Solving Combinatorial Exchanges: Optimality via a Few Partial Bids

Anshul Kothari
Department of Computer Science
University of California
Santa Barbara, CA 93106
kothari@cs.ucsb.edu

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
sandholm@cs.cmu.edu

Subhash Suri
Department of Computer Science
University of California
Santa Barbara, CA 93106
suri@cs.ucsb.edu

Abstract

We investigate the problem of matching buyers and sellers in a multi-item multi-unit combinatorial exchange so as to maximize either the surplus (revenue minus cost) or the trading volume (number of units traded). In such an exchange, participants can place bids to buy or sell *bundles* of goods. While even highly specialized cases of this problem are both NP-Complete and inapproximable, we show that optimal surplus or trade volume can be achieved in polynomial time if some bids can be satisfied *partially*. Using theory of linear programming, we show that in exchanges trading multiple units of k distinct items, maximum surplus is possible with at most k partial bids, and maximum trade volume is possible with at most $k+1$ partial bids. These bounds on the number of partial bids are the best possible in the worst case. For the simple but important case of single item exchanges, we also develop fast combinatorial algorithms for optimal matching. The bidding language that simply allows bidders to bid on alternative (nonexclusive) combinations enables bidders to express complementarity between items, but not substitutability (the value of a bundle being less than the sum of the parts). We show that if the bidding language is enriched with XOR-constraints between bids (a common, necessary and sufficient enrichment that achieves full expressiveness by the bidders), then allowing for partial acceptance of bids does not help—even for single item exchanges, computing the optimal matching remains NP-Complete no matter how many partial bids are allowed.

Introduction

Auctions have been studied in economics and game theory for a long time as important resource allocation mechanisms in distributed environments. In recent years, their role has grown with the emergence of Internet and electronic commerce, as businesses and corporations leverage the new medium to streamline their procurement process. Many businesses are moving to an auction-based purchase method where they issue a request for quotes for the goods and services needed, and let the suppliers bid for a piece of the business. (See, for example, the web sites of companies such as FreeMarkets, i2 Technologies, or CommerceOne.) Driven by these fundamentals, auctions—and algorithms related to them—have become important and popular research topics in computer science, especially in AI.

An exchange generalizes the auction mechanism to the setting with multiple buyers and sellers. Some familiar examples are the exchanges for equities and commodities, transportation, electricity, and the business-to-business exchanges. A *combinatorial exchange* is an exchange where buyers and sellers can bid on bundles (subsets) of goods. Combinatorial markets are desirable both for the bidders and the overall economic efficiency because items often have complementarity, and combinatorial bidding minimizes bidders' risk of getting stuck with only a partial subset. (Some examples of complementarity are FCC's spectrum auctions, where a bidder may highly value licenses in some neighboring regions or wish to get a national footprint, or direct-material procurement, where a company must acquire all the raw materials necessary to produce its goods.)

Although combinatorial markets were proposed 20 years ago (Rassenti, Smith, & Bulfin 1982), the last few years have seen a surge of research, no doubt by their relevance to the Internet and the electronic commerce (Rothkopf, Pekeč, & Harstad 1998; Sandholm 2002a; Fujishima, Leyton-Brown, & Shoham 1999; Sandholm & Suri 2000; Nisan 2000; Sandholm *et al.* 2001; 2002). Given a set of buyers and sellers, determining which set of trades leads to the highest surplus (profit) is NP-Complete (Rothkopf, Pekeč, & Harstad 1998) and inapproximable to a ratio better than $\Omega(n^{1-\varepsilon})$, where n is the number of bids (Sandholm 2002a). Typical approaches to solving combinatorial auctions fall into three categories: search algorithms (such as branch and bound) (Sandholm 2002a; Fujishima, Leyton-Brown, & Shoham 1999; Sandholm & Suri 2000; Nisan 2000; Sandholm *et al.* 2001; 2002), approximation algorithms (Lehmann, O'Callaghan, & Shoham 1999; Hoos & Boutilier 2000), and restricted auctions where bids on only some subsets are allowed (Rothkopf, Pekeč, & Harstad 1998; Sandholm & Suri 2000; Tennenholtz 2000; Penn & Tennenholtz 2000). The first approach obtains optimal solutions but can be exponential in the worst-case; the second approach is typically polynomial time but usually can't provide any decent guarantee on the quality of solution; the third approach suffers from the same economic inefficiencies as noncombinatorial markets because the bidders might not be able to bid on the bundles that they really want.

All-Or-Nothing vs. Partial Acceptance of Bids

Most research on combinatorial markets has focused on the binary case where each bid must be either fully accepted or rejected, which makes the problem NP-Complete. In many real combinatorial markets, bids could be accepted partially, but it is more desirable to accept them entirely or not at all. For instance, spectrum licenses can be refined but at an additional cost; transportation carriers prefer full truck loads but will carry partial truckloads; in business procurement, suppliers would rather have full contracts, but given a choice between no contract or a partial one, may prefer the partial contract. In other possible applications, such as computational ressource markets (e.g. CPU, disk, bandwidth), there is no inherent reason to not accept partial bids. However, since a combinatorial market maker's reputation is linked to its ability to allocate full bundles, there is an incentive to minimize the number of partial bids. Allowing bids to be accepted partially has two advantages. First, a better economic allocation is obtained: the relaxed linear program has a better (or same) solution than the integer program. Second, the computational complexity of the solving the exchange might be reduced significantly.

In this paper we study markets where bids can be accepted partially, but it is desirable to only accept a small number of them partially. We show how many bids have to be accepted partially to obtain a solution of equal value as that where all bids can be accepted partially. We also determine the computational complexity of finding such a solution.

Clearing Objectives

We consider two natural objective functions in solving combinatorial exchanges: *surplus* and *trade volume*. The former is the *net* monetary gain (profit) realized by trading the goods: the difference between the revenue collected from the buyers and the amount paid to the sellers. Maximizing the surplus is also equivalent to maximizing the *social welfare* in an economy because it puts the goods in the hands of the people who value them most. (When we maximize surplus, we actually maximize *revealed surplus*, that is, surplus as revealed by the bidders' bids. If the bidders bid truthfully, then this corresponds to actual surplus. We do not address the question of motivating the bidders to bid truthfully. There is a large literature on this topic in game theory, and our algorithms can be used in conjunction with those mechanisms.) The second objective function is the total number of *units* traded over all goods *without external subsidy*. That is, the goal is to maximize the trade volume under the constraint that the net surplus is non-negative—with external subsidy, all the units can be traded trivially.[1]

Our Results

We state our results for combinatorial exchanges; since auctions and reverse auctions are special cases, our results apply to those markets as well. In an exchange, each buyer may be matched with multiple sellers and vice versa. A seller places a *sell bid* (also called an *ask*), which specifies the goods he wishes to sell and at what price. Similarly, a buyer places a *buy bid* (simply called a *bid*), which specifies the goods he wishes to buy and at what price.[2]

Suppose there are n participants, labeled 1 through n, of which $S = \{1, 2, \ldots, s\}$ is the set of sellers, and $B = \{s+1, \ldots, n\}$ is the set of buyers. Let P^* denote the maximum possible surplus (profit) in any matching of buyers B and sellers S *allowing any number of bids to be partially accepted*. Similarly, let V^* denote the maximum possible trading volume with any number of partial bids. In these optimal solutions, however, a rather large number of buy or sell bids can be fractional. A natural question is how much surplus or trade volume is sacrificed if we put an *upper bound* on the number of bids that can be partially accepted. Suppose the market has k distinct items, each with any number of identical units, and each bid can be an arbitrary combinatorial bid, (G, p), where G is the bundle of items, specifying how many units of each item, and p is the price attached to the bid. We show that theory of linear programming easily implies that that the maximum surplus P^* is achievable with at most k partial bids. Similarly, when the objective is to maximize the trade volume then V^* is achievable with at most $(k+1)$ partial bids. These bounds are the best possible in the worst-case.

Next, we consider the special case of $k = 1$—that is, markets that trade multiple identical units of one item. We derive fast combinatorial algorithm for maximizing the surplus or trade volume while allowing at most partial bid for the surplus and at most two partial bids for the volume. (Note that, even in this restricted setting, the problem is intractable—both the knapsack and the subset sum are special cases.) These result are sharp: if no partial bid is allowed, then determining the exact surplus is NP-Complete, and in a worst-case the surplus can go to zero. Similarly, if no partial bid is allowed, then the trading volume can go to zero from V^*, and with one partial bid the best guarantee on the trade volume is $\frac{1}{2}V^*$.

Next, we consider the case of a more expressive bidding language, where each bidder can submit multiple bids that are XOR'ed together. A bid has the form $(q_1, p_1) \oplus (q_2, p_2) \oplus \ldots \oplus (q_j, p_j)$, where at most one (q_i, p_i) is to be accepted. Such bids are more expressive because they can encode substitutability among items and decreasing marginal utility in units. They are also used in practice for expressing *volume discounts*. At most one bid from a bidder can be accepted, but that bid can be accepted partially. We now ask: is there a polynomial time algorithm for maximizing the surplus or trade volume with XOR-bids where some fixed number ℓ of bids can be partially fulfilled. Surprisingly, the answer is negative. We show that computing the optimal surplus or

[1] The subsidy provider buys everything from the sellers, paying them their ask prices, and then sells everything to the buyers, collecting their bid prices; any shortfall comes out of his subsidy. We assume here that demand equals supply for each item.

[2] For simplicity, we assume that each participant is either a buyer or a seller; a participant who is both selling and buying in the same exchange can be modeled by two separate proxies, one seller and one buyer. Our exchange algorithms can also handle the case where a bid includes both the buy and sell component, but to simplify the discussion, we assume that each bid is either a buy or a sell bid.

trade volume is still NP-Complete, even for single-item exchanges (actually even auctions), irrespective of how many partial bids are allowed.

Thus, we have a sharp division: markets where valuation functions are *superadditive* can be optimally solved using a small number of partial bids (depending on the number of distinct item types). But if the valuation functions are *subadditive*, then even accepting bids partially does not help clear the market.

Multi-Item Exchanges

We begin with the general multi-item multi-unit combinatorial exchange, where each buyer or seller has a bid on a *bundle* of goods. Suppose the exchange deals with k item types (or goods). We use the notation, (q_i, p_i) to denote the bid of participant i, where q_i is a *vector* of length k, and p_i is the price for the bundle. The number of units of item j in this bid is denoted by q_i^j. For example, a seller's bid $(\langle 3, 0, 5\rangle, \$100)$ is a bid to sell three units of item 1 and five units of item 2 for \$100.

Given a set of n combinatorial bids, determining the surplus maximizing matching is NP-Complete. The set packing problem is a special case of the combinatorial exchange problem. We argue in this section that it is possible to find a surplus or trade maximizing matching in polynomial time where the number of bids accepted partially depends only on k, the number of items, independent of the number of units and the number of bidders. In particular, for surplus maximization, k partial bids suffice; for trade maximization, $k+1$ partial bids suffice. We are implicitly assuming here that the items are divisible (such as oil, gas, electricity, pollution and logging rights) or roundable (commodities traded in large quantities, such as equities, direct and indirect materials in business procurement) so that a partially satisfied bid is viable. We begin by formulating the combinatorial exchange problem as an integer program. Let x_i denote the binary variable, which represents whether the bid i is accepted ($x_i = 1$) or rejected ($x_i = 0$). (Recall that B and S are the sets denoting buyers and sellers, respectively.) The following program maximizes the surplus.

$$\begin{aligned} \max \quad & \sum_{i\in B} p_i x_i \;-\; \sum_{i \in S} p_i x_i & (1)\\ \text{s.t.} \quad & \sum_{i\in B} q_i^j x_i \;=\; \sum_{i\in S} q_i^j x_i \quad \text{for } j = 1, 2, \ldots, k & (2)\\ & x_i \in \{0, 1\} & (3) \end{aligned}$$

The objective function maximizes the difference between the revenue from the buyers and the payment to the sellers. The first constraint group ensures that demand equals supply for each item, and the second constraint group enforces binary decision variables. (The equality in the first constraint group can be relaxed to $\geq$ inequality if extra items can be freely disposed.)

The LP relaxation of this program changes the last constraint to $0 \leq x_i \leq 1$. To show that a surplus maximizing matching exists with at most k partial bids, we argue that every vertex of this LP polytope corresponds to a solution in which at most k variables are fractional. The argument is quite straightforward: the linear program has n variables (the bids), and $2n + k$ constraints. (In addition to the k constraints of Eq. (3), there are 2 constraints of type $x_i \geq 0$ and $x_i \leq 1$ for each variable.) Each of these constraints is defined by a hyperplane, and a vertex of the LP polytope corresponds to the intersection of n hyperplanes. Since there are only k constraints of type Eq. (3), any vertex must involve at least $n-k$ hyperplanes of the type $x_i = 0$ or $x_i = 1$. Since the latter hyperplanes correspond to binary decisions on bids, at least $n-k$ coordinates of any vertex are integral. Thus, the number of fractional coordinates, which correspond to partial bids, is at most k. Thus, any vertex of the LP, and hence its optimal solution involves at most k partially satisfied bids.

The integer program for maximizing the trade volume is

$$\begin{aligned} \max \quad & \sum_{i\in B} \sum_{j=1}^{k} q_i^j x_i \\ \text{s.t.} \quad & \sum_{i\in B} q_i^j x_i \;=\; \sum_{i\in S} q_i^j x_i \qquad \text{for } j = 1, 2, \ldots, k \\ & \sum_{i\in B} p_i x_i \;-\; \sum_{i\in S} p_i x_i \;\geq\; 0 \\ & x_i \in \{0, 1\} \end{aligned}$$

The objective function maximizes the number of units bought by all the buyers (which equals the number of units sold by sellers). The first constraint group is the same as before: matching demand and supply for each item. The second constraint is a new one: it ensures that total net surplus is non-negative. Arguing as before, it follows easily that any vertex of the LP polytope can have at most $k+1$ fractional x_i's, and so the trade-maximizing matching has at most $k+1$ partially satisfied bids. (It is easy to see that our formulation and the proof holds even if a bid contains both the buy and sell components in it.) We conclude with the following simple result.

Theorem 1 *Consider a multi-unit combinatorial exchange with k items. Given any set of n combinatorial buy or sell bids, there is a surplus-maximizing matching with at most k partially satisfied bids, and a trade-maximizing matching with at most $k+1$ partially satisfied bids. These matchings can be found by solving a linear program with n variables and $2n+k$ constraints.*

These bounds are close to the best possible. The following example shows that the surplus can drop from kL to zero if the number of partial bids is reduced below k. Consider a single buyer, who places a combinatorial bid $(\langle 1, 1, \ldots, 1\rangle, \$kL)$ to buy one unit of each of the k items, for a total price of kL. There are k sellers, where seller i has 2 units of item i, for price 2ϵ. A surplus of $k(L-\epsilon)$ is possible by satisfying each seller partially. But if fewer than k partial bids are allowed, then no trading is possible, and the surplus is reduced to zero. We can extend this example to show similar behavior for trade volume.

Next, we consider the special case of single item exchanges, and develop fast combinatorial algorithms for it.

This is an important setting, since many markets have dedicated sub-markets for individual good types (e.g. stocks, commodities).

Single-Item Exchanges

Each bid specifies the number of units and the price for the bid. It will be more convenient to use the *per unit* price in bids, so that a bid (q, p) means that the bidder wants to buy or sell q units at the price of $\$p$ per unit. (If fewer than q units are traded, then this bid is considered partially accepted.)

The exchange has s sellers and $n - s$ buyers, and our goal is to determine an optimal trading among them maximizing either the total surplus or the volume (number of units traded). We use the notation (q_i, p_i) to denote the bid of participant i; we will use the index sets S and B to distinguish between sellers and buyers. We also use the notation (q'_i, p_i) to denote the portion of the bid (q_i, p_i) that was accepted. A *matching* between the buyers and sellers is defined as a *bipartite graph*, with nodes $S \cup B$, and an edge (s, b) if a non-zero trade occurs between the seller s and buyer b. Each buyer may be matched with multiple sellers and vice versa. We use the notation $exch(s, b)$ to denote the number of units traded between s and b. Observe that, for any seller s, $\sum_{b \in B} exch(s, b) = q'_s \leq q_s$ and, for any buyer b, $\sum_{s \in S} exch(s, b) = q'_b \leq q_b$.

Suppose M is a matching between buyers and sellers (with any number of partial bids), then the surplus generated by M is written as

$$P_M \;=\; \sum_{i \in B} q'_i \; p_i \;-\; \sum_{j \in S} q'_j \; p_j .$$

The trade volume generated by M is written as

$$V_M \;=\; \sum_{i \in B} q'_i \;=\; \sum_{j \in S} q'_j .$$

The next two subsections separately deal with the objectives of maximizing the surplus and the trade volume.

Surplus Maximization

Our algorithm is quite simple and based on sorting the buyer bids in descending order of per unit price, and sorting seller bids in ascending order of per unit price. For technical convenience, we will assume that the bids of all buyers have distinct unit prices, and similarly for the sellers. Otherwise, an appropriate tie-breaking rule can enforce a unique ordering. The algorithm is described below.

Single-Item-Max-Surplus

1. Sort the buyers in descending and sellers in ascending order of per unit price.
2. Initially all buyers and sellers are *unmarked*.
3. While there is at least one unmarked buyer and one unmarked seller, repeat the following steps:
   - Let b and s be the first unmarked buyer and seller.
   - If $p_b \leq p_s$, then terminate.
   - Else if $q_b \leq q_s$, then mark b, set $q_s = q_s - q_b$, and $exch(s, b) = q_b$.
   - Else if $q_b > q_s$, then mark s, set $q_b = q_b - q_s$, and $exch(s, b) = q_s$.

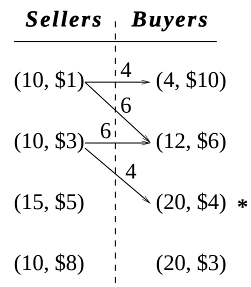

Figure 1: Algorithm **Single-Item-Max-Surplus**. Asterisk marks the buyer with partial bid.

Fig. 1 illustrates the algorithm. We now prove that this algorithm produces maximum possible surplus and results in at most one partial bid. Assume, without loss of generality, that the descending price order of buyers is $b_1, b_2, \ldots$ and the ascending price order of sellers is $s_1, s_2, \ldots$. We begin by introducing the definition of bid contiguity, which says that an optimal matching can't skip bids.

Let M be a matching between buyers B and sellers S, and suppose that b_k is the lowest price buyer with non-zero trade, and s_ℓ is the highest price seller with non-zero trade. We say that matching M satisfies the *bid contiguity* conditions if the bids of all buyers $b_1, b_2, \ldots, b_{k-1}$ and all sellers $s_1, s_2, \ldots, s_{\ell-1}$ are fully satisfied.

Lemma 1 *Any surplus-maximizing matching M_{opt} between B and S satisfies bid contiguity.*

PROOF. The proof is by a simple swapping argument. Assume that b_k is the lowest price buyer with non-zero trade, and s_ℓ is the highest price seller with non-zero trade. Suppose there is a buyer b_i, with $p_i > p_k$, whose bid is partially fulfilled. Let s_j be a seller with $exch(s_j, b_k) \neq 0$. Since $p_{s_j} < p_{b_k}$, it follows that $p_{s_j} < p_{b_i}$, and so we can transfer one unit of trade from (s_j, b_k) to (s_j, b_i), which improves the surplus, contradicting the optimality of M_{opt}. Thus, b_i's bid must be fully satisfied. An identical argument works for the sellers. □

Lemma 2 *In a surplus-maximizing matching M_{opt} between B and S, suppose b_k is the lowest price buyer with non-zero trade, and s_ℓ is the highest price seller with non-zero trade. Then, $p_{b_k} > p_{s_\ell}$.*

PROOF. By contradiction. We can write the surplus generated by M_{opt} as

$$P_{opt} \;=\; \sum_{i=1}^{k} q'_i \; p_i \;-\; \sum_{j=1}^{l} q'_j \; p_j .$$

The remaining buyers and sellers do not participate in trading, and therefore generate no surplus. Suppose, without loss of generality, that $q_{b_k} \geq q_{s_\ell}$. We rearrange the terms in the above equation as

$$P_{opt} = \sum_{i=1}^{k-1} q_i \; p_i - \sum_{j=1}^{\ell-1} q_j \; p_j + p_{b_k}(q_{b_k} - q_{s_\ell}) + q_{s_\ell}(p_{b_k} - p_{s_\ell}).$$

This follows because, by the preceding lemma, $q'_i = q_i$ for the first $k-1$ buyers and first $\ell-1$ sellers. Now, if $p_{b_k} \leq p_{s_\ell}$ holds, then we can improve the surplus by modifying the matching to one where s_ℓ's trade is reduced to zero, and b_k's trade is reduced to $q_{b_k} - q_{s_\ell}$. But that contradicts the optimality of M_{opt}. □

Next we introduce the following definition of *price inversion.* We say that a matching M, where b_k is the lowest price buyer with non-zero trade, and s_ℓ is the highest price seller with non-zero trade, satisfies the *price inversion* if

- If buyer b_k's bid is partially satisfied, then $p_{b_k} < p_{s_{\ell+1}}$.
- If seller s_ℓ's bid is partially satisfied, then $p_{b_{k+1}} < p_{s_\ell}$.
- If neither b_k nor s_ℓ is partial, then $p_{b_{k+1}} < p_{s_{\ell+1}}$.

We have the following result.

Lemma 3 *Any matching M between buyers B and seller S that satisfies the bid contiguity and price inversion conditions is a surplus-maximizing matching, and it has at most one partially satisfied bid.*

PROOF. We begin by showing that M maximizes surplus. Let M_{opt} be an optimal matching, with maximum surplus, and suppose b_i is the lowest price buyer with non-zero trade, and s_j is the highest price seller with non-zero trade in M_{opt}. Since M_{opt} maximizes surplus, by Lemma 1, it satisfies bid contiguity. The matching M satisfies bid contiguity by assumption. Thus, the bids of buyers $b_1, \ldots, b_{k-1}$ and sellers $s_1, \ldots, s_{\ell-1}$ are fully satisfied in M, and, similarly, the bids of buyers $b_1, \ldots, b_{i-1}$ and sellers $s_1, \ldots, s_{j-1}$ are fully satisfied in M_{opt}. Since M_{opt} is optimal, it must be the case that $i \geq k$ and $j \geq l$, which by the sorted order implies that

$$p_{b_i} \leq p_{b_k} \qquad \text{and} \qquad p_{s_j} \geq p_{s_\ell}.$$

Of course, if the bids of both b_k and s_ℓ in M are fully satisfied, then it must be the case that $i = k$ and $j = l$, meaning that $M = M_{opt}$; this follows because price inversion $p_{b_{k+1}} < p_{s_{\ell+1}}$ occurs, and so there can't be any other trades past (s_ℓ, b_k). So, without loss of generality, assume that the bid of seller s_ℓ is partially satisfied. We get from the price inversion that $p_{b_{k+1}} < p_{s_\ell}$. If $b_i \neq b_k$, then we must have $p_{s_\ell} > p_{b_i}$, because $i > k$ and prices are descending in the buyer sequence. Next, since $j \geq l$ and the prices are increasing in the seller sequence, we must have that $p_{s_j} > p_{b_i}$. But this contradicts the claim of Lemma 2. Therefore, we must have $b_i = b_k$ and $s_j = s_\ell$, which immediately implies that the matching M is optimal.

Finally, it remains to show that M has at most one partial bid. By Lemma 1, the only candidates for partial bids are seller s_ℓ and buyer p_k. By Lemma 2, however, the buyer b_k's unit price is bigger than the seller s_ℓ's unit price, and if both were partial, additional trade would be possible. Therefore, at most one of s_ℓ and b_k can be partial. This completes the proof. □

We now show that the matching output by the algorithm **Single-Item-Max-Surplus** satisfies both the bid contiguity and price inversion, as follows. The bid contiguity follows because we scan the bids in order (descending for buyers, ascending for sellers), and a new bid is considered only if the previous bid is fully satisfied. The price inversion follows because as long as the seller price is smaller than the buyer price, namely, $p_s < p_b$, we set up exchange between s and b. The algorithm stops only when one of the following three conditions occurs: (i) b_k is fully satisfied but $p_{s_\ell} > p_{b_{k+1}}$, so next buyer is ineligible; or (ii) s_ℓ is fully satisfied but $p_{s_{\ell+1}} > p_{b_k}$, so next seller is ineligible; or (iii) b_k and s_ℓ are both fully satisfied and $p_{s_{\ell+1}} > p_{b_{k+1}}$, and so no more exchange among remaining buyers and sellers is possible. But these conditions precisely state the price inversion condition. We therefore have the following result.

Theorem 2 *Given n buyers and sellers in a single-item, multi-unit exchange, we can determine a surplus maximizing matching where at most one bid is partially satisfied in time $O(n \log n)$.*

It is also easy to see that the partially satisfied bid in our matching belongs to either the buyer with the lowest (per unit) bid price, or to the seller with the highest (per unit) ask price.

Trade Volume Maximization

We now consider a matching between buyers and sellers that maximizes the number of units traded. We require that the matching must produce a non-negative surplus; otherwise, the problem is trivial. Recall that the trade volume generated by a matching M is $V_M = \sum_{i \in B} q'_i$, where q'_i is the number of units bought by buyer i. We begin with a simple algorithm that achieves optimal trade volume *but it may result in an unbounded number of partial bids,* as shown in Fig. 2. We then show how to use a sequence of swaps that reduces the number of partial bids to two without reducing the trade volume.

Single-Item-Max-Trade

1. Sort the buyers and sellers in descending order of per unit price.
2. Initially all buyers and sellers are *unmarked*.
3. While there is at least one unmarked buyer and one unmarked seller, repeat the following steps:
   - Let b and s be the first unmarked buyer and seller.
   - If $p_b < p_s$, then mark seller s.
   - Else if $q_b \leq q_s$, then mark b, set $q_s = q_s - q_b$, and $exch(s, b) = q_b$.
   - Else if $q_b > q_s$, then mark s, set $q_b = q_b - q_s$, and $exch(s, b) = q_s$.

It is easy to see that this algorithm produces at most one buyer with partially satisfied bid, but the number of *partial sellers can be arbitrarily large,* as shown by the example in Figure 2. In order to prove that this algorithm realizes the maximum possible trade volume, we first show that it is

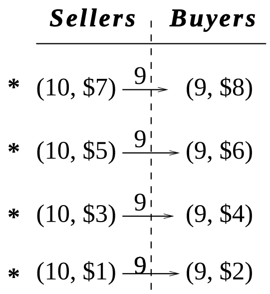

Figure 2: Illustration of the algorithm **Single-Item-Max-Trade**. All the sellers are partially satisfied.

sufficient to restrict our attention to "non-crossing" matchings. More precisely, consider an embedded bipartite graph, where the seller nodes are listed in descending price order along the line $x = 0$, and the buyer nodes are listed in descending price order along the line $x = 1$. Given a matching m, put an edge between buyer b and seller s if $exch(b, s) \neq 0$. We say that M is non-crossing if the bipartite graph of M does not contain any crossing edges.

Lemma 4 *There exists a trade-maximizing non-crossing matching.*

PROOF. Given a matching M, let the cross-trade of M be the number of units corresponding to edges that are crossed by some other edge. We want to show that there is a matching with cross-trade zero. For the sake of contradiction, assume that M is a trade-maximizing matching with the smallest cross-trade. Consider the highest price buyer b_1 that has a crossing edge incident to it in the bipartite graph. Consider two trading pairs (b_1, s_1) and (b_2, s_2) that are crossing in the bipartite graph of M. Assume, without loss of generality, that $p_{b_1} \geq p_{b_2}$, and therefore $p_{s_2} \geq p_{s_1}$. By definition of feasible trade, we have $p_{b_2} \geq p_{s_2}$, and therefore $p_{b_1} \geq p_{s_2}$. Similarly, $p_{b_2} \geq p_{s_1}$. We can "uncross" one unit of trade, where s_1 sells to b_2 and s_2 sells to b_1. This does not decrease the total trade volume, but reduces the cross-trade by at least one unit, since the unit traded between b_1 and s_2 is not crossed by any other edge. This contradicts the choice of M. Thus, M must be non-crossing. □

Lemma 5 *The matching found by* **Single-Item-Max-Trade** *is trade-maximizing.*

PROOF. By the preceding lemma, it is sufficient to show that among all non-crossing matchings, the algorithm find the one with maximum trade. We prove this by contradiction. Suppose there is a matching M' with $V_{M'} > V_M$, where M is the output of our algorithm. Let s_i be the smallest indexed seller such that the total number of units sold by $\{s_1, s_2, \ldots, s_i\}$ in M is smaller than the units sold in M'. Thus, the bid of s_i must be partially satisfied. But this can only happen in our algorithm if all buyers b_j with unit price $p_{b_j} \geq p_{s_i}$ are completely satisfied. Since these are the only buyers that can trade with sellers in $\{s_1, s_2, \ldots, s_i\}$, we have a contradiction that sellers in $\{s_1, s_2, \ldots, s_i\}$ realize a larger trade in M' than in M. □

We now show how to execute a sequence of swaps on the matching returned by the algorithm **Single-Item-Max-Trade** so that the trade volume remains the same but the number of partial bids can be reduced to at most two. We pointed out earlier that the number of partial buyer bids is at most one, and we will show that a suitable sequence of swaps can reduce the number of partial seller bids to at most one as well.

Theorem 3 *Given n buyers and sellers in a single-item, multi-unit exchange, we can determine a trade-maximizing matching where at most one seller bid and one buyer is partially satisfied in time $O(n \log n)$.*

PROOF. We first compute the trade-maximizing matching using the algorithm **Single-Item-Max-Trade**. This takes $O(n \log n)$ time, and the matching has at most one buyer partial bid, but possibly many seller partial bids. Suppose there are at least two partially satisfied sellers. In the descending price sorted order of sellers, let s_ℓ be the last seller whose bid is partially satisfied, and let s_i be the seller with partial bid before s_ℓ. We use the notation q_s and q'_s to denote the quantity bid and quantity traded by seller s. Thus, $q'_{s_\ell} < q_{s_\ell}$ and $q'_{s_i} < q_{s_i}$.

Suppose b_j is a buyer such that $exch(s_i, b_j) \neq 0$. Let $\delta_\ell = q_{s_\ell} - q'_{s_\ell}$ be the number of unfulfilled units for the seller s_ℓ. If $exch(s_i, b_j) > \delta_\ell$, then set $exch(s_i, b_j) = exch(s_i, b_j) - \delta_\ell$, and $exch(s_\ell, b_j) = exch(s_\ell, b_j) + \delta_\ell$. Otherwise, set $exch(s_\ell, b_j) = exch(s_\ell, b_j) + exch(s_i, b_j)$, and set $exch(s_i, b_j) = 0$. In the former case, seller s_ℓ's bid is fully satisfied, and in the latter, the seller s_i's bid is no longer accepted, thus the number of partial bids reduces by one, while the trade volume remains unchanged. We repeat this process until at most seller bid is partially satisfied.

The time complexity of the algorithm is $O(n)$ after the initial sorting of buyers and sellers, because in each step at least one partial bid is eliminated in constant time. □

Exchanges with XOR-Bids

Combinatorial exchanges where items have *substitutability* (*negative complementarity*) or where units have *decreasing marginal utility* require a more expressive form of bidding. For instance, if a buyer values object A at \$10, object B at \$5, but the bundle $\langle A, B \rangle$ at only \$12, then with straightforward method of placing all three bids is inadequate. A seller can sell items A and B to a buyer and expect the payment of \$15, arguing he sold the items separately. A more expressive form of bidding using XORs is needed (Sandholm 2002a; Fujishima, Leyton-Brown, & Shoham 1999; Sandholm 2002b; Nisan 2000). In our example, the buyer will like to bid $(A, \$10) \oplus (B, \$5) \oplus (\langle A, B \rangle, \$12)$.

The XOR-bids are also used for volume discount in multi-unit exchanges, such as a bid $(10, \$10) \oplus (25, \$8) \oplus (100, \$5)$, indicating a desire to buy larger quantities as the per unit price drops. Of all the bids that are linked by XORs, at most one can be accepted. Even for single-item multi-unit exchanges, with or without XORs, the problem of maximizing the surplus or trade volume is NP-Complete; for instance, both the knapsack and subset sum are special cases. But can allowing some partial bids help, as we found to be the case without XORs? Surprisingly, the answer turns out to be negative. We show that even for a single-item exchange, determining the maximum surplus matching is NP-Complete even allowing (any number of) partial bids. Our reduction is from the *subset sum problem*.

Consider an instance of the subset sum problem: $(a_1, a_2, \ldots, a_n, Z)$, where a_i's and Z are positive integers, and the problem is to decide if there is a subset of a_i's that sum to exactly Z. We create an instance of the one-sided single-item exchange, with a single seller (the auctioneer) who has $n + Z$ units to sell, and n bidders. The bidder i places an XOR-bid of the form

$$(1, \$A) \oplus \left(a_i + 1,\ \$\left(\frac{A + a_i}{a_i + 1}\right)\right),$$

where $A \geq \max\{a_1, a_2, \ldots, a_n\}$. (Recall that the bid price is *per unit* price.) We assume that the auctioneer's selling price is zero, and thus maximizing the surplus is the same as maximizing the auctioneer's revenue. A feasible matching is one where at most one bid in each XOR pair is accepted, but it can be accepted partially.

Our reduction will show that this exchange yields a revenue $\geq\ nA + Z$ if and only if the subset sum problem has a solution. Furthermore, we show that for this particular exchange, partial bids do not help, in fact, *they worsen* the revenue. In order to argue this more formally, let us call a solution of the exchange problem *strictly integral* if it contains no partially satisfied bid, and *fractional* if at least one bid is accepted partially. We prove the following theorem. For convenience, for each bidder i, let us call the $(1, \$A)$ bid his *high bid*, and the $(a_i + 1,\ \$(A + a_i)/(a_i + 1)$ bid his *low bid*.

Theorem 4 *Suppose that the subset sum problem $(a_1, a_2, \ldots, a_n, Z)$ has a solution. Then, the associated exchange problem has a strictly integral solution with surplus $\geq\ nA + Z$, while every fractional solution of it has surplus $<\ nA + Z$. If the subset sum problem does not have a solution, then every solution of the exchange problem (integral or fractional) has surplus $<\ nA + Z$.*

PROOF. If the subset sum problem has a solution, then let I be the index set of elements that sum to Z. We obtain a strictly integral solution with surplus $nA + Z$, as follows. The auctioneer sells $a_i + 1$ units to each bidder $i \in I$, and one unit to each remaining bidder. The total number of units sold is $\sum_{i \in I}(a_i + 1) + (n - I) = Z + I + n - I = n + Z$. The total surplus is $\sum_{i \in I}(A + a_i) + (n - I)A = nA + Z$. We now show that every fractional solution is worse.

The argument depends on two observations: (i) in an optimal matching, there can be at most one partial bid; and (ii) if buyer i's low bid is *fully* satisfied, then we can interpret that bid as "\$A for for the first unit, and \$1 for each of the remaining a_i units." (Note that this interpretation is *incorrect* if the low bid is *partially satisfied*.) An argument for the first observation is as follows. In an optimal matching, all high bids must be fully satisfied, because there are $n + Z$ units, and only n high bids, and the per unit price of any high bid is strictly larger than that of any low bid. Secondly, if there two partially satisfied low bids, then we can improve the solution by transferring units from the buyer with lower per unit price to the buyer with higher per unit price. The second observation follows from elementary algebra.

Now, let us consider a fractional solution of the exchange. Let i be a buyer whose bid is partially satisfied. We first argue that the partial bid must be the low bid of i. Otherwise, we can raise the surplus by increasing the partial high bid of i by ϵ and reducing any other buyer's *low* bid by the same amount. Since the unit price of a high bid is strictly greater than the unit price of every low bid, this increases the surplus. Observe that there must be at least one buyer whole low bid is satisfied, since otherwise the total surplus can't exceed nA.

Thus, the partial bid is a low bid, and suppose that buyer i received $m > 1$ units for this bid. By Observation (i) above, the bids of the remaining $n - 1$ buyers are fully satisfied. Thus, each of the remaining $n - 1$ buyers gets at least one unit. In addition, these buyers have at most $(n + Z) - (n - 1) - m = Z - m + 1$ additional units among them. Any bidder that receives strictly more than one unit has his *low* bid satisfied. Now, by Observation (ii), these $Z - m + 1$ units generate a surplus of \$1 each. Thus, the total surplus by the $n - 1$ buyers is at most

$$(n - 1)A + (Z - m + 1).$$

The ith bidder's bid is partially satisfied, and therefore he generates the surplus of $m(A + a_i)/(a_i + 1)$. Thus, for the fractional solution under consideration, the surplus is at most

$$\begin{aligned} & (n-1)A + (Z - m + 1) + m\left(\frac{A + a_i}{a_i + 1}\right) \\ \leq\ & (n-1)A + Z + \left((1 - m) + m\left(\frac{A + a_i}{a_i + 1}\right)\right) \\ <\ & (n-1)A + Z + A \\ =\ & nA + Z \end{aligned}$$

This proves that every fractional solution of the exchange has surplus $<\ nA + Z$. Finally, if the subset sum does not have a solution, then it is easy to see that both an integral and a fractional solution must have value $<\ nA + Z$. □

The preceding theorem shows that determining a surplus-maximizing matching in a single-item exchange is NP-Complete *even if (any number of) partial bids are permissible.*

Conclusions

We investigated the problem of matching buyers and sellers in a combinatorial exchange with the objective of maximizing either the surplus or the number of units traded. These problems are NP-Complete, and prior work has focused on tree search algorithms, approximation schemes, and restricted classes of exchanges. In this paper, we took a new approach, and showed that even multi-item multi-unit exchanges can be solved optimally if we allow a small number of bids to be *partially* satisfied. Our optimal solution has value equal to the optimal solution that is obtainable if all bids were partially acceptable, which is generally higher (and never lower) than the value when bids have to be entirely accepted or rejected. The polynomial solution time should be contrasted with the fact that, in general, combinatorial exchanges are at least as hard as weighted set packing, which cannot even be approximated within a factor better than $\Omega(n^{1-\varepsilon})$ (Håstad 1999).

The bidding language that simply allows bidders to bid on multiple (nonexclusive) combinations enables bidders to express complementarity between items, but not substitutability (the value of a bundle being less than the sum of the parts). We show that if the bidding language is enriched with XOR-constraints between bids (a common, necessary and sufficient enrichment that achieves full expressiveness by the bidders), then allowing for partial acceptance of bids does not help—even for single item exchanges, computing the optimal matching remains NP-Complete no matter how many partial bids are allowed.

To address this difficulty, one interesting avenue of future research includes studying whether a market can be cleared if only restricted substitutability (and unrestricted complementarity through bundle bidding) is allowed. For example, to achieve polynomial-time clearability, how many bids need to be accepted partially if substitutability only comes in the form of a capacity constraint from each seller and/or a budget constraint from each buyer? Our LP hyperplane-based analysis approach can easily be extended to answer questions of this type.

Acknowledgments

This work was supported in part by the National Science Foundation under ITR grant IIS-0121678. Kothari and Suri were also supported in part by National Science Foundation grants IIS-0121562 and CCR-9901958. Sandholm was also supported in part by National Science Foundation CAREER Award IRI-9703122, and grants IIS-9800994 and ITR IIS-0081246.

References

Fujishima, Y.; Leyton-Brown, K.; and Shoham, Y. 1999. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, 548–553.

Håstad, J. 1999. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica* 105–142.

Hoos, H., and Boutilier, C. 2000. Solving combinatorial auctions using stochastic local search. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 22–29.

Lehmann, D.; O'Callaghan, L. I.; and Shoham, Y. 1999. Truth revelation in rapid, approximately efficient combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, 96–102.

Nisan, N. 2000. Bidding and allocation in combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, 1–12.

Penn, M., and Tennenholtz, M. 2000. Constrained multi-object auctions and b-matching. *Information Processing Letters* 75(1–2):29–34.

Rassenti, S. J.; Smith, V. L.; and Bulfin, R. L. 1982. A combinatorial auction mechanism for airport time slot allocation. *Bell J. of Economics* 13:402–417.

Rothkopf, M. H.; Pekeč, A.; and Harstad, R. M. 1998. Computationally manageable combinatorial auctions. *Management Science* 44(8):1131–1147.

Sandholm, T., and Suri, S. 2000. Improved algorithms for optimal winner determination in combinatorial auctions and generalizations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 90–97.

Sandholm, T.; Suri, S.; Gilpin, A.; and Levine, D. 2001. CABOB: A fast optimal algorithm for combinatorial auctions. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, 1102–1108.

Sandholm, T.; Suri, S.; Gilpin, A.; and Levine, D. 2002. Winner determination in combinatorial auction generalizations. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. Early version appeared at the AGENTS-01 Workshop on Agent-Based Approaches to B2B, pp. 35–41, Montreal, Canada, May 2001.

Sandholm, T. 2002a. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence* 135:1–54. First appeared as an invited talk at the First International Conference on Information and Computation Economies, Charleston, SC, Oct. 25–28, 1998. Extended version appeared as Washington Univ., Dept. of Computer Science, tech report WUCS-99-01, January 28th, 1999. Conference version appeared at the International Joint Conference on Artificial Intelligence (IJCAI), pp. 542–547, Stockholm, Sweden, 1999.

Sandholm, T. 2002b. eMediator: A next generation electronic commerce server. *Computational Intelligence*. Special issue on Agent Technology for Electronic Commerce. To appear. Early versions appeared in the Conference on Autonomous Agents (AGENTS-00), pp. 73–96, 2000; AAAI-99 Workshop on AI in Electronic Commerce, Orlando, FL, pp. 46–55, July 1999; and as a Washington University, St. Louis, Dept. of Computer Science technical report WU-CS-99-02, Jan. 1999.

Tennenholtz, M. 2000. Some tractable combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.