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. 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.

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 uniform distributions over prespecified intervals. For the sake of simplicity, we considered *purely quadratic* buyer utility functions $U_j(\cdot), j \in \mathcal{J}$ and seller cost functions $C_l(\cdot), l \in \mathcal{L}$. This implies no loss of generality, since a simple transformation suggested in Bourbeau *et al.* [22] makes it possible 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 6.1 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 ($\Delta_m = \underline{M} - \bar{m}$, where \underline{M} designates the minimum value M_j^k can take, and \bar{m} the maximum value of m_j^k). These two parameters are important since they directly impact the number and forcefulness of constraints (6.30) in the market-clearing formulation. Every series of problems consisted of 10 randomly generated instances.

We have set up four auction processes. The first three are based on variants of the subgradient algorithm, while the fourth relies on Dantzig-Wolfe decomposition. More specifically, the subgradient variants used are :

1. The “basic” subgradient method : $\lambda^{(k+1)} = \lambda^{(k)} - \xi^{(k)} g^{(k)}$, with stepsize formula $\xi^{(k)} = \rho^{(k)} \frac{\Theta(\lambda^{(k)}) - \Theta^{(k)}}{\|g^{(k)}\|^2}$. We have used the simple estimate $\bar{\Theta}^{(k)} = 0.5 \cdot \Theta_*^{(k)}$, where

Problem series	Problem description			
	# buyers	# lots	# products	$\Delta_m = M - \bar{m}$ (%)
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	300	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	300	10	10

TAB. 6.1 – Characteristics of problem instances

$\Theta_*^{(k)}$ denotes the best value of the lagrangian dual function $\Theta(\cdot)$ 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 (Camerini, Fratta, and Maffioli [26]). 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 stepsize formula used is similar to that of the basic subgradient, that is $\xi^{(k)} = \rho^{(k)} \frac{\Theta(\lambda^{(k)}) - \Theta^{(k)}}{g^{(k)}d^{(k)}}$.

3. The subgradient method with the modified Camerini-Fratta-Maffioli rule. Instead of using a “fixed” value of parameter μ , this variant uses the rule : $\mu = -\frac{\|g^{(k)}\|\|d^{(k-1)}\|}{g^{(k)}d^{(k-1)}}$ (Crainic, Frangioni, and Gendron [33]).

In order to compare the different auction mechanisms, we fixed the maximum number of rounds to 1000 for the subgradient-based auctions, and to 200 for the DW auction. The following metrics are used :

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

1. the gap $GAP_{lr} = (Z_{lr} - Z_{cent})/Z_{cent}$, where Z_{lr} is the best upper bound obtained by the corresponding subgradient methods, and Z_{cent} is the optimal surplus achieved by the centralized allocation model (6.26-6.32) ;
2. the gap $GAP_{lr}^P = (Z_{cent} - \tilde{Z})/Z_{cent}$, where \tilde{Z} is the economic surplus achieved by the “closest” feasible allocation to the primal solution $q^{(k_*)}$, obtained by projecting the latter on the feasible domain of model (6.26-6.32) ;
3. the gap GAP_{lr}^E 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 6.4.1 on the domain of feasible allocations, that is $GAP_{lr}^E = (Z_{cent} - \bar{Z})/Z_{cent}$, where \bar{Z} is the economic surplus achieved by q^E ;
4. the Euclidean norms $\|g^{(k_*)}\|$ and $\|g^{(K)}\|$ of the trivial sub-gradients corresponding to allocations $q^{(k_*)}$ and $q^{(K)}$, respectively.

(b) For the DW-based auction process, we evaluated the gap $GAP_{dw} = (Z_{cent} - Z_{dw})/Z_{cent}$ between Z_{cent} and the lower bound obtained by Dantzig-Wolfe auction process Z_{dw} .

6.6.2 Numerical results

Table 6.2 displays the results obtained by the basic subgradient, the Camerini-Fratta-Maffioli, and the modified Camerini-Fratta-Maffioli methods. The first figure (GAP_{lr}) indicates the average gaps corresponding to the best upper bound achieved at the last (1000th) bidding round. We only retained the best gaps (with respect to the five possible values of parameter $\rho^{(0)}$), and we listed the best $\rho^{(0)}$ s in the table. The second and the third figures of the table are the average gaps of the allocations obtained by projecting $q^{(k_*)}$ and $q^{(K)}$, respectively, on the feasible allocation domain. The

fourth and fifth figures indicate the average norms of the subgradients corresponding to $q^{(k_*)}$ and $q^{(K)}$.

Several interesting observations can be made regarding these results. First, while both the basic subgradient and the CFM methods display quite small average gaps, consistently converging to within 3% of the optimal solution of the centralized market-clearing formulation, the convergence of the modified CFM is a more mixed bag : on some series ($S - 01$ and $S - 06$, for instance), the average gaps are comparable to those of the two other methods ; on others (especially $S - 03$, $S - 07$, and $S - 08$) the gaps are much larger. In that regard, the poor performance of the modified CFM method seems likely to be a consequence of its relatively greater sensitiveness to the choice of the initial scaling factor $\rho^{(0)}$, rather than an inherent lack of effectiveness. The feasibility of the primal solutions returned by the three methods is a major source of concern, though, as fairly large subgradients (with respect of the magnitude of the randomly generated quantities $Q^j, j \in \mathcal{J}$ and $Q^l, l \in \mathcal{L}$) are displayed, and the gaps associated with the projection of the primal solution on the feasible allocation domain are large too. It is noticeable that, despite significantly lowering infeasibility, the ergodic sequence of Larsson, Patriksson, and Strömberg [81] does not solve this issue in a satisfactory manner, as their primal convergence seems to be extremely slow, and the projection of the terms of the sequence on the feasible domain yields even larger gaps than the original allocations.

In order to gain more insight into the behavior of the three relaxation-based auctions, we have taken the third instance of series $S - 01$ and we mapped out in Figure 6.3 the best upper bound of each auction as the latter evolves in time. The figure shows quite clearly the relatively superior speed of convergence of the two CFM methods, which was expected : whereas barely 189 and 183 rounds of bidding are enough for the CFM and the modified CFM methods, respectively, to attain a less than 1% gap, the basic subgradient method needs 463 rounds to achieve that same level. Another issue which is interesting to examine in detail is the sensitiveness of the subgradient methods to the initial scaling factor $\rho^{(0)}$. We have thus considered another problem in series $S - 01$ and plotted in Figure 6.4 the upper bound of the three methods with respect to different values of parameter $\rho^{(0)}$. The resulting plots confirm the greater

Series	Basic subgradient					
	GAP_{lr} (%)	GAP_{lr}^P (%)	GAP_{lr}^E (%)	$\ g^{(k_*)}\ $	$\ g^{(K)}\ $	$\rho_*^{(0)}$
$S - 01$	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_{lr}^P (%)	GAP_{lr}^E (%)	$\ g^{(k_*)}\ $	$\ g^{(K)}\ $	$\rho_*^{(0)}$
$S - 01$	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_{lr}^P (%)	GAP_{lr}^E (%)	$\ g^{(k_*)}\ $	$\ g^{(K)}\ $	$\rho_*^{(0)}$
$S - 01$	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

TAB. 6.2 – Behavior of the subgradient, the CFM, and the modified CFM methods

sensitiveness to $\rho^{(0)}$ of the modified CFM, which might perform very poorly and fail to converge within an acceptable number of rounds if the “wrong” value of the initial scaling factor is chosen (e.g., $\rho^{(0)} = 1.0$ for the instance of Figure 6.4).

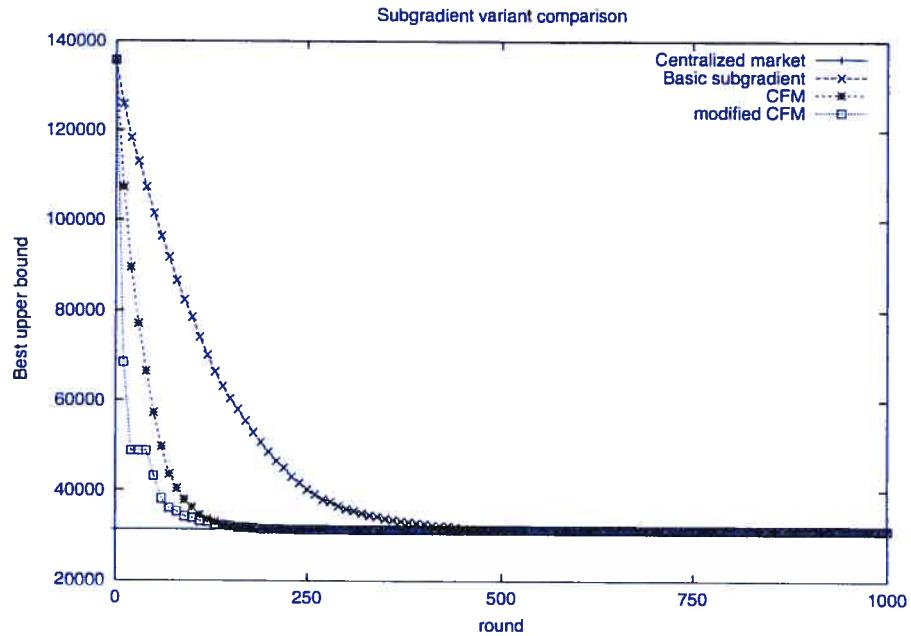
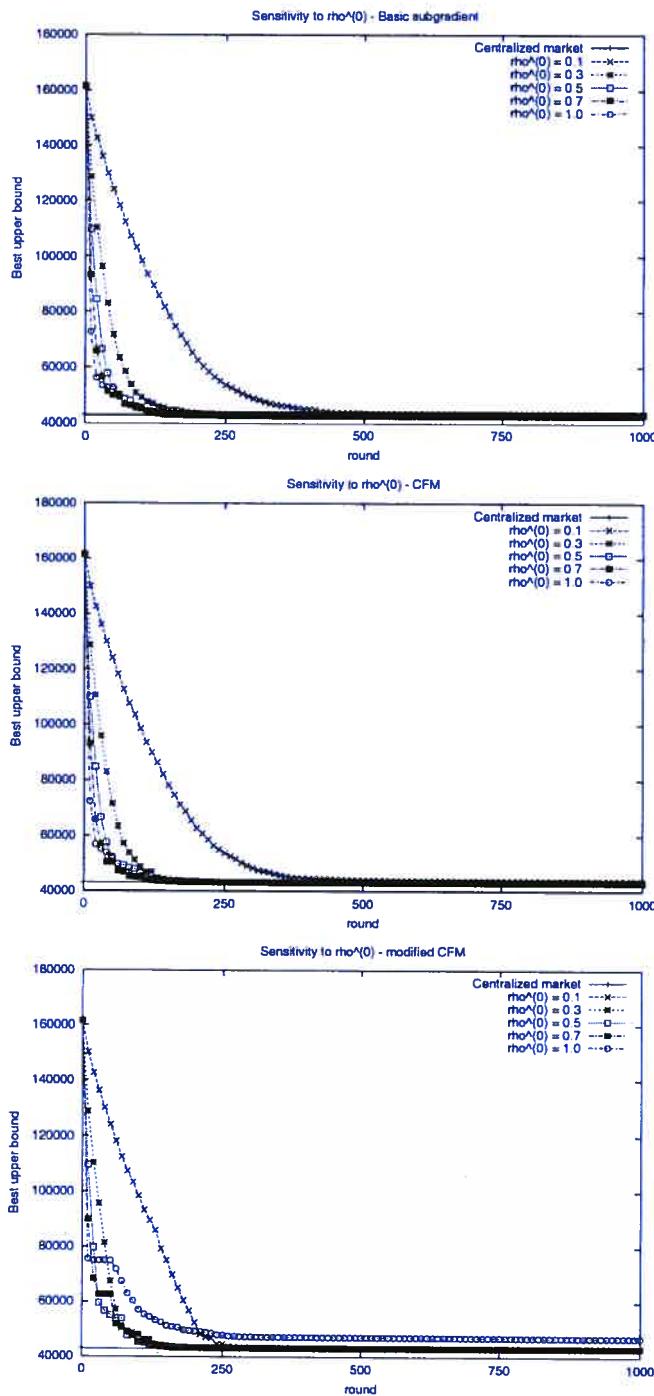


FIG. 6.3 – Evolution of the best upper bound for the three subgradient methods.

Figure 6.5 presents the evolution of GAP_{dw} , the gap between the lower bound at each round of the DW-based auction and the optimal surplus of the centralized market-clearing model, for instances of series $S - 01$, $S - 04$, and $S - 09$. The initialization phase consisted in announcing a 0 price vector, then a “large” price vector, and collecting the bids made by the sellers and the buyers at these prices. The figure indicates that the DW process has each time succeeded, in a fairly small number of rounds, in constructing very close approximations of the actual valuation and cost functions of the participants. One could also remark that the smaller the instance ($S - 01$ vs $S - 04$) and the fewer products it involves ($S - 04$ vs $S - 09$), the quicker the convergence of the iterative process to the centralized market-clearing outcome.

FIG. 6.4 – Sensitivity of the subgradient methods to $\rho^{(0)}$.

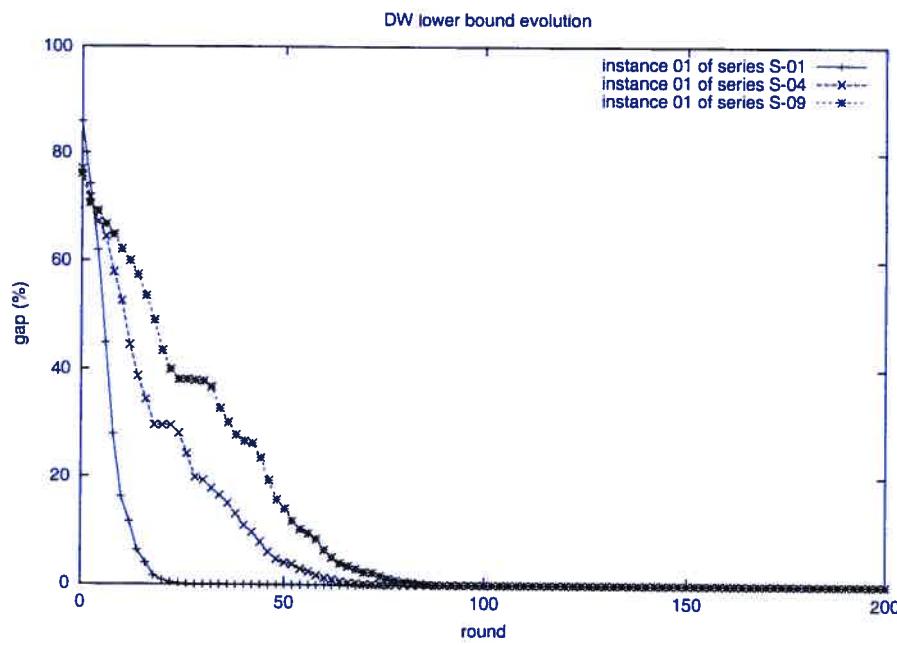


FIG. 6.5 – Evolution of the gap associated with the DW-based auction for instances of $S - 01$, $S - 04$, and $S - 09$.

6.7 Conclusions

This paper has presented a new perspective of mathematical decomposition methods as iterative auctions in combinatorial exchanges of interdependent goods. We have put the emphasis on two well-known approaches, the Lagrangian relaxation and Dantzig-Wolfe decomposition, and we have shown that they can be interpreted as auction processes 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. We have furthermore established that subtle yet important differences exist between the two kinds of approaches in the structure of information that the participants in the market need to disclose and how that information is exploited, as well as in the assumption made regarding the strategic behavior of the participants.

The numerical results have shown that the different variants of the subgradient method often converge in the dual space to the optimal market surplus. Nevertheless, they generally fail to produce feasible allocations. This is a major impediment to

the acceptance and usability of the auction processes associated with these methods. On the other hand, Dantzig-Wolfe decomposition maintains primal feasibility but requires the participants reveal parts of their private preferences, which may not be acceptable to them. The bundle methods, with their more sophisticated price update rules, could help get the best of both worlds, that is, to resolve the primal infeasibility issue while having minimal preference revelation requirements.

Several interesting research issues are before us. 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 inviting are : (a) the exploration of extended formulations of the market-clearing allocation problem (see Bikhchandani and Ostrov [20] for the CAP) ; and (b) pricing schemes based on “approximated” linear prices, which sacrifice either dual feasibility or complementary slackness (e.g. DeMartini *et al.* [41]). 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. 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 who helped in coding the numerical methods and in realizing the experimental study. This project has received financial support from the Bell Canada University Laboratories (LUB), the Natural Sciences and Engineering Research Council of Canada through their Basic Research and Collaborative Industry-University Research grant programs, and the FCAR research development fund of the Province of Québec.

Chapitre 7

Conclusion

Un grand nombre de contextes de marché impliquent la vente simultanée d'objets hétérogènes. En raison d'effets de complémentarité ou de substituabilité, la préférence d'un participant pour un objet donné peut dépendre de l'achat ou de la vente d'autres objets. Les enchères combinatoires, où il est possible de placer des mises sur des paquets d'objets, constituent une alternative intéressante aux enchères simultanées ascendantes. En permettant aux participants de communiquer explicitement leurs préférences exactes pour les paquets d'objets désirés, les enchères combinatoires réduisent considérablement l'ampleur de plusieurs problèmes communs aux enchères basées sur des mises sur des objets individuels, notamment la réduction de l'efficacité économique du marché due aux participants choisissant de miser "prudemment" pour éviter d'obtenir un sous-ensemble d'objets du paquet désiré pour un prix supérieur à leur préférence pour le sous-ensemble.

La conception d'enchères combinatoires est une problématique fondamentalement complexe impliquant des défis importants de modélisation et d'élaboration d'approches de résolution efficaces pour les problèmes de décision rencontrés par l'enchanteur et les participants à l'enquête. Cette thèse a été consacrée à l'étude de quelques aspects importants de la conception de mécanismes d'enquête combinatoire. Dans ce qui suit, nous présentons un résumé des principales contributions de la thèse et nous suggérons des avenues de recherche prometteuses que nous comptons explorer dans le cadre de nos travaux futurs.

7.1 Principales contributions

Nous avons présenté dans le chapitre 2 une revue de littérature qui se veut une synthèse des développements les plus importants reliés à la conception de mécanismes d'enchère combinatoire. Après avoir exposé quelques concepts fondamentaux concernant les cadres plus simples d'enchères d'objets uniques et d'enchères multi-unités, nous avons présenté les modèles d'équilibre pertinents dans le cas d'enchères de biens divisibles, les principales formulations de base d'enchère combinatoire d'objets indivisibles, ainsi que quelques applications d'enchère combinatoire dans une variété de contextes. Nous avons par ailleurs répertorié les principaux mécanismes progressifs d'enchère combinatoire proposés dans la littérature, ainsi que les différents formalismes d'expression de la préférence. Enfin, nous avons présenté les travaux reliés à la mise au point de mécanismes d'enchère combinatoire incitatifs, où la déclaration vérifique de la préférence est une stratégie dominante des participants.

Le chapitre 3 a porté sur l'étude plus spécifique de quatre facettes importantes de la problématique générale de la conception de mécanismes d'enchère combinatoire, qui forment le cadre de travail pour la suite de la thèse. Dans un premier temps, nous avons procédé à une classification multidimensionnelle des marchés et des mécanismes d'enchère combinatoire qui nous a permis, en particulier, de dériver plusieurs formulations du problème fondamental de la détermination des mises gagnantes dans une enchère combinatoire. Nous avons ensuite discuté le besoin de langages de mise permettant aux participants d'exprimer de manière succincte leurs préférences pour les objets transigés sur le marché. Le troisième volet du chapitre est, quant à lui, consacré aux mécanismes itératifs d'enchère combinatoire, où l'activité de mise et l'allocation des objets se déroulent sur l'espace de plusieurs rondes. Enfin, nous avons abordé les problèmes de décision du point de vue des participants à une enchère combinatoire et le besoin d'outils d'aide à la décision permettant d'assister ces derniers dans la mise au point de stratégies de mise profitables.

Dans le chapitre 4, nous avons présenté un nouveau cadre formel pour la définition de langages de mises pour les enchères combinatoires. Le cadre proposé répond au besoin de rassembler, au sein d'un seul et même formalisme, les outils syntaxiques et sémantiques nécessaires à la définition et l'expression de mises combinées complexes

dans des enchères d'objets divisibles et indivisibles. Ainsi, nous faisons la remarque fondamentale qu'un langage de mise dans le cas divisible doit permettre aux participants, *en plus* d'exprimer des conditions logiques reliées à l'exécution ou non d'une mise, de définir des contraintes sur les *proportions* d'exécution des mises. Ce constat se traduit naturellement dans le cadre que nous proposons par une définition à deux niveaux d'une mise combinée et l'extension de la notion classique d'opérateur de mise. La grande généralité de notre cadre formel constitue un atout important, en égard à l'émergence récente d'un grand nombre de marchés transigeant des biens divisibles (télécommunications, énergie, matières premières, etc.). Un autre attrait du cadre formel suggéré, que l'analyse a permis de démontrer, est son remarquable niveau d'expressivité. Enfin, nous avons examiné l'impact du formalisme sur le problème d'allocation (quelles mises exécuter ? dans quelles proportions ?) et évalué empiriquement la complexité des modèles d'optimisation correspondants.

Dans le chapitre 5, nous avons proposé un nouveau modèle de marché financier basé sur des ordres *composites* consolidant l'exécution (dans les mêmes proportions) de plusieurs ordres simples d'achat et de vente de valeurs financières. Notre modèle apporte plusieurs améliorations aux formulations précédemment suggérées dans la littérature, en permettant notamment aux participants de spécifier des bornes sur les proportions d'exécution de leurs ordres et de définir des relations d'exclusion mutuelle (XOR) portant sur l'exécution d'ordres "équivalents". L'étude expérimentale que nous avons réalisée a eu pour objectif d'évaluer l'influence de la consolidation caractérisant le modèle sur la liquidité du marché, ainsi que de mesurer l'impact des extensions suggérées sur l'efficacité économique et la complexité numérique des modèles correspondants. Par ailleurs, nous avons mis au point des procédures de discrimination des allocations et des prix optimaux - par rapport aux proportions d'exécution et aux paiements à effectuer ou à recevoir - sur la base d'un critère éthique simple, soit le temps de soumission des ordres sur le marché. Ces procédures, qui combinent une lacune des mécanismes actuels, constituent à notre avis une des contributions pratiques importantes de la thèse.

Le chapitre 6 présente un point de vue nouveau sur les méthodes de décomposition en programmation mathématique. En effet, en exploitant le potentiel de prise de

décision decentralisée de ces méthodes et le fait qu'elles ont une interprétation économique “naturelle”, nous avons pu établir que ces méthodes, prises dans un contexte de marché, sont à la base de mécanismes itératifs d'enchère combinatoire. Plus précisément, nous avons considéré une économie générale composée de producteurs et de consommateurs de biens, pour laquelle nous avons formulé le problème (centralisé) consistant à déterminer l'allocation efficace maximisant le surplus social du marché. Nous avons ensuite démontré que, sous certaines hypothèses, tout aussi bien des approches “duales” basées sur la relaxation lagrangienne que la décomposition (“primale”) de Dantzig-Wolfe sont à même de déterminer l'allocation efficace sans avoir *directement* accès aux préférences des participants, en autant que ces derniers adoptent un comportement *compétitif*, en misant selon les niveaux de production et de consommation qui maximisent leur propre surplus compte tenu des prix courants. À la lumière de notre analyse, nous avons par ailleurs réussi à identifier certaines différences fondamentales qui existent entre les processus d'enchères correspondant respectivement à la relaxation lagrangienne et à la décomposition de Dantzig-Wolfe. En particulier, en ce qui concerne l'exploitation de l'information contenue dans les mises des participants, nous avons noté que le sous-gradient constituait un processus “ponctuel”, qui ne tient compte que des dernières mises placées, alors que la décomposition de Dantzig-Wolfe était une approche “cumulative” qui utilise l'historique complet des mises.

7.2 Avenues de recherche

Un certain nombre de perspectives de recherche s'ouvrent suite à ce travail. Parmi les avenues qui nous semblent les plus prometteuses, mentionnons ce qui suit :

- Nous avons bâti notre interprétation des approches de décomposition comme étant des processus itératifs d'enchère combinatoire sur la non-existence d'une marge duale pour le problème d'allocation, rendue possible par les hypothèses de convexité et de stabilité du problème. Si la stabilité ne pose pas de problème particulier, la condition de convexité est assez contraignante. Notamment, le fait que le domaine réalisable du problème d'allocation doive être convexe exclut a priori les économies d'objets indivisibles, tout comme la possibilité que les

participants formulent des mises à l'aide de certains opérateurs “discrets” (OU exclusif, par exemple). Toutefois, il convient de noter qu'une formulation *de base* du problème d'allocation a été utilisée. À cet égard, il serait utile d'explorer la possibilité de dériver des formulations *étendues* du problème d'allocation possédant la propriété d'intégralité. Deux points de repère possibles sont : (i) le travail de Bikhchandani et Ostroy [20], qui ont appliqué avec succès une telle approche au problème classique de l'allocation combinatoire (CAP) ; et (ii) les approches heuristiques basées sur l'utilisation de prix linéaires “approximatifs” (par exemple, DeMartini *et al.* [41]). Nous croyons que le défi majeur résidera dans la généralisation des idées et des techniques mises en œuvre dans ces travaux tout en sauvegardant l'essence des approches de décomposition.

- Tout au long de cette thèse, nous nous sommes limités à l'étude d'enchères combinatoires impliquant plusieurs objets transigés *simultanément* sur le marché. Il existe néanmoins des situations où des enchères *séquentielles* (les objets sont transigés l'un à la suite de l'autre), ou *hybrides* séquentielles et simultanées sont plus pertinentes. C'est le cas, par exemple, de certaines enchères pour la pro-curement de biens dans le domaine industriel où l'acheteur *impose* aux vendeurs que certains “lots” soient procurés avant d'autres, ou encore, tout simplement, de participants désirant miser sur des objets interdépendants transigés sur des places de marchés différentes. Les participants à des enchères séquentielles ou hybrides sont confrontés à des problèmes de décision très complexes : sur quels objets miser à un moment donné ? quels prix miser ? quand faut-il arrêter de miser sur un objet ? etc. La difficulté de déterminer des stratégies de mises profitables est d'autant plus grande que certains résultats fondamentaux ne sont plus valables dès que l'on passe des enchères simultanées aux enchères séquentielles ou hybrides. À titre d'exemple, le fait de miser sa vraie préférence pour les objets transigés quand un mécanisme “second-prix” est utilisé pour vendre chacun des objets dans une enchère séquentielle n'est pas forcément une stratégie dominante d'un participant. Quoique la recherche concernant les enchères séquentielles ait connu récemment des développements importants (Matsumoto et Fujita [98], Chang, Crainic et Gendreau [29], Elmaghriby [43]),

ces travaux n'ont pas considéré le cas combinatoire qui demeure un problème ouvert.

- Les marchés électroniques de fret catalysent de plus en plus les échanges entre les expéditeurs et les transporteurs dans l'industrie de transport terrestre de marchandises. Nous nous intéressons en particulier à l'impact sur le processus décisionnel d'un transporteur de sa participation à une ou plusieurs enchères électroniques. Face à des commandes de transport émanant des expéditeurs, un transporteur (évoluant selon le modèle de transport des "charges complètes") doit typiquement décider quand, sur quelles commandes, et combien miser compte tenu de l'état de sa flotte (camions, conteneurs, personnel de conducteurs, etc.) et de la dynamique du marché. Nous comptons formuler ce problème en adaptant une des approches de programmation stochastique (voir Powell [119], par exemple). Nous devrons tenir compte dans notre approche de : (i) la complexité d'évaluer les combinaisons intéressantes de commandes de transport et les stratégies les plus profitables d'un participant en fonction de son profil et de la connaissance qu'il possède de ses compétiteurs, (ii) le caractère stochastique du service de transport (bris d'équipements, accidents, congestion, etc.), et (iii) le fait que la prise de décision doive se faire dans un environnement "temps-réel", où les temps de réponse alloués aux transporteurs peuvent être très courts. Ce dernier point constitue un réel défi. À cet égard, nous comptons accorder un soin particulier à l'étude des interactions avec les systèmes d'aquisition de données, à l'analyse du compromis entre la prise des décisions en temps réel et la précision de ces dernières, et à l'intégration des outils au sein d'un système de transport intelligent.

Bibliographie

- [1] Abrache, J., Bourbeau, B., Crainic, T.G., and Gendreau, M. A New Bidding Framework for Combinatorial E-Auctions. Publication CRT-2002-19, Centre de recherche sur les transports, Université de Montréal, Canada, 2002.
- [2] Abrache, J., Crainic, T.G., and Gendreau, M. Design Issues for Combinatorial Auctions. In *Proceedings of the 4th International Conference on Electronic Commerce Research (ICECR-4)*. Dallas, TX, 2001.
- [3] Abrache, J., Crainic, T.G., and Gendreau, M. Combinatorial Auctions for Portfolio Bundle Trading. Working paper, Centre de recherche sur les transports, Université de Montréal, Canada, 2002.
- [4] Ahuja, R.K., Magnanti, T.L., and Orlin, J.B. *Network Flows : Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- [5] Andersson, A., Tenhunen, M., and Ygge, F. Integer Programming for Combinatorial Auction Winner Determination. In *Proceedings of the 4th International Conference on Multi-Agent Systems (ICMAS-00)*, pages 39–46, 2000.
- [6] Arrow, K. and Debreu, G. Existence of Equilibrium for a Competitive Economy. *Econometrica*, 22 :265–290, 1954.
- [7] Arrow, K. and Hahn, F.H. *General Competitive Analysis*. Holden-Day, San Francisco, 1971.
- [8] Ausubel, L.M. An Efficient Ascending-Bid Auction for Multiple Objects. Working Paper 97-06, Department of Economics, University of Maryland, 1997.
- [9] Ausubel, L.M. An Efficient Dynamic Auction for Heterogeneous Commodities. Technical report, Department of Economics, University of Maryland, 2002.

- [10] Ausubel, L.M. and Cramton, P. Demand Reduction and Inefficiency in Multi-Unit Auctions. Working Paper 96-07, Department of Economics, University of Maryland, 1996.
- [11] Bachem, A., Hochstättler, W., and Malich, M. Simulated Trading : A New Approach For Solving Vehicle Routing Problems. Working Paper 93.139, Mathematisches Institut der Universität zu Köln, 1992.
- [12] Balas, E. and Padberg, M.W. Set Partitioning : A Survey. *SIAM Review*, 18 :710–760, 1976.
- [13] Banks, J.S., Ledyard, J.O., and Porter, D.P. Allocating uncertain and unresponsive resources : An experimental approach. *The RAND Journal of Economics*, 20(1) :1–25, 1989.
- [14] Benyoucef, M., Alj, H., Vézeau, M., and Keller, R.K. Combined Negotiations in E-Commerce : Concepts and Architecture. *Electronic Commerce Research*, 1(3) :277–299, 2001.
- [15] Bertsekas, D. A New Algorithm for the Assignment Problem. *Mathematical Programming*, 21 :152–171, 1981.
- [16] Bertsekas, D. *Linear Network Optimization*. The MIT Press, Cambridge, 1991.
- [17] Bertsekas, D. Auctions Algorithms for Network Flow Problems : A Tutorial Introduction. *Computational Optimization and Applications*, 1 :7–66, 1992.
- [18] Bikhchandani, S., De Vries, S., Schummer, J., and Vohra, R.V. Linear Programming and Vickrey Auctions. In *Mathematics of the Internet : E-Auctions and Markets*. Springer-Verlag, Berlin, 2001.
- [19] Bikhchandani, S. and Mamer, J. Competitive Equilibrium in an Exchange Economy with Indivisibilities. *Journal of Economic Theory*, 74(2) :385–413, 1997.
- [20] Bikhchandani, S. and Ostroy, J.M. The Package Assignment Problem. Technical report, Anderson School of Management and Department of Economics, UCLA, 2000.

- [21] Bossaerts, P., Fine, L., and Ledyard, J. Inducing Liquidity in Thin Financial Markets through Combined-Value Trading Mechanisms. Working paper, Caltech, 2000.
- [22] Bourbeau, B., Crainic, T.G., Gendreau, M., and Robert, J. Design for Optimized Multi-lateral Multi-commodity Markets. Working paper, Centre de recherche sur les transports, Université de Montréal, Canada, 2002.
- [23] Boutilier, C. and Hoos, H.H. Bidding Languages for Combinatorial Auctions. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, pages 1211–1217, 2001.
- [24] Brewer, P.J. and Plott, C.R. A Binary Conflict Ascending Price (BICAP) Mechanism for The Decentralized Allocation of The Right to Use Railroads Tracks. *International Journal of Industrial Organization*, 14 :857–886, 1996.
- [25] Bykowsky, M.M., Cull, R.J., and Ledyard, J.O. Mutually Destructive Bidding : The FCC Auction Design Problem. *Journal of Regulatory Economics*, 3(3) :205–228, 2000.
- [26] Camerini, P.M., Fratta, L., and Maffioli, F. On improving relaxation methods by modified gradient techniques. *Mathematical Programming Study*, 3 :26–34, 1975.
- [27] Chandra, B. and Halldórsson, M. Greedy local improvement and weighted set packing approximation. In *SODA-99*, pages 169–176. Baltimore, 1999.
- [28] Chang, T.-S., Crainic, T.G., and Gendreau, M. Dynamic Advisors for Freight Carriers for Bidding in Combinatorial Auctions. In *Proceedings of the 5th International Conference on Electronic Commerce Research (ICECR-5)*, CDROM. Centre de recherche sur les transports, Montréal, CANADA, 2002.
- [29] Chang, T.-S., Crainic, T.G., and Gendreau, M. Dynamic Bidding Strategies in On-line Sequential Auctions. In *Proceedings of the 5th International Conference on Electronic Commerce Research (ICECR-5)*, CDROM. Centre de recherche sur les transports, Montréal, CANADA, 2002.

- [30] Cheng, J.Q. and Wellman, M.P. The WALRAS Algorithm : A Convergent Distributed Implementation of General Equilibrium Outcomes. *Computational Economics*, 12 :1–24, 1998.
- [31] Clarke, E.H. Multipart Pricing of Public Goods. *Public Choice*, 11 :17–33, 1971.
- [32] Conen, W. and Sandholm, T. Preference Elicitation in Combinatorial Auctions : Extended Abstract. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 256–259. Tampa, FL, 2001.
- [33] Crainic, T.G., Frangioni, A., and Gendron, B. Bundle-based relaxation methods for multicommodity capacitated fixed charge network design. *Discrete Applied Mathematics*, 112(1-3) :73–99, 2001.
- [34] Cramton, P. The FCC Spectrum Auctions : An Early Assessment. *Journal of Economics and Management Strategy*, 6 :431–495, 1997.
- [35] Cramton, P. Ascending Auctions. *European Economic Review*, 42(3-5) :745–756, 1998.
- [36] Dantzig, G.B. *Linear Programming and Extensions*. Princeton University Press, Princeton, 1963.
- [37] Davenport, A.J. and Kalagnanam, J.R. Price Negotiations for Procurement of Direct Inputs. In *Mathematics of the Internet : E-Auctions and Markets*. Springer-Verlag, Berlin, 2001.
- [38] Davis, R. and Smith, R.G. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20 :63–109, 1983.
- [39] de Vries, S. and Vohra, S. Combinatorial Auctions : A Survey. *INFORMS Journal on Computing*, 15(3) :284–309, 2003.
- [40] Demange, G., Gale, D., and Sotomayor, M. Multi-item auctions. *Journal of Political Economy*, 94 :863–872, 1986.
- [41] DeMartini, C., Kwasnica, A.M., Ledyard, J.O., and Porter, D.P. A New and Improved Design for Multi-object Iterative Auctions. Social Science Working Paper SSWP 1054, Caltech, 1999.

- [42] Domowitz, I. The Mechanics of Automated Trade Systems. *Journal of Financial Intermediation*, 1 :167–194, 1990.
- [43] Elmaghribi, W. The Importance of Ordering in Sequential Auctions. *Management Science*, 49(5) :673–682, 2003.
- [44] Engelbrecht-Wiggans, R. Auctions and Bidding Models : A Survey. *Management Science*, 26 :119–142, 1980.
- [45] Engelbrecht-Wiggans, R. and Kahn, C.M. Multi-Unit Auctions with Uniform Prices. Working paper, University of Illinois, 1995.
- [46] Epstein, R., Henríquez, L., Catalán, J., Weintraub, G.Y., and Martínez, C. A Combinatorial Auction Improves School Meals in Chile. *Interfaces*, 32(6) :1–14, 2002.
- [47] Eso, M. An Iterative Auction for Online Seats. In *Mathematics of the Internet : E-Auctions and Markets*. Springer-Verlag, Berlin, 2001.
- [48] Everett, H. Generalized Lagrange Multiplier Method for Solving Problems of Optimum Allocation of Resources. *Operations Research*, 11 :399–417, 1963.
- [49] Fan, M., Srinivasan, S., Stallaert, J., and Whinston, A.B. *Electronic Commerce and the Revolution in Financial Markets*. South-Western Thomson Learning, 2002.
- [50] Fan, M., Stallaert, J., and Whinston, A.B. The Design and Development of a Financial Cybermarket Based on a Bundle Trading Mechanisms. Working paper, Center for Research on Electronic Commerce, The University of Texas, Austin, TX, 1998.
- [51] Fan, M., Stallaert, J., and Whinston, A.B. A Web-based Financial Trading System. *IEEE Computer*, 32(4) :64–70, 1999.
- [52] Fan, M., Stallaert, J., and Whinston, A.B. The Internet and the Future of Financial Markets. *Communications of the ACM*, 43(11) :83–88, 2000.
- [53] Figliozi, M., Mahmassani, H.S., and Jaillet, P. Modeling Online Transportation Market Performance : Carrier Competition under Asymmetric Fleet Assignment Technology. In *Proceedings of the 5th International Conference on*

- Electronic Commerce Research (ICECR-5), CDROM.* Centre de recherche sur les transports, Montréal, CANADA, 2002.
- [54] Fischer, K., Müller, J.P., and Pischel, M. Cooperative Transportation Scheduling : an Application Domain for DAI. DFKI GmbH, D-66123, Saarbrücken, 1996.
 - [55] Fujushima, Y., Leyton-Brown, K., and Shoham, Y. Taming the Computational Complexity of Combinatorial Auctions. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, pages 548–553, 1999.
 - [56] Garfinkel, R. and Nemhauser, G.L. The Set Partitionning Problem : Set Covering with Equality Constraints. *Operations Research*, 17(5) :848–856, 1969.
 - [57] Geoffrion, A.M. Elements of Large-Scale Mathematical Programming. *Management Science*, 16(11) :652–685, 1970.
 - [58] Geoffrion, A.M. Duality in Nonlinear Programming : A Simplified Applications-Oriented Development. *SIAM Review*, 13(1) :1–37, 1971.
 - [59] Gonen, R. and Lehmann, D. Optimal Solutions for Multi-Unit Combinatorial Auctions : Branch and Bound Heuristics. In *Proceedings of the ACM Conference on Electronic Commerce (EC'00)*, pages 13–20. ACM 2000, Minneapolis, MN, 2000.
 - [60] Graves, R.L., Schrage, L., and Sankaran, J. An Auction Method for course registration. *Interfaces*, 23(5) :81–92, 1993.
 - [61] Groves, T. Incentives in Teams. *Econometrica*, 41 :617–631, 1973.
 - [62] Gul, F. and Stacchetti, E. Walrasian Equilibrium with Gross Substitutes. *Journal of Economic Theory*, 87 :95–124, 1999.
 - [63] Gul, F. and Stacchetti, E. The English Auction with Differentiated Commodities. *Journal of Economic Theory*, 92 :66–95, 2000.
 - [64] Harris, M. and Raviv, A. Allocation Mechanisms and the Design of Auctions. *Econometrica*, 49 :1477–1499, 1981.

- [65] Hiriart-Urruty, J. and Lemaréchal, C. *Convex Analysis and Minimization Algorithms II : Advanced Theory and Bundle Methods*. A Series of Comprehensive Studies in Mathematics. Springer, New York, 1991.
- [66] Hoos, H.H. and Boutilier, C. Solving combinatorial auctions using stochastic local search. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 22–29, 2000.
- [67] Hudson, B. and Sandholm, T. Effectiveness of Preference Elicitation in Combinatorial Auction. In *Agent-Mediated Electronic Commerce (AMEC) workshop at AAMAS-02*. Bologna, Italy, 2002.
- [68] Jones, J.L. and Koehler, G.J. Combinatorial auctions using rule-based bids. *Decision Support Systems*, 34 :59–74, 2002.
- [69] Kalagnanam, J.R., Davenport, A.J., and Lee, H.S. Computational Aspects of Clearing Continuous Call Double Auctions with Assignment Constraints and Indivisible Demand. *Electronic Commerce Research Journal*, 1(3) :221–237, 2001.
- [70] Kalagnanam, J.R. and Parkes, D.C. Auctions, Bidding and Exchange Design. In Simchi-Levi, D., Wu, D.S., and Chen, M.Z., editors, *Supply Chain Analysis in the eBusiness Era*. Kluwer Academic Publishers, 2003. forthcoming.
- [71] Kelley, J.E. The cutting-plane method for solving convex programs. *J. SIAM*, 8 :703–712, 1960.
- [72] Kelly, F. and Steinberg, R. A Combinatorial Auction with Multiple Winners for Universal Service. *Management Science*, 46(4) :586–596, 2000.
- [73] Kfir-Dahav, N., Monderer, D., and Tennenholtz, M. Mechanism design for resource bounded agents. In *International Conference on Multi-Agent Systems, ICMAS 2001*, 2001.
- [74] Kim, K.H., Chung, W.J., Huang, H., and Chang, S.K. An Electronic Brokering Process for Truckload Freight. Technical report, Department of Industrial Engineer, Pusan National University, 2002.
- [75] Kiwiel, K.C. Proximity Control in Bundle Methods for Convex Nondifferentiable Optimization. *Mathematical Programming*, 46 :105–122, 1990.

- [76] Klemperer, P.D. Auction theory : A guide to the literature. *Journal of Economic Surveys*, 13 :227–286, 1999.
- [77] Kothari, A., Sandholm, S., and Suri, S. Solving Combinatorial Exchanges : Optimality via a Few Partial Bids. In *AAAI-02 workshop on AI for Intelligent Business*. Edmonton, Canada, 2002.
- [78] Krishna, V. and Perry, M. Efficient Mechanism Design. Technical report, Department of Economics, Pennsylvania State University, 1997.
- [79] Kurose, J.F. and Simha, R. Microeconomic Optimal Resource Allocation in Distributed Computer Systems. *IEEE Transactions on Computers*, 38(5) :705–717, 1989.
- [80] Kutanoglu, E. and Wu, D. On Combinatorial Auction and Lagrangian Relaxation for Distributed Resource Scheduling. *IEEE Transactions*, 31 :813–826, 1999.
- [81] Larsson, T., Patriksson, M., and Strömberg, A-B. Ergodic Primal Convergence in Dual Subgradient Schemes for Convex Programming. *Mathematical Programming*, 86 :283–312, 1999.
- [82] Lasdon, L.S. *Optimization Theory for Large Systems*. Macmillan, London, 1970.
- [83] Ledyard, J.O., Olson, M., Porter, D., Swanson, J.A., and Torma, D.P. The First Use of a Combined Value Auction for Transportation Services. Social Science Working Paper SSWP 1093, Caltech, 2000.
- [84] Lehmann, D. and Gonen, R. Linear programming helps solving large multi-unit combinatorial auctions. In *Proceedings of the Electronic Market Design Workshop*. Maastricht, The Netherlands, 2001.
- [85] Lehmann, D., O’Callaghan, L.I., and Shoham, Y. Truth revelation in approximately efficient combinatorial auctions. In *Proc. of the first ACM Conference on Electronic Commerce*, pages 96–102, 1999.
- [86] Lemaréchal, C. 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.

- [87] Lemaréchal, C. Lagrangian Relaxation. In Juenger, M. and Naddef, D., editors, *Computational Combinatorial Optimization*. Springer Verlag, 2001. Lecture Notes in Computer Science 2241.
- [88] Leonard, H.B. Elicitation of Honest Preferences for the Assignment of Individuals to Positions. *Journal of Political Economy*, 91(3) :461–479, 1983.
- [89] Leyton-Brown, K., Shoham, Y., and Tennenholz, M. An Algorithm for Multi-unit Combinatorial Auctions. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 56–61, 2000.
- [90] Madhavan, A. Trading Mechanisms in Securities Markets. *Journal of Finance*, 47(2) :607–641, 1992.
- [91] Malone, T.W., Yates, J., and Benjamin, R.I. Electronic Markets and Electronic Hierarchies. *Communications of the ACM*, 30(6) :484–497, 1987.
- [92] Mansini, R. and Speranza, M.G. Multi-unit Combinatorial Auctions : An Exact Approach. In *Proceedings of the 5th International Conference on Electronic Commerce Research (ICECR-5)*, CDROM. Centre de recherche sur les transports, Montréal, CANADA, 2002.
- [93] Markowitz, H. Portfolio Selection. *Journal of Finance*, 7(1) :77–91, 1992.
- [94] Martello, S. and Toth, P. *Knapsack Problems : Algorithms and Computer Implementation*. Wiley & Sons, 1990.
- [95] Mas-Colell, A., Whinston, M.D., and Green, J.R. *Microeconomic Theory*. Oxford University Press, 1995.
- [96] Maskin, E.S. and Riley, J.G. Optimal Auctions with Risk-Averse Buyers. *Econometrica*, 52(6) :1473–1518, 1984.
- [97] Mathiesen, L. Computational Experience in Solving Equilibrium Models by a Sequence of Linear Complementarity Problems. *Operations Research*, 33 :1225–1250, 1985.
- [98] Matsumoto, Y. and Fujita, S. An Auction Agent for Bidding on Combinations of Items. In *Proceedings of the 5th International Conference on Autonomous Agents*. Montréal, Canada, 2001.

- [99] McMillan, J. Selling Spectrum Rights. *Journal of Economic Perspectives*, 8 :145–162, 1994.
- [100] Milgrom, P. Auctionning the Radio Spectrum. In *Auction Theory for Privatization*. Cambridge University Press, Cambridge, UK, 1995.
- [101] Milgrom, P.R. A Convergence Theorem for Competitive Bidding with Differential Information. *Econometrica*, 47 :679–688, 1979.
- [102] Milgrom, P.R. Putting auction theory to work : the simultaneous ascending auction. Technical Report 98-0002, Department of Economics, Stanford University, 1997.
- [103] Myerson, R.B. Optimal Auction Design. *Mathematics of Operations Research*, 6(1) :58–73, 1981.
- [104] Nagurney, A. *Network Economics : a Variational Inequality Approach*. Kluwer Academic Publishers, 1993.
- [105] Nisan, N. Bidding and Allocation in Combinatorial Auctions. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, pages 1–12. ACM 2000, Minneapolis, MN, 2000.
- [106] Nisan, N. The Communication Complexity of Combinatorial Auctions. Draft, Institute of Computer Science, The Hebrew University of Jerusalem, 2001. Available at : <http://www.cs.huji.ac.il/~noam/mkts.html>.
- [107] Nisan, N. and Ronen, A. Computationally Feasible VCG Mechanisms. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, pages 242–252. ACM 2000, Minneapolis, MN, 2000.
- [108] Noussair, C. Equilibria in a Multi-Object Uniform Price Sealed Bid Auction with Multi-Unit Demands. *Economic Theory*, 5 :337–351, 1995.
- [109] Oren, M.E. and Williams, A.C. On Competitive Bidding. *Operations Research*, 23(6) :1072–1079, 1975.
- [110] Ortega-Reichert, A. *Models for Competitive Bidding Under Uncertainty*. PhD thesis, Stanford University, 1968.

- [111] Parkes, D.C. iBundle : An Efficient Ascending Price Bundle Auction. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 148–157. ACM 2000, Minneapolis, MN, 2000.
- [112] Parkes, D.C. An Iterative Generalized Vickrey Auction : Strategy-proofness Without Complete Revelation. In AAAI Press, editor, *Proc. AAAI Spring Symposium on Game Theoretic and Decision Theoretic Agents*, pages 78–87, 2001.
- [113] Parkes, D.C. *Iterative Combinatorial Auctions : Achieving Economic and Computational Efficiency*. PhD thesis, Department of Computer and Information Science, University of Pennsylvania, 2001.
- [114] Parkes, D.C., Kalagnanam, J., and Eso, M. Achieving Budget-Balance with Vickrey-Based Payment Schemes in Exchanges. In *Proc. 17th International Joint Conference on Artificial Intelligence*, pages 1161–1168, 2001.
- [115] Parkes, D.C. and Ungar, L. Iterative Combinatorial Auctions : Theory and Practice. In *Proceedings of the 17th National Conference on Artificial Intelligence*, pages 74–81, 2000.
- [116] Parkes, D.C. and Ungar, L. An ascending-price generalized vickrey auction. Technical report, Division of Engineering and Applied Science, Harvard University, 2002.
- [117] Polk, C.W. and Schulman, E. Enhancing the Liquidity of Bond Trading. In *The Handbook of Fixed-Income Technology*. Summit Group Press, New York, 2000.
- [118] Popper, M. The Ins and Outs of Portfolio Trading. *Global Investor*, pages 18–20, 1995.
- [119] Powell, W.B. Dynamic Models of Transportation Operations. In Graves, S. and de Tok, T.A.G., editors, *Handbook in Operations Research and Management Science : Supply Chain Management*. North-Holland, 2003. To appear.
- [120] Rassenti, S.J., Smith, V.L., and Bulfin, R.L. A Combinatorial Auction Mechanism for Airport Time Slot Allocation. *Bell Journal Of Economics*, 13(2) :402–417, 1982.

- [121] Reece, D.K. Competitive Bidding for Offshore Petroleum Leases. *Bell Journal of Economics*, 9 :369–384, 1978.
- [122] Riley, J.G. and Samuelson, W.F. Optimal Auctions. *American Economic Review*, 71 :381–392, 1981.
- [123] Rothkopf, M.H. A Model of Rational Competitive Bidding. *Management Science*, 15 :362–373, 1969.
- [124] Rothkopf, M.H. and Park, S. An Elementary Introduction to Auctions. *Interfaces*, 31(6) :83–97, 2001.
- [125] Rothkopf, M.H., Pekeč, A., and Harstad, R.M. Computationally manageable combinatorial auctions. *Management Science*, 44 :1131–1147, 1998.
- [126] Rothkopf, M.H., Teisberg, T.J., and Kahn, E.P. Why are Vickrey Auctions Rare ? *Journal of Political Economy*, 98 :94–109, 1990.
- [127] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the 12th International Workshop on Distributed Artificial Intelligence*, pages 295–308, Hidden Valley, Pennsylvania, 1993.
- [128] Sandholm, S. and Suri, S. Side Constraints and Non-Price Attributes in Markets. In *International Joint Conference on Artificial Intelligence (IJCAI), Workshop on Distributed Constraint Reasoning*, pages 55–61. Seattle, WA, 2001.
- [129] Sandholm, S. and Suri, S. BOB : Improved Winner Determination in Combinatorial Auctions and Generalizations. *Artificial Intelligence*, 145 :33–58, 2003.
- [130] Sandholm, T. An Algorithm for Optimal Winner Determination in Combinatorial Auctions. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, pages 542–547, 1999.
- [131] Sandholm, T. e-Mediator : a next-generation electronic commerce server. In *Proceedings of the 4th International Conference on Autonomous Agents*, pages 341–348. Barcelona, 2000.
- [132] Sandholm, T., Suri, S., Gilpin, A., and Levine, D. CABOB : A Fast Optimal Algorithm for Combinatorial Auctions. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, pages 1102–1108, 2001.

- [133] Sandholm, T., Suri, S., Gilpin, A., and Levine, D. Winner determination in combinatorial auction generalizations. In *First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 69–76. Bologna, Italy, 2002.
- [134] Scarf, H.E. *Computation of Economic Equilibria*. Yale University Press, New Haven, CT, 1973.
- [135] Schramm, H. and Zowe, J. 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.
- [136] Sharpe, W.F. Capital Asset Prices : A Theory of Market Equilibrium Under Conditions of Risk. *Journal of Finance*, 19(3) :425–442, 1964.
- [137] Smith, T., Sandholm, T., and Simmons, R. Constructing and Clearing Combinatorial Exchanges Using Preference Elicitation. In *Proceedings of AAAI-02 workshop on Preferences in AI and CP : Symbolic Approaches*, pages 87–93, 2002.
- [138] Srinivasan, S., Stallaert, J., and Whinston, A.B. Portfolio Trading and Electronic Networks. CREC working paper, University of Texas at Austin, 1998.
- [139] Tan, J.C. and Harker, P.T. Designing workflow coordination : centralized versus market-based mechanisms. *Information Systems Research*, 10 :328–342, 1999.
- [140] Vickrey, W. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16 :8–37, 1961.
- [141] Walras, L. *Eléments d'économie politique pure*. L. Corbaz, Lausanne, 1874.
- [142] Walsh, W.E, Wellman, M.P., and Ygge, F. Combinatorial Auctions for Supply Chain Formation. In *Proc. ACM Conference on Electronic Commerce*, pages 260–269, 2000.
- [143] M.P. Wellman. Market-oriented programming : Some early lessons. In Scott Clearwater, editor, *Market-Based Control : A Paradigm for Distributed Resource Allocation*. World Scientific, River Edge, New Jersey, 1996.

- [144] Wellman, M.P. A market-oriented programming environment and its application to multi-commodity flow problems. *Journal of Artificial Intelligence Research*, 1 :1–22, 1993.
- [145] Wellman, M.P., Walsh, W.E., Wurman, P.R., and MacKie-Mason, J.K. Auction Protocols for Decentralized Scheduling. *Games and Economic Behavior*, 35 :271–303, 2001.
- [146] Wilson, R.B. Competitive Bidding With Disparate Information. *Management Science*, 15(7) :446–448, 1969.
- [147] Wolfe, P. A Method of Conjugate Subgradients for Minimizing Nondifferentiable Convex Functions. *Mathematical Programming Study*, 3 :145–173, 1975.
- [148] Wurman, P.R. and Wellman, M.P. AkBA : A Progressive Anonymous-price Combinatorial Auction. In *Second ACM Conference on Electronic Commerce*, pages 21–29, 1999.
- [149] Wurman, P.R. and Wellman, M.P. Equilibrium Prices in Bundle Auctions. In *AAAI-99 Workshop on Artificial Intelligence for Electronic Commerce*, pages 56–61, 1999.
- [150] Wurman, P.R., Wellman, M.P., and Walsh, W.E. A Parametrization of the Auction Design Space. *Games and Economic Behavior*, 35 :304–338, 2001.
- [151] Xia, M., Koehler, G.J., and Whinston, A.B. Pricing Combinatorial Auctions. *European Journal of Operations Research*, 2003. Forthcoming.
- [152] Yamaki, H., Wellman, M.P., and Ishida, T. A market-based approach to allocating QoS for multimedia applications. In *Second International Conference on Multi-Agent Systems*, pages 385–392. Kyoto, 1996.
- [153] Ygge, F. *Market-Oriented Programming and its Application to Power Load Management*. PhD thesis, Department of Computer Science, Lund University, 1998.
- [154] Ygge, F. and Akkermans, H. Resource-Oriented Multi-Commodity Market Algorithms. *Autonomous Agents and Multi-Agent Systems Journal*, 3(1) :53–71, 2000.

