

Truth Revelation in Approximately Efficient Combinatorial Auctions

DANIEL LEHMANN

Hebrew University, Jerusalem, Israel

AND

LIADAN ITA O'CALLAGHAN AND YOAV SHOHAM

Stanford University, Stanford, California

Abstract. Some important classical mechanisms considered in Microeconomics and Game Theory require the solution of a difficult optimization problem. This is true of mechanisms for combinatorial auctions, which have in recent years assumed practical importance, and in particular of the gold standard for combinatorial auctions, the Generalized Vickrey Auction (GVA). Traditional analysis of these mechanisms—in particular, their truth revelation properties—assumes that the optimization problems are solved precisely. In reality, these optimization problems can usually be solved only in an approximate fashion. We investigate the impact on such mechanisms of replacing exact solutions by approximate ones. Specifically, we look at a particular greedy optimization method. We show that the GVA payment scheme does not provide for a truth revealing mechanism. We introduce another scheme that does guarantee truthfulness for a restricted class of players. We demonstrate the latter property by identifying natural properties for combinatorial auctions and showing that, for our restricted class of players, they imply that truthful strategies are dominant. Those properties have applicability beyond the specific auction studied.

Categories and Subject Descriptors: F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems; I.2.m [**Artificial Intelligence**]: Miscellaneous; J.4 [**Social and Behavioral Sciences**]: *economics*

General Terms: Algorithms, Design, Theory

Additional Key Words and Phrases: Combinatorial auctions, computational complexity, mechanism design

This work was partially supported by the Jean and Helene Alfassa Fund for Research in Artificial Intelligence.

Authors' addresses: D. Lehmann, School of Engineering and Computer Science, Hebrew University, Jerusalem, 91904 Israel, e-mail: lehmann@cs.huji.ac.il; L. I. O'Callaghan and Y. Shoham, Computer Science Department, 353 Serra Mall, Stanford University, Stanford, CA 94305-9025, e-mail: {loc;shoham}@cs.stanford.edu.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@ansi.org.

© 2002 ACM 0004-5411/02/0900-0001 \$5.00

1. Introduction

This article concerns combinatorial auctions (also called combinatorial), that is, auctions in which multiple goods are available and in which bidders can post bids for subsets (i.e., bundles) of the goods. Such auctions have become the object of increased interest recently, in part because of the general interest in auctions, and in part because of specific auctions in which combinatorial bidding would seem natural, such as the series of the FCC spectrum auctions [McMillan 1994; Cramton 1997; Milgrom 2000].¹

Combinatorial auctions (henceforth CAs) typically require the solution of one or more difficult optimization problems. The computational complexity of these problems threatens to render the traditional auction designs a mere theoretical construct. One approach to meeting this threat is to replace the exact optimization by an approximate one. This, however, gives rise to a new challenge: traditional analysis of established CA mechanisms relies strongly on the fact that the goods are allocated in an optimal manner, and the properties guaranteed by the mechanism (such as truthful bidding, to be defined later), disappear if the allocation is anything less than optimal. This is true in particular of the Generalized Vickrey Auction (GVA), also defined later, which is widely taken to be the gold standard for CAs. The primary focus of this article is to present a simple approximate optimization method for CAs that possesses two attractive properties:

- the method performs a reasonably effective optimization, and
- there exists a novel payment scheme that, when coupled with the approximate optimization method, makes for a combinatorial auction in which truth-telling is a dominant strategy.

In order to show the latter property we identify several axioms which are sufficient to ensure truth-telling for a restricted class of players, in any combinatorial auction; these axioms are interesting in their own right, as they can be applied to auctions other than the one discussed here. They were, in fact, shown to ensure truth-telling for any *one parameter* agents in Archer and Tardos [2001].

Note: Since we aim to make this article easily accessible to both computer scientists and game theorists, we include some rather basic material.

2. A Brief Introduction to Combinatorial Auctions

In this section, we briefly cover the notions of complementarity and substitutability, as motivating CAs; the two degrees of freedom in a sealed-bid CA, namely allocation and payment policies; and why one needs to be careful when applying the desiderata of efficiency and revenue maximization to CAs.

2.1. COMPLEMENTARITY AND SUBSTITUTABILITY. Throughout this article, we shall consider single-side CAs with a single seller and multiple buyers. The reverse situation with a single buyer and multiple sellers is symmetric; the two-sided case, with multiple buyers and sellers, is more complex. Let us assume, then, that an auctioneer is selling a number of different goods. In such a situation, a bidder may be willing to pay more for the whole than the sum of what he is willing to pay for

¹ Up until now the FCC auctions have not in fact been combinatorial, due in part to the complexity problem discussed below. However, the FCC is currently actively considering a combinatorial auction.

the parts: this is the case if the parts complement each other well, for example, a left shoe and a right shoe. This phenomenon is called *complementarity*. In other cases, a bidder may be willing to pay for the whole only less than the sum of what he is willing to pay for the parts, maybe only as much as one of the parts. This is especially the case if the bidder has a limited budget or if the goods are similar, or interchangeable, for example, two tickets to the same performance. This phenomenon is called *substitutability*. In general, complementarity and substitutability can both play heavily in the same auction.

In the absence of complementarity and substitutability, that is, if every participant values a set of goods at the sum of the values of its elements, one should organize the multiple auction as a set of independent simple auctions, but, in the presence of complementarity, organizing the multiple auction as a set or even a sequence of simple auctions will lead to less than optimal results: for example, a participant ending up with a left shoe and another one with the right shoe, or the left shoe auctioned for almost nothing because bidders fear not to be able to get the right shoe and the right shoe then auctioned for nothing to the buyer of the left shoe since no one is interested in just a right shoe. The problem is particularly acute when the complementarity and substitutability relations vary among the various bidders.

2.2. SPECIFYING A COMBINATORIAL AUCTION. Several auction designs have been proposed to deal with complementarity and substitutability. For example, the Simultaneous Ascending Auction was devised in connection with the FCC Spectrum Auction mentioned above. In this article, we shall consider only what is perhaps the most obvious approach, which is to allow combinatorial bidding. For the history of combinatorial auctions, see Rothkopf [1983]. What does it take to specify a CA? In general, any auction must specify three elements: the bidding rules (i.e., what one is allowed to bid for and when), the market clearing rules (i.e., when is it decided who bought what and who pays what), and the information disclosure rules (i.e., what information about the bid state is disclosed to whom and when).

We consider only one-stage, sealed-bid CAs; in these, each bidder submits zero or more bids, the auction clears, and the results are announced. The third element of the specification is thus straightforward: no information is released about other bidders' bids prior to the close of the auction.

The first element of the specification is almost as straightforward: each bidder may submit one or more bids, each of which mentions a subset of the goods and a price. One has to be precise, however, about the semantics of the collection of bids submitted by a single bidder; if I bid \$5 for a and \$7 for b , what does it mean about my willingness to pay for $\{a, b\}$? If I bid \$10 for $\{a, b\}$ and \$20 for $\{b, c\}$, what does it mean about my willingness to pay for $\{a, b, c\}$? This is not a mysterious issue, but one needs to be precise about it. We shall return to this issue later when we discuss the notion of a bidder's *type*.

The scheme above allows one to express complementarity. Bidding for \$5 for a , \$7 for b and \$15 for $\{a, b\}$ clearly indicates complementarity. On the face of it, though, substitutability cannot be expressed, since bidding \$8 for $\{a, b\}$, \$5 for a and \$7 for b does not preclude, under the usual market clearing rules, one being allocated a and b separately. However, a simple encoding trick presented in Fujishima et al. [1999] allows the expression of substitutability, at least partially. More elaborate proposals for bidding languages may be found in Sandholm [2000], Nisan [2000], and Boutilier and Hoos [2001].

Thus far, the designer of a combinatorial auction has no discretion. Only the second element of specification, the clearing policy, provides choices. There are two choices to be made here: which goods does every bidder receive, and how much does every bidder pay? We address these below.

2.3. MAXIMIZING EFFICIENCY AND REVENUE. The standard yardsticks for auction design, which are sometimes at odds with one another, are guaranteeing efficiency and maximizing (in our case, the seller's) revenue. We shall be concentrating primarily on efficiency in this article, but a very preliminary study of revenue is found in Section 13. Efficiency means that the allocation (of goods and money) resulting from the auction is Pareto optimal: no further trade among the buyers can improve the situation of some trader without hurting any of them. This is typically achieved by ensuring that the clearing rules maximize the sum of the values the various bidders place on the actual allocation decided on by the auctioneer. On the whole, one can expect that an efficient auction, after which the participants are globally satisfied, allows the seller to extract a higher revenue than an inefficient auction after which the level of social satisfaction is lesser. Efficiency, therefore, which may be a goal in itself, may also be a step in the direction of revenue maximization. In fact, this correlation holds only in part, and auctions that are maximizing revenue are not always efficient [Myerson 1981]. Nevertheless, we shall seek efficient, at least approximately, auction mechanisms.

Note four problems here. We have already mentioned that bidders specify bids, not their profile of preferences over bundles. This does not pose a real challenge, so long as one is clear about the meaning of those bids. The second one is that those profiles of preferences over bundles do not allow for a full specification of preferences about the outcomes of the auction, that is, the resulting allocation. A bidder cannot express *externalities*, for example, that he would prefer, if he does not get a specific good, this good to be allocated to bidder *X* and not to bidder *Y*. Third, we have an optimization problem on our hands; as it turns out, it is an NP-hard optimization problem that cannot be even approximated in a feasible way, in the worst case. This means that, for all practical purposes, there does not exist a polynomial-time algorithm for computing the optimal allocation, or even for computing an allocation that is guaranteed to be off from optimal by at most a constant, any given constant. The fourth and deepest problem is that the optimization is supposed to happen over the bidder's true valuations, as opposed to merely their bid amounts, but that information is not available to the auctioneer and the bidder will reveal this information only if it is in his/her best interest.

An ingenious method, discussed in the next section, has been developed in game theory to overcome the fourth problem. The problem is that not only does it not address the second problem, it actually mildly exacerbates it by requiring that the optimization be performed once for each bidder. The primary goal of this article is to devise a method that promises good (albeit suboptimal) efficiency, while being computationally feasible. In a nutshell, the goal is to simultaneously ensure economic and computational efficiency.

3. Mechanism Design for CA

In this section, we consider the design of combinatorial auctions as a problem of designing a game of incomplete information for which the weakly dominant

strategies present a *good* way of allocating the goods and paying for them. The general setting is that of economic mechanism design: see Mas-Colell et al. [1995, Chap. 23], for example, for an introduction to the field and Varian [1995] for a description of auctions in this framework. Contrary to the latter, we shall restrict our description to combinatorial auctions in which no externalities can be expressed. Informally, each bidder sends a message describing (truthfully or not) his preferences. The auctioneer, then, computes the resulting allocation of the goods and the payments, based on the bidders' messages but according to rules known in advance. The mechanism is a *truthful* one if it is in the best interest of the bidders to send messages that truthfully reveal their preferences.

Formally, we consider a set P of n bidders. The indices $i, j, 1 \leq i, j \leq n$, will range over the bidders. Bidders are selfish, but rational, and trying to maximize their utility in the final outcome. A bidder knows his own utility function (i.e., his *type*), but this information is private and neither the auctioneer nor the other players have access to it. The final result of an auction consists of two elements: an allocation of the goods and a vector of payments from the bidders to the auctioneer, both of which are functions of the bidders' declarations (i.e., bids). Formally, we have a finite set G of k goods and an allocation is a *partial* function from G to P , that is, a function $a : G \rightarrow P'$, with $P' = P \cup \{\text{unallocated}\}$, since we do not insist that all goods be allocated. Notice that the allocations produced by the Generalized Vickrey Auctions of Section 4 and by our Greedy algorithm of Section 7 are not always total. The set of outcomes (i.e., allocations) is $\mathcal{O} = P'^G$, the set of partial functions from G to P . Since we do not allow for externalities, the set Θ_i of the possible types for bidder i is $\mathbf{R}_+^{2^G}$, where \mathbf{R}_+ is the set of all nonnegative real numbers. Notice that such a type allows for both complementarity and substitutability, but not for externalities. Since the set Θ_i does not depend on i , we shall write Θ . An element of Θ assigns a real nonnegative valuation to every possible bundle. The set Θ is also the set of messages that bidder i may send. A bidder may send any element of Θ , irrespective of his (true) type, that is, a bidder may lie. We shall typically use t to denote a (true) type, d to denote a message, T or D to denote vectors of n types and P for a payment vector, that is, a vector of n nonnegative numbers.

Since we assume the Independent Value Model and Quasi-Linear utilities, fairly standard assumptions in the field, the utility, for a bidder of type t , of bundle $s \subseteq G$ and payment x is:

$$u = t(s) - x. \quad (1)$$

Definition 3.1. A (*direct mechanism for combinatorial auctions*) consists of

- an allocation algorithm f that picks, for each vector D (D is a vector of declared types), an allocation $f(D)$,
- a payment scheme p that determines, for each vector D a payment vector $p(D)$: $p_i(D)$ is paid by bidder i to the auctioneer.

Let us denote the bundle obtained by i as:

$$g_i(D) = f(D)^{-1}(i). \quad (2)$$

Notation. In general, g_i depends on the allocation algorithm f , but when f is clear from the context, we shall abuse the notation and treat g_i as a direct function of the bid vector, D . Equation (1) implies that if bidder i has (true) type t , his utility

from the mechanism is:

$$u_i = t(g_i(D)) - p_i(D), \quad (3)$$

where $D = \langle d_1, \dots, d_n \rangle$ is the vector of declarations.

The first term of this sum, $t(g_i(D))$ is often called the valuation of i : $v_i(f(D), t)$. The game-theoretic solution concept used throughout this article is that of a weakly dominant strategy, that is, a strategy that is, as good as any other for a given player, no matter what other players do. This is in contrast with the weaker and more common notion of Nash equilibria. The particular property we would like to ensure for our mechanism is that the dominant strategy for each player is to bid his true valuation; in other words, no bidder can be better off by lying, no matter how other bidders behave. A mechanism is truthful if no bidder can be better off by lying, even if other bidders lie. This is a very strong requirement, making for a very sturdy mechanism.

Definition 3.2. A mechanism $\langle f, p \rangle$ is *truthful* if and only if for every $i \in P$, $t \in \Theta$ and any vector D of declarations, if D' is the vector obtained from D by replacing the i th coordinate d_i by t , then: $t(g_i(D')) - p_i(D') \geq t(g_i(D)) - p_i(D)$.

In the definition above, t is the true type of bidder i and D is a vector of declared types. The term, $t(g_i(D))$, represents the true satisfaction i receives from the allocation resulting from declarations D and $t(g_i(D'))$ represents his true satisfaction from the allocation that would have been obtained had i been truthful.

4. The Generalized Vickrey Auction

A very general method for designing truthful mechanisms has been devised by Clarke [1971] and Groves [1973]. Applied to combinatorial auctions, it generalizes the second price auctions of Vickrey [1961]. We shall now describe those generalized Vickrey auctions, prove that the mechanism described is truthful and then discuss the complexity issues that render those auctions infeasible when k , the number of goods, is large. Generalized Vickrey Auctions (GVAs) appear to be part of the folklore of mechanism design. A description of a more general type may be found in MacKie-Mason and Varian [1994] and Varian [1995]; we adopt a special case of it, one which does not allow for externalities.

In a GVA, the allocation chosen maximizes the sum of the declared valuations of the bidders, each bidder receives a monetary amount that equals the sum of the declared valuations of all other bidders, and pays the auctioneer the sum of such valuations that would have been obtained if he had not participated in the auction. A way to describe such an auction, in which i does not participate, is to consider the auction in which bidder i declares a zero valuation for all possible bundles. A bidder with zero valuation for all bundles has no influence on the outcome.

Formally, given a vector D of declarations, the generalized Vickrey auction defines the allocation and payment policies as follows (notice that $a^{-1}(i)$ is the bundle allocated to i by allocation a , and that g_i is defined in Eq. (2)):

$$f(D) = \operatorname{argmax}_{a \in \mathcal{O}} \sum_{i=1}^n d_i(a^{-1}(i)), \quad (4)$$

$$p_j(D) = - \sum_{i=1, i \neq j}^n d_i(g_i(D)) + \sum_{i=1, i \neq j}^n d_i(g_i(Z)), \quad (5)$$

where $Z_i = D_i$ for any $i \neq j$ and $Z_j(s) = 0$ for any bundle $s \subseteq G$. Since $d_j(g_j(Z)) = 0$, we may as well have written:

$$p_j(D) = - \sum_{i=1, i \neq j}^n d_i(g_i(D)) + \sum_{i=1}^n d_i(g_i(Z)). \quad (6)$$

A proof of the truthfulness of the Clarke–Groves–Vickrey mechanism may be found, for example, in Mas-Colell et al. [1995, Proposition 23.C.4]. We include the proof here only to stress how easy it is.

THEOREM 4.1. *The generalized Vickrey auction is a truthful mechanism.*

PROOF. Assume $j \in P$, $t \in \Theta$, D is a vector of declarations, and $D'_i = D_i$ for any $i \neq j$ and $D'_j = t$. By Eq. (4),

$$\sum_{i=1}^n d'_i(g_i(D')) \geq \sum_{i=1}^n d'_i(g_i(D)).$$

But, for $E = D, D'$, we have:

$$d'_i(g_i(E)) = d_i(g_i(E)), \quad \text{if } i \neq j \text{ and } d'_j(g_j(E)) = t(g_j(E)).$$

Therefore,

$$t(g_j(D')) - p_j(D') + \sum_{i=1, i \neq j}^n d_i(g_i(Z)) \geq t(g_j(D)) - p_j(D) + \sum_{i=1, i \neq j}^n d_i(g_i(Z))$$

and $t(g_j(D')) - p_j(D') \geq t(g_j(D)) - p_j(D)$. \square

Notice that the second term in the payment of j does not depend on j 's declaration and is therefore irrelevant to his decision on what to declare. A feature of the GVA is that no truthful bidder's utility can be negative.

PROPOSITION 4.2. *If j is truthful, his utility u_j in the GVA is nonnegative.*

PROOF. By Eq. (3), since j is truthful, by Eq. (6), and finally by Eq. (4):

$$\begin{aligned} u_j &= d_j(g_j(D)) + \sum_{i=1, i \neq j}^n d_i(g_i(D)) - \sum_{i=1}^n d_i(g_i(Z)) \\ &= \sum_{i=1}^n d_i(g_i(D)) - \sum_{i=1}^n d_i(g_i(Z)) \geq 0. \end{aligned} \quad \square$$

Since bidders truthfully declare their type and the allocation maximizes the sum of the declared utilities, in a GVA, the allocation maximizes the sum of the true valuations of the bidders, that is, the social welfare. In a quasilinear setting, this is equivalent to Pareto optimality. Therefore, a GVA is Pareto optimal. The mechanism to be presented in Section 10 only approximately maximizes the sum of the true valuations of the bidders, and is not Pareto optimal.

As we discuss in the following sections, it is known that algorithmic complexity considerations imply that Pareto optimality cannot be feasibly attained. Specifically, ensuring Pareto efficiency requires solving an intractable optimization problem. This is true even if we restrict the class of bidders severely, as we propose in Section 5.

5. Single-Minded Bidders

As is customary, we shall consider that any algorithm whose running-time is polynomial in $k + n$ is feasible, but any algorithm whose running-time is not polynomial in $k + n$ is infeasible. The size of the set \mathcal{O} of allocations is exponential in k , if there are at least two bidders, and the set Θ of possible types is doubly exponential in k . Since, in a direct mechanism (we consider no others), the message that a bidder sends describes one specific element (type) of Θ , a bidder needs an exponential number of bits to describe his type: the length of the messages sent in any such mechanism, and in a generalized Vickrey auction, is exponential in k , therefore infeasible. The design of a feasible version of the GVA must begin, therefore, by reducing the set of possible types to some set of singly exponential size. Most implementations of auctions assume that the bidders express their preferences by a small set of *bids*. Other restrictions about the types of the bidders that facilitate their description have been considered in Lehmann et al. [2001]. We shall start with a most sweeping restriction. In Section 11, we shall consider relaxing this restriction and an impossibility result will be presented in Section 12.

We shall assume that bidders are single-minded and care only about one specific (bidder-dependent) set of goods. If they do not get this set they value the outcome at the lowest possible value 0. In other words, our bidders are restricted to one single bid.

Definition 5.1. Bidder i is *single-minded* if and only if there is a set $s \subset G$ of goods and a value $v \in \mathbf{R}_+$ such that its type t can be described as: $t(s') = v$ if $s \subseteq s'$ and $t(s') = 0$ otherwise.

Note that single-minded bidders are not *one parameter* agents in the sense of Archer and Tardos [2001], since they have the freedom of deciding which bundle they bid for on top on the amount they are willing to pay. In Mu'alem and Nisan [2002], truthfulness is shown to be attainable for allocation algorithms providing better approximations if agents are assumed to be single-minded and *one parameter*. We shall denote by $\langle s, v \rangle$ the type just described. Note that a single-minded bidder enjoys *free disposal*. We shall assume, in most of this article, that all bidders are single-minded, that is, there are sets of goods s_i and nonnegative real numbers v_i such that bidder i is of type $\langle s_i, v_i \rangle$. We shall denote by Σ the set of all single-minded types. The size of the set Σ , contrary to the size of Θ , is singly exponential in k . A string of polynomial size will be enough to code the declarations of the bidders: it will describe a set of goods and a value. In this setting, we identify bids and bidders. Note that in a simple auction, that is, $k = 1$, assuming there are no externalities as we do, any agent is single-minded.

We shall design a feasible truthful mechanism for combinatorial auctions among single-minded bidders. At least anecdotal evidence suggests that single-mindedness is not an uncommon situation. Indeed, R. Wilson (private communication) reports

that, in the GVA used for selling timber harvesting rights in New Zealand, the bidders were almost single-minded: they were typically interested in all of the locations in a specific geographical area. There are situations in which one may realistically assume that bidders are single-minded. Consider, for example, an auction of pollution rights, such as the ones considered in DeMartini et al. [1999]. Some authority is selling rights to emit different pollutants during a fixed time-frame. Bidders are typically chemical plants whose production process cannot be changed at short notice and therefore bidders either buy all the rights corresponding to the chemicals emitted by their plant or have to close the plant for the given time-frame. The bidders are therefore single-minded and even fit the restrictions of Mu'alem and Nisan [2002]. If the plant owner has the possibility of operating its plant at less than full capacity, the situation is more complex. Other situations in which bidders are single-minded would be: the auction of communication links in a tree, or the auction of auto parts to buyers desiring a specific model. We are indebted to Noam Nisan for those last examples. Consider, indeed, communication links in a tree auctioned to bidders who want a path between two specific nodes. Since there is only one path between any pair of nodes, the bidders are single-minded. Or consider an auto manufacturer capable of assembling a number of different models out of its stock of auto parts. The buyers have set their mind on the model they want and are therefore single-minded: they need exactly all the parts included in their model.

It might also seem that this restriction does away with the computational issue; however, as we see in the next section, GVAs are infeasible even with the restriction to single-minded bidders. In Section 11, we shall discuss the generalization of our results to larger families of bidders.

6. Infeasibility of the GVA

Let us now assume that all bidders are single-minded, that is, the set of all possible types is now Σ . It follows easily from Proposition 4.2 that, in a GVA, a single-minded bidder of type $\langle s, v \rangle$ never pays more than v and pays nothing if he is not allocated the whole set s .

In a GVA, the allocation is the one defined in Eq. (4). Computing this allocation requires optimizing $\sum_{i=1}^n d_i(a)$ over all a 's in the set \mathcal{O} that is of exponential size. One may suspect that this is an infeasible task. Indeed, the problem of finding the allocation of Eq. (4) has been shown to be NP-hard in Rothkopf et al. [1998]. We remark that the restriction to single-minded bidders does nothing to alleviate the problem. Not only is it infeasible to solve exactly the optimization problem, but it turns out to be also infeasible to guarantee any nontrivial approximation to the optimal allocation. This follows easily from a celebrated result of Håstad [1999] and was, independently, noticed by Sandholm to whom priority is due: final version in Sandholm [2002].

THEOREM 6.1. *Let a single-minded type $d_i = \langle s_i, v_i \rangle$, $s_i \subseteq G$, $v_i \in \mathbf{R}_+$ be given for each bidder $i \in P$. Let $|G| = k$ and $|P| = n$. The problem of finding an allocation a that maximizes $\sum_{i=1}^n d_i(a)$ is NP-hard in $k + n$. Moreover, the existence of a polynomial time algorithm guaranteed to find an allocation whose value is at least $k^{-1/2+\epsilon}$ times the value of the optimal solution would imply that $NP = ZPP$.*

A short note on complexity classes: in the above, NP is the class of sets for which membership can be decided nondeterministically in polynomial time and ZPP is the subclass of NP consisting of those sets for which there is some constant c and a probabilistic Turing machine M that on input x runs in expected time $O(|x|^c)$ and outputs 1 if and only if $x \in L$. The question of whether $\text{NP} = \text{ZPP}$ is a deep open question in theoretical computer science, related with the famous $\text{P} = \text{NP}$ question. $\text{NP} = \text{ZPP}$ is not known to imply $\text{P} = \text{NP}$, but does imply $\text{NP} = \text{RP} = \text{co-RP} = \text{co-NP}$. $\text{P} = \text{NP}$ obviously implies $\text{NP} = \text{ZPP}$. RP is the class of sets for which membership can be decided in polynomial time by a randomizing algorithm. The class co-RP is the class of sets whose complements are in RP: nonmembership can be decided polynomially by a randomizing algorithm. Similarly for co-NP. ZPP is the intersection of RP and co-RP. (end of short note).

PROOF. The problem at hand may be described as the weighted version of the Set Packing problem of Karp [1972]. Karp shows that Set Packing is NP-hard by reducing the Clique problem to it. The k used in this reduction is of the order of n^2 . A direct reduction of Clique to our allocation problem is obtained in the following way. Given a graph G , let the goods be the edges and the bids be the vertices. Each vertex requests the edges it is adjacent to for a price of 1. An optimal allocation is a maximal independent vertex set. Håstad [1999] has shown that Clique cannot be approximated within $|V|^{1-\epsilon}$ unless $\text{NP} = \text{ZPP}$. The reduction mentioned above shows our claim. \square

Notwithstanding Theorem 6.1, if $n \leq \log k$, an optimal allocation may be found in time linear in k and if $k \leq \log n$, then dynamic programming provides an optimal allocation in time quadratic in n , as shown in Rothkopf et al. [1998].

Let us now consider the significance of Theorem 6.1. Even if agents declare their type truthfully, we cannot always attain an efficient allocation. This negative result holds even if agents are single-minded. Global restrictions on the structure of the set of bidders are considered in Rothkopf et al. [1998] and shown to allow a polynomial search for the efficient allocation. They severely restrict the possible types of the bidders to a subset of Σ , based on some inherent structure of G . Those restrictions are rarely met in practice.

If the number of goods is large, we may either find an algorithm that computes the efficient allocation but may, in the worst cases, never terminate (for all practical purposes) or settle for an algorithm that provides a subefficient allocation. Both ideas have been proposed in Fujishima et al. [1999] and Sandholm [1999]. Recent results on the impact of such an approximation on the quality of the mechanism, that is, its truthfulness may be found in Nisan and Ronen [2000] for combinatorial auctions and in Roughgarden and Tardos [2002] and Jain and Vazirani [2001] for other combinatorial optimization problems. A pioneering study of the properties of approximate mechanisms, but not for combinatorial auctions, may be found in Nisan [1999] and Nisan and Ronen [1999]. The recent [Bartal et al. 2002] provides a truthful approximate mechanism for a multi-unit combinatorial auction but is not applicable to the single-unit combinatorial auctions considered here. In Section 7, we shall provide a feasible approximation algorithm that appears to be very effective in practice and, in Section 10, we shall describe a payment scheme, different from the GVA's, that guarantees truthfulness. The payment scheme is carefully tailored to the specific approximation algorithm.

7. The Greedy Allocation

Since an efficient solution seems out of reach, we shall look for an approximately efficient solution. We shall propose a family of algorithms that provide such an approximation. Each of those algorithms runs in almost linear time in n , the number of single-minded bidders. One algorithm of the family guarantees an approximation ratio of $k^{-1/2}$.

A single-minded bidder declaring $\langle s, a \rangle$, with $s \subseteq G$ and $a \in \mathbf{R}_+$ will be said to put out a bid $b = \langle s, a \rangle$. We shall use $s(b)$ and $a(b)$ to denote the components of b and call $a(b)$ the *amount* of the bid b . As explained in Section 5, we identify bids and bidders. Two bids $b = \langle s, a \rangle$ and $b' = \langle s', a' \rangle$ conflict if $s \cap s' \neq \emptyset$.

The algorithms we consider execute in two phases.

- In the first phase, the bids are sorted by some criterion. The algorithms of the family are distinguished by the different criteria they use. Since there are n bids, this phase takes time of the order of $n \log n$. We assume a criterion, that is, a norm is defined and the bids are sorted in decreasing order following this norm. Since we shall have, in Theorem 10.2, to compare the sorted lists of bids of slightly different auctions, we also assume a consistent treatment of ties, that is, bids with equal norms. Formally, we shall assume that no two different bids have the same norm, that is, there are no ties.
- In the second phase, a greedy algorithm generates an allocation. Let L be the list of sorted bids obtained in the first phase. The first bid of L , say $b = \langle s, a \rangle$ is granted, that is, the set s will be allocated to b and then the algorithm examines each bid of L , in order, and grants it if it does not conflict with any of the bids previously granted. If it does, it denies (i.e., does not grant) the bid. This phase requires time linear in n .

The use of such a greedy scheme is very straightforward and speedy. We shall now discuss its efficiency: how efficient is the allocation generated? The efficiency of the allocation generated depends obviously both on the criterion used in the first phase and on the types of the bidders, or on the distribution with which the bidders are generated. It is clear that, to obtain allocations close to efficiency, one should use a norm that pushes bids that have a good chance to be part of an efficient allocation toward the beginning of the list L . The amount of a bid is a good criterion in this respect: we want bids with higher amounts to have a larger norm than bids with lower amounts, at least when the bids are for the same set of goods. Similarly, leaving the amount of a bid unchanged but making its bundle a smaller set (inclusion-wise), should also increase the norm. We shall require that changing s to s' with $s' \subset s$ or changing v to v' with $v' > v$ increase the norm of a bid. Let us call this property *bid-monotonicity*. This is the only requirement we shall make. Many criteria satisfy it.

In real-life situations, one can typically find a suitable natural norm related to the economic parameters of the bundle that measures the a-priori attractiveness of the bid (for the auctioneer). In the FCC auction, goods (licenses) are characterized by the population they cover. The (inverse of the) sum of those populations is a good indicator. In the abstract, if we know nothing concrete about the goods, our best bet is to use the size of the set of goods mentioned in a bid. We shall look in particular at the average-amount-per-good measure.

Definition 7.1. The average amount per good of a bid $b = \langle s, a \rangle$ is $\frac{a}{|s|}$.

Sorting the list L by descending average amount per good is a very reasonable idea. But many other possibilities may be considered. Sorting L by descending amounts for example, or, more generally sorting L by a criterion of the form $a/|s|^l$ for some number l , $l \geq 0$, possibly depending on k . All such criteria satisfy bid-monotonicity.

How good is the greedy allocation in comparison with the optimal one? For $l = 1$, the worst case may be analyzed without much difficulty. The ratio between the total value of the optimal allocation and that of the allocation found by the greedy algorithm cannot be larger than k , and this bound is tight. As usual in this sort of situations, on the average, on realistic distributions of bids, the performance of the greedy allocation scheme is much better than the lower bound above. We have been able to perform a full analysis of the worst case performance of those norms for different l 's and found out that $l = 1/2$ is best: it guarantees an approximation ratio of at least \sqrt{k} and, by Theorem 6.1, this is, up to a multiplicative constant, essentially, the best approximation ratio one can hope for a polynomial-time algorithm. The \sqrt{k} upper-bound improves on the previously best known result of Halldórsson [1999] by a factor of 2. Better results may now be found in Halldórsson [2000]. The following has since been generalized to multi-unit combinatorial auctions in Gonen and Lehmann [2000].

THEOREM 7.2. *The greedy allocation scheme with norm $a/|s|^{1/2}$ approximates the optimal allocation within a factor of \sqrt{k} .*

PROOF. Assume the bids (i.e., bidders) are $\langle s_i, a_i \rangle$ for $i = 1, \dots, n$. Let $w_i = |s_i|$. Our norm is: $r_i = a_i/\sqrt{w_i}$. Let OP be the optimal solution, that is, the set of bids contained in the optimal solution. The value of the optimal solution is $\alpha = \sum_{i \in OP} a_i$. Let GR be the solution obtained by the greedy allocation and β its value: $\beta = \sum_{i \in GR} a_i$. We want to show that:

$$\alpha \leq \beta \sqrt{k}. \quad (7)$$

Notice, first, that we may, without loss of generality, assume that the sets OP and GR have no bid in common. Indeed, if they have, one considers the problem in which the common bids and all the units they request have been removed. The greedy and optimal solutions of the new problem are similar to the old ones and the inequality for the new smaller problem implies the same for the original problem.

Let us consider β . By elementary algebraic considerations:

$$\beta = \sum_{i \in GR} a_i \geq \sqrt{\sum_{i \in GR} a_i^2} = \sqrt{\sum_{i \in GR} r_i^2 w_i}.$$

Consider α . By the Cauchy–Schwarz inequality:

$$\alpha = \sum_{i \in OP} r_i \sqrt{w_i} \leq \sqrt{\sum_{i \in OP} r_i^2} \sqrt{\sum_{i \in OP} w_i}.$$

The expression $\sum_{i \in OP} w_i$ represents the total number of goods allocated in the optimal allocation OP and is therefore bounded from above by k , the number of

goods available. We conclude that:

$$\alpha \leq \sqrt{\sum_{i \in OP} r_i^2} \sqrt{k}.$$

To prove (7), it will be enough, then, to prove that:

$$\sum_{i \in OP} r_i^2 \leq \sum_{i \in GR} r_i^2 w_i.$$

Consider the optimal solution OP . By assumption, the bids of OP did not enter the greedy solution GR . This means that, at the time any such bid i is considered during the execution of the greedy algorithm, it cannot be entered in the partial allocation already built. This implies that there is a good $l \in s_i$ that has already been allocated in the partial greedy solution, that is, there is a bid j in GR , with $r_j \geq r_i$ and $l \in s_j$.

A number of different bids from OP may, in this way, be associated with the same bid j of GR , but at most w_j different bids of OP may be associated with bid j of GR , since the sets of goods requested by two different bids of OP have an empty intersection. If OP_j is the set of bids of OP that are associated with bid j :

$$\sum_{i \in OP_j} r_i^2 \leq r_j^2 w_j. \quad \square$$

In other words, the greedy scheme does not guarantee any fixed ratio of approximation, but guarantees the best achievable ratio (assuming $NP \neq ZPP$). Experiments reported about in Gonen and Lehmann [2000] have confirmed that, on average for a specific distribution, the greedy algorithm using the norm of Theorem 7.2 performs extremely well, much better than the lower bound described in the theorem. More experiments are necessary to study the average case performance of different norms. In the sequel, all examples will use the average amount per good criterion but it is not difficult to find similar examples for any criterion of the form $a/|s|^l$.

Example 7.3. Assume there are two goods a and b and three bidders Red, Green and Blue. Red bids 10 for a , Green bids 19 for the set $\{a, b\}$ and Blue bids 8 for b . We sort the bids by decreasing average amount and obtain: Red's bid for a (average 10), Green's bid for $\{a, b\}$ (average 9.5) and Blue's bid for b (average 8). The greedy algorithm grants Red's bid for a , denies Green's bid for $\{a, b\}$ since it conflicts with Red's and grants Blue's bid for b . The allocation is not efficient. The efficient allocation grants Green's bid for $\{a, b\}$ and denies both other bids.

Our goal is to devise truthful mechanisms for combinatorial auctions among single-minded bidders. Given a suitable greedy allocation, can one find a payment scheme that makes the pair a truthful mechanism?

8. Greedy Allocation and Clarke's Payment Scheme Do Not Make a Truthful Mechanism, Even for Single-Minded Bidders

In Section 10, a mechanism based on the greedy allocation will be built and shown to be truthful if all bidders are single-minded. In this section, we show that the use of Clarke's payment scheme, used in the GVA and described in Eq. (5), in conjunction with the greedy allocation does *not* make for a truthful mechanism, even if bidders

are single-minded. In other terms, if the greedy allocation and Clarke's payment scheme are used, a bidder may have an incentive to lie about his valuation. The payment scheme used in the truthful mechanism of Section 10 is different from Clarke's. This is in stark contrast with the almost universal use of Clarke's scheme for devising mechanisms that are truthful in dominant strategies. Even in Nisan and Ronen [1999] where approximate mechanisms are shown to be truthful for certain situations, the payment schemes are Clarke's scheme. In Nisan and Ronen [2000] it is shown that for combinatorial auctions and other allocation problems any truthful mechanism whose payments are the VCG payments must find the optimal allocation. For us, a very simple example will suffice.

Example 8.1. As in Example 7.3, there are two goods a and b and three bidders Red, Green, and Blue. Red bids 10 for a , Green bids 19 for the set $\{a, b\}$ and Blue bids 8 for b . The greedy algorithm grants Red's and Blue's bids and denies Green's bid, that is, $f(D)(a) = \text{Red}$ and $f(D)(b) = \text{Blue}$. We shall now compute Red's payment. For this allocation, we have the following declared valuations: $v_{\text{Blue}} = 8$ and $v_{\text{Green}} = 0$. If Red had bid zero, the greedy algorithm would have granted Green's bid and denied Blue's bid. Therefore, the allocation $f(Z)$ is defined by: $f(Z)(a) = f(Z)(b) = \text{Green}$, where $v_{\text{Blue}} = 0$ and $v_{\text{Green}} = 19$. Clarke's payment scheme gives to Red: $8 - 0$ for Blue and $0 - 19$ for Green, that is, Red pays 11. Red ends up paying more than the amount he declared. If Red has been truthful and his valuation is indeed 10, his utility is -1 . He would have been better off lying, under-bidding at, say 9, or 0. In such a case, the greedy algorithm would have granted Green's bid and denied Blue's and Red's bids and the payment to Red would have been zero, making his utility 0, better than -1 .

Since this example is very simple and can be embedded in many more complex situations, we may conclude that, typically, the use of a method that is only approximately efficient is incompatible with the use of a Clarke's payment scheme. The next sections present a positive result: there is a payment scheme (necessarily different from Clarke's) that makes truth-telling a dominant strategy.

9. A Sufficient Condition for a Truthful Mechanism for Single-Minded Bidders

We shall describe in this section a number of properties of allocation schemes and of payment schemes for combinatorial auctions. Those properties seem natural properties to expect from a truthful mechanism and they are satisfied by the GVA. We shall then show that, in any mechanism that satisfies those properties, telling the truth is a dominant strategy. The literature concerned with strategic aspects of multi-item auctions has so far been interested in efficient mechanisms, that is, mechanisms that allocate the goods in an optimal way (see Krishna and Perry [1998] and Monderer and Tennenholtz [2000], for example). The conditions presented here are remarkable in that they apply to nonefficient mechanisms too. Those conditions have been shown in Archer and Tardos [2001] to apply to any situation in which agents are *one-parameter*. Single-minded bidders are not *one-parameter*. In Section 10, we shall describe a payment scheme and show that the greedy allocation scheme, together with this new payment scheme, satisfy those properties. The properties we are about to describe concern combinatorial auctions among single-minded bidders. The question of generalizing those conditions to a

more general setting is an intriguing one. Independently of this work, such a setting has been proposed in Monderer et al. [2000]. Their setting is rich enough to encompass combinatorial auctions among single-minded bidders, but not among arbitrary bidders. Our mechanism does not satisfy their Axiom 2 and its payment scheme is not of the Clarke's type they propose. The properties below are sufficient conditions for truthfulness and we do not claim they are necessary. Some of them are obviously not necessary. Nevertheless many of those properties can be shown to be necessary in the presence of others and for some others one can show that given any truthful mechanism one can easily describe another similar truthful mechanism that satisfies them. We leave to further work the exact characterization of truthful mechanisms for combinatorial auctions among single-minded bidders.

The general structure of the properties of interest is that they consider a given set of single-minded types and vary *one* of those types. They restrict the changes that can appear in the allocation or the payments as a result of such a change. Let declarations be fixed, but arbitrary, for all bidders except j . Consider two possible declarations for j : $\langle s, v \rangle$ and $\langle s', v' \rangle$. Given an allocation scheme f and a payment scheme p , we shall consider the allocations and payments generated by both declarations of j . Let g_i be the set of goods obtained by bidder i if j declares $\langle s, v \rangle$, and g'_i the set he obtains if j declares $\langle s', v' \rangle$. Similarly denote by p_i and p'_i the payments of i .

Our first property requires that the allocation, among single-minded bidders, be exact, that is, a single-minded bidder either gets exactly the set of goods he desires, nothing added, or he gets nothing. He never gets only part of what he requested. This is a very natural property, when dealing with single-minded bidders: the valuation of the bidder does not increase by giving him part of what he requested instead of nothing or by giving him more than what he requested instead of just the bundle he requested.

Exactness Either $g_j = s$ or $g_j = \emptyset$.

In an exact allocation, we shall say that j 's bid is *granted* in the first case, and *denied* in the second case. In such a scheme, the allocation may be viewed as a set of bids (or bidders) that is conflict-free, that is, the s coordinates have pairwise empty intersections. A GVA, as we defined it, does not in fact always satisfy Exactness. If nobody is interested in a , an optimal allocation could still allocate it to one of the bidders. An obvious modification of the GVA for single-minded bidders can ensure Exactness.

Our next property, Monotonicity, also concerns only the allocation scheme. It requires that, if j 's bid is granted if he declares $\langle s, v \rangle$, it is also granted if he declares $\langle s', v' \rangle$ for any $s' \subseteq s$, $v' \geq v$. In other words, proposing more money for fewer goods cannot cause a bidder to lose his bid. It follows that, similarly, offering less money for more goods cannot cause a lost bid to win. Formally:

Monotonicity $s \subseteq g_j$, $s' \subseteq s$, $v' \geq v \Rightarrow s' \subseteq g'_j$.

The GVA's allocation scheme picks the efficient allocation, that is, the allocation that maximizes the sum of the amounts of a conflict-free subset of bids. If a bid is included in the optimal allocation and its amount increases then the same allocation's total amount increases by the same amount and therefore stays optimal. Similarly, if the amount stays unchanged but the set of goods requested becomes smaller (inclusion-wise), the previous allocation, after the obvious change, is still

conflict-free and its total amount has not changed. Any allocation not containing the new bid was a suitable allocation before the change and therefore is not better. Similarly, if a bid is denied and its amount decreases, the optimal allocation's value stays fixed but the value of any allocation including the bid decreases, and similarly when varying the set s . We conclude that, assuming that there is a unique optimal allocation, the GVA's allocation scheme satisfies Monotonicity. In general, when many allocations could be tied for optimality, a GVA scheme may not be monotonic, but one may modify the GVA scheme to ensure Monotonicity.

We must immediately consider the consequences of Monotonicity, since we shall need them in stating the upcoming Critical property.

LEMMA 9.1. *In a mechanism that satisfies Exactness and Monotonicity, given a bidder j , a set s of goods and declarations for all other bidders, there exists a critical value v_c such that*

$$\begin{aligned}\forall v, v < v_c &\Rightarrow g_j = \emptyset, \\ \forall v, v > v_c &\Rightarrow g_j = s,\end{aligned}$$

We allow v_c to be infinite if $f(A^{s,v})^{-1}(j) = \emptyset$ for every v . Note that we do not know whether j 's bid is granted or not in case $v = v_c$.

PROOF. By Monotonicity, the set of v 's for which $g_j = \emptyset$ is empty (in which case take $v_c = 0$), a semi-open set of the form $[0, v_c[$ or a closed set of the form $[0, v_c]$ or equal to \mathcal{R}_+ . \square

Our third property deals with a satisfied bidder: a satisfied bidder pays exactly the critical value of Lemma 9.1, that is, the lowest value he could have declared and still be allocated the goods he desires.

$$\mathbf{Critical} \quad s \subseteq g_j \Rightarrow p_j = v_c$$

Notice that Critical says, first, that the payment for a bid that is granted does not depend on the amount of the bid, it depends only on the other bids. Then it says that it is exactly equal to the critical value below which the bid would have lost.

Critical is a necessary property for a truthful mechanism that satisfies Exactness, Monotonicity and the Participation property below. If the payment p is smaller than v_c , any bidder with real value between p and v_c loses if he declares truthfully but wins and pays less than his true value if he declares just above v_c . If the payment p is larger than v_c , any bidder with real value between v_c and p wins but gets negative utility if he declares truthfully and would be better off declaring a value below v_c and losing. Since a GVA is truthful and satisfies Exactness, Monotonicity and Participation, it also satisfies Critical.

Our last property concerns the payment scheme. Together with Critical, it implies that the utility of no truthful bidder is negative. It concerns unsatisfied bidders, bids that are denied. We require that an unsatisfied bidder pay zero. The utility of an unsatisfied bidder is then zero. This is simply tuning the utility scales of the different bidders, or, ensuring that bidders may not lose by participating in the auction.

$$\mathbf{Participation} \quad s \not\subseteq g_j \Rightarrow p_j = 0$$

A GVA satisfies Participation. In fact, the second term of Eq. (5) is precisely tuned to satisfy Participation.

Any mechanism that satisfies the conditions above is truthful. A number of preliminary lemmas are needed.

LEMMA 9.2. *In a mechanism that satisfies Exactness and Participation, a bidder whose bid is denied has utility zero.*

PROOF. By Exactness, the bidder gets nothing and his valuation is zero. By Participation, his payment is zero. \square

LEMMA 9.3. *In a mechanism that satisfies Exactness, Monotonicity, Participation and Critical a truthful bidder's utility is nonnegative.*

PROOF. If j 's bid is denied, we conclude by Lemma 9.2. Assume j 's bid is granted and his type is $\langle s, v \rangle$. Since he is truthful, his declaration is $d_j = \langle s, v \rangle$. We conclude that j is allocated s and his valuation is v . By Lemma 9.1, since j 's bid is granted, $v \geq v_c$. By Critical, j 's payment is v_c , and his utility is $v - v_c \geq 0$. \square

The next lemma shows that a bidder cannot benefit from lying just about his value (he truthfully declares the set of goods he is interested in).

LEMMA 9.4. *In a mechanism that satisfies Exactness, Monotonicity, Participation and Critical, a bidder j of type $\langle s, v \rangle$ is never better off declaring $\langle s, v' \rangle$ for some $v' \neq v$ than by being truthful.*

PROOF. Compare the case j bids, truthfully, $\langle s, v \rangle$ and the case he bids $\langle s, v' \rangle$. Let g_j be the bundle he gets in the first case and g'_j the bundle he gets in the second case. If j 's bid is denied in the second case (i.e., if $g'_j \neq s$), then, by Lemma 9.2, his utility is zero in the second case and, by Lemma 9.3, his utility in the first case is nonnegative. The claim holds.

Assume therefore that $g'_j = s$. If both bids are granted, j has the same valuation (v) and pays the same payment, v_c (by Critical). If $g'_j = s$ but $g_j = \emptyset$, it must be the case that $v \leq v_c \leq v'$. Being truthful gives j , by Lemma 9.2, zero utility. Lying gives him utility $v - v_c \leq 0$. \square

LEMMA 9.5. *In a mechanism that satisfies Exactness, Monotonicity and Critical, a bidder j declaring type $\langle s, v \rangle$ whose bid is granted, that is, $g_j = s$, pays a price p_j that is at least the price p'_j that he would have paid had he declared his type as $\langle s', v \rangle$ for any $s' \subseteq s$.*

PROOF. By Monotonicity, the bid $\langle s', v \rangle$ would have been granted and by Critical, the price p'_j paid for such a bid satisfies: for any $x < p'_j$ the bid $\langle s', x \rangle$ would not have been granted. By Monotonicity, for any such x the bid $\langle s, x \rangle$ would not have been granted. By Critical, for any x such that $x > p_j$, the bid $\langle s, x \rangle$ would have been granted. We conclude that $p'_j \leq p_j$. \square

Finally, we may prove a central result.

THEOREM 9.6. *If a mechanism satisfies Exactness, Monotonicity, Participation and Critical, then it is a truthful mechanism.*

PROOF. Suppose j 's type is $\langle s, v \rangle$. Could j have any interest in declaring his type as $\langle s', v' \rangle$? By Lemma 9.3, the only case we have to consider is when declaring $\langle s', v' \rangle$ j gets a positive utility, and by Lemma 9.2 this means that j 's bid is granted. Assume, therefore that $g'_j = s'$. If $s \not\subseteq s'$, the valuation of j is zero. Since, by

Critical, his payment is non-negative, his utility cannot be positive. Assume then $s \subseteq s'$. Since j 's valuation for s' is the same as for s , Lemma 9.5 implies that, instead of declaring $\langle s', v' \rangle$, j would not have been worse off by declaring $\langle s, v' \rangle$. Lemma 9.4 implies that declaring $\langle s, v' \rangle$ cannot be better than being truthful. \square

10. A Truthful Mechanism with Greedy Allocation

We shall now describe the payment mechanism that we propose to be used in conjunction with the greedy allocation of Section 7. The description of the payments is tightly linked with that of the greedy algorithm. The computation of the payment is performed in parallel with the execution of the greedy algorithm and takes time linear in the number of bidders for each payment. On the whole, computing the allocation and the payments takes time at most quadratic in the number of bids.

We assume that the criterion used is average amount per good, the adaptation to most other suitable greedy allocations is obvious. Informally, a bidder pays, per good, the average price proposed by the first bid in the list L that is denied because of this bid. Consider a bid j in L . Let $c(j)$ be the average amount per good of j . We shall denote by $n(j)$ the first bid following j (bids are sorted in decreasing order, that is, $c(j) \geq c(n(j))$) that has been denied but would have been granted were it not for the presence of j . Assume that such a bid exists. Notice that such a bid necessarily conflicts with j , and therefore:

$$n(j) = \min\{i \mid j < i, s(j) \cap s(i) \neq \emptyset, \forall l, l < i, l \neq j, l \text{ granted} \Rightarrow s(l) \cap s(i) = \emptyset\}.$$

Definition 10.1 Greedy Payment Schéme. Let L be the sorted list obtained in the first phase.

- j pays zero if his bid is denied or if there is no bid $n(j)$,
- if there is an $n(j)$ and j 's bid $\langle s, v \rangle$ is granted he pays $|s| \times c(n(j))$.

We may now state the main result of this article.

THEOREM 10.2. *The mechanism composed of the greedy allocation and payment schemes is truthful for single-minded bidders.*

PROOF. We shall prove that greedy mechanism satisfies Exactness, Monotonicity, Participation and Critical and use Theorem 9.6. The description of the greedy allocation scheme makes it clear that every bid is either granted or denied. The greedy allocation satisfies Exactness. For Monotonicity, assume that $s \subseteq s'$ and that $v \geq v'$ and let c be the norm of $\langle s, v \rangle$ and c' the norm of $\langle s', v' \rangle$. By our assumption concerning norms we have $c \geq c'$. If we compare the list L and L' obtained, respectively, we see that, since there are no ties by assumption, they differ only in that j 's bid may have been moved backwards by the change from $\langle s, v \rangle$ to $\langle s', v' \rangle$. The greedy allocation algorithm performs, that is, grants or denies bids, in exactly the same way on L and L' until it gets to j 's bid in L . Assume j 's bid is denied in L : there is some bid that conflict with it that has been granted already. The same bid also conflicts with j 's bid in L' since $s \subseteq s'$ and this bid will also be denied. Similarly, if j 's bid in L' is granted, no bid granted before conflicts with it and therefore no bid granted before j 's in L conflicts with it either and j 's bid is also granted in L . We have shown that the greedy allocation satisfies Monotonicity. It is

clear from the first part of Definition 10.1 that it satisfies Participation. For Critical, notice that the second part of Definition 10.1 defines the payment for a bid granted at exactly the minimal declared value that would have allowed it to be granted, v_c . Any declared value above $|s| \times c(n(j))$ leaves j before $n(j)$. If there was a bid i , $j < i < n(j)$ that would prevent the granting of j displaced in such a way, i would have to be granted and conflict with j . It is therefore a bid denied in the original allocation, that would have been granted were it not for j , contradicting the fact that $n(j)$ is the first such bid. Any declared value below $|s| \times c(n(j))$ guarantees the denial of j because $n(j)$ is granted. \square

Let us now describe this payment scheme on two examples.

Example 10.3. Consider the bidders of Example 8.1. The goods are a and b and the bidders are Red, Green, and Blue. Red bids 10 for a , Green bids 19 for the set $\{a, b\}$ and Blue bids 8 for b . We have seen that Red's and Blue's bids are granted, Green's bid is denied. This is not the efficient solution. If Red had not participated, Green's bid would have been the one with highest average price and would have been granted. Red pays Green's average price. Red pays 9.5. Green pays 0, since his bid is denied. Blue pays 0 since he is not keeping any other bid from being granted. Note that a GVA would have allocated both goods to Green and made him pay 18.

Example 10.4. Assume, as usual, two goods and three bidders. Red bids 20 for a , Green bids 15 for b and Blue bids 20 for the set $\{a, b\}$. Red's and Green's bids are granted. Blue's bid is denied. If Red had not participated, Blue's bid would still have been denied, because of Green's. Therefore, Red pays zero. Similarly, Green pays zero. Notice that, in this case, the allocation is the efficient allocation, as in a GVA, but the GVA's payments are different: Red pays 5 and Green pays 0.

11. Complex Bidders

Our assumption that bidders are single minded seems very restrictive; is there a way to extend our results to more complex players? Why not view a player as a collection of single-minded agents, or, equivalently, view the type of a player as a collection of bids? In such a setting, the game played becomes much richer in strategies and players may be better-off lying on some of their bids to obtain an advantage on others. Our discussion will, by necessity, be sketchy.

In Section 5, we presented single-minded bidders as an answer to the combinatorial explosion in bidders' types triggered by a growth in the number of goods, k . The set of types is doubly exponential in k , but the set of single-minded types is only exponential in k . In trying to overcome the limitation to single-minded bidders, one could consider any super-set of the single-minded types that grows only exponentially with k . A very natural idea is to consider players that send off single-minded agent bidders to do their work. The agents play rationally, but individually, and bring the goods and the payments due to the player. In the final analysis, a player gets all the goods obtained by each of his agents and pays all the payments imposed on each of his agents. A player's strategy is then a *small* (i.e., polynomial in k) set of single-minded agents (i.e., bids).

Our proposal is a formalization of this idea and enables us to raise fundamental game-theoretic questions about this setting. This setting is by no means a trivial

restriction. Notice, for example, that, even though a GVA may be described in terms of bids placed by the players, a player placing one bid for each subset of the goods, the allocation and payment schemes require knowledge of the identity of the player who placed the bid: a player can have at most one of his bids granted and his payment is not a function only of the bids but also of their owners.

One may ask the following questions: given a type t , not necessarily single-minded, what is a truthful description of t as a small set of single-minded bidders? For which types is there such a description? Given a mechanism, what is the declaration, that is, small set of single-minded bidders, that a player of type t should use to get the most out of the mechanism? Is there a mechanism for which a truthful declaration is a dominant strategy? The sequel will show that, if the mechanism uses any reasonable variation on the greedy allocation, the answer is negative for any reasonable definition of a truthful description.

First, let us remark that one positive result has been obtained. Theorem 10.2 shows that a single-minded bidder has, in our mechanism, a weakly dominant strategy that is to tell the truth, *even if the other players are complex players represented by a collection of single-minded agents*. But what is the optimal strategy of a complex player, that is, which agents should he send off?

It is not clear what are the mechanisms we should consider in this setting. One could assume a blind mechanism, in which the auctioneer has to allocate the goods between the single-minded agents without knowing which agents are owned by the same player. But one could also provide the auctioneer with this information. This would allow him, for example, to avoid making the payment for a bid depend on another bid from the same player, which is certainly a step toward truthfulness. One could also require the auctioneer does not grant more than one bid from each bidder, but the literature does not seem to favor this policy. As noticed in Section 2.1, a player may naturally express complementarity by the bids he puts out, but expressing substitutability is more difficult. To this effect, one could allow the players to declare not only a set of bids but also an incompatibility list describing which of his bids may not be granted simultaneously. This is the policy proposed in Fujishima et al. [1999] under the name *dummy goods*.

A further discussion of these issues can be left for a future article since our result, concerning the greedy allocation's properties, is negative and based on a simple situation that can be embedded in any of the proposals above. In Section 12, a strong result will be presented but it is necessarily formal, and general. Here, we shall present a concrete example.

Example 11.1. The mechanism we consider is the greedy mechanism. Red is a single-minded bidder and his type is $\langle \{a\}, 12 \rangle$, that is, he bids 12 for a alone. Green is a complex bidder. His type t_G is described by: $t_G(\{a\}) = 10$, $t_G(\{b\}) = 10$ and $t_G(\{a, b\}) = 30$. Notice that Green exhibits complementarity: he values the set $\{a, b\}$ at more than the sum of his values for a and b . Whatever stance one takes about the way a set of single-minded bidders can, in general, represent a type, in this case, the set of three bids: $\langle \{a\}, 10 \rangle$, $\langle \{b\}, 10 \rangle$ and $\langle \{a, b\}, 30 \rangle$ is a truthful representation of Green's type. Even if the rules of the auction allow the auctioneer to grant Green both his bid for a and his bid for b , Green cannot complain, in such a case, about his bid for the set $\{a, b\}$ being denied since he will, under any reasonable payment scheme, pay less for a and b separately than for his bid for the whole set. Suppose Green bids truthfully. The greedy mechanism grants Green's bid for the

set $\{a, b\}$ and denies all other bids. Green pays 24 (in a GVA, he would pay only 12), and therefore his utility is 6. To eliminate all doubts about the legitimacy of the payment scheme here, notice that Green's payment is determined by Red's bid, not by Green's other bids.

But consider what would have happened if Green had under-bid and declared: $\langle\{a\}, 10\rangle$, $\langle\{b\}, 10\rangle$ and $\langle\{a, b\}, 23\rangle$. The greedy mechanism now allocates a to Red (he pays 11.5) and b to Green. Green pays zero. His utility is 10. Green is better off lying. Notice that, by lying on his valuation for the set $\{a, b\}$, Green loses (6) on this bid: considered in isolation, this bid had no incentive to lie, but this lie favors the bid for b which happens to be Green's also.

Example 11.1 above exhibits a situation in which a gang of single-minded players may be globally better off under-bidding and losing utility on one of its bids, in order to have another of the gang's bid granted and making up for the loss, and more. A similar situation can arise in which a gang may be better off over-bidding on a bid b_1 to ensure that it is granted, even at a loss, in order to keep another bidder from getting goods included in another bid of the gang.

The greedy mechanism is not truthful for complex players. In the next section, we shall show that the fault does not lie with the payment scheme: no payment scheme can make the greedy allocation algorithm truthful. The problem lies with the allocation scheme. Nevertheless, the greedy mechanism has some truthfulness in it. If a player's bidding is decided in a myopic way by his single-minded agents they will bid truthfully. It is only global considerations that can induce a society of agents to require its agents not be truthful. We think we have here some kind of myopic, limited or bounded truthfulness that may be a very useful ingredient in certain types of mechanisms. Situations in which the players have too little information and too few resources to be able to analyze intelligently the global strategic situation may induce them to delegate their strategy to myopic agents. In such situations, one may be content with a mechanism that exhibit this kind of limited truthfulness.

12. No Payment Scheme Makes the Greedy Allocation a Truthful Mechanism for Complex Bidders

In Section 11, we showed that the greedy scheme, that is, greedy allocation + greedy payment, cannot be extended to a truthful mechanism for complex players. We shall now show that no payment scheme can complement the greedy allocation.

If a bidder is not single-minded, but double-minded (i.e., interested in two different sets of two goods), there may be no payment scheme that, combined with the greedy allocation algorithm, will make for a truthful mechanism. We shall consider a very simple situation: two goods, two bidders, one of them single-minded, the other one double-minded. The search for a family of bidders that is significantly larger than the single-minded ones and a suitable payment scheme is open, but it starts with a negative result. Notice the result does not only show that our payment scheme is unsuitable, it shows that no payment scheme exists (to be used in conjunction with the greedy allocation scheme).

Assume there are two goods a and b and two bidders Green and Red. Red is single-minded and his type is $\langle\{a\}, 10\rangle$. Red truthfully declares his type. Green is interested in both b and the set $\{a, b\}$. His valuation is 0 for any allocation in which

he does not get b . It is v_b for any allocation in which he gets b but not a , and it is v_{ab} if he gets both a and b , $v_{ab} > v_b$. Green's declaration is 0 for all allocations that do not give him b , g_b for all allocations that give him b but not a and g_{ab} for the allocation in which he gets both a and b . Notice that four parameters describe the auction: g_{ab} , g_b , v_{ab} and v_b . Assume, furthermore, that $0 \leq g_b < 10$. We reason by contradiction and assume there is a payment scheme that makes truth-telling a dominant strategy for Green. Let us consider two cases.

First, assume that $g_{ab} > 20$. In this case, the greedy algorithm will allocate both goods to Green. The payment mechanism will make Green pay p_{ab} . Notice that this payment p_{ab} cannot depend on:

- g_{ab} (as long as $g_{ab} > 20$): otherwise, Green would have an interest in declaring the g_{ab} most favorable to him, irrespective of his v_{ab} ,
- g_b : otherwise, similarly, Green would have an interest in declaring the g_b most favorable to him irrespective of his v_b ,
- v_{ab} : since payments cannot depend on private values,
- v_b : similarly.

Therefore, p_{ab} is simply a number. The utility of Green, in this first case, is: $v_{ab} - p_{ab}$.

Consider, now, a second case: $g_{ab} < 20$. In this case, the greedy algorithm allocates a to Red and b to Green. Let us denote by p_b the payment of Green. For the same reasons as above, p_b cannot be a function of any of the parameters. The utility of Green, in this second case, is: $v_b - p_b$.

Assume that, in fact, Green is bidding his true valuation on b (i.e., $g_b = v_b$). Since truth-telling is a dominant strategy for Green, it must be the case that,

- if $v_{ab} > 20$, Green gets from case 1 not less than from case 2 (i.e., $v_{ab} - p_{ab} \geq v_b - p_b = g_b - p_b$);
- if $v_{ab} < 20$, Green gets from case 2 not less than from case 1 (i.e., $g_b - p_b = v_b - p_b \geq v_{ab} - p_{ab}$).

By considering the case v_{ab} is just greater than 20 and g_b is just less than 10, the first inequality gives us $20 - p_{ab} \geq 10 - p_b$ (i.e., $p_{ab} - p_b \leq 10$). By considering the case v_{ab} is just less than 20 and g_b is 0, the second inequality gives us $-p_b \geq 20 - p_{ab}$ (i.e., $p_{ab} - p_b \geq 20$). A contradiction.

Let us try, now, to discuss the reasons for the negative result just presented. Why is there a scheme for single-minded bidders and no scheme for more complex bidders? The impossibility to devise a truth-conducting payment scheme around the greedy allocation stems from the richness of the strategic possibilities offered to a complex bidder. Let us explain why the obvious extension of our payment scheme does not work. Bidder i , in order to get good a against the competition of another bidder interested in $\{a, b, c\}$, may have an interest in over-bidding on c and get it at a loss, just to keep his opponent from getting the whole set. Similarly, i underbidding on a and losing it, may give $\{a, b\}$ to another bidder, which in turn may keep a third bidder from getting $\{b, c\}$ and cause i to get much coveted $\{c\}$.

The discussion just above is, in fact, very similar to Vickrey's discussion in Section 5 of Vickrey [1961] of the reasons why his scheme for an auction of identical objects is truth-revealing only if one assumes buyers of a very simple type: interested in at most one item.

13. Revenue Considerations

We have described a feasible mechanism for combinatorial auctions that is truthful when bidders are single-minded. Should a seller use it for selling goods? It is very difficult to say anything general about the revenue generated by this mechanism. We shall compare the revenue generated by our mechanism to the revenue generated by a GVA. Since a GVA allocates the goods in an efficient way but our mechanism does not, one can fear that the revenue generated by our mechanism will be significantly smaller in all those cases in which the allocation is not efficient. This does not seem to be the case. There are cases in which our algorithm generates a higher revenue than a GVA and there are cases in which a GVA is preferable. The comparison does not seem to be tightly correlated to the relative efficiency of the allocations. We shall present four simple situations. All examples assume single-minded bidders Green, Red, Black and sometimes Blue. The first two examples are typical of purely combinatorial situations.

Example 13.1. Assume there are four goods, a , b , c and d . Green is interested in $\{a, b\}$, Red in $\{c, d\}$ and Black in $\{a, c\}$. All bids are for the same amount: 1.

Let us first consider a GVA. A GVA allocates the efficient way: Green gets $\{a, b\}$ and Red gets $\{c, d\}$. Green and Red pay nothing: if they had not participated only one bid could have been granted. The revenue generated by a GVA is zero.

Because of the tie our greedy scheme may end in one of three possible situations, up to symmetry between Green and Red. First scenario: the order is Green, Red, Black. The allocation is efficient and nobody pays anything, as for a GVA. Second scenario: the order is Green, Black, Red. The allocation is efficient, but this time Green pays 1, Red pays nothing. Third scenario: the order is Black, Green, Red. The allocation is *not* efficient: Black gets $\{a, c\}$ and b and d are unallocated. Black pays 1.

In this case, our scheme generates, on average, $2/3$, whereas a GVA generates 0.

Example 13.2. Four goods, a , b , c and d . Green is interested in $\{a, b\}$, Red in $\{c, d\}$, Black in $\{a, c\}$ and Blue in $\{b, d\}$. All bids are for the same amount: 1.

A GVA allocates the efficient way: either to Green and Red or to Black and Blue. In any case, each successful bidder pays 1: the revenue is 2 and the full surplus is extracted.

Because of the tie our greedy scheme may end in one of three possible situations, up to symmetry. First scenario: the order is Green, Red, Black, Blue. The allocation is efficient (to Green and Red) and nobody pays anything. Second scenario: the order is Green, Black, Red, Blue. The allocation is efficient (to Green and Red), but this time Green pays 1, Blue pays nothing. Third scenario: the order is Green, Black, Blue, Red. The allocation is efficient (to Green and Red). Green pays 1 and Red pays nothing.

In this case, our scheme generates, on average, $2/3$, whereas a GVA generates 2.

Our next example is typical of strong complementarity.

Example 13.3. Red bids 20 for the set $\{a, b\}$, Green bids 9 for a and Black bids 1 for b . Both our greedy algorithm and a GVA allocate a and b to Red. With our scheme, Red pays 18, with a GVA, he pays 10.

Example 13.4. Green bids 20 for a , Red bids 37 for the set $\{a, b\}$ and Black bids 18 for b . Both our greedy algorithm and the efficient allocation of the GVA give a to Green and b to Black. With us, Green pays 18.5 and Black pays nothing. With a GVA, Green pays 19 and Black pays 17. Our mechanism generates 18.5 to the GVA's 36.

Example 13.5. Green bids 10 for a and Red bids 19 for the set $\{a, b\}$. Our greedy scheme allocates a to Green and leaves b unallocated. The efficient allocation of the GVA gives both a and b to Red. In our scheme Green pays 9.5. In a GVA, Red pays 10.

More work is needed to assess the revenue generated by the mechanism proposed.

14. Conclusion and Future Work

To overcome the complexity of computing, the efficient allocation in combinatorial auctions, we propose to use a greedy approximation together with a payment scheme tailored to fit it. The combination provides a truthful mechanism tailored to a specific nonoptimal approximation scheme that is not a member of the VCG family. This mechanism admits dominant strategies and is therefore very sturdy. Between the presentation of a first version of this work [Lehmann et al. 1999a, 1999b] and the publication of this final version, a number of both impossibility and possibility results [Nisan and Ronen 2000; Archer and Tardos 2001; Lehmann et al. 2001; Bartal et al. 2002; Mu'alem and Nisan 2002] on the topic of feasible approximations and truthful mechanisms for different sorts of combinatorial auctions.

A number of additions, modifications or extensions should be considered. Reserve prices are a necessary feature of real-life auctions. Adding reserve prices to our scheme poses no problem: reserve prices are bids put out by the auctioneer and truthfulness is still a dominant strategy for the buyers. The optimal strategy of the seller, that is, the optimal reserve price, requires a different study. In a combinatorial auction, the reserve prices can, very naturally, express the complementarity of the seller. In particular, a seller who does not want to sell too large sets of goods to the same buyer, to avoid monopolies for example, will put high reserve prices for large sets of goods.

Before one can apply the ideas presented here to auctions of identical items, and to such double auctions, those ideas need to be adapted to this setting. This is the topic of further research.

A combinatorial auction that features a number of different types of goods, a number of items of each type of goods being for sale, represent the ultimate combinatorial auction. The ideas presented in this paper may provide a computationally feasible solution for such auctions.

The revenue generated by the mechanism proposed should be studied in depth.

The approximation scheme presented in this paper: greedy, is quite rudimentary. Even though it attains the theoretically optimal (worst-case) ratio, it should, probably, in practice, be either iterated with different criteria or be included in some more complex scheme with some sort of backtracking. The main avenue for further research is probably the consideration such better approximation schemes and the design of suitable payment schemes. The properties described in Section 9 are a clear guide on how to do that. Note, in particular, that Critical leaves no freedom in the design of the payment scheme.

The properties of Section 9 are sufficient for truthfulness, among single-minded bidders, but some of them also seem to be necessary, at least in the presence of others. A full characterization of truthful schemes for combinatorial auctions should be attempted.

ACKNOWLEDGMENTS. We want to thank Noam Nisan, Moshe Tennenholtz and Robert Wilson for their insightful remarks. Anonymous watchful referees helped us improve the presentation and clarify a few points, thanks to them too.

REFERENCES

- ARCHER, A. F., AND TARDOS, É. 2001. Truthful mechanisms for one-parameter agents. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science (Las Vegas, Nev.)*. IEEE Computer Society, Los Alamitos, Calif., pp. 482–491.
- BARTAL, Y., GONEN, R., AND NISAN, N. 2002. Incentive compatible multi unit combinatorial auctions. Tech. Rep. TR-2002-36, Leibniz Center for Research in Computer Science, School of Computer Science and Engineering, Hebrew University, Jerusalem, Israel. June. (Presented at Dagstuhl workshop on Electronic Market Design.)
- BOUTILIER, C., AND HOOS, H. H. 2001. Bidding languages for combinatorial auctions. In *Proceedings of IJCAI'01 (Seattle, Wash.)*. Morgan-Kaufmann, San Mateo, Calif.
- CLARKE, E. H. 1971. Multipart pricing of public goods. *Public Choice* 11, 17–33.
- CRAMTON, P. C. 1997. The FCC spectrum auction: An early assessment. *J. Econ. Manage. Strat.* 6, 3, 431–495.
- DEMARTINI, C., KWASNICA, A. M., LEDYARD, J. O., AND PORTER, D. 1999. A new and improved design for multi-object iterative auctions. Social Science Working Paper 1054, California Institute of Technology, Pasadena, Calif.
- FUJISHIMA, Y., LEYTON-BROWN, K., AND SHOHAM, Y. 1999. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Proceedings of IJCAI'99 (Stockholm, Sweden)*. Morgan-Kaufmann, San Mateo, Calif.
- GONEN, R., AND LEHMANN, D. 2000. Optimal solutions for multi-unit combinatorial auctions: Branch and bound heuristics. In *Proceedings of the 2nd Conference on Electronic Commerce (EC'00) (Minneapolis, Minn.)*. ACM, New York, pp. 13–20.
- GROVES, T. 1973. Incentives in teams. *Econometrica* 41, 617–631.
- HALLDÓRSSON, M. M. 1999. Approximation of weighted independent set and hereditary subset problems. In *Proceedings of COCOON'99*. Lecture Notes in Computer Science, vol. 1627. Springer-Verlag, New York.
- HALLDÓRSSON, M. M. 2000. Approximation of weighted independent set and hereditary subset problems. *J. Graph Algori. Appl.* 4, 1, 1–16.
- HÅSTAD, J. 1999. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Math.* 182, 105–142.
- JAIN, K., AND VAZIRANI, V. V. 2001. Applications of approximation algorithms to cooperative games. In *Proceedings of 33rd Annual ACM Symposium on Theory of Computing (Crete, Greece)*. ACM, New York, pp. 364–372.
- KARP, R. M. 1972. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, R. E. Miller and J. W. Thatcher, Eds. Plenum Press, New York, pp. 85–103.
- KRISHNA, V., AND PERRY, M. 1998. Efficient mechanism design. Available at <http://www.ma.huji.ac.il/~motty>.
- LEHMANN, B., LEHMANN, D., AND NISAN, N. 2001. Combinatorial auctions with decreasing marginal utilities. In *Proceedings of the 3rd Conference on Electronic Commerce EC'01 (Tampa, Fla.)*. Y. Shoham, Ed. ACM, New York.
- LEHMANN, D., O'CALLAGHAN, L. I., AND SHOHAM, Y. 1999a. Truth revelation in rapid, approximately efficient combinatorial auctions. Tech. Note STAN-CS-TN-99-88, Stanford University, Computer Science, Robotics Lab., Stanford, Calif.
- LEHMANN, D., O'CALLAGHAN, L. I., AND SHOHAM, Y. 1999b. Truth revelation in rapid, approximately efficient combinatorial auctions. In *Proceedings of the 1st ACM Conference on Electronic Commerce (EC'99) (Denver, Col.)*. ACM, New York, pp. 96–102.
- MACKIE-MASON, J. K., AND VARIAN, H. R. 1994. Generalized Vickrey auctions. Working paper, Univ. Michigan. July.

- MAS-COLELL, A., WHINSTON, M. D., AND GREEN, J. R. 1995. *Microeconomic Theory*. Oxford University Press, New York, Oxford.
- MC MILLAN, J. 1994. Selling spectrum rights. *J. Econ. Pers.* 68, 145–162.
- MILGROM, P. 2000. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy* 108, 2, 245–272.
- MONDERER, D., KFIR-DAHAV, N. E., AND TENNENHOLTZ, M. 2000. Mechanism design for resource bounded agents. In *Proceedings of ICMAS-2000*.
- MONDERER, D., AND TENNENHOLTZ, M. 2000. Asymptotically optimal multi-object auctions for risk-averse agents. Available at <http://ie.technion.ac.il/dov.phtml>.
- MU'ALEM, A., AND NISAN, N. 2002. Truthful approximation mechanisms for restricted combinatorial auctions. In *Proceedings 18th National Conference on Artificial Intelligence, AAAI'02*. Edmonton, Canada. (Also presented at Dagstuhl workshop on Electronic Market Design), 379–384.
- MYERSON, R. B. 1981. Optimal auction design. *Math. Oper. Res.* 6, 58–73.
- NISAN, N. 1999. Algorithms for selfish agents. In *Symposium on Theoretical Aspects in Computer Science* (Trier, Germany), 1–15.
- NISAN, N. 2000. Bidding and allocation in combinatorial auctions. In *Proceedings of the 2nd ACM Conference on Electronic Commerce (EC'00)* (Minneapolis, Minn.). ACM, New York, pp. 1–12.
- NISAN, N., AND RONEN, A. 1999. Algorithmic mechanism design. In *Proceedings of the 31st Annual Symposium on Theory of Computing*. ACM, 129–140.
- NISAN, N., AND RONEN, A. 2000. Computationally feasible VCG mechanisms. In *Proceedings of the 2nd ACM Conference on Electronic Commerce (EC'00)* (Minneapolis, Minn.). ACM, New York.
- ROTHKOPF, M. H. 1983. Bidding theory: the phenomena to be modeled. In *Auctions, Bidding and Contracting: Uses and Theory*, R. Engelbrecht-Wiggans, M. Shubik, and R. Stark, Eds. New York University Press, New York, pp. 105–120.
- ROTHKOPF, M. H., PEKEČ, A., AND HARSTAD, R. M. 1998. Computationally manageable combinatorial auctions. *Manage. Sci.* 44, 8, 1131–1147.
- ROUGHGARDEN, T., AND TARDOS, E. 2002. How bad is selfish routing? *J. ACM* 49, 2 (Mar.), 236–259.
- SANDHOLM, T. 1999. An algorithm for optimal winner determination in combinatorial auctions. In *Proceedings of IJCAI'99* (Stockholm, Sweden). Morgan-Kaufmann, San Mateo, Calif., pp. 542–547.
- SANDHOLM, T. 2000. eMediator: A next generation electronic commerce server. In the *International Conference on Autonomous Agents* (Barcelona, Spain). 73–96. (First appeared as a Washington University, St. Louis, Dept of Computer Science technical report WUCS-99-02, January 1999.)
- SANDHOLM, T. 2002. Algorithm for optimal winner determination in combinatorial auctions. *Artif. Intell.* 135, 1–54.
- VARIAN, H. R. 1995. Economic mechanism design for computerized agents. In *Proceedings of the 1st Usenix Conference on Electronic Commerce*. New York.
- VICKREY, W. S. 1961. Counterspeculation, auctions and competitive sealed tenders. *J. Finance* 16, 8–37.

RECEIVED MARCH 2001; REVISED MARCH 2001; ACCEPTED MAY 2002