CIRRELT

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

Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation

Combinatorial Auctions

Jawad Abrache
Teodor Gabriel Crainic
Michel Gendreau
Monia Rekik

January 2007

CIRRELT-2007-02

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

Bureaux de Québec :
Université Laval
Pavillon Palasis-Prince, local 2642
Québec (Québec)
Canada G1K 7P4
Téléphone : 418 656-2073
Télécopie : 418 656-2624

www.cirrelt.ca

Université Laval · UQÀM · HEC Montréal · École Polytechnique Montréal · Université de Montréal

Combinatorial Auctions

Jawad Abrache[1], Teodor Gabriel Crainic[2,3*], Michel Gendreau[2,4], Monia Rekik[2]

[1] School of Business Administration, Al Akhawayn University, P.O. Box 104, Ifrane 5300, Morroco

[2] Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT), Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada H3C 3J7

[3] Chaire de recherche industrielle du CRSNG en management logistique, ESG, Université du Québec à Montréal, C.P. 8888, succursale Centre-ville, Montréal, Canada H3C 3P8

[4] Département d'informatique et de recherche opérationnelle, 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.

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 FQRNT research development fund of the Province of Québec.

---

* Corresponding author: Teodor-Gabriel.Crainic@cirrelt.ca

This document is also published as Publication #1293 by the Department of Computer Science and Operations Research (DIRO) of the Université de Montréal.

Dépôt légal – Bibliothèque nationale du Québec,
Bibliothèque nationale du Canada, 2007

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-business 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, 1994) 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 et al., 1982; Ball et al., 2006), course registration (Graves et al., 1993), private and public procurement (Davenport and Kalagnanam, 2001), sale of online seats (Eso, 2001), distribution routes (Ledyard et al., 2002), job shop scheduling (Wellman et al., 2001), electricity markets (Ausubel and Cramton, 2004; **?**), trading (Abrache et al., 2005), 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 airport time slots as an example. A take-off slot associated to the origin airport of a flight and a corresponding landing slot at 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 reaches its preference for the collection. But since the auctions of the different items run independently of each other, the participant may find itself 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 (see, for example, (Pekeč and Rothkopf, 2003). 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 et al., 1998) and even difficult to approximate (Sandholm, 2002).

The challenge of *mechanism design* (Mas-Colell et al., 1995) 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 participants, 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, 2001), some others are indeed fundamental. These issues concern, for instance, the decision to make bidding in the auction one-shot or progressive, and 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, 1981), 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 et al., 1989), the RAD mechanism (DeMartini et al., 1999), the PAUSE mechanism (Kelly and Steinberg, 2000), the AkBA family of auctions (Wurman and Wellman, 1999a), the iBundle mechanism (Parkes, 1999), the ascending proxy auctions (Ausubel and Milgrom, 2002), and the clock-proxy auctions (Ausubel et al., 2006).

As pointed out in (Rothkopf and Park, 2001), 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, 1999). By putting game theory into application, they have built models that describe the strategic behavior of the participants in many auction types. Among main issues relevant to auctions, they have shaped powerful theories for economic efficiency, pricing, incentives, and collusive behavior. 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 simulation platforms for the evaluation of auction mechanisms in controlled artificial environments. As for operations research, it already provides tools for efficiently addressing winner determination and pricing issues. Operations research 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 emphasize 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. This paper is an updated version of (Abrache et al., 2004b) and includes the new research on combinatorial auctions published in the last two years.

The paper is organized as follows. In Section 2, we present an elementary taxonomy of auctions and survey several important formulations of the winner determination problem. In Section 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 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 5 with a general discussion about the role of advisors to participants in combinatorial auctions.

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, 1980) and (Wurman et al., 2001). 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 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 (the one-to-many case), or (b) many sellers and one buyer (the many-to-one case). Multilateral auctions, often designated by the name *exchanges*, involve many sellers and many buyers (the many-to-many case). 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 et al., 2001) 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.

- **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.

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 $(M1)$:

$$\max \sum_{1 \leq j \leq n} \sum_{S \subseteq G} p_{j,S} x_{j,S} \tag{1}$$

$$s.t. \sum_{1 \leq j \leq n} \sum_{S \subseteq G} \delta_{i,S} x_{j,S} \leq 1, \forall i \in G, \tag{2}$$

$$\sum_{S \subseteq G} x_{j,S} \leq 1, \forall j, 1 \leq j \leq n, \tag{3}$$

$$x_{j,S} \in \{0,1\}, \forall S \subseteq G, \forall j, 1 \leq j \leq n, \tag{4}$$

where $\delta_{i,S} = 1$ if $i \in S$, and 0 otherwise. Constraints (2) establish that no single item is allocated to more than one buyer, while constraints (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, 2003). The classical account of the problem by (Rothkopf et al., 1998) 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). However, restricting the bundles on which bids are submitted may deprive participants of bidding on the bundles they desire, which leads to the same inefficiency and exposure issues encountered in non-combinatorial auctions. (Park and Rothkopf, 2005) recently proposed a sealed-bid auction in which the decision on what is biddable is delegated to the bidders themselves. Bidders submit a prioritized list of combinations upon which they want to bid at the beginning of the auction and the auctioneer uses as many of these combinations as is computationally possible. More recently, (Müller, 2006) provided a survey of tractable winner determination problems based on both subset restrictions and preference-type restrictions.

By opposition to the worst-case analysis of (Rothkopf et al., 1998), the search algorithms that have been proposed in the literature (e.g., (Fujushima et al., 1999; Sandholm, 2002; Sandholm et al., 2001; Hoos and Boutilier, 2000)) 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, 2002) 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 et al., 1999) 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 $r^\star_C$ that can be achieved from the sale of $C \subseteq G$. Consider a partial feasible allocation of the subset $F \subseteq G$ at a given step of the search such that $G \setminus F \subseteq C$. Then, if $r^\star_C + r_F$ 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., 2001) calls upon additional techniques that include pruning with upper and lower bounds, decomposition of the bid graph, and dynamic branching heuristics. (Andersson et al., 2000) 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 concluded that the overall performance of CPLEX is actually very good compared to that of the search algorithms. Recently, (Sandholm et al., 2005) compared the performance of the newest development of their search algorithm, CABOB, to CPLEX 8.0 using different combinatorial auction benchmark distributions. CPLEX was faster than CABOB for most distributions. However, CABOB required less CPU time for the so-called *Components* distribution (i.e., a distribution in which the problem can be decomposed into smaller independent sub-problems) thanks to the decomposition techniques it incorporates. (Günlük et al., 2005) presented a Branch-and-Price algorithm for solving the winner determination problem for combinatorial auctions with bids having a XOR-of-OR structure (see Section 3.1 for the definition of XOR-of-OR bidding languages). The algorithm is applied to a formulation of the WDP which uses considerably more variables than the classical formulation, but which is proved to yield tighter linear programming relaxations. The proposed Branch-and-Price algorithm is compared to CPLEX 8.0 (applied to the classical formulation) for a large set of generated instances inspired from FCC spectrum auctions. Although the reported results give a computational advantage to the Branch-and-Price procedure over CPLEX, no conclusive results can be made on the relative performance of one method to the other for general instances. In fact, the authors reported that CPLEX performed better than their Branch-and-Bound procedure for other instances derived from pre-existing generators for combinatorial auctions.

While many approximate methods for the general set packing problem have been suggested in the literature (Chandra and Halldórsson, 2001), the Casanova algorithm by (Hoos and Boutilier, 2000) is, to the best of our knowledge, the first 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. (Sakurai et al., 2000) presented an approximate algorithm based on limited discrepancy search (LDS). The LDS algorithm limits the search efforts to the region that is likely to contain good solutions. When compared to the optimal IDA^* search algorithm of (Sandholm, 2002), the LDS algorithm yielded good solutions (within 5% of optimal) with smaller computing times (1% of the running time required by the IDA^* algorithm). Near-optimal solutions (usually within 1% of optimal) were also obtained by the two-phase heuristic proposed by (Zürel and Nisan, 2001). The first phase of the heuristic approximates the LP relaxation of the set packing problem. This fractional approximate solution determines a first order on the bids that is adopted by a greedy algorithm in the second phase. The resulting allocation is improved through local changes in the ordering of the bids.

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 $x_b = 1$ if bid b wins, and 0 if it loses, $\forall b \in B$. The winner determination problem can be written in this case as model ($M2$):

$$\max \quad \sum_{b \in B} p_b x_b \tag{5}$$

$$s.t. \quad \sum_{b \in B} a_{b,i} x_b \leq M_i, \forall i \in G, \tag{6}$$

$$x_b \in \{0, 1\}, \forall b \in B. \tag{7}$$

Model ($M2$) is an 0-1 multidimensional knapsack problem for which exact and heuristic solution methods have been designed and implemented (Martello and Toth, 1997). In the context of combinatorial auctions, many search algorithms have been recently proposed. Among the important contributions, (Leyton-Brown et al., 2000) and (Leyton-Brown, 2003) present the CAMUS ("Combinatorial Auction Multi-Unit Search") algorithm is 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 context of the Branch-and-Bound algorithms of (Gonen and Lehmann, 2000) and (Lehmann and Gonen, 2001). Finally, the observation in (Mansini and Speranza, 2002) 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 improvements in performance in comparison with CPLEX 7.0, particularly for large problems.

2.2 Many-to-one combinatorial auctions

In a many-to-one combinatorial auction configuration (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. 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 $(M3)$:

$$\min \quad \sum_{b \in B} p_b x_b \tag{8}$$

$$s.t. \quad \sum_{b \in B} \delta_{i,S_b} x_b \geq 1, \forall i \in G, \tag{9}$$

$$x_b \in \{0,1\}, \forall b \in B, \tag{10}$$

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. It is important to note that an implicit *free disposal* assumption is made in model $(M3)$; that is, the buyer tolerates 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 (9) need to be changed to equalities. The corresponding set partitioning problem proves to be relatively more difficult to tackle (Sandholm et al., 2002).

In practice, reverse auctions are especially useful as market mechanisms for the procurement of goods and services. Among the many applications that have invoked reverse combinatorial auctions in recent years, one may cite trucking service acquisition for Sears Logistical Services (Ledyard et al., 2002), procurement of direct inputs for a food manufacturer (Davenport and Kalagnanam, 2001; **?**), and assignment of nationwide meal service contracts by the Chilean government (Epstein et al., 2002). In a typical procurement application, side constraints are often needed to ensure additional conditions on a valid assignment are satisfied. For instance, constraints that set limits on the number of winning bids and the volumes of goods received from each of the bidders appear in (Davenport and Kalagnanam, 2001) and (Hohner et al., 2003), while (Epstein et al., 2002) set limits on (a) the total number of winning bidders, (b) the number of territorial units (geographic regions) assigned to a bidder, and (c) the number of bidders that will operate in a given region. In a recent discussion on the use of combinatorial auctions for transportation services procurement, (Caplice and Sheffi, 2006) classify the side constraints commonly encountered in the United States truckload markets into three categories: (1) business guarantee constraints, which set bounds on the amount of traffic (in terms of loads won or total estimated dollar value) a carrier or a set of carriers wins, (2) carrier base size constraints, which restrict the total number of winning carriers in the system, the region or the lane, and (3) If-Then constraints, which express the fact that if a carrier is awarded any business, then it has to be of a certain minimum level. The authors illustrate how these constraints can be mathematically modeled.

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 network flow algorithms can help in finding the optimal allocations more efficiently. 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 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 $O_2 - I_2 - D_2$.

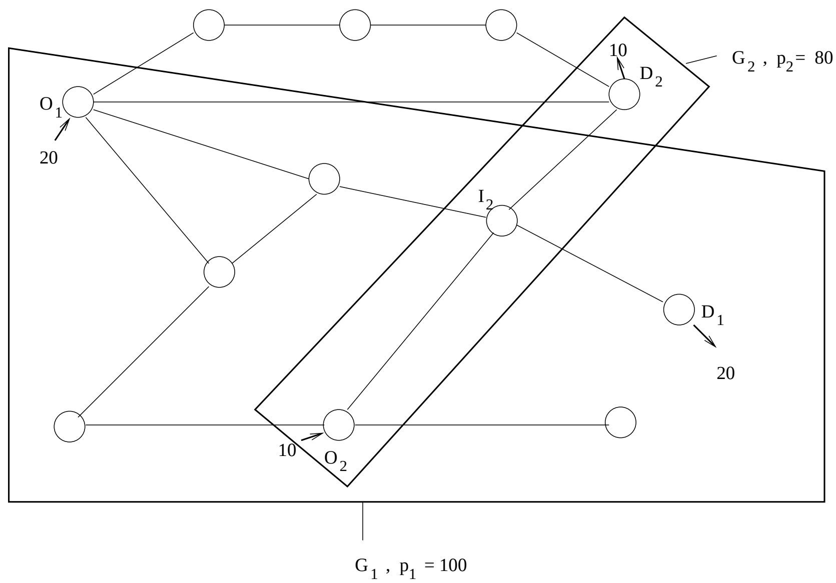

Figure 1: Combinatorial bids on capacity

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 $k \in K_j$.

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

$$\max \quad \sum_{j \in N} x_j p_j \tag{11}$$

$$s.t. \quad \sum_{k \in K_j} h_k = c_j x_j, \forall j \in N, \tag{12}$$

$$\sum_{j \in N} \sum_{k \in K_j} \delta_{a,k} h_k \leq v_a, \forall a \in A, \tag{13}$$

$$x_j \in \{0, 1\}, \forall k \in K_j, \forall j \in N, \tag{14}$$

$$h_k \geq 0, \forall k \in K_j, \forall j \in N, \tag{15}$$

where $\delta_{a,k} = 1$ if $a \in k$, 0 otherwise. Constraints (12) state that the capacity allocated to a winning bid must be within the bid's sub-network, while constraints (13) correspond to capacity availability on links.

When paths are completely specified by buyers, i.e., G_j is limited to a single path between O_j and D_j, $\forall j \in N$, and single units 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 do 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 et al., 1993) can therefore be used instead of plain simplex in a Branch-and-Bound procedure, for example.

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 $(M5-a)$:

$$\max \quad \sum_{b\in B} p_b x_b \tag{16}$$

$$s.t. \quad \sum_{b\in B} q_{b,i} x_b \leq 0, \forall i \in G, \tag{17}$$

$$x_b \in \{0,1\}, \forall b \in B. \tag{18}$$

Model $(M5-a)$ maximizes the total surplus of the market under the constraint that sales should cover buys. Notice again that inequalities in constraints (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 $(M5-b)$:

$$\max \quad \sum_{b\in B} p_b(x_b) \tag{19}$$

$$s.t. \quad \sum_{b\in B} q_{b,i} x_b \leq 0, \forall i \in G, \tag{20}$$

$$0 \leq x_b \leq 1, \forall b \in B. \tag{21}$$

While model $(M5-b)$ is generally easy to solve, especially when the price mappings $p_b(.)$ 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, 2003) 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 et al., 2002) 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 et al., 2000; **?**)), supply chain formation and coordination (Walsh et al., 2000), and market clearing in process industries (Kalagnanam et al., 2001). 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.

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 (Sandholm and Suri, 2001; **?**). 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, 2001). 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, 2001).

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.

3.1 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 for the execution of these orders. In order to permit an optimal usage 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 et al., 1999), as well as in (Sandholm, 2002), one may find the expression of exclusive OR (XOR) conditions through the introduction of *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 $A \cup \{l\}$, $B \cup \{l\}$, ..., $E \cup \{l\}$, and submit five unrelated bids on these new bundles. The OR and XOR logics have been combined for the first time in (Hoos and Boutilier, 2000). 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, 2000) 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, 2001) 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 she 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 her 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_2 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 she 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).

(Abrache et al., 2004a) pointed out that all the languages previously considered in the literature were formalized for one-sided combinatorial auctions of indivisible single-unit items. They present a new bidding language framework that is independent of the physical nature of the items traded and their divisibility. The proposed language is general enough to be applied to either unilateral or multilateral markets. It also allows participants to submit complex bidding definitions, requirements and conditions in a succinct way. We describe this framework in more details in Section 3.2.

(Cavallo et al., 2005) presented a new logical tree-based bidding language (TBBL) that allows participants to express preferences for both buying and selling goods in the same structure. The leaves of the tree-bid correspond to individual item trades, either a buy or a sell, and the internal nodes describe the relationship between lower-level nodes. More precisely, an internal node is an "interval-choose" (IC_x^y) operator which defines a range of children nodes (at least x and at most y) that need to be satisfied. This class of IC_x^y operators include the standard XOR, OR, AND and $k-of$ operators. TBBL also imposes that a given node in the tree can be satisfied only if its parent node is satisfied. This latter rule gives the internal operator nodes the possibility to act as constraints on what allocations are acceptable. When compared to OR^*, a TBBL language allowing a node to have multiple parents is proved to be more concise than OR^*. This increase in conciseness is also observed with regard to the $\mathcal{L}_{GB}$ language of (Boutilier and Hoos, 2001) in some but not all settings. For example, TBBL is more concise when valuations are expressed on an interval range of satisfied leaves.

TBBL also supports a partial value revelation. That is, bidders are allowed to specify upper and lower bounds on their values for trades. Such a functionality is interesting in iterative mechanisms since bidders are given the possibility to refine their valuations at each round of the auction as a response to the information feedback. (Parkes et al., 2005) recently proposed an iterative combinatorial exchange mechanism that incorporates the TBBL language.

(Day and Raghavan, 2006) introduced a matrix bidding language and showed that it is as expressive as that proposed by (Boutilier and Hoos, 2001) with the $k-of$ operator. A bidder expresses its preferences through a matrix $\mathcal{B}$ in which a column j represents a fictive "unit-demand" agent and a row i corresponds to an auctioned item. A bidder must specify a total and strict ordering of the auctioned items so that the highest ranked item corresponds to the first row of the matrix. A value b_{ij} in that matrix corresponds to the bid offered by the bidder for item i given that it is the j^{th} best item it receives. It is also assumed that any "unit-demand" agent (except the first one) cannot receive an item unless the next more highly ranked agent receives a more highly ranked item. Matrix bidding has the advantage to be particularly compact in expressing some preferences (for example, preferences with capacity constraints and/or cost structure) that could have required several $k - of$ expressions. However, the value an item brings to a bidder is determined only by its ranking with respect to the other items in the bundle and is assumed independent of the identity of the items already received. In other words, the incremental value an item provides to a bundle is assumed the same for all the bundles containing higher ranked items.

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 (Abrache et al., 2004a) 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 the 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., 2004a). In the following, we present a brief survey of the important concepts of the framework.

3.2.1 The inner level

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

Definition 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 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., 2004a), 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-INNER** 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:

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

3.2.2 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 (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*
   - $\Omega_j = \{\Theta_{\bar{\jmath}}\}_{\bar{\jmath} \in J_j}$ *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 $\Omega_j = \{\Theta_{\bar{\jmath}}\}_{\bar{\jmath} \in J_j}$. Denote by $\Psi_j = \{\bar{\jmath} \in J_j : \Theta_{\bar{\jmath}} \text{ is executed}\}$ 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

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

Here N^l and N^u are integer parameters such that $0 \le N^l \le N^u \le |J_j|$. 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 Θ_1 AND $\Theta_2 \equiv S_{2,2}(\{\Theta_1, \Theta_2\})$, Θ_1 OR $\Theta_2 \equiv S_{1,2}(\{\Theta_1, \Theta_2\})$, and Θ_1 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., 2004a), 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 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:

$$\max \quad \sum_{l \in L} \sum_{j \in J_l} p_{lj} x_{lj} \tag{22}$$

$$s.t. \quad \sum_{l \in L} \sum_{j \in J_l} q_{ljr} x_{lj} = 0, \quad r \in G, \tag{23}$$

$$0 \leq x_{lj} \leq 1, \quad l \in L, j \in J_l. \tag{24}$$

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 considers 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 of 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:

$$\max \quad \sum_{l \in L} \sum_{j \in J_l} p_{lj} x_{lj} \tag{25}$$

$$s.t. \quad \sum_{l \in L} \sum_{j \in J_l} q_{ljr} x_{lj} = 0, \quad r \in G, \tag{26}$$

$$b_{lj} y_{lj} \leq x_{lj} \leq y_{lj}, \quad l \in L, j \in J_l, \tag{27}$$

$$\sum_{j \in \mathcal{X}} y_{lj} \leq 1, \quad \mathcal{X} \in \mathcal{X}_l, l \in L. \tag{28}$$

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 et al., 2005).

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., 2005). 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 software 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, 1975)). 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, 2001) and implies that optimal outcomes in combinatorial auctions cannot be achieved, in the general case, with sub-exponential data communication.

Iterative auctions (Cramton, 1998) 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 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 own 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.

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

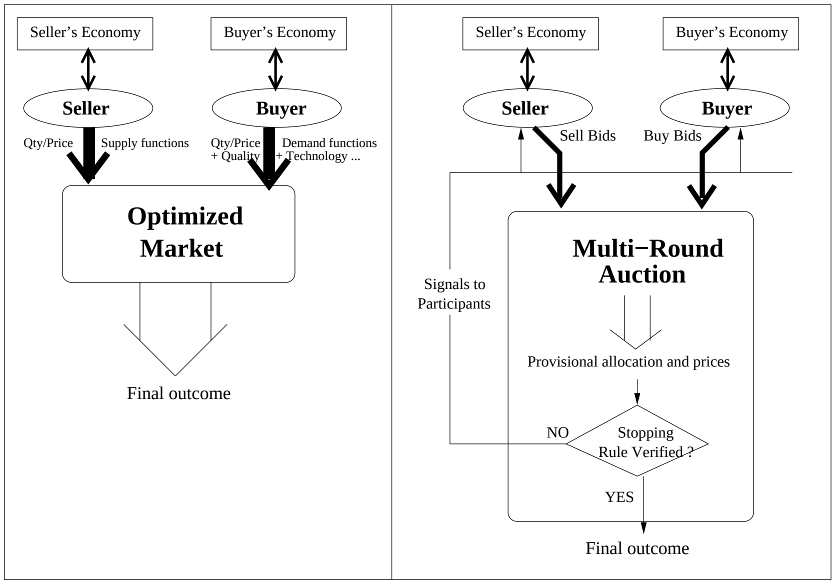

Figure 2: Direct revelation mechanisms vs multi-round auctions

schemes in price-directed auctions, and the incentive-compatibility properties of an iterative combinatorial auction.

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 threw 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.

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 in the auction. Each phase may consequently be characterized by its own rules and "mechanism", and produce a provisional outcome supplied as 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.

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. See (Xia et al., 2004) for a recent survey on pricing combinatorial auctions.

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 $M = [M_1, \ldots, M_m]$ of the goods. Each buyer j has a preference $v_j(x_j)$ for bundle $x_j = [x_{j,1}, \ldots, x_{j,m}]$. A *socially-efficient* allocation $x^\star$ 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 $x^\star$ maximizes the payoff of each buyer; that is

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

with the usual quasi-linear utility assumption.

A classical result of the general equilibrium theory (Arrow and Debreu, 1954) establishes that Walrasian equilibrium prices exist under conditions of continuity, monotony, and concavity of preference functions $v_j(.)$. 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, for example, (Arrow and Hahn, 1971)).

The literature on iterative auction design in presence of indivisibilities has mainly focused on the Combinatorial Allocation Problem (CAP) (de Vries and Vohra, 2003). 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 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):

$$\max \quad \sum_{1 \le j \le n} \sum_{S \subseteq G} v_j(S) x_{j,S} \tag{29}$$

$$s.t. \quad \sum_{1 \le j \le n} \sum_{S \subseteq G} \delta_{i,S} x_{j,S} \le 1, \forall i \in G, \tag{30}$$

$$\sum_{S \subseteq G} x_{j,S} \le 1, \forall j, 1 \le j \le n, \tag{31}$$

$$x_{j,S} \in \{0, 1\}, \forall S \subseteq G, \forall j, 1 \le j \le n. \tag{32}$$

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, 1997): optimal solutions $\{p^{\star}{}_i\}_{i \in G}$ and $\{\tau^{\star}{}_j\}_{1 \le j \le n}$ of

$$\min \quad \sum_{1 \le j \le n} \tau_j + \sum_{i \in G} p_i \tag{33}$$

$$s.t. \quad \sum_{i \in S} p_i + \tau_j \ge v_j(S), \forall j \in N, \forall S \subseteq G, \tag{34}$$

$$p_i \ge 0, \forall i \in G, \tag{35}$$

$$\tau_j \ge 0, \forall j, 1 \le j \le n, \tag{36}$$

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, 1999) 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, 2002) 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, 1999) 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. More recently, (O'Neill et al., 2005) presented a primal-dual method to construct a set of linear prices that support a Walrasian competitive equilibrium in markets with non-convexities. These prices are proved to correspond to the dual prices of an "augmented" linear program. The latter is obtained by adding to the linear relaxation of the original MIP, a set of equality constraints forcing the variables to take their optimal integer values. Such a method assumes however that optimal integer solutions can be efficiently computed.

Many other iterative auctions based on different price-adjustment schemes have been suggested. For instance, the experimental RAD mechanism of (DeMartini et al., 1999) 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 to the next round, new bids by the participants need to beat the current provisional prices by a certain increment. Although the RAD mechanism achieves high level of ex post efficiency in experiments (Kwasnica et al., 2005), theoretical efficiency results are not available. (Wurman and Wellman, 1999b) 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 (Leonard, 1983). Given that this problem has multiple solutions in general, the authors suggest 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 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, 1999) for an application to distributed job shop scheduling). (de Vries and Vohra, 2003) recently emphasized the connections between the duality theory of optimization problems and the design of auctions. They showed that some of the auction mechanisms already proposed, such as the RAD, and the iBundle auctions can be given a Lagrangian relaxation interpretation. In a subsequent paper, (de Vries et al., 2005) showed that many of the well-known ascending auctions can be derived from either primal-dual or subgradient algorithms. (Abrache et al., 2003) were the first to analyse mathematical programming decomposition methods for combinatorial multilateral auctions, in which many sellers and buyers interact. They investigate iterative auction mechanisms based on Lagrangian relaxation and Dantzig-Wolfe decomposition for a general combinatorial exchange economy in which participants trade heterogeneous divisible commodities. They show that, under appropriate assumptions, the application of these decomposition techniques to the centralized allocation problem leads to indirect mechanisms that can achieve social efficiency without requiring complete information revelation from the participants.

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, 1961; Clarke, 1971; Groves, 1973) is known to be an economically efficient, strategy-proof mechanism. Given preferences $\{\tilde{v}_j(.)\}_{j \in J}$ reported by participants, the VCG's allocation and payment rules are:

- Return an allocation $x^\star \in \arg\max_{x \in \mathcal{D}} \sum_{j \in J} \tilde{v}_j(x_j)$ that maximizes total value given the reported valuations;
- Participant j pays $V^\star_{-j} - (V^\star - \tilde{v}_j(x^\star_j))$, where $V^\star = \max_{x \in \mathcal{D}} \sum_{j \in J} \tilde{v}_j(x_j)$ and $V^\star_{-j} = \max_{x \in \mathcal{D}} \sum_{k \in J - \{j\}} \tilde{v}_k(x_k)$ (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, 2000) 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., 2001) give a primal-dual interpretation to Gul and Stacchetti's auction. (Ausubel, 2006) 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, 2001; Parkes and Ungar, 2002) 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. (de Vries et al., 2005) and (Mishra and Parkes, 2005) recently observed that deriving VCG payments with an ascending-price primal-dual algorithm is not always guaranteed. (de Vries et al., 2005) prove that a stronger condition on the *submodularity of coalition values* needs to be satisfied to achieve the VCG outcomes with an ascending price auction. (Mishra and Parkes, 2005) generalize the auction in (de Vries et al., 2005) and the iBundle auction in (Parkes and Ungar, 2000) from the restricted-valuation class of buyer-submodularity to general valuations. They introduce the concept of *universal competitive equilibrium prices* (these are competitive equilibrium prices for the main economy as well as for every marginal economy, i.e., an economy where one agent is excluded) and show that these prices are necessary and sufficient to yield VCG outcomes in an ascending-price auction. Truthful bidding is shown to be an ex post Nash equilibrium in the auctions proposed.

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 et al., 1990; **?**; **?**; **?**), 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. Moreover, (Ausubel and Milgrom, 2006) showed that the VCG auction loses its dominant-strategy property when bidders face effective budget constraints. Hence, some recent researches turned towards the design of alternative auction mechanisms that overcome some of the VCG auction drawbacks. (Parkes et al., 2001) proposed a Threshold payment rule that determines budget-balanced payments that minimize the maximal error to the VCG outcome across all agents. (Parkes et al., 2005) extended this threshold payment rule from one-sided to exchange markets with partial revelation. (Ausubel and Milgrom, 2002) proposed an ascending proxy auction and proved that the final payoffs it yields are always in the core. That is, there is no coalition of bidders that can trade among themselves in a way that generates strictly more revenue for the seller (the market maker) and equally or more preferred outcomes for all the bidders of the coalition. (Ausubel and Milgrom, 2006) showed that the Vickrey auction leads to such core allocations only when the goods-are-substitutes condition holds for all bidders and bidders' budgets are unlimited.

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 bidding 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 et al., 2002; Figliozzi et al., 2002). 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 may be used 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 realizing 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 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 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., 2001), 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.

(Das et al., 2001) undertook to compare the performances of software agents against humans in a continuous double auction where multiple units of a single hypothetical commodity could be bought or sold. The reported results show that software agents consistently obtained significantly larger trade gains than their human counterparts. However, convergence to equilibrium was generally slower than in prior all-software agent or all-human traders auctions. (Ausubel and Milgrom, 2002) introduced the concept of "proxy agents" in the ascending auction they propose. A proxy agent is a software agent that acts on behalf the bidder. More specifically, each bidder reports his values to the corresponding proxy agent for all the packages that it is interested in. Based on these values, the proxy submit bids on behalf the bidder so as to maximize its utility. (Ausubel and Milgrom, 2002) proved that ascending proxy auctions help bidders focus on the packages they are more likely to win given the bids made by competitors in earlier rounds. Moreover, they prove that, under some assumptions, using such proxy agents yields core allocations. (Hoffman et al., 2004) focused on the package determination problem in combinatorial FCC spectrum auctions. They describe an optimization-based bidder aid-tool support to simultaneously generate and valuate the optimal set of biddable packages at the beginning of the auction and dynamically before each round.

(Song and Regan, 2005) recently proposed an optimization-based approximation algorithm that helps trucking companies construct optimal or near-optimal bids for the procurement of freight transportation contracts. This carrier bid construction problem is modeled as a vehicle routing problem in which each route represents a candidate bid for the carrier and the objective function minimizes the total empty movement cost. The reported results are promising. (Lee et al., 2006) also considered vehicle routing models to help bidders identify sets of origin-destination pairs that maximize their profits. Column generation and Lagrangian based techniques are employed for solving this carrier optimization model.

(Kwon et al., 2005) studied a more general bid construction problem in the context of an iterative ascending framework (Parkes and Ungar, 2000). They propose a bidding mechanism that enables bidders to identify new valuable packages during the course of the auction. New bids are identified by the use of single-item prices that maintain the ascending property of the auction. Thus, at each round of the auction, a procedure similar to that proposed in the RAD mechanism (DeMartini et al., 1999) computes approximate single-item prices based on the provisional allocation. Given this set of prices, *best-response* packages are determined, i.e., packages that maximize the quasi-linear utility of the bidder, and conveniently priced. The authors prove that the efficiency yielded by an ascending mechanism with the proposed endogenous bidding can be greater than the efficiency produced by other ascending mechanisms where bidding is restricted to a fixed set of packages determined before the start of the auction.

Advisors have been considered in the literature to guide not only bidders in constructing profitable bids but also auctioneers in determining efficient allocations in partial-revelation mechanisms. In such mechanisms, bidders need not to bid on all the bundles nor to specify exact valuations for the packages they desire, circumventing thus the complexity of the communication and valuation tasks encountered in combinatorial auctions. Besides, (Nisan and Segal, 2006) showed that, in the worst case, identifying an optimal allocation requires an exponential amount of communication in the number of auctioned items. The role of the advisor, also called "*elicitor*" in this case, is to help the auctioneer elicit the bidders'preferences by cleverly querying them, i.e., asking them the "right" questions at the "right" times, without requesting bids on all bundles. *Incremental* querying enables thus the auctioneer to collect the information needed to derive optimal allocations. (Conen and Sandholm, 2001) used three types of queries in the *preference elicitation* framework they proposed: order, value and rank queries. In an order query, the elicitor asks the bidder to order two given bundles. In a value query, the elicitor presents a bundle and the bidder responds with either exact or approximate (i.e., upper and lower bounds) values on this bundle. A rank query asks a bidder on the rank of a given bundle, the bundle having a giving rank, etc. The gathered information is progressively stored in particular data structures (constraint networks) and is used by the proposed search algorithms to tighten the search space. In a subsequent paper, (Hudson and Sandholm, 2002) showed the merits of using such elicitation procedures in reducing the size of the information communicated by bidders. (Zinkevich et al., 2003) and (Santi et al., 2004) studied special classes of preferences for which they prove that elicitation can be achieved in a polynomial number of value queries (polynomial with regard to the number of auctioned items). Their proof relies on some known results of learning theory.

(Blum et al., 2004) and (Lahaie and Parkes, 2004) explored the links between the preference elicitation problem in combinatorial auctions with general valuations and the problem of learning an unknown function from computational learning theory. They considered learning algorithms with membership queries (i.e., a query in which the learner asks the oracle on the values of the unknown function at some points) and equivalence queries (i.e., a query in which the learner presents its current estimate of the unknown function and the oracle either agrees or returns a point at which the values of the estimate function and the true one are different).

(Lahaie and Parkes, 2004) proved that any exact learning algorithm with membership and equivalence queries can be converted into a preference elicitation algorithm with value and demand queries. In a demand query, the elicitor presents a vector of non-negative prices over all the possible bundles as well as a specified bundle. The bidder either confirms that the proposed bundle is the most preferred at the specified prices or indicates a better one. (Lahaie and Parkes, 2004) also presented an elicitation algorithm that guarantees elicitation in a polynomial number of value and demand queries (polynomial in the number of items, agents and the size of the agent's valuation functions). Later, (Lahaie et al., 2005) extended the work of (Lahaie and

Parkes, 2004) by providing a preference elicitation scheme for a wide class of the so-called *atomic* languages (i.e., languages in which a bid is expressed through a set of items and the corresponding price) such as OR and XOR ones. They also addressed the problem of incentive compatibility by introducing a new type of demand query, namely the *universal demand query*. They reported that using this new type of query together with value and traditional demand queries in the preference elicitation scheme yields to universal competitive equilibrium prices and thus to VCG outcomes (Mishra and Parkes, 2005). In a universal demand query, the elicitor presents to a bidder i a set of prices on all bundles and a set of n bundles (where n denote the number of bidders). The bidder agrees if every bundle is a best-response at the proposed prices or disagrees by pointing out the "bad" bundle and proposing a better one. We refer the reader to (Sandholm and Boutilier, 2006)'s survey for more details on preference elicitation in combinatorial auctions.

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. Some excellent surveys have recently discussed combinatorial auction design in general (Pekeč and Rothkopf, 2003; de Vries and Vohra, 2003), or highlighted a more specific issue (Xia et al., 2004; **?**; **?**; **?**). Nevertheless, some other issues, such as advising participants in combinatorial auctions, still need more attention.

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.

Acknowledgment 4 *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 FQRNT research development fund of the Province of Québec.*

References

Abrache, J., Bourbeau, B., Crainic, T., and Gendreau, M. (2004a). A New Bidding Framework for Combinatorial E-Auctions. *Computers & Operations Research*, 31:1177–1203.

Abrache, J., Crainic, T., and Gendreau, M. (2003). Decomposition Methods and Iterative Combinatorial Auctions. Publication CRT-2003-04, Centre de recherche sur les transports, Université de Montréal, Canada.

Abrache, J., Crainic, T., and Gendreau, M. (2004b). Design Issues for Combinatorial Auctions. *4OR*, 2:1–33.

Abrache, J., Crainic, T., and Gendreau, M. (2005). Models for Bundle Trading in Financial Markets. *European Journal of Operational Research*, 160:88–105.

Ahuja, R., Magnanti, T., and Orlin, J. (1993). *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall.

Andersson, A., Tenhunen, M., and Ygge, F. (2000). Integer Programming for Combinatorial Auction Winner Determination. In *Proceedings of the 4th International Conference on Multi-Agent Systems (ICMAS-00)*, pages 39–46.

Arrow, K. and Debreu, G. (1954). Existence of Equilibrium for a Competitive Economy. *Econometrica*, 22:265–290.

Arrow, K. and Hahn, F. (1971). *General Competitive Analysis*. Holden-Day, San Francisco.

Ausubel, L. (2006). An Efficient Dynamic Auction for Heterogeneous Commodities. *The American Economic Review*, 96(3).

Ausubel, L. and Cramton, P. (2004). Auctioning Many Divisible Goods. *Journal of the European Economic Association*, 2(2-3):480–493.

Ausubel, L., Cramton, P., and Milgrom, P. (2006). The Clock-Proxy Auction: A Practical Combinatorial Auction Design. In P. Cramton, Y. Shoham, R. S., editor, *Combinatorial Auctions*, chapter 5. MIT Press.

Ausubel, L. and Milgrom, P. (2002). Ascending Auctions with Package Bidding. *Frontiers of Theoretical Economics*, 1:1–45.

Ausubel, L. and Milgrom, P. (2006). The lovely but Lonely Vickrey Auction. In *Combinatorial Auctions*, chapter 1, pages 17–40. MIT Press.

Ball, M., Donohue, G., and Hoffman, K. (2006). Auctions for the Safe, Efficient and Equitable allocation of National Airspace Systems Resources. In P. Cramton, Y. Shoham, R. S., editor, *Combinatorial Auctions*, chapter 22. MIT Press.

Banks, J., Ledyard, J., and Porter, D. (1989). Allocating uncertain and unresponsive resources: An experimental approach. *The RAND Journal of Economics*, 20(1):1–25.

Benyoucef, M., Alj, H., Vézeau, M., and Keller, R. (2001). Combined Negotiations in E-Commerce: Concepts and Architecture. *Electronic Commerce Research*, 1(3):277–299.

Beraldi, P., Conforti, D., Triki, C., and Violi, A. (2004). A New Algorithm for the Assignment Problem. *4OR*, 2(1):35–51.

Bikhchandani, S., De Vries, S., Schummer, J., and Vohra, R. (2001). Linear Programming and Vickrey Auctions. In *Mathematics of the Internet: E-Auctions and Markets*. Springer-Verlag, Berlin.

Bikhchandani, S. and Mamer, J. (1997). Competitive Equilibrium in an Exchange Economy with Indivisibilities. *Journal of Economic Theory*, 74(2):385–413.

Bikhchandani, S. and Ostroy, J. (2002). The Package Assignment Model. *Journal of Economic Theory*, 107(2):377–406.

Blum, A., jackson, J., Sandholm, T., and Zinkevich, M. (2004). Preference Elicitation and Query Learning. *Journal of Machine Learning Research (JMLR)*, 5:649–667.

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

Boutilier, C. and Hoos, H. (2001). Bidding Languages for Combinatorial Auctions. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, pages 1211–1217.

Caplice, C. and Sheffi, Y. (2006). Combinatorial Auctions for Truckload Transportation. In *Combinatorial Auctions*, pages 539–571. MIT Press.

Cavallo, R., Parkes, D., Juda, A., Kirsh, A., Kulesza, A., Lahaie, S., Lubin, B., Michael, L., and Shneidman, J. (2005). TBBL: A Tree-Based Bidding Language for Iterative Combinatorial Exchanges. In *In The Multidisciplinary Workshop on Advance in Preference Handling (IJCAI)*.

Chandra, B. and Halldórsson, M. (2001). Greedy local improvement and weighted set packing approximation. *Journal of Algorithms*, 39(2):223–240.

Chang, T.-S., Crainic, T., and Gendreau, M. (2002). 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.

Clarke, E. (1971). Multipart Pricing of Public Goods. *Public Choice*, 11:17–33.

Conen, W. and Sandholm, T. (2001). Preference Elicitation in Combinatorial Auctions: Extended Abstract. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 256–259. Tampa, FL.

Cramton, P. (1998). Ascending Auctions. *European Economic Review*, 42(3-5):745–756.

Das, R., Hanson, J., Kephart, J., and Tesauro, G. (2001). Agent-Human Interactions in the Continuous Double Auction. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4–10.

Davenport, A. and Kalagnanam, J. (2001). Price Negotiations for Procurement of Direct Inputs. In *Mathematics of the Internet: E-Auctions and Markets*. Springer-Verlag, Berlin.

Day, R. and Raghavan, S. (2006). Matrix Bidding in Combinatorial Auctions. Working paper.

de Vries, S., Schummer, J., and Vohra, S. (2005). On Ascending Vickrey Auctions for Heterogeneous Objects. *Journal of Economic Theory*. Forthcoming.

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

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

EngelBrecht-Wiggans, R. (1980). Auctions and Bidding Models: A Survey. *Management Science*, 26:119–142.

Epstein, R., Henrìquez, L., Catalán, J., Weintraub, G., and Martinez, C. (2002). A Combinatorial Auction Improves School Meals in Chile. *Interfaces*, 32(6):1–14.

Eso, M. (2001). An Iterative Auction for Online Seats. In *Mathematics of the Internet: E-Auctions and Markets*. Springer-Verlag, Berlin.

Fan, M., Srinivasan, S., Stallaert, J., and Whinston, A. (2002). *Electronic Commerce and the Revolution in Financial Markets.* South-western Thomson Learning.

Fan, M., Stallaert, J., and Whinston, A. (2000). The Internet and the Future of Financial Markets. *Communications of the ACM*, 43(11):83–88.

Figliozzi, M., Mahmassani, H., and Jaillet, P. (2002). 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.

Fujushima, Y., Leyton-Brown, K., and Shoham, Y. (1999). Taming the Computational Complexity of Combinatorial Auctions. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, pages 548–553.

Gonen, R. and Lehmann, D. (2000). 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.

Graves, R., Schrage, L., and Sankaran, J. (1993). An Auction Method for course registration. *Interfaces*, 23(5):81–92.

Groves, T. (1973). Incentives in Teams. *Econometrica*, 41:617–631.

Gul, F. and Stacchetti, E. (1999). Walrasian Equilibrium with Gross Substitutes. *Journal of Economic Theory*, 87:95–124.

Gul, F. and Stacchetti, E. (2000). The English Auction with Differentiated Commodities. *Journal of Economic Theory*, 92:66–95.

Günlük, O., Lászlo, L., and Devries, S. (2005). A Branch-and-Price Algorithm and New Test Problems for Spectrum Auctions. *Management Science*, 51(3):391–406.

Hoffman, K., Menon, D., and van den Heever, S. (2004). A Bidder Aid Tool for Dynamic Package Creation in the FCC Spectrum Auctions. Technical report, Departement of Systems Engineering and Operations Research, George Masson University. Working paper.

Hohner, G., Rich, J., Ng, E., Reid, G., Davenport, A., Kalagnanam, J., Lee, H., and An, C. (2003). Combinatorial and Quantity-Discount Procurement Auctions Benefit Mars, Incorporated and Its Suppliers. *Interfaces*, 33(1):23–35.

Hoos, H. and Boutilier, C. (2000). Solving Combinatorial Auctions Using Stochastic Local Search. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 22–29.

Hudson, B. and Sandholm, T. (2002). Effectiveness of Preference Elicitation in Combinatorial Auction. In *Agent-Mediated Electronic Commerce (AMEC) workshop at AAMAS-02*. Bologna, Italy.

Kalagnanam, J., Davenport, A., and Lee, H. (2001). Computational Aspects of Clearing Continuous Call Double Auctions with Assignment Constraints and Indivisible Demand. *Electronic Commerce Research Journal*, 1(3):221–237.

Kelly, F. and Steinberg, R. (2000). A Combinatorial Auction with Multiple Winners for Universal Service. *Management Science*, 46(4):586–596.

Klemperer, P. (1999). Auction theory: A guide to the literature. *Journal of Economic Surveys*, 13:227–286.

Kothari, A., Sandholm, S., and Suri, S. (2002). Solving Combinatorial Exchanges: Optimality via a Few Partial Bids. In *AAAI-02 workshop on AI for Intelligent Business*. Edmonton, Canada.

Kutanoglu, E. and Wu, D. (1999). On Combinatorial Auction and Lagrangian Relaxation for Distributed Resource Scheduling. *IEEE Transactions*, 31:813–826.

Kwasnica, A., Ledyard, J., Porter, D., and DeMartini, C. (2005). A New and Improved Design for Multiobject Iterative Auctions. *Management Science*, 51(3):419–434.

Kwon, R., Anandalingam, G., and Ungar, L. (2005). Iterative Combinatorial Auctions with Bidder-Determined Combinations. *Management Science*, 51(3):407–418.

Lahaie, S., Constantin, F., and Parkes, D. (2005). More on the Power of Demand Queries in Combinatorial Auctions: Learning Atomic Languages and Handling Incentives. In *Proc. 19th International Joint Conference on Artificial Intelligence (IJCAI'05)*.

Lahaie, S. and Parkes, D. (2004). Applying Learning Algorithms to Preference Elicitation. In *Proc. of the ACM Conference on Electronic Commerce*, pages 180–188.

Ledyard, J., Olson, M., Porter, D., Swanson, J., and Torma, D. (2002). The First Use of a Combined Value Auction for Transportation Services. *Interfaces*, 32(5):2–14.

Lee, C.-G., Kwon, R., and Ma, Z. (2006). A Carrier's Optimal Bid Generation Problem in Combinatorial Auctions for Transportation Procurement. *Transportation Research Part E*. Forthcoming.

Lehmann, D. and Gonen, R. (2001). Linear programming helps solving large multi-unit combinatorial auctions. In *Proceedings of the Electronic Market Design Workshop*. Maastricht, The Netherlands.

Leonard, H. (1983). Elicitation of Honest Preferences for the Assignment of Individuals to Positions. *Journal of Political Economy*, 91(3):461–479.

Leyton-Brown, K. (2003). *Resource Allocation in Competitive Multiagent Systems*. Doctoral dissertation, Standford University.

Leyton-Brown, K., Shoham, Y., and Tennenholz, M. (2000). An Algorithm for Multi-unit Combinatorial Auctions. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 56–61.

Mansini, R. and Speranza, M. (2002). 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.

Martello, S. and Toth, P. (1997). *Knapsack Problems: Algorithms and Computer Implementation*. Wiley & Sons.

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

McMillan, J. (1994). Selling Spectrum Rights. *Journal of Economic Perspectives*, 8:145–162.

Mishra, D. and Parkes, D. (2005). Ascending Price Vickrey Auctions for General Valuations. *Journal of Economic Theory*. Forthcoming.

Müller, R. (2006). Tractable Cases of the Winner Determination Problem. In P. Cramton, Y. Shoham, R. S., editor, *Combinatorial Auctions*, chapter 13. MIT Press.

Myerson, R. (1981). Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73.

Nisan, N. (2000). Bidding and Allocation in Combinatorial Auctions. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, pages 1–12. ACM 2000, Minneapolis, MN.

Nisan, N. (2001). The Communication Complexity of Combinatorial Auctions. Draft, Institute of Computer Science, The Hebrew University of Jerusalem. Available at: http://www.cs.huji.ac.il/~noam/mkts.html.

Nisan, N. (2006). Bidding languages. In P. Cramton, Y. Shoham, R. S., editor, *Combinatorial Auctions*, chapter 1. MIT Press.

Nisan, N. and Segal, I. (2006). The Communication Requirement of Efficient allocations and supporting prices. *JOURNAL OF Economic Theory*, 129(1):192–224.

O'Neill, R., Stokiewicz, P., Hobbs, B., Rothkopf, M., and Stewart, W. J. (2005). Efficient Market-Clearing Prices in Markets with Nonconvexities. *European Journal of Operational Research*, 164:269–285.

Oren, M. and Williams, A. (1975). On Competitive Bidding. *Operations Research*, 23(6):1072–1079.

Park, S. and Rothkopf, M. (2005). Auctions with Bidder-Determined Allowable Combinations. *European Journal of Operational Research*, 161:399–415.

Parkes, D. (1999). iBundle: An Efficient Ascending Price Bundle Auction. In *Procedings of the 1st ACM Conference on Electronic Commerce*, pages 148–157. ACM Press, New York, USA.

Parkes, D. (2001). An Iterative Generalized Vickrey Auction: Strategy-proofness Without Complete Revelation. In Press, A., editor, *Proc. AAAI Spring Symposium on Game Theoretic and Decision Theoretic Agents*, pages 78–87.

Parkes, D. (2006). Iterative Combinatorial Auctions. In P. Cramton, Y. Shoham, R. S., editor, *Combinatorial Auctions*, chapter 2. MIT Press.

Parkes, D., Cavallo, R., Elprin, N., Juda, A., Lahaie, S., Lubin, B., Michael, L., Shneidman, J., and Sultan, H. (2005). ICE: An Iterative Combinatrial Exchange. In *Proceedings of ACM Conference on Electronic Commerce (EC'05)*.

Parkes, D. and Kalagnanam, J. (2005). Models for Iterative Multiattribute Procurement Auctions. *Management Science*, 51(3):435–451.

Parkes, D., Kalagnanam, J., and Eso, M. (2001). Achieving Budget-Balance with Vickrey-Based Payment Schemes in Exchanges. In *Proc. 17th International Joint Conference on Artificial Intelligence*, pages 1161–1168.

Parkes, D. and Ungar, L. (2000). Iterative Combinatorial Auctions: Theory and Practice. In *Proceedings of the 17th National Conference on Artificial Intelligence*, pages 74–81.

Parkes, D. and Ungar, L. (2002). An Ascending-price Generalized Vickrey Auction. Technical report, Harvard University.

Pekeč, A. and Rothkopf, M. (2003). Combinatorial Auction Design. *Management Science*, 49(11):1485–1503.

Rassenti, S., Smith, V., and Bulfin, R. (1982). A Combinatorial Auction Mechanism for Airport Time Slot Allocation. *Bell Journal Of Economics*, 13(2):402–417.

Rothkopf, M. and Park, S. (2001). An Elementary Introduction to Auctions. *Interfaces*, 31(6):83–97.

Rothkopf, M., Pekeč, A., and Harstad, R. (1998). Computationally manageable combinatorial auctions. *Management Science*, 44:1131–1147.

Rothkopf, M., Teisberg, T., and Kahn, E. (1990). Why are Vickrey Auctions Rare? *Journal of Political Economy*, 98:94–109.

Sakurai, Y., Yokoo, M., and Kamei, K. (2000). An Efficient Approximate Algorithm for Winner Determination in Combinatorial Auctions. In ACM, editor, *EC'00*, pages 30–37.

Sandholm, S. and Suri, S. (2001). 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.

Sandholm, T. (2002). Algorithm for Optimal Winner Determination in Combinatorial Auctions. *Artificial Intelligence*, 135:1–54.

Sandholm, T. and Boutilier, C. (2006). Preference Elicitation in Combinatorial Auctions. In *Combinatorial Auctions*, chapter 10, pages 233–264. MIT Press.

Sandholm, T. and Suri, S. (2003). BOB: Improved Winner Determination in Combinatorial Auctions and Generalizations. *Artificial Intelligence*, 145:33–58.

Sandholm, T., Suri, S., Gilpin, A., and Levine, D. (2001). CABOB: A Fast Optimal Algorithm for Combinatorial Auctions. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, pages 1102–1108.

Sandholm, T., Suri, S., Gilpin, A., and Levine, D. (2002). Winner determination in combinatorial auction generalizations. In *First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 69–76. Bologna, Italy.

Sandholm, T., Suri, S., Gilpin, A., and Levine, D. (2005). CABOB: A Fast Optimal Algorithm for Winner Determination in Combinatorial Auctions. *Management Science*, 51(3):374–390.

Santi, P., Conitzer, V., and Sandholm, T. (2004). Towards a Characterization of Polynomial Preferenec elicitation with Value Queries in Combinatorial Auctions. In *Conference on Computational Learning Theory (COLT)*, pages 1–16.

Song, J. and Regan, A. (2005). Approximation Algorithms for the Bid Construction Problem in Combinatorial Auctions for the Procurement of Fraight Transportation Contracts. *Transportation Research Part B*, 39:914–933.

Vickrey, W. (1961). Conterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16:8–37.

Walsh, W., Wellman, M., and Ygge, F. (2000). Combinatorial Auctions for Supply Chain Formation. In *Proc. ACM Conference on Electronic Commerce*, pages 260–269.

Wellman, M., Walsh, W., Wurman, P., and MacKie-Mason, J. (2001). Auction Protocols for Decentralized Scheduling. *Games and Economic Behavior*, 35:271–303.

Wurman, P. and Wellman, M. (1999a). AkBA: A Progressive Anonymous-price Combinatorial Auction. In *Second ACM Conference on Electronic Commerce*, pages 21–29.

Wurman, P. and Wellman, M. (1999b). Equilibrium Prices in Bundle Auctions. In *AAAI-99 Workshop on Artificial Intelligence for Electronic Commerce*, pages 56–61.

Wurman, P., Wellman, M., and Walsh, W. (2001). A Parametrization of the Auction Design Space. *Games and Economic Behavior*, 35:304–338.

Xia, M., Koehler, G., and Whinston, A. (2004). Pricing Combinatorial Auctions. *European Journal of Operational Research*, 154:251–270.

Zinkevich, M., Blum, A., and Sandholm, T. (2003). On polynomial-time preference elicitation with value queries. In *ACM Conference on Electronic Commerce*, pages 176–185. ACM Press.

Zürel, E. and Nisan, N. (2001). An efficient approximate allocation algorithm for combinatorial auctions. In *ACM Conference on Electronic Commerce*, pages 125–136. ACM Press.