Market Clearability*

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

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

Abstract

Market mechanisms play a central role in AI as a coordination tool in multiagent systems and as an application area for algorithm design. Mechanisms where buyers are directly cleared with sellers, and thus do not require an external liquidity provider, are highly desirable for electronic marketplaces for several reasons. In this paper we study the inherent complexity of, and design algorithms for, clearing auctions and reverse auctions with multiple indistinguishable units for sale. We consider settings where bidders express their preferences via price-quantity *curves*, and settings where the bids are price-quantity *pairs*. We show that markets with piecewise linear supply/demand curves and *non-discriminatory pricing* can always be cleared in polynomial time. Surprisingly, if *discriminatory pricing* is used to clear the market, the problem becomes $\mathcal{NP}$-Complete (even for step function curves). If the price-quantity curves are all *linear*, then, in most variants, the problem admits a poly-time solution even for discriminatory pricing. When bidders express their preferences with price-quantity pairs, the problem is $\mathcal{NP}$-Complete, but solvable in pseudo-polynomial time. With free disposal, the problem admits a poly-time approximation scheme, but no such approximation scheme is possible without free disposal. We also present pseudo-polynomial algorithms for XOR bids and OR-of-XORs bids, and analyze the approximability.

1 Introduction

Market mechanisms play a central role in AI for several reasons. First, they provide a tool for resource and task allocation in multiagent systems where the agents may be self-interested. Second, AI techniques can be used to clear markets. For example, there has been a recent surge of research in the AI community on search algorithms [Sandholm, 1999; Fujishima *et al.*, 1999; Sandholm and Suri, 2000] and special-case polynomial algorithms [Tennenholtz, 2000] for clearing combinatorial auctions. Third, recent electronic commerce server prototypes such as eMediator [Sandholm, 2000] and AuctionBot [Wurman *et al.*, 1998] from academic AI groups have led to the uncovering of a need for fast clearing algorithms for a vast space of market designs.

In this paper we analyze the inherent complexity of, and design algorithms for, clearing auctions and reverse auctions in the ubiquitous setting where there are multiple indistinguishable units of an item for sale.

In the largest securities markets where there are multiple units of each item (e.g., stock) for sale, there usually are liquidity providers (*market makers* on the NASDAQ and a *specialist* on the NYSE) that carry inventory, and guarantee that trades are possible at essentially any quantity. Therefore, direct matching between buyers and sellers is not absolutely necessary. However, we argue that if technically possible, it would be highly desirable to construct market clearing algorithms that directly match buyers and sellers, and do not rely on an external liquidity provider, for several reasons. First, most ecommerce marketplaces do not have external liquidity providers. Second, liquidity providers incur operating cost, and need to be compensated. This compensation tends to be paid (implicitly) by the market participants. Third, if the items that are being traded are not securities, the SEC does not impose or monitor rules on liquidity-provisioning parties. Finally, even in markets where the liquidity providers are regulated, they frequently violate the regulations (the most famous recent cases are from the NASDAQ).

We study the possibility of algorithms for accomplishing this in the context of price-quantity curves first, and price-quantity pairs second. We analyze markets with and without free disposal of units. We also uncover the complexity implications of non-discriminatory vs. discriminatory pricing.

2 Auctions with Demand Curve Bids

We consider the auction setting where each bidder submits a *demand curve* indicating the quantity $q(p)$ he will accept at each *unit price* p. If his bid is cleared at price p, he receives $q(p)$ units, for a total price of $p \cdot q(p)$. Recently, several computational markets have been built that use piecewise linear [Sandholm, 2000] or step function [Lupien and Rickard, 1997; Sandholm, 2000] demand curves.

We focus mainly on piecewise linear curves because they can approximate any curve arbitrarily closely, and because their complexity (in the sense of measuring the length of the input) can be characterized systematically. In order to keep our discussion simple, we use a single parameter k to denote the complexity of the piecewise linear curves. That is, k is the *largest* number of pieces in any bidder's demand curve.

We begin with an elementary lemma, which will be used repeatedly in the following discussion. For now, in order to keep the discussion simple, let us assume that the seller has an infinite supply of units to sell. When the number of units available is finite, the solution follows as an easy corollary of Lemmata 2.1 and 2.2, as is discussed after those lemmata.

Lemma 2.1 *Consider a linear demand curve* $q = ap + b$ *subject to* $p, q \geq 0$.

*This work was funded by, and conducted at, CombineNet, Inc., 311 S. Craig St., Pittsburgh, PA 15213.

- *If $a \geq 0$, then an infinite revenue is achievable.*
- *If $a < 0$, then the revenue is maximized at price $p^* = -\frac{b}{2a}$. The corresponding quantity sold at this price is $q^* = \frac{b}{2}$, and the revenue is $-\frac{b^2}{4a}$.*

PROOF. If the slope is non-negative, the revenue is maximized by setting $p^* = \infty$. The second case is more interesting. The revenue at price p equals $p(ap + b) = ap^2 + bp$. Setting the first derivative with respect to p to zero, we get $2ap = -b$, which yields $p^* = -\frac{b}{2a}$. The second derivative is negative (since $a < 0$), so the revenue is maximized at this p. The values of q and $p \cdot q$ follow easily. □

The demand curve could also have boundaries, meaning that the bidder expresses his preference by specifying the linear demand curve $q = ap + b$ but with explicit bounds on the price. If the curve is restricted to the price range $[p_1, p_2]$, where $p_1 < p_2$, we call it a *bounded linear demand curve*.

Lemma 2.2 *Consider a bounded linear demand curve, $q = ap + b$, restricted to the price range $[p_1, p_2]$.*

- *If $a \geq 0$, then the revenue is maximized at price $p^* = p_2$.*
- *If $a < 0$, then the revenue is maximized either at $p^* = -\frac{b}{2a}$ provided $p_1 \leq -\frac{b}{2a} \leq p_2$, or at that endpoint of the range $[p_1, p_2]$ which is closer to $-\frac{b}{2a}$.*

PROOF. Let us first consider the case of non-negative slope, $a \geq 0$. When $a \geq 0$, we have $q' > q''$ whenever $p' > p''$, and so $p'q' > p''q''$. Thus, the maximum revenue is achieved at the highest permissible price, which is p_2.

When the demand curve has negative slope, Lemma 2.1 tells us that the optimal solution without price boundaries is $p^* = -b/2a$. If this optimal price is within the range $[p_1, p_2]$, then it obviously maximizes the revenue, and we are done. So, let us now assume that $p^* \notin [p_1, p_2]$.

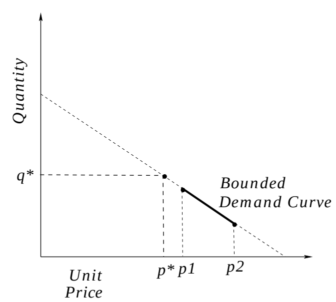

Figure 1: Revenue maximization for a linear demand curve with boundaries.

We consider the effect on revenue of changing the price by an amount ε from the unconstrained optimum $p^* = -b/2a$. (See Figure 1.) Let (p', q'), where $q' = q(p)$, be an arbitrary point on the linear curve $q = ap+b$. Since a is the slope of the demand curve, we note that $\frac{q'-q^*}{p'-p^*} = a$. Let $\varepsilon = p'-p^*$ be the change in the price, and let $\delta = q' - q^*$ be the corresponding change in the quantity. Then, it follows that $\delta = a\varepsilon$. The revenue at point (p', q') is

$$\begin{aligned}
p'q' &= (p^* + \varepsilon) \cdot (q^* + \delta) \\
&= p^*q^* + \varepsilon q^* + \delta p^* + \varepsilon\delta \\
&= p^*q^* + \varepsilon q^* + a\varepsilon p^* + a\varepsilon^2 \\
&= p^*q^* + \varepsilon(b/2) + a\varepsilon(-b/2a) + a\varepsilon^2 \\
&= p^*q^* + a\varepsilon^2
\end{aligned}$$

In line 4 we used the fact that $p^* = -\frac{b}{2a}$ and $q^* = \frac{b}{2}$. Since the demand curve slope a is negative, this shows that a change of ε from the unconstrained optimal price $p^* = -\frac{b}{2a}$ *reduces* the revenue by $|a|\varepsilon^2$. Thus, if the price curve is bound to the range $[p_1, p_2]$, and $p^* \notin [p_1, p_2]$, the maximum revenue is achieved at either p_1 or p_2, whichever is *closer* to p^*. □

In the preceding discussion, we assumed that the seller has an unlimited supply of units. When this supply is bounded by some quantity Q, the optimal solution can be derived easily, as follows. Of course, a feasible solution exists if and only if $Q \geq \min\{q_1, q_2\}$. When the slope is positive, the seller sells $\min\{q_2, Q\}$ units. When the slope is negative and $Q \geq b/2$, revenue is maximized at $q^* = b/2$. But if $Q < b/2$, then the revenue is maximized at $q^* = Q$.

A *piecewise linear* demand curve consists of one or more bounded linear curves. We do not require the demand curve to be continuous (i.e., the quantity can "jump" between pieces). Given a single piecewise linear demand curve bid, we can use Lemma 2.2 on each linear piece separately to determine the revenue-maximizing allocation. To lay the groundwork for our auction-clearing algorithm, we next discuss the problem of *aggregating* the demand over multiple piecewise linear curves.

Demand Curve Aggregation: Consider a set of piecewise linear curves $f_1, f_2, \ldots, f_n$. Their *aggregate curve* is a piecewise linear function $f : R^+ \to R^+$ such that $f(p)$ is the total demand at unit price p. That is, $f(p) = f_1(p) + f_2(p) + \ldots + f_n(p)$, where $f_i(p)$ is the demand by curve i at unit price p. For instance, if the demand curves are linear functions $q = a_i p + b_i$, $i = 1, 2, \ldots, n$, then their aggregate curve is easily shown to be the linear function $q = (\sum_i a_i)p + \sum_i b_i$.

Breakpoints of Aggregate Curve: The aggregate curve f changes only when one of the component curve changes; that is, the breakpoints of f are the union of the breakpoints of the component curves. Thus, given a set of n piecewise linear curves each of which has at most k pieces, their aggregate curve f has at most nk breakpoints.

Given n piecewise linear curves, we can compute their aggregate curve in time $O(nk \log(nk))$, as follows, where k is maximum number of pieces in any curve. Let $z_1, z_2, \ldots, z_L$, where $L \leq nk$, denote the breakpoints of all the component curves, in sorted order. We scan these breakpoints in right to left order (decreasing order of price), and determine the linear aggregate curve between two consecutive breakpoints.

Initially, we compute the linear aggregate function in the range (z_L, ∞), in $O(n)$ time. Next, as we move to the next breakpoint, at most one linear piece changes—one piece may end and another may begin. (If multiple curves begin or end at the same point, we can enforce an artificial order among those, and consider them one at a time.) We can update the linear aggregate by deleting the coefficients of the leaving curve and adding those of the entering curve, and so each update takes $O(1)$ time. Thus, the complete aggregate curve can be determined in time $O(nk)$, after an initial sorting cost of $O(nk \log(nk))$.

While the construction just described should be sufficiently fast for most practical applications, we can build the aggregate curve even faster if only part of the curve is needed.

Lemma 2.3 *Given n piecewise linear demand curves, we can construct the m rightmost pieces of their aggregate curve in time $O(m \log n + n)$, where $m = O(nk)$, and each curve has at most k pieces.*

PROOF. We maintain a priority queue that stores the "next" breakpoint (end, begin, or change) of each component curve. We process events in the order presented by the priority queue: when we delete a breakpoint from the queue, we insert the next linear piece (if any) of the same curve. We initially compute the rightmost piece of the aggregate curve in $O(n)$ time. After that the aggregate curve is updated at each event point in $O(1)$ time. Inserting or deleting an event from the priority queue takes $O(\log n)$ time, and so the total cost to construct m rightmost pieces of the aggregate curve is $O(m \log n + n)$. □

2.1 Auctions with Non-Discriminatory Pricing

Say that the seller wants to auction off Q indistinguishable units of an item. Each of the n bidders submits a piecewise linear demand curve bid. In a *non-discriminatory* auction, the seller determines an optimal price p^* to maximize his revenue, and every buyer pays the same unit price p^*. (The number of items received by bidder i is computed using his demand curve, evaluated at price p^*.)

As Lemma 2.2 shows, if the seller wants to maximize his revenue, he might not sell all the units. We therefore consider the auction both with and without *free disposal*. With free disposal, the seller may choose to keep some units (because he can dispose of them for free), but without free disposal, he must sell all the units.

Theorem 2.4 *Consider a single-item, multi-unit auction with n bidders, each with a piecewise linear demand curve. Under non-discriminatory pricing, the auction can be cleared so as to maximize the seller's revenue in time $O(nk \log(nk))$ with or without free disposal, where k is the maximum number of pieces in any bidder's demand curve.*

PROOF.

- [*Without Free Disposal.*] We construct the aggregate demand curve, incrementally from the right, as described in Lemma 2.3. For each linear piece of the aggregate curve, we check to see if it intersects the supply line $q = Q$. If there is an intersection, then the intersection point is a feasible solution. Since the goal is to maximize seller's revenue, we want the rightmost (highest price) intersection. Thus, we can stop the algorithm as soon as we find an intersection. From this price we can determine the quantities sold to each bidder, using their demand curves. (The problem is clearly infeasible when there is no intersection between the line and the aggregate curve.)
- [*With Free Disposal.*] In this case, we compute the entire aggregate curve, since we cannot stop at the rightmost feasible solution. For each linear piece of the aggregate curve, we compute the maximum feasible revenue and keep track of the optimum found so far.
  1. If the piece lies entirely above the line $q = Q$, no feasible solution exists for this piece.
  2. If the piece is entirely below the line, we take the solution given by Lemma 2.2. In other words, we compute the unconstrained optimum p^* for this linear curve. If p^* is within the price bounds of the piece, we take that solution; otherwise, we choose the endpoint of the linear piece whose price is closer to p^*.
  3. If the piece intersects the line $q = Q$, we compute the unconstrained optimum q^*. If $Q > q^*$, we take the unconstrained solution; otherwise, we sell Q units.

  Each of the three cases takes constant time to evaluate. Thus the complexity is dominated by the time to build the aggregate curve, which is $O(nk \log(nk))$. □

2.2 Auctions with Discriminatory Pricing

In a *discriminatory* price auction, the seller determines for each buyer j a distinct unit price p_j with the objective of maximizing his revenue $\sum_j p_j q_j$ subject to the supply constraint $\sum_j q_j \leq Q$. The quantity q_j sold to buyer j is determined from j's demand curve at price p_j. For a fixed set of demand curve bids, the seller's revenue in a discriminatory auction is generally higher (and never lower) than in a non-discriminatory auction, but the latter offers a stronger notion of fairness among bidders. Discriminatory price auctions however do offer a weak form of *ex ante* fairness: they are anonymous in the sense that had two players swapped their bids, their allocations would also have been swapped.

Intractability under Piecewise Linear Demand Curves
In sharp contrast to a non-discriminatory auction, we show that clearing a discriminatory auction with piecewise linear demand curves is $\mathcal{NP}$-Complete. In fact, this complexity jump occurs even for the simplest piecewise linear demand curve, a *step function*.

Step Function Demand Curve: A step function demand curve is defined by a tuple (p_i, q_i), indicating a buyer's willingness to buy q_i units at or below the unit price p_i; the buyer is not willing to buy any units at price strictly greater than p_i.

Theorem 2.5 *Consider a single-item, multi-unit auction with n bidders, each making a step function demand curve bid. Then the problem of determining a revenue-maximizing allocation using discriminatory pricing, is $\mathcal{NP}$-Complete. This holds with or without free disposal.*

PROOF.

- [*With Free Disposal.*]

  We reduce the **knapsack** problem to our auction problem. Let $\{(s_1, v_1), (s_2, v_2), \ldots, (s_n, v_n),\ K\}$ be an instance of the knapsack problem—K is the knapsack capacity, s_i and v_i, respectively, are the size and value of item i. The goal is to choose a subset of items of maximum value with total size at most K. We create an instance of the single-item multi-unit auction using step function demand curves, as follows. Bidder i places a step function bid $(v_i/s_i, s_i)$, meaning he is willing to buy s_i units at lot price v_i (or maximum unit price v_i/s_i), and no units for a higher price. The total number of units available is K. Since we are using discriminatory pricing, the goal is to choose a subset of bids maximizing the total revenue subject to the total quantity constraint K. Now, it is easy to see that any solution to the auction problem is a solution to the knapsack, and vice versa.

- [*Without Free Disposal.*]

  We reduce the **subset sum** problem [Garey and Johnson, 1979] to the auction problem. In the subset sum problem, we are given a set of integers $X = \{x_1, x_2, \ldots, x_n\}$, an integer K, and the goal is to choose a subset of X whose elements sum to exactly K. We create n bids, where the bidder i places a step function demand curve bid $(\$1, x_i)$—that is, the buyer is willing to pay one dollar per unit for x_i units, but does not accept any other quantity. (Actually, the price is immaterial in this transformation, so we use a default value of \$1.) The total number of units available is K. It is easy to see that the auction without free disposal has a feasible solution if and only if the original subset sum problem has a solution. □

Thus, we conclude that even with step function demand curve bids, or more generally piecewise linear curve bids, the discriminatory auction becomes intractable.

Polynomial Algorithm for Linear Demand Curves

There is an important case of the discriminatory price auction for which we can clear the market in polynomial time. This is the case where all bids are *downward sloping linear demand curves*, that is, in the demand curve $q = ap+b$, we have $a < 0$ and $b \geq 0$. In other words, the buyer's demand decreases linearly as the price increases. We begin by describing the algorithm for auctions with free disposal. Let us suppose that the auctioneer has Q units of the item for sale. Then, the algorithm has the following steps.

1. Let $S = \{1, 2, \ldots, n\}$ denote the index set of bids.
2. Compute the unconstrained optimal solution for each bid independently: $(p_j, q_j) = \left(\frac{-b_j}{2a_j}, \frac{b_j}{2}\right)$.
3. If $\sum_{j \in S} q_j \leq Q$, then these unconstrained price-quantity pairs are the optimal solution.
4. Otherwise, let $\delta = \min_{j \in S} \{p_j\}$, and let ℓ be the index of the bid that achieves this minimum.
5. Set $p'_j = p_j + p_\ell$, and $q'_j = a_j p'_j + b_j$, for $j \in S$. That is, increase each bid's unit price by δ, and determine the new quantity. (Note that $q'_j < q_j$ because a_j is negative.)
6. If $\sum_{j \in S} q'_j \leq Q$, then set $q^*_j = q_j - a_j C$, where $C = (\sum_{j \in S} b_j - 2Q)/(2 \sum_{j \in S} a_j)$, and stop. This is the optimal quantity sold to buyer j, at the corresponding price $p^*_j = p_j - C$. (Note that both a_j and C are negative, and so $q^*_j < q_j$ and $p^*_j > p_j$.)
7. Otherwise, set $S = S - \{\ell\}$, and go to step 4.

The correctness of the algorithm depends on three facts: (a) if the unconstrained solution is quantity-infeasible, then one must increase the unit price *uniformly* for all bids, until either the solution becomes feasible or the price becomes infeasible for some bid, (b) when the price becomes infeasible for some bid, that bid receives zero units in the optimal solution, and (c) in the price-interval where the set of feasible bids remains constant, the formula of line (6) gives the optimal solution. Due to lack of space, we omit the proofs of these claims. Instead, we briefly discuss the time complexity of the algorithm. Step (4) repeatedly chooses a bid with the next smallest unconstrained optimal price p_j. In order to facilitate this choice, we initially sort the bids in increasing order of p_j's. The key step is to determine whether $\sum_{j \in S} q'_j \leq Q$. However, it is easy to see that $\sum_{j \in S} q'_j = \sum_{j \in S} q_j + (\sum_{j \in S} a_j) p_l$. Thus, we can calculate $\sum_{j \in S} q'_j$ in $O(1)$ time, by simply keeping track of the sums $\sum_{j \in S} a_j$ and $\sum_{j \in S} q_j$. These sums only change when a bid is removed from S, and can be maintained in $O(1)$ time per update to S. In summary, once the bids have been sorted, the total cost of running the algorithm is $O(n)$. We summarize this result in the following theorem.

Theorem 2.6 *Consider a single-item, multi-unit auction with n bidders, each making a downward sloping linear demand curve bid. Under discriminatory pricing, the auction can be cleared so as to maximize revenue in $O(n \log n)$ time, with or without free disposal.* [1]

Remark. For completeness, we now discuss the complexity of clearing absurd types of linear demand curves. If all the demand curves are upward sloping (the higher the unit price, the more the bidder will buy), then the optimal solution is easily obtained (with or without free disposal) by selling all Q units to the bidder whose demand line intersects the constant line $q = Q$ at the *highest* unit price. This can be determined in $O(n)$ time by calculating the intersection point for every demand line. Next we consider the case where all demand curves are constant ($a = 0$), i.e., the bidders do not care about price. With free disposal, the seller's revenue can be maximized in $O(n)$ time by choosing any bidder (whose quantity is positive, i.e., $b > 0$), and charging an infinite price. Without free disposal, finding a feasible solution (finding a combination of the constant curves whose quantities sum up to exactly Q) is $\mathcal{NP}$-Complete because that is equivalent to the *subset sum* problem. If a feasible solution is found, the seller can charge an infinite price.

3 Reverse Auctions with Supply Curve Bids

In this section, we consider reverse (or, buyer-based) auctions using an analog of demand curve bidding. These types of auctions are frequently used for Requests for Proposals (RFPs) and Requests for Quotes (RFQs). The buyer posts the items (goods, services, etc.) she wants to purchase, and sellers compete for the business by bidding to sell the items. We assume that a buyer wishes to acquire Q indistinguishable units of a certain item, and each seller submits a *supply curve bid*.

We start by establishing simple properties of *linear supply curves*. As in auctions, we consider two settings: free disposal and no free disposal. Under free disposal, the buyer is willing to accept more than the desired Q units if it leads to lower cost (at worst he can dispose of the extra units for

[1] When free disposal is not allowed, the only change to the algorithm is this: if the initial unconstrained optimal solution is quantity-feasible, meaning $\sum_{j \in S} q_j \leq Q$, instead of stopping, we *decrease* the price uniformly over all bids, until either the total demand reaches Q, or some bid becomes infeasible (its price reaches zero). When a bid becomes infeasible, we remove it from S and continue.

free). Without free disposal, the buyer wants exactly Q units (or none if the solution is infeasible).

Lemma 3.1 *Consider a linear supply curve $q = ap + b$ subject to $p, q \geq 0$. Suppose that the buyer wishes to acquire Q units of the item.*

- *If $a \leq 0$, then there is a feasible solution if and only if $Q \leq b$. If $Q \leq b$, the cost to the buyer is 0 with free disposal, and $Q(Q-b)/a$ without free disposal.*
- *If $a > 0$, then the cost of acquiring the items is $\max\{0,\ Q(Q-b)/a\}$.*

PROOF. When the supply curve has non-positive slope, the maximum number of items that can be purchased is b. Thus, the trade is feasible if and only if $Q \leq b$. If free disposal is allowed, the buyer buys b units at price 0; without free disposal, the buyer pays $(Q-b)/a$ per unit.

If the supply curve has positive slope, the price per unit is increasing with the number of units. If $b > Q$, then the buyer can buy Q units at zero cost. Otherwise, the buyer purchases exactly Q units at the unit cost of $(Q-b)/a$. □

Lemma 3.2 *Consider a bounded linear supply curve, $q = ap + b$, restricted to the price range $[p_1, p_2]$, where $p_1 < p_2$. Suppose the buyer wishes to acquire Q units of the item.*

- *Without free disposal, a feasible solution exists iff $q_1 \geq Q \geq q_2$, and if feasible, the solution has cost $Q(Q-b)/a$.*
- *With free disposal and $a \leq 0$, a feasible solution exists iff $Q \leq q_1$. Set $q_2' = \max\{q_2, Q\}$. If feasible, the solution has cost $\min\{p_1 q_1, p_2' q_2'\}$, where p_2' is the price corresponding to q_2'.*
- *With free disposal and $a > 0$, a feasible solution exists iff $Q \leq q_2$, and in that case the solution has cost $\max\{p_1 q_1, Q(Q-b)/a\}$.*

PROOF. See Figure 2 for illustration. The first case follows easily, since the boundaries of the supply curve dictate that the quantity supplied is in the range $[q_1, q_2]$. If Q lies in this range, the cost equals $Q(Q-b)/a$.

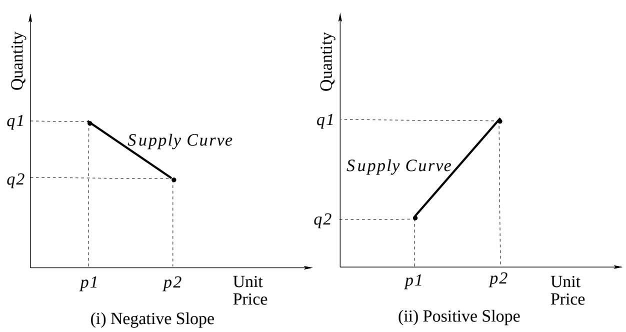

Figure 2: Bounded linear supply curves.

In the second case, since the slope is negative, the maximum quantity supplied is q_1, and so for feasibility, we must have $Q \leq q_1$. Since the buyer admits free disposal, we can purchase any quantity between Q and q_1, in order to minimize the cost. We define $q_2' = \max\{q_2, Q\}$ so that the feasible range of quantity is $[q_1, q_2']$. We now invoke the algebra of Lemma 2.2 to observe that quantity $p \cdot q(p)$ decreases the farther the price deviates from $p^* = -\frac{b}{2a}$. (Recall that we have $a \leq 0$ in this case.) Thus, the cost of acquiring at least Q units in the range $[q_1, q_2']$ is minimized at one of the endpoints of the range. So, we take the smaller of the two values.

Finally, in the third case, the largest quantity supplied by the curve is q_2, so the feasibility condition checks whether $Q \leq q_2$. If $Q \leq q_2$, then we can set $q_1' = \max\{Q, q_1\}$, and our cost is $p_1' q_1'$, where p_1' is the price corresponding to q_1'. □

3.1 Reverse Auctions with Non-Discriminatory Pricing

Consider a reverse auction where a buyer wants Q units of an item. Each of the n sellers submits a piecewise linear supply curve. In a *non-discriminatory* reverse auction, the buyer determines an optimum price p^* to minimize his cost, and sellers supply their share of units at price p^*. (The number of items purchased from seller i is computed using his supply curve, evaluated at price p^*.) The proof of the following theorem is similar to the proof of Theorem 2.4.

Theorem 3.3 *Consider a single-item, multi-unit reverse auctions with n bidders, each making a k-piece supply curve bid. Under non-discriminatory pricing, the auction can be solved so as to minimize cost in time $O(nk \log(nk))$ with or without free disposal.*

3.2 Reverse Auctions with Discriminatory Pricing

In a *discriminatory* reverse auction, the buyer determines for each seller j a distinct price p_j with the objective of minimizing the total cost $\sum_j p_j q_j$ subject to the supply constraint $\sum_j q_j \geq Q$. (Without free disposal, the quantity constraint is an equality.) The problem of clearing such an auction with piecewise linear supply curves again turns out to be $\mathcal{NP}$-Complete.

Theorem 3.4 *A single-item, multi-unit reverse auction with discriminatory pricing and piecewise linear (specifically step function) supply curve bids is $\mathcal{NP}$-Complete to clear so as to minimize cost (with or without free disposal).*

PROOF. The proof *without free disposal* is the same as in Theorem 2.5. The details of the *free disposal* case are different, so we discuss that here. We reduce the *knapsack* problem to the reverse auction, but since knapsack is a value-maximization problem, while the reverse auction is a cost-minimization problem, we need a minor transformation in between.

Let $\{(s_1, v_1), (s_2, v_2), \ldots, (s_n, v_n),\ K\}$ be an instance of the knapsack problem—the knapsack has capacity K, and item i has size s_i and value v_i. Let us create an instance of the single-item multi-unit reverse auction, as follows.

Bidder i places a step function bid (v_i, s_i), meaning he is willing to supply s_i units at lot price v_i (or at any price higher than v_i), and no units for a price less than v_i. Let $T = \sum_i s_i$ be the total number of units in all the bids. We set up the auction so that the buyer wishes to purchase at least $T - K$ units, at minimum possible price. Let S' be the set of bids that are winning bids in the reverse auction. Then, we claim that the remaining bids (those *not* in S') form a solution to the knapsack problem. First, since the bids in S' have a total of at least $T - K$ units, the remaining units are at most K, and so the solution is knapsack feasible. Second, since the bids in S' provide $T - K$ units at least possible price, the total price of the remaining bids is largest possible subject to the knapsack constraint. Thus, the auction problem is $\mathcal{NP}$-Complete. □

In fact, the reverse auction without free disposal is $\mathcal{NP}$-Complete *even with linear downward sloping supply curves*, unlike the seller auction with linear demand curves which is polytime solvable (except in the absurd case of constant demand lines and no free disposal), cf. Theorem 2.6.

Theorem 3.5 *A single-item, multi-unit reverse auction with discriminatory pricing and no free disposal is $\mathcal{NP}$-Complete to clear so as to minimize cost with downward sloping linear supply curve bids.*

PROOF. Consider an instance of the subset sum problem: set of non-negative integers $X = \{x_1, x_2, \ldots, x_n\}$, and an integer K. Corresponding to x_i, we create a linear supply curve $q = -p + x_i$. The subset sum problem has a solution if and only if the reverse auction without free disposal has a solution for quantity Q and cost 0. □

Remark. In the preceding theorem, the price zero is only for convenience. We can easily truncate the supply curves, for example at $p = 1$, to enforce a minimum unit price of one. In that case, the subset sum has solution if and only if the buyer can acquire exactly K units at total price K.

Remark. With free disposal and downward sloping or constant supply lines (that start at $p = 0$), the optimal solution is obtained in $O(n)$ time by accepting all supply lines at $p = 0$. If the aggregate quantity is at least Q, then the solution is feasible and optimal. Otherwise, no solution is feasible.

Remark. With constant supply lines and no free disposal, finding a feasible solution to the reverse auction is $\mathcal{NP}$-Complete because that corresponds to the *subset sum* problem. If a feasible solution exists, and the lines start at $p = 0$, then the optimal solution has zero cost (accept at $p = 0$ each one of the lines that are part of the feasible solution).

Remark. In reverse auctions, downward sloping supply corresponds to quantity discounts. A more common case would be *upward sloping supply*, i.e., the higher the price, the more the supplier is willing to sell. We analyze upward sloping supply lines in another paper [Sandholm and Suri, 2001], showing clearing complexity of $O(n \log n)$—the same complexity as *auctions* with *downward* sloping demand lines.

4 Price-Quantity Pair Bids

In this section we consider auctions where the auctioneer has multiple indistinguishable units of one item to sell, and bidders express their preferences via *price-quantity pairs*.

Definition 1 *In a* price-quantity pair *bid (p, q), the bidder states a price $p \in Z^+$ that he is willing to pay for $q \in Z^+$ units.*[2] *The bid is atomic, meaning it must be accepted as a whole or rejected—it cannot be accepted fractionally.*

Theorem 4.1 (Known) *If the seller has Q units, and each buyer submits one price-quantity bid, then the problem of maximizing revenue is equivalent to the $\mathcal{NP}$-Complete knapsack problem. It is solvable in pseudo-polynomial time $O(n^2V)$ [Garey and Johnson, 1979], where n is the number of bids and V is the maximum price of any bid. It is approximable within a $(1-\varepsilon)$ factor of the optimum in polynomial time $O(n \log \frac{1}{\varepsilon} + \frac{1}{\varepsilon^4})$ [Lawler, 1976].*

[2] In this section we consider the case where p is the price for the entire lot q. If a unit price is stated instead, it can be trivially converted to a lot price by multiplying by q.

If there is *free disposal*, the auctioneer can always sell fewer than Q units because, at worst, he can dispose of the extra units for free. On the other hand, for many real goods there is no free disposal.

Theorem 4.2 *If the seller has to sell exactly Q units (or none if the instance is infeasible), and each buyer submits one price-quantity bid, then even finding a feasible solution is $\mathcal{NP}$-Complete. The problem of maximizing revenue can be solved optimally in pseudo-polynomial time.*

PROOF. Finding a feasible solution is $\mathcal{NP}$-Complete because the special case where $p_i = q_i$ for all i is equivalent to the *subset sum* problem, which is $\mathcal{NP}$-Complete. The problem can be solved in pseudo-polynomial time using a straightforward dynamic program. We omit it due to limited space. □

In many settings, a bidder could accept alternative quantities at different prices. This can be enabled by allowing each bidder to submit multiple price-quantity bids which are combined with XOR.

Theorem 4.3 *If a seller can sell at most Q units, and each buyer submits a set of alternative price-quantity bids (i.e., XOR bids), then the problem is $\mathcal{NP}$-Complete. It can be solved in pseudo-polynomial time $O(n^2VL)$, where n is the number of bidders, L is the maximum number of alternative bids by any buyer, and V is the maximum price of any bid. The problem can also be approximated within a $(1-\varepsilon)$ factor of the optimum in $O(nL \log \frac{1}{\varepsilon} + \frac{1}{\varepsilon^4})$ time.*

PROOF. $\mathcal{NP}$-Completeness follows from Theorem 4.1 since each bidder might only submit one price-quantity pair.

We now devise a pseudo-polynomial algorithm. We label the bidders (arbitrarily), 1 through n. Observe that nV is an upper bound on the revenue. Let $A(i, v)$ denote the smallest number of units that can be sold to bidders in the set $\{1, 2, \ldots, i\}$ with total revenue exactly v.

1. [**Initialize:**] Let the alternative price-quantity bids of buyer 1 be (q_1, p_1) XOR (q_2, p_2) XOR ... (q_j, p_j). Set
   - $A(1, p_t) = q_t$, for $t = 1, 2, \ldots, j$.
   - $A(1, v) = \infty$, for all other values of v.
2. **for** $i = 2$ to n
   for $v = 1$ to nV
   $A(i, v) = \infty$
   if buyer i has bid (q_1, p_1) XOR (q_2, p_2) XOR ... (q_j, p_j), then
   for $t = 1$ to j
   if $p_t \le v$ **then**
   $$A(i, v) = \min \left\{ \begin{array}{l} A(i, v), \\ A(i-1, v), \\ q_t + A(i-1, v-p_t) \end{array} \right\}$$
   else $A(i, v) = A(i-1, v)$.

This algorithm computes a solution in $O(n^2VL)$ time. This pseudo-polynomial algorithm can be converted to an ε-approximation algorithm with running time $O(nL \log \frac{1}{\varepsilon} + \frac{1}{\varepsilon^4})$ using the scheme of Lawler [Lawler, 1976]. □

While XOR bids are fully expressive in the sense that they allow a bidder to express any valuation (mapping from the number of units to a price), in many cases, a more compact (and never less compact) representation of the same valuation can be obtained using the OR-of-XORs bidding language[3]. An

[3] XOR bids and OR-of-XORs bids were originally introduced for combinatorial auctions where there are multiple distinguish-

Market type	Linear curves: Upward sloping	Linear curves: constant	Linear curves: downward sloping	Piecewise linear curves
Nondiscriminatory auction	$O(n)$	$O(n)$	$O(n \log n)$	$O(nk \log(nk))$ *
Discriminatory auction	$O(n)$	fd: $O(n)$ nfd: $\mathcal{NP}$-Complete	$O(n \log n)$ *	$\mathcal{NP}$-Complete *
Nondiscriminatory reverse auction	$O(n \log n)$	$O(n)$	$O(n)$	$O(nk \log(nk))$ *
Discriminatory reverse auction	$O(n \log n)$ * [Sandholm and Suri, 2001]	fd: $O(n)$ nfd: $\mathcal{NP}$-Complete	fd: $O(n)$ nfd: $\mathcal{NP}$-Complete	$\mathcal{NP}$-Complete *

Table 1: Summary of our results on clearing supply/demand curves (fd = free disposal, nfd = no free disposal). The nontrivial results are marked with a "*".

OR-of-XORs bid is a bid where multiple XOR price-quantity bids are offered and any number of these can be accepted (subject to honoring the overall quantity Q):

$$[(q_1, p_1) \text{ XOR } (q_2, p_2) \text{ XOR } \ldots (q_i, p_i)] \text{ OR }$$
$$[(q_{i+1}, p_{i+1}) \text{ XOR } (q_{i+2}, p_{i+2}) \text{ XOR } \ldots (q_j, p_j)] \text{ OR} \ldots$$
$$[(q_k, p_k) \text{ XOR } (q_{k+1}, p_{k+1}) \text{ XOR } \ldots (q_l, p_l)].$$

Theorem 4.4 *If the auctioneer can sell at most Q units, and each bidder submits an OR-of-XORs bid, then the problem can be solved and approximated with the time complexities stated in Theorem 4.3, where n now is the number of XOR-disjuncts submitted overall, and L is the maximum number of bids within any XOR-disjunct.*

PROOF. We can treat different XOR bids from the same bidder as coming from different bidders. Since each bidder can be awarded any number of OR bids, this transformation is sound. We can thus assume that each bidder has submitted only one XOR bid, and solve the problem using the algorithm described in the proof of Theorem 4.3. □

Now, consider XOR bids and OR-of-XORs bids in settings where the auctioneer has to sell exactly Q units. It follows from Theorem 4.2 that finding a feasible solution (and therefore also approximation) is $\mathcal{NP}$-Complete (unlike in the free disposal setting). The problem can be solved in pseudo-polynomial time using a dynamic program akin to the ones above where the auctioneer can keep any of the units (we omit these due to limited space).

5 Conclusions

Market mechanisms play a central role in AI as a coordination tool in multiagent systems, and as an application area for algorithm design. Market mechanisms where buyers are directly cleared with sellers, and thus do not require an external liquidity provider, are highly desirable for electronic marketplaces for several reasons. In this paper we studied the inherent complexity of, and designed algorithms for, clearing auctions and reverse auctions with multiple indistinguishable units for sale.

Table 1 summarizes our results on market clearability under demand/supply curves. Note that in non-discriminatory settings, even when each bidder's curve is linear, the aggregate curve might not be linear but piecewise linear (because the bidder's curves have to be ignored below zero quantity).

In addition to the results summarized in the table, we believe that one of the most surprising result of our paper is the following property of discriminatory auctions: at the unconstrained optimum, each bidder generally gets a different price, but interestingly, to accommodate the constraint of limited supply, each bidder's price is incremented *equally* from the unconstrained optimum.

When bidders express their preferences with price-quantity pairs, the market clearing problem is essentially equivalent to the knapsack problem, and therefore $\mathcal{NP}$-Complete but solvable in pseudo-polynomial time. With free disposal, the problem admits a polynomial-time approximation scheme, but no such approximation scheme is possible without free disposal. We also describe pseudo-polynomial algorithms for XOR bids and OR-of-XORs bids, and their polynomial approximability when free disposal is allowed.

Our $\mathcal{NP}$-completeness results carry over to **exchanges** (where the objective is to maximize surplus, i.e., sum of accepted bids minus sum of accepted asks) directly since auctions and reverse auctions are special cases of exchanges.

Our algorithms help pave the way toward automated electronic markets without external liquidity providers, but at the same time, our $\mathcal{NP}$-Completeness results curtail the space of automated market designs that are computationally tractable.

References

[Fujishima *et al.*, 1999] Yuzo Fujishima, Kevin Leyton-Brown, and Yoav Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *IJCAI*, pages 548–553.

[Garey and Johnson, 1979] Michael R Garey and David S Johnson. *Computers and Intractability*. W. H. Freeman and Company.

[Lawler, 1976] E. L. Lawler. *Combinatorial Optimization: Networks and Matroids*. Holt, Rinehart and Winston.

[Lupien and Rickard, 1997] William A Lupien and John T Rickard. Crossing network utilizing optimal mutual satisfaction density profile. US Patent 5,689,652, granted Nov. 18 to Optimark Technologies.

[Sandholm and Suri, 2000] Tuomas Sandholm and Subhash Suri. Improved algorithms for optimal winner determination in combinatorial auctions and generalizations. In *AAAI*, pages 90–97.

[Sandholm and Suri, 2001] Tuomas Sandholm and Subhash Suri. Computational complexity of clearing exchanges. Technical report, 2001. In progress.

[Sandholm, 1999] Tuomas Sandholm. An algorithm for optimal winner determination in combinatorial auctions. In *IJCAI*, pages 542–547. First appeared: Washington Univ., Dept. of Computer Science WUCS-99-01, Jan. 28th.

[Sandholm, 2000] Tuomas Sandholm. eMediator: A next generation electronic commerce server. In *AGENTS*, pages 73–96, Barcelona, Spain, June 2000. Early version: AAAI-99 Workshop on AI in Electronic Commerce, pp. 46–55, July 1999, and Washington Univ., St. Louis, Dept. of Computer Science WU-CS-99-02, Jan. 1999.

[Tennenholtz, 2000] Moshe Tennenholtz. Some tractable combinatorial auctions. *AAAI*.

[Wurman *et al.*, 1998] Peter R Wurman, Michael P Wellman, and William E Walsh. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. In *AGENTS*, pages 301–308.

able items for sale, and bids can be submitted on combinations of items [Sandholm, 1999; 2000]