

eMediator: A Next Generation Electronic Commerce Server*

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh PA 15213
sandholm@cs.cmu.edu

May 5, 2002

Abstract

This paper presents *eMediator*, an electronic commerce server prototype that demonstrates ways in which algorithmic support and game-theoretic incentive engineering can jointly improve the efficiency of ecommerce. *eAuctionHouse*, the configurable auction server, includes a variety of generalized combinatorial auctions and exchanges, pricing schemes, bidding languages, mobile agents, and user support for choosing an auction type. We introduce two new logical bidding languages for combinatorial markets: the XOR bidding language and the OR-of-XORs bidding language. Unlike the traditional OR bidding language, these are fully expressive. They therefore enable the use of the Clarke-Groves pricing mechanism for motivating the bidders to bid truthfully. *eAuctionHouse* also supports supply/demand curve bidding. *eCommitter*, the leveled commitment contract optimizer, determines the optimal contract price and decommitting penalties for a variety of leveled commitment contracting mechanisms, taking into account that rational agents will decommit strategically in Nash equilibrium. It also determines the optimal decommitting strategies for any given leveled commitment contract. *eExchangeHouse*, the safe exchange planner, enables unenforced anonymous exchanges by dividing the exchange into chunks and sequencing those chunks to be delivered safely in alternation between the buyer and the seller.

1 Introduction

Electronic commerce is taking off rapidly, but the full power of algorithmic support and game-theoretic tools has not been harnessed to improve its efficiency. This paper presents *eMediator*, a next generation electronic commerce server that demonstrates ways in which these techniques can improve ecommerce both in terms of processes and outcomes. The server is designed to be used by both human and automated agents. The result of our 4-year design and implementation effort is now available for use on the web at <http://www.cs.cmu.edu/~amem/eMediator>. The three components of *eMediator* are an auction house, a leveled commitment contract optimizer, and a safe exchange planner. These are discussed in Sections 2, 3, and 4 respectively. Each one exhibits interesting interplay between algorithms (such as search) and game-theoretic incentive engineering.

*Part of this work was supported by the National Science Foundation under CAREER Award IRI-9703122 and Grants IRI-9610122, IIS-9800994, ITR IIS-0081246, and ITR IIS-0121678. Early versions of this paper appeared at the International Conference on Autonomous Agents (AGENTS-00), pp. 73–96, Barcelona, Spain; in the AAAI-99 Workshop on AI in Electronic Commerce, Orlando, FL, pp. 46–55, 1999; and as a Washington University, St. Louis, Dept. of Computer Science technical report WU-CS-99-02, Jan. 1999.

2 *eAuctionHouse*: A configurable auction

One of the services that *eMediator* provides is a free-to-use Internet auction and exchange prototype called *eAuctionHouse*. Several successful commercial Internet auction sites already exist—such as eBay and Yahoo—and interesting academic auction houses have recently appeared on the Internet [25, 45]. Our motivation in developing an auction server was to prototype novel next generation features, and test their feasibility both computationally and in terms of user comfort. *eAuctionHouse* allows users from across the Internet to buy and sell goods as well as to set up markets. It is a third party site, so both sellers and buyers can trust that it executes the market mechanisms as stated. It is implemented in Java, with some of the computationally intensive algorithms in C++. The information about the markets is stored in a relational database to increase reliability. *eAuctionHouse* has been in operation on the web since December 1998, and to our knowledge, was the first Internet auction to support combinatorial auctions, bidding via graphically drawn price-quantity graphs, and by mobile agents. *eAuctionHouse* also offers a wide range of market types to be chosen from, and supports the user in choosing the appropriate market type for the setting at hand. These features are now discussed in order.

2.1 Combinatorial auctions

In a *sequential auction*, items are auctioned one at a time. If a bidder has preferences over bundles, that is, combinations of items (as is often the case for example in electricity markets, equities trading, bandwidth auctions [16, 17], and transportation exchanges [27]), then bidding in such auctions is difficult. To determine her valuation for an item, the bidder needs to guess what items she will receive in later auctions. This requires speculation on what the others will bid in the future because that affects what items she will receive. Furthermore, what the others bid in the future depends on what they believe others will bid, etc. This counterspeculation introduces computational cost and other wasteful overhead. Moreover, in auctions with a reasonable number of items, such lookahead in the game tree is intractable, and then there is no known way to bid rationally. Bidding rationally would involve optimally trading off the cost of lookahead against the gains it provides, but that would again depend on how others strike that tradeoff. Furthermore, even if lookahead were computationally manageable, usually uncertainty remains about the others' bids because agents do not have exact information about each others' preferences. This often leads to inefficient allocations where bidders fail to get the combinations they want and get ones they do not [2, 29].

In a *parallel auction* the items are open for auction simultaneously and bidders may place their bids during a certain time period. This has the advantage that the others' bids partially signal to the bidder what the others' bids will end up being for the different items, so the uncertainty and the need for lookahead is not as drastic as in a sequential auction. However, the same problems prevail as in sequential auctions, albeit in a mitigated form.

Combinatorial auctions can be used to overcome the need for lookahead and the inefficiencies that stem from the related uncertainties [17, 24, 26, 27]. In a combinatorial auction bidders may place bids on combinations of items. This allows the bidders to express complementarities between items instead of having to speculate into an item's valuation the impact of possibly getting other, complementary items. This capability is particularly important in illiquid, highly volatile, or non-commoditized markets where it is unsure whether one can acquire the items of a desired bundle one at a time (for the anticipated price).

While combinatorial auctions can avoid the need for lookahead by the bidders and tend to therefore lead to more efficient allocations, they impose significant complexity on the auctioneer because the auctioneer needs to determine the winners. This is a nontrivial task. For example, the Federal Communications Commission (FCC) saw the desirability of combinatorial bidding in their bandwidth auctions, but it was not allowed due to perceived intractability of winner determination (among other reasons). (This is about to change given the speed improvement in winner determination algorithms. For one, the FCC is planning to hold its first combinatorial spectrum auction in summer 2002.)

Our auction server supports a variety of combinatorial auctions. The following subsections discuss some of the main variants.

2.1.1 The OR bidding language

In the combinatorial auction setting that has been most commonly discussed [26], each bidder can bid on combinations of indivisible items, and her bids are implicitly joined with non-exclusive OR, meaning that any number of her bids can be accepted. Winner determination in this case is the problem of deciding which bids win so as to maximize the sum of the bid prices, under the constraint that every item is allocated to at most one bid. This cannot be done in general in polynomial time in the size of the input (unless $\mathcal{P} = \mathcal{NP}$) because the problem is \mathcal{NP} -complete [26] (it is analogous to weighted set packing which is \mathcal{NP} -complete [11]).

Recently, we proved that even approximate winner determination is hard if one is interested in worst case guarantees. Specifically, no polynomial time algorithm can guarantee an allocation within a bound $n^{1-\epsilon}$ from optimum for any $\epsilon > 0$ (unless $\mathcal{NP} = \mathcal{ZPP}$) [30].¹ If the bids exhibit special structure, better approximations can be achieved in polynomial time [3, 7, 8, 10, 35], but even these guarantees are so far from optimum that they are irrelevant for auctions in practice [30].

Polynomial time winner determination can be achieved by restricting the combinations on which the agents are allowed to bid [21, 26, 35, 42]. However, because the agents may then not be able to bid on the combinations they want, similar economic inefficiencies prevail as in non-combinatorial auctions.

We recently generated another approach to optimal winner determination [30]. The motivation was to

- allow bidding on all combinations.
- strive for the optimal allocation.
- capitalize heavily on the sparseness of bids. In practice the space of bids is extremely sparsely populated. For example, if there are 100 items, there are $2^{100} - 1$ combinations, and it would take longer than the life of the universe to bid on all of them even if every person in the world submitted a bid per second. Sparseness of bids implies sparseness of the allocations that need to be checked. Our algorithm constructively checks each allocation that has positive value exactly once, and does not construct the other allocations. Therefore, the algorithm only generates those parts of the search space which are actually populated by bids. The disadvantage then is that the run time depends on the bids received.

We achieve these goals by a tree search algorithm that does provably sufficient selective generation of children in the search and by using a method for fast child generation, heuristics that are accurate and optimized for speed, and four methods for preprocessing the search space. While the worst case complexity is exponential in the number of items for sale, it is polynomial in the number of bids. This is desirable because the auctioneer can control the number of items that are for sale in a given combinatorial auction, but usually cannot (and should not) restrict the number of bids submitted. The algorithm scales up well in practice. For details and experimental results, see [30]. The algorithm has been further refined by others [5], and we then designed a completely different search algorithm that promised to be drastically faster [35]. Recently, we developed yet a faster algorithm, CABOB, for the problem. Experimental results show that it is currently the fastest algorithm for optimal winner determination, scaling to hundreds of items and thousands or tens of thousands of bids [38].

2.1.2 The XOR bidding language

The OR bidding language implicitly assumes that the bids are superadditive: $b_i(S \cup S') \geq b_i(S) + b_i(S')$ where $b_i(S)$ is the bid of agent i on combination S . But what would happen if

¹We proved this via an approximation-preserving polynomial reduction from maximum clique (which is known to be inapproximable [9]).

agent 1 bid $b_1(\{1\}) = \$5$, $b_1(\{2\}) = \$4$, and $b_1(\{1, 2\}) = \$7$, and there were no other bidders? The auctioneer could allocate items 1 and 2 to agent 1 separately, and that agent’s bid for the combination would value at $\$5 + \$4 = \$9$ instead of $\$7$. So, the OR bidding language focuses on settings where combinatorial bids are introduced to capture synergies (complementarities) among items. On the other hand, in many real world settings local subadditivities (substitutability) can occur as well. For example, when bidding for a landing slot for a plane, the bidder is willing to take any one of a host of slots, but does not want more than one.

To address this, we introduced the *XOR bidding language*, which has become popular recently (e.g. [20]). In this language, each bidder can submit one *XOR disjunct*, that is, each bidder can submit multiple bids (each bid including a set of items and a price), but only one bid from any given bidder can be accepted. *eAuctionHouse* supports the XOR bidding language. This allows the bidders to express general preferences with both positive and negative complementarities. In other words, the language is expressive enough to represent any mapping from combinations of items to prices.

Optimal winner determination in the XOR bidding language is at least as hard as in the OR bidding language. The reason is that any problem instance from the OR bidding language can be converted into the XOR bidding language (simply have each bid come from a separate virtual bidder; this way there would be no XOR constraints). Therefore, the negative results, \mathcal{NP} -hardness and inapproximability, apply to the XOR bidding language as well. Our winner determination algorithms extend to the XOR bidding language by inserting the extra constraints that no two bids from the same bidder can be accepted [30]. These extra constraints actually *reduce* the size of the search space. However, this does not imply that the search becomes faster.

To enable an economically efficient allocation to be reached, it would be desirable to extract truthful valuation revelations as the bids from the bidders. This is an issue because in traditional “pay-your-winning-bid” auctions, the bidders have incentive to underbid strategically. Bidding truthfully can be made incentive compatible (a weakly dominant strategy) by using the *Groves-Clarke mechanism* [4, 6] (aka. *generalized Vickrey auction* [15]). This means that each bidder is motivated to bid truthfully irrespective of how the others bid. This renders counterspeculation unnecessary. With unrestricted bidder preferences, the Groves-Clarke mechanism requires a fully expressive bidding language. Therefore, our XOR bidding language enables incentive-compatible combinatorial auctions, while the traditional OR bidding language does not.

With our XOR bidding language, the Groves-Clarke mechanism can be applied to the combinatorial auction setting as follows. Winning bids are determined so as to maximize the auctioneer’s revenue under the constraint that each item can be allocated to at most one bid (while honoring the XOR constraints). The amount that an agent needs to pay is the sum of the others’ winning bids had the agent not submitted any bids, minus the sum of the others’ winning bids in the actual optimal allocation. Therefore, the winner determination problem has to be solved once overall, and once per winning agent without any of that agent’s bids. This makes fast winner determination even more crucial. Note that for example just removing one winning bid at a time would not constitute an incentive compatible mechanism. Incentive compatibility can also be lost if either winner determination or price determination is done only approximately.²

2.1.3 The OR-of-XORs bidding language

In addition to full expressiveness, the convenience of using combinatorial auctions is another important issue. While the XOR bidding language is fully expressive, in the worst case expressing one’s preferences requires one to submit a bid on each of the $2^m - 1$ bundles of items (where m is the number of items). A shorter representation of preferences without loss of expressive power could be possible by allowing a richer input language.³ To this end, we introduced the

²A recent paper shows that if each bidder is interested in only one bundle, then incentive compatibility can be guaranteed even under approximate winner determination [12].

³One idea toward this direction was presented early on by Rassenti et al [24]. They allowed each bidder to place combinatorial bids and to state the maximal number of bids that could be accepted from that bidder. Our XOR bidding language could be viewed as a special case of this where that number is one.

OR-of-XORs bidding language. Each bidder can submit multiple XOR disjuncts, where these disjuncts are combined with OR. For example, a bidder could offer

$$\begin{aligned} & [(\{1, 2\}, \$30) \text{ XOR } (\{1, 3\}, \$25)] \\ \text{OR } & [(\{4, 5\}, \$75) \text{ XOR } (\{3, 5\}, \$55) \text{ XOR } (\{6\}, \$25)] \\ \text{OR } & [(\{7, 8\}, \$25)] \end{aligned}$$

Such *OR-of-XORs* expressions maintain full expressive power, but lead to shorter (at least no longer) input descriptions than in the XOR bidding language (because the latter is a special case). We also believe that *OR-of-XORs* expressions are a more natural way to input preferences. The OR connectors represent independence. For example, if a company bids for inputs to its production, and has three plants, it could submit 3 XOR disjuncts (one for each plant, and the different alternative inputs would be XOR'ed) to represent the fact that the plants' production plans are independent. This avoids the combinatorial listing of *combinations* of bundles (as would be the case with XOR bids) in cases where the values of the bundles are independent.

In *eAuctionHouse*, the bidder can submit such *OR-of-XORs* expressions in table form. Each row in the table is a combinatorial bid (that lists the set of items and a price), and the rows are combined with XOR. These XOR disjuncts are combined together with a non-exclusive OR. We do this by allowing the user to submit multiple tables; an intuitive way of representing independence across XOR disjuncts (tables). Winners can be determined with the same algorithm as in the case of XOR bids by using exclusivity constraints only between bids that are in the same XOR disjunct (table).

Other enrichments to the language are also possible. Since even the XOR bidding language is fully expressive, expressiveness should be viewed as a necessary, but not sufficient, condition in designing bidding languages. Between fully expressive input languages the appropriate comparison criterion is the convenience of their use in the particular application domain in question.⁴

2.1.4 Other combinatorial auction generalizations

eAuctionHouse also supports *multi-unit* combinatorial markets where there can be multiple indistinguishable units of each item for sale. Furthermore, it supports combinatorial *exchanges*, that is, combinatorial double auctions where there can be multiple buyers and multiple sellers. Some of the buyers may be sellers as well. Even in a single bid, a bidder can buy some items while selling others. Figure 1 illustrates the table-based bidding interface in a multi-unit combinatorial exchange that uses our OR-of-XORs bidding language (the number of units is specified in each cell of the table). Recent research has focused on developing even faster winner determination algorithms for multi-unit combinatorial auctions [13, 35], and for single- and multi-unit combinatorial exchanges [35, 39].

2.2 Bidding via price-quantity graphs

For markets where there is one item for sale, but multiple indistinguishable units of it, *eAuctionHouse* supports price-quantity graphs (that is, supply curves and demand curves) so that bidders can express continuous preferences, see Figure 2. For example when a bidder buys a larger quantity, she might only accept a lower unit price. Naturally, a bidder accepts anything below the curve as well (automatically colored region) because she will get the same quantity as on the curve, but at a lower price. Similarly, a seller would accept anything above her curve.⁵

⁴A recent paper proposes a generalization of the OR-of-XORs language called OR*. In that language, a bidder can submit bids as nodes in a graph, and can submit edges into the graph corresponding to arbitrary XOR constraints. Because the OR* language is a generalization of the OR-of-XORs language, it is fully expressive and never less compact. Furthermore, it was shown that there exist preferences for which it is exponentially more compact than the OR-of-XORs language [19]. We believe, however, that the OR-of-XORs language is more natural.

⁵Optimark Technologies, Inc. does a fuzzier match where each bidder specifies how much she prefers different regions of the price-quantity plane [14].

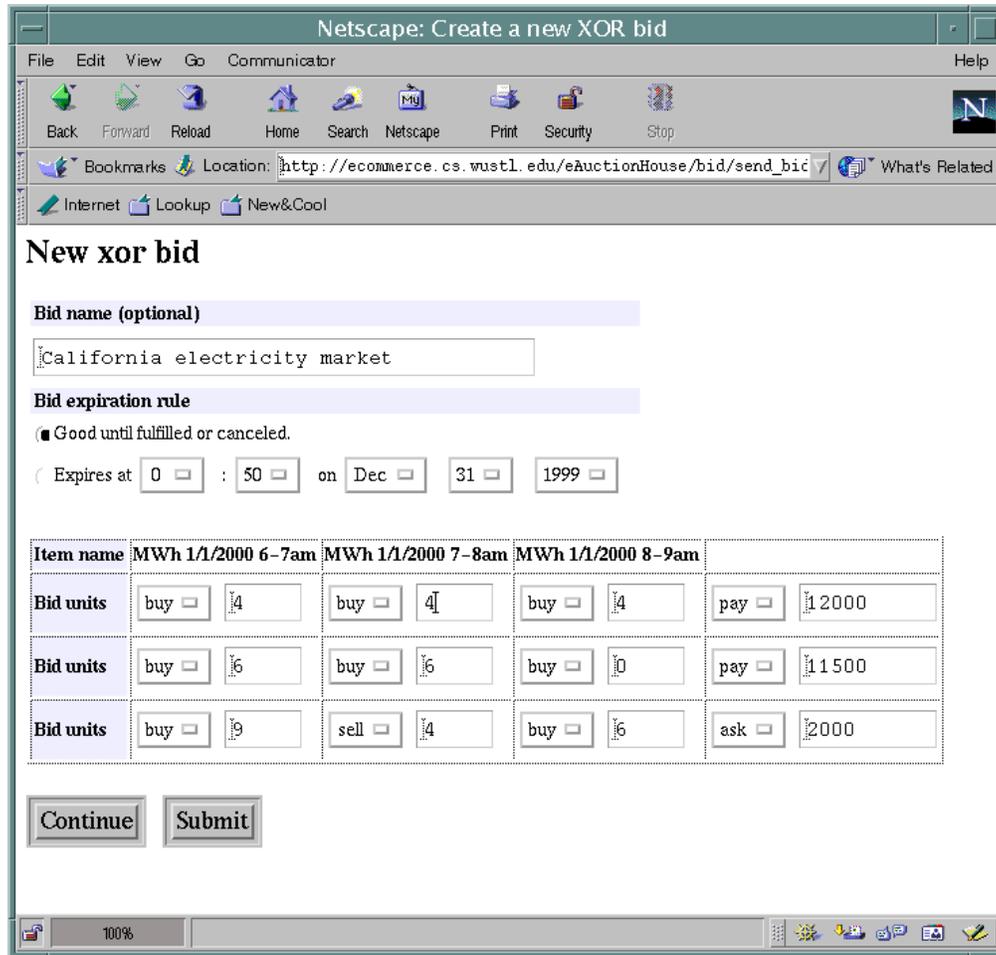


Figure 1: An XOR disjunct in a multi-item, multi-unit combinatorial exchange. The example is from an electricity market scenario where the agents can bid for combinations of electricity for different hours of the day, and for multiple Mega Watt hours for each hour of the day. In this example, a refinery operator wants consecutive hours of electricity for her plant. She prefers to run the plant for three hours at a rate of 4 MWh per hour, but running for two hours at 6 MWh per hour is also feasible. In the last alternative she acts both as a buyer and as a seller of different items in one combination.

In our implementation, the curves are piecewise linear both for the bidders' drawing convenience and for the convenience of winner determination. In single-sided auctions with price-quantity graph bidding, the winner determination algorithm works as follows. It sums the demand for every unit price (it does not loop though prices but uses the endpoints of each linear piece of the curve to do this). Then, it picks the unit price that maximizes the aggregate price under the constraint that not more is demanded than is available. Each bidder then gets the amount that she bid at that unit price. In double auctions, both supply and demand curves are separately aggregated, and any one of the points where supply meets demand is chosen. If the curves were noncontinuous, it would be possible that no match exists. This holds both for single and double auctions. To prevent this, we use continuous curves, that is, the slope of each linear segment is finite. We do not assume that the curves are monotonic; for example, the curve in Figure 2 is nonmonotonic.

If the objective is to maximize economic efficiency (welfare among the buyers and sellers), then the best way to clear a fixed set of supply/demand curves is to pick a solution that maximizes

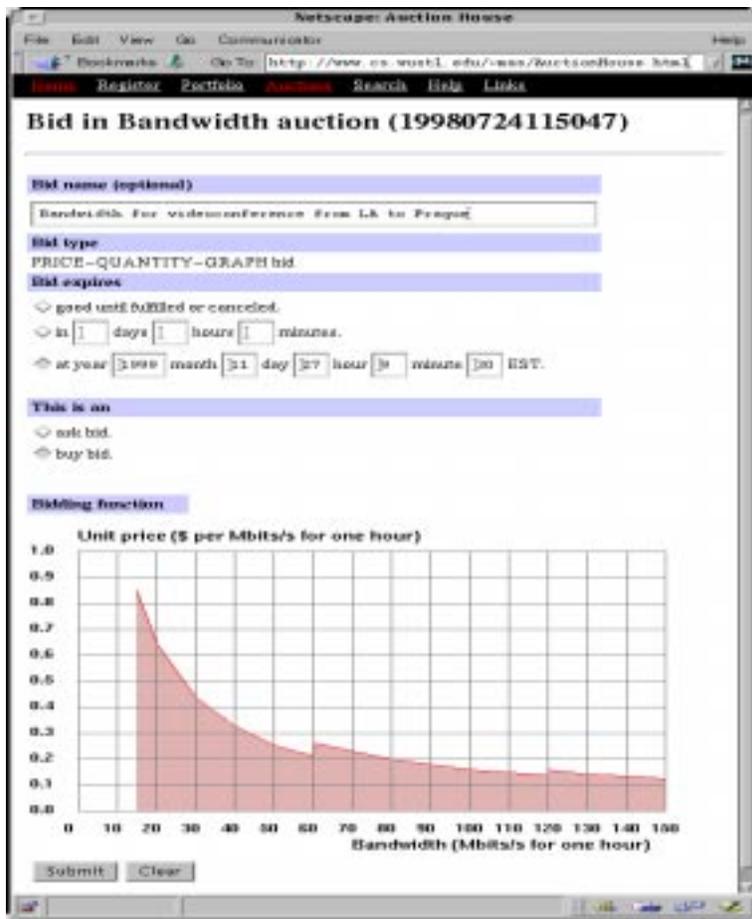


Figure 2: A price-quantity graph allows the user to express continuous preferences in the auction server of *eMediator*. This figure corresponds to the user being able to hold a video-conference at three alternative picture resolutions requiring a bandwidth of 15, 60, or 120 Mbps. The auctioned item could be a prebundled combination of items. For example, the virtual circuit from LA to Prague can use several network links owned by different backbone providers.

the amount of trade under the constraint that supply meets demand. On the other hand, in a recent paper, Sandholm and Suri studied the complexity of clearing supply/demand curves so as to maximize the profit of the party who runs the market (prices paid by buyers minus prices paid to sellers). Under that objective, it is usually best to reduce the amount of trade. It turns out that with piecewise-linear curves, *nondiscriminatory clearing* (where each buyer is cleared at the same price, and each seller is cleared at the same price—which is generally different from the buyers' price) can be done optimally in polynomial time. However, optimal *discriminatory clearing* (where each party is cleared at a potentially different price) is \mathcal{NP} -complete. If the curves are linear, then even discriminatory clearing can be done in polynomial time. These results hold both for auctions [36] and for exchanges [37].

2.3 Support for choosing an auction type

The auction server, *eAuctionHouse*, supports a wide variety of auction types. The user that sets up a given auction (she may be a buyer, a seller, or a third party facilitator) decides the auction type. However, since the space of different auction types is enormous, the auction server helps the user in making the choice. First, only choices that are sensible based on game-theoretic

analyses or economics experiments are provided as alternatives. Furthermore, there is an expert system that restricts the choice of auction types given the auction setting (Figure 3). For any

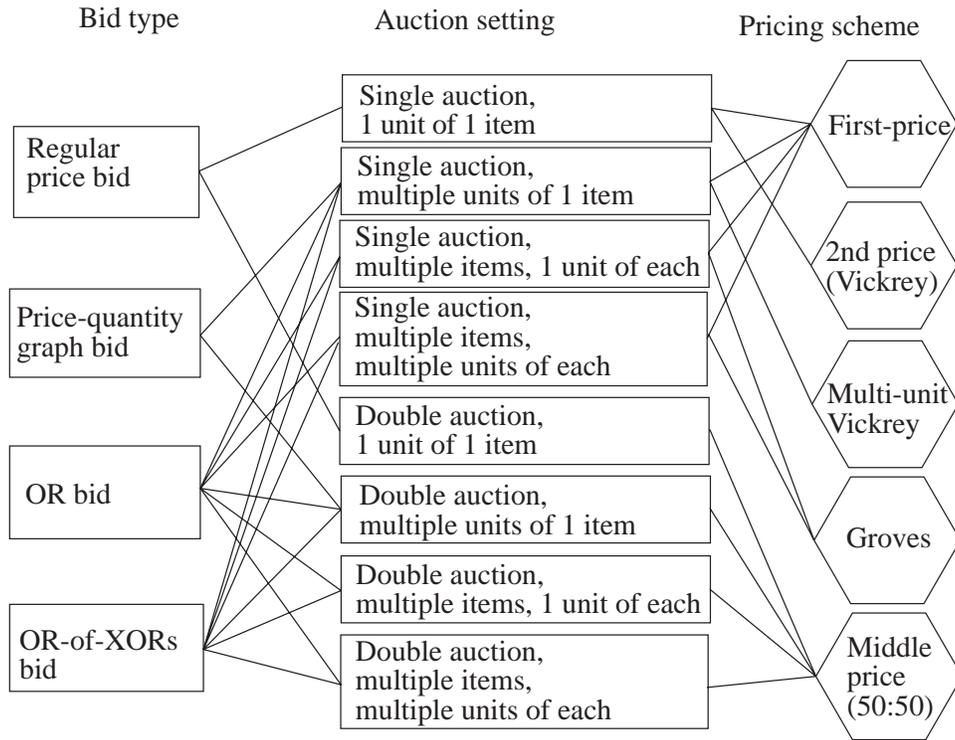


Figure 3: *Expertise showing valid combinations of choices of some of the parameters of eAuctionHouse.*

given auction setting, it tells the user what kinds of bids can be accepted, and what price determination schemes should be used. The auction setting differs based on whether it is a single or double auction, whether there is one or multiple items, and whether there is one or multiple units of each item. Furthermore, the units can be divisible or indivisible. The bid types include a regular price bid where the bidder specifies the price for a good; a price-quantity graph bid (Figure 2); an OR bid; and an OR-of-XORs bid.

The first-price pricing scheme charges the buyer the price of her bid. This scheme leads to strategic underbidding. In single-unit ascending open-cry auctions, each bidder's dominant strategy is to bid a small increment more than the current price, and stop when her valuation is reached. In sealed bid auctions with common knowledge assumptions about the priors from which the bidders' private valuations are drawn, a Nash equilibrium analysis can be conducted to determine how much each agent should underbid as a function of her valuation (e.g. [44]).

The second-price (Vickrey) auction charges the winning bidder the price of the second highest bid. Under certain restrictions [29], it is each bidder's dominant strategy to bid her true valuation [43].

The multi-unit Vickrey auction is a generalization of the Vickrey auction to settings with multiple units of an item, or in other words, multiple indistinguishable goods. The bidders can use any of the following fully expressive bidding languages: 1) the XOR bidding language, where each bid includes a price and a quantity, 2) the OR-of-XORs bidding languages, where each bid includes a price and a quantity, or 3) a price-quantity curve as described above. (Even in this restricted setting, the OR bidding language would not be fully expressive. For example, it would not enable an agent to state that the value of 5 units is lower than the value of 2 units plus the value of 3 units.) The units are assigned to the bidders so as to maximize the sum of the accepted bid prices. Incentive-compatible pricing is achieved, again, using the Clarke-Groves mechanism.

Specifically, the amount a bidder has to pay is the sum of the prices of the others' accepted bids had the bidder not participated, minus the sum of the prices of the others' accepted bids. Another possible generalization of the Vickrey auction to the multi-unit setting is to charge every bidder the price of the highest bid that just did not win (aka. the *uniform-price auction*). Our auction server uses the former method because the latter is not incentive compatible: it falls prey to *demand reduction lies* by the bidders [1].

For multi-item auctions (with one or more units per item), the Groves-Clarke mechanism, as discussed earlier, is the appropriate generalization of the Vickrey auction. Each bidder's dominant strategy is to bid truthfully.

In double auctions (exchanges), our server splits the gains equally in the standard way. The price is half way between the bid and the ask. In combinatorial double auctions (exchanges), we maximize the *surplus* (prices of accepted bids minus prices of accepted asks). The surplus can be divided arbitrarily between the buyers, the sellers, and the exchange—as long as no buyer pays more than the prices of her accepted bids (this guarantees that each bidder is motivated to participate, even in hindsight).

The auctions in *eAuctionHouse* also have several other parameters, including:

- Whether or not matches have to be exact in multi-unit auctions.
- Tie-breaking rule: random, older bid overrides, or newer bid overrides.
- When to clear the auction: when a specific time is reached, every time a bid is received, periodically, or when no bids have been received for a specified time.
- When the auction permanently closes: when it is cleared, when no bids have been received for a specified time, or when the auction's owner cancels it.
- Whether or not bid retraction is allowed (possibly for a penalty) before winners are determined.
- Whether or not bid retraction is allowed (possibly for a penalty) after winners are determined.
- What information is revealed to the bidders during bidding: all bids, highest bids, or none.
- What information is revealed to the bidders after clearing: all bids, winning bids, or none.

2.4 *NOMAD*: Mobile proxy agents

Our auction house supports mobile agents so that a user can have her agent actively participating in the auction while she is disconnected. For example, the user can launch her agent over the phone from an airplane using a laptop, and then disconnect. Mobile agents that execute on the agent dock which is on (or near) the host machine of the auction server also reduce the network latency—an issue of key importance in time-critical bidding. The Michigan Internet AuctionBot [45] provides a TCP/IP-level message protocol via which agents could participate in their auction. Their auction server differs from ours in that they do not provide support for mobile agents. Our auction server uses the commercial Concordia agent dock from Mitsubishi to provide mobile agents a safe execution platform from where they can observe what is transpiring in the auctions, bid, set up auctions, move to other hosts, etc. The user has the full flexibility of Java programming at her disposal when designing her mobile agent. We also provide an easy-to-use HTML interface for non-programmers where the user can specify what she wants her agent to do, and our system automatically generates the Java code for the corresponding mobile agent, and launches it. The following parameterizable mobile agent templates are currently available:

1. The *information agent* goes to an auction and sends email to the user when specified events occur. Using this agent, the user does not have to poll the auction, and gets notified of important events immediately.

2. The *incrementor agent* implements the dominant strategy on the user’s behalf in single-item single-unit ascending open-cry first-price private-value auctions. It bids a small increment more than the current highest price, and stops if the user’s reservation price is reached. With this agent the user does not have to follow the auction, and her dominant strategy in these settings is to report her valuation truthfully to the agent.
3. The *N-agent* underbids optimally on the user’s behalf in single-item single-unit sealed-bid first-price auctions where the number of bidders, N , is known, and the bidders’ private valuations are independently drawn from a uniform distribution. Specifically, the symmetric Nash equilibrium strategy is to bid the user’s valuation times $(N-1)/N$ [23]. The user is then motivated to reveal her true valuation to the agent.
4. The *control agent* goes to an auction and submits very low noncompetitive bids. It is a speculator’s tool to artificially increase the number, N , of bidders in an auction to mislead others, e.g. the N -agent. For example, it is in the seller’s interest to submit control agents so that N -agents would bid higher.
5. The *discover agent* computes the expected gain from bidding a small amount more than the current highest price according to the agent’s current distribution of her valuation. This is intended for settings where the user does not know her exact valuation for the item, but only a probability distribution on it. In the future, the probability distribution could be updated based on new information, or in non-private-value auctions, based on what others have bid.

Unlike current electronic negotiation servers which usually only provide an auction house, *eMediator* provides other types of services for facilitating ecommerce in addition, such as a leveled commitment contract optimizer, and a safe exchange planner. These are discussed in Sections 3 and 4 respectively.

3 *eCommitter*: A leveled commitment contract optimizer

Normal full commitment contracts are unable to take advantage of the possibilities that unknown future events provide. Once an agent agrees to a contract, she has to follow through no matter how future events unravel. Although a contract may be profitable to an agent when viewed *ex ante*, it need not be profitable when viewed after some future events have occurred. Similarly, a contract may have too low expected payoff *ex ante*, but in some realizations of the future events, it may be desirable.

Contingency contracts have been suggested for utilizing the potential provided by future events among self-interested agents [22]. The contract obligations are made contingent on future events. In some games this increases the expected payoff to both parties compared to any full commitment contract. However, contingency contracts are often impractical because the space of combinations of future events may be large and unknown. Also, when events are not mutually observable, the observing agent can lie about what transpired.

Leveled commitment contracts are another method for capitalizing on future events [33]. Instead of conditioning the contract on future events, a mechanism is built into the contract that allows unilateral decommitting. This is achieved by specifying in the contract the level of commitment by decommitment penalties, one for each agent. If an agent wants to decommit—that is, to be freed from the contract obligations—it can do so simply by paying the decommitment penalty to the other party. The method requires no explicit conditioning on future events: each agent can do her own conditioning dynamically. No event verification mechanism against lying is required either. The decommitment possibility increases each agent’s expected payoff under very general assumptions [33].

eMediator includes an optimizer for leveled commitment contracts. We call it *eCommitter*. We analyze contracting situations from the perspective of two risk neutral agents each of which attempts to maximize his own expected payoff: the *contractor* who pays to get a task done,

and the *contractee* who gets paid for handling the task. The framework can be interpreted as modeling other types of settings than task allocation also, for example general allocation of rights and obligations where the agents' costs and gains of the rights and obligations may change. In what follows, we word the results in the context of task allocation.

The contractor tries to minimize the contract price ρ that he has to pay to get the task handled. The contractee tries to maximize the payoff ρ that she receives from the contractor for handling the task. We study a setting where the future of the agents involves uncertainty. Specifically, the agents might receive outside offers.⁶ The contractor's best outside offer \check{a} is only probabilistically known *ex ante* by both agents, and is characterized by a probability density function $f(\check{a})$. If the contractor does not receive an outside offer, \check{a} corresponds to its best outstanding outside offer or its fall-back payoff, i.e. payoff that it receives if no contract is made. The contractee's best outside offer \check{b} is also only probabilistically known *ex ante*, and is characterized by a probability density function $g(\check{b})$. If the contractee does not receive an outside offer, \check{b} corresponds to its best outstanding outside offer or its fall-back payoff.⁷ The variables \check{a} and \check{b} are assumed statistically independent, and f and g are assumed to be common knowledge.

The contractor's options are either to make a contract with the contractee or to wait for \check{a} . Similarly, the contractee's options are either to make a contract with the contractor or to wait for \check{b} . The two agents could make a full commitment contract at some price. Alternatively, they can make a leveled commitment contract which is specified by the contract price, ρ , the contractor's decommitment penalty, a , and the contractee's decommitment penalty, b . We restrict our attention to contracts where $a \geq 0$ and $b \geq 0$, that is, agents do not get paid for decommitting.⁸ The contractor has to decide on decommitting when he knows his outside offer \check{a} but does not know the contractee's outside offer \check{b} . Similarly, the contractee has to decide on decommitting when she knows her outside offer \check{b} but does not know the contractor's. This seems realistic from a practical contracting perspective.

3.1 Nash equilibria for a given contract

One concern is that a rational self-interested agent is reluctant in decommitting because there is a chance that the other party will decommit, in which case the former agent gets freed from the contract, does not have to pay a penalty, and collects a penalty from the breacher. Sandholm and Lesser showed that despite such insincere decommitting the leveled commitment feature increases each contract party's expected payoff, and enables contracts in settings where no full commitment contract is beneficial to all parties [33].

The contractor decommits if he gets a low enough outside offer, e.g., he can get his task handled at a low cost. We denote his decommitting threshold by \check{a}^* , so his decommitting probability is

$$p_a = \int_{-\infty}^{\check{a}^*} f(\check{a}) d\check{a} \tag{1}$$

The contractee decommits if she gets a high enough outside offer, e.g., gets paid for handling a task. We denote her decommitting threshold by \check{b}^* , so her decommitting probability is

$$p_b = \int_{\check{b}^*}^{\infty} g(\check{b}) d\check{b} \tag{2}$$

3.1.1 Sequential decommitting mechanisms

In our sequential decommitting mechanism, one agent has to reveal her decommitting decision before knowing whether the other party decommits. While our implementation analyzes both orders of decommitting, here we only discuss the setting where the contractee has to decide first. The case where the contractor decides first is analogous. There are two alternative leveled commitment contracts that differ on whether or not the agents have to pay the penalties if both decommit.

⁶The framework can also be interpreted to model situations where the agents' cost structures for handling tasks and for getting tasks handled change e.g. due to resources going off-line or becoming back on-line.

⁷Games where at least one agent's future is certain, are a subset of these games. In such games all of the probability mass of $f(\check{a})$ and/or $g(\check{b})$ is on one point.

⁸A recent paper shows that this restriction does not reduce the welfare of the contract parties [41].

If the contractee has decommitted, the contractor's best move is not to decommit because $-\check{a} - a + b \leq -\check{a} + b$ (because $a \geq 0$). This also holds for a contract where neither agent has to pay a decommitment penalty if both decommit since $-\check{a} \leq -\check{a} + b$. In the subgame where the contractee has not decommitted, the contractor's best move is to decommit if $-\check{a} - a > -\rho$, i.e. $\check{a}^* = \rho - a$ (3a)

The contractee gets $\check{b} - b$ if she decommits, $\check{b} + a$ if she does not but the contractor does, and ρ if neither decommits. Thus the contractee decommits if $\check{b} - b > p_a(\check{b} + a) + (1 - p_a)\rho$. If $p_a = 1$, this is equivalent to $-b > a$ which is false because $a \geq 0$ and $b \geq 0$. In other words, if the contractee surely decommits, the contractor does not. On the other hand, the above is equivalent to

$$\check{b} > \rho + \frac{b + ap_a}{1 - p_a} \stackrel{\text{def}}{=} \check{b}^* \text{ when } p_a < 1 \quad (4a)$$

3.1.2 Simultaneous decommitting mechanisms

In our simultaneous decommitting mechanisms, agents have to reveal their decommitment decisions simultaneously. We first discuss the variant where both have to pay the penalties if both decommit. The contractor decommits if $p_b \cdot (-\check{a} + b - a) + (1 - p_b)(-\check{a} - a) > p_b \cdot (-\check{a} + b) + (1 - p_b)(-\rho)$. If $p_b = 1$, this equates to $a < 0$, but we already ruled out contracts where an agent gets paid for decommitting. On the other hand, this equates to

$$\check{a} < \rho - \frac{a}{1 - p_b} \stackrel{\text{def}}{=} \check{a}^* \text{ when } p_b < 1 \quad (3b)$$

The contractee decommits if $(1 - p_a)(\check{b} - b) + p_a(\check{b} - b + a) > (1 - p_a)\rho + p_a(\check{b} + a)$. If $p_a = 1$, this equates to $b < 0$, but we ruled out contracts where an agent gets paid for decommitting. However, this equates to

$$\check{b} > \rho + \frac{b}{1 - p_a} \stackrel{\text{def}}{=} \check{b}^* \text{ when } p_a < 1 \quad (4b)$$

In another type of simultaneous decommitting mechanism, neither agent has to pay if both decommit. The contractor decommits if $p_b \cdot (-\check{a}) + (1 - p_b)(-\check{a} - a) > p_b \cdot (-\check{a} + b) + (1 - p_b)(-\rho)$. If $p_b = 1$, this equates to $b < 0$, but we already ruled out contracts where an agent gets paid for decommitting. On the other hand, this equates to

$$\check{a} < \rho - a - \frac{bp_b}{1 - p_b} \stackrel{\text{def}}{=} \check{a}^* \text{ when } p_b < 1 \quad (3c)$$

The contractee decommits if $(1 - p_a)(\check{b} - b) + p_a\check{b} > (1 - p_a)\rho + p_a(\check{b} + a)$. If $p_a = 1$, this equates to $a < 0$, but we ruled out contracts where an agent gets paid for decommitting. However, this equates to

$$\check{b} > \rho + b - \frac{ap_a}{1 - p_a} \stackrel{\text{def}}{=} \check{b}^* \text{ when } p_a < 1 \quad (4c)$$

3.1.3 Using *eCommitter* to find the Nash equilibria for a given contract

To use the *eCommitter* optimizer, the user inputs the probability distributions f and g graphically or textually. The user also inputs the contract price ρ and the decommitment penalties (a and b). For each one of the decommitting mechanisms (which vary based on who has to reveal her decommitting decision first—the simultaneous mechanisms are also considered—and whether or not the agents have to pay the penalties if both decommit), *eCommitter* finds all the Nash equilibria. In other words, *eCommitter* determines decommitting threshold pairs. *eCommitter* does this by solving the simultaneous equations (3) and (4) which use (1) and (2).⁹ Given an equilibrium, it is easy to compute the agents' expected payoffs under the contract. Finally, *eCommitter* presents all the equilibria and the associated expected payoffs to the user.

⁹While solving this equation group is complex in general, we have the user input piecewise linear f and g . This allows the problem to be decomposed into rectangles (corresponding to one piece of f and one of g), where f and g are linear within each rectangle. In that case, the equation group can be easily solved in constant time. Our optimizer solves the problem in this way for each rectangle separately. For details, see [34].

3.2 Using *eCommitter* to find optimal contracts

In addition to determining how agents should play under a given leveled commitment contract, *eCommitter* also determines the optimal leveled commitment contracts. Again, the user inputs the distributions f and g graphically or textually. However, the user does not input the contract price or decommitment penalties. Rather, *eCommitter* determines the values of these (ρ , a , b) in a way that maximizes the sum of the contract parties' expected payoffs [34]. (Again, that optimization algorithm takes into account that the agents decommit insincerely in Nash equilibrium according to Equations (1), (2), (3), and (4).) The expected benefit from the contract is the sum of the agents' expected payoffs under the contract, minus the sum of the agents' (expected) payoffs under the *status quo* where they do not make a contract.

The contract is optimized separately for each one of the decommitting mechanisms, which vary based on who has to reveal her decommitting decision first—the simultaneous mechanisms are also considered—and whether or not the agents have to pay the penalties if both decommit. The optimal contracts for each mechanism are then presented to the user. For each mechanism, there is a continuum of optimal contracts that vary based on how the expected benefit is divided between the agents. *eCommitter* visualizes how each agent's payoff changes as a function of the contract price ρ , and shows how the decommitment penalties (a and b) should be tailored to the contract price so as to maximize the expected benefit.¹⁰ *eCommitter* also points out, for each of the mechanisms, the fair optimal contract, that is, the contract that maximizes the expected benefit and divides it equally between the agents.

4 *eExchangeHouse*: A safe exchange planner

Contract execution is more difficult in electronic commerce than physical commerce because the parties may be anonymous and can disappear easily. The problem will be exacerbated in the future by software agents for automated negotiation. Such an agent can vanish by simply killing its process, and litigation is infeasible unless the other contract party knows which real-world entity the agent represented. Another problem is the lack of uniform laws on electronic commerce—and particularly agent-mediated commerce—in different countries.

An important aspect of contract execution is making sure that the seller gets paid, and that the buyer gets the goods. The risk is that once one party has received the item, he may be motivated to vanish without delivering his part of the contract. In fact, 6% of US people with online buying experience have reported nondelivery [18].

Nondelivery could be avoided by a trusted third party that takes the payment and goods, and carries out the transaction only after all parties have delivered their part to the intermediary. Today's electronic commerce implements a coarse one-sided variant of this where the third party takes the payment into escrow, and releases it to the seller only after the buyer has verified receipt of the goods. A disadvantage of these third party escrow companies like *Tradenable* and *Tradesafe* is the cost of running such an intermediary, which is recovered as fees—currently about 3-5% of the contract price—from the contract parties.

A method for tackling this problem without third parties was developed in [32]. The exchange is divided into chunks where each party delivers a small amount at a time, and the exchange proceeds with such alternation. The method targets settings where dividing the goods into chunks is relatively inexpensive, such as is often the case for example with information goods and computational services. A sequence is called safe if each party is motivated to follow the exchange at every step in anticipation of the profit from the rest of the exchange instead of vanishing with what the other party has delivered so far. (Specifically, a safe sequence can be executed in subgame perfect Nash equilibrium.) Some chunkings allow a safe sequence while

¹⁰The expected benefit can be maximized while dividing the benefit in any way—as long as each party gets a nonnegative portion of the benefit [34, 41]. In other words, here integrative bargaining (increasing the expected benefit) does not hinder distributive bargaining (dividing the expected benefit). However, this independence only holds if the decommitment penalties are tailored to the chosen contract price.

others do not [28, 32]. Similarly, some sequences of delivering given chunks are safe while others are not.

As part of *eMediator*, we built a safe exchange planner called *eExchangeHouse*. As input, the planner takes the following information:

- the contract price,
- how the seller’s cost accrues as a function of how much he has delivered so far in the transaction,
- how the buyer’s valuation accrues as a function of how much he has received in the transaction,
- how much gain the seller is willing to forego to avoid possible reputation costs from defecting in the exchange (or costs associated with the risk of getting caught after defecting), and
- how much gain the buyer is willing to forego for the same reasons.

The form of the cost and valuation accrual information depends on the setting. The setting can have one item to be exchanged, multiple independent items, or multiple dependent items. Also, there can be one unit of each item, multiple independent units of each item, or multiple dependent units of each item. Finally, the units can be divisible or indivisible.

For example, in the case where a single divisible item is exchanged, the user inputs a graph of how the buyer’s valuation accrues as a function of how much has been delivered, and another graph that shows how the seller’s cost accrues. As another example, in the case of multiple distinguishable items, the user lists the items and how many units of each item are to be exchanged. For each possible *state* of the exchange (units of item 1 delivered \times units of item 2 delivered \times ...) the user inputs the the buyer’s valuation and the seller’s cost. Combining indistinguishable goods into units of a single item significantly reduces the state space because within each good, only the number of units delivered matters, not which specific units have been delivered.

Based on this input, the planner finds a safe chunking that minimizes the number of chunks and a safe chunk sequence if they exist. If they do not exist, the user is alerted of this. The planner can determine the optimal safe exchange plan under any of the following four alternative objectives: 1) minimizing the number of deliveries, 2) minimizing the number of payments, 3) minimizing the number of payments plus the number of deliveries, and 4) minimizing the number of steps, where each step may include both a payment and a delivery. The chunking algorithms and the chunk sequencing algorithms are highly nontrivial; they are described and analyzed in detail in other papers [28, 31, 32].

Generally, a user (buyer or seller) only knows the desired inputs regarding himself. He can use bounds on the other party’s values and the exchange plan will still be safe. An upper bound can be used for how the buyer’s value accrues and a lower bound on how the seller’s cost accrues. Similarly, an upper bound can be used for the seller’s cost of the completed contract and a lower bound can be used on the buyer’s valuation of the completed contract. Finally, lower bounds can be used on the amounts that the agents are willing to forego in order to avoid defection related averse reputation costs or the chance of getting caught after nondelivery. Although the planned exchange will be safe when such bounds are used as the input, the planning problem may be infeasible if the bounds are too far off. In that case, *eExchangeHouse* reports that no safe exchange plan exists given the inputs.

Each trade party could have a software agent execute its sequence of delivery chunks. This would provide a level of isolation between the user and the exchange so the user need not even know that the trade is occurring in chunks instead of all at once.

4.1 Example of exchanging independent units of independent items

In this example the agents exchange multiple indivisible independent units of multiple independent items. Specifically, they exchange five stocks (multiple units of each): Microsoft (MSFT),

Hewlett-Packard (HWP), Dell (DELL), Ebay (EBAY) and Procter & Gamble (PG). The stocks have different values for the supplier and the demander (for example, because the supplier and the demander have different portfolios, and the value of a stock in a portfolio depends on how correlated the stock is with the rest of the portfolio statistically).

The input to the planner includes the contract price (\$25864.60) and the amounts the agents are willing to forego to avoid reputation-related problems (the seller is ready to forego \$1000, and the buyer is ready to forego \$1500). The input also includes the information in Table 1. The user actually submits this information by filling such a table in *eExchangeHouse*. Note that the seller’s cost per unit ($\frac{\Delta v_s}{\Delta x_i}$) can be greater than the demander’s valuation per unit ($\frac{\Delta v_d}{\Delta x_i}$) for some items i —but not all of them—to allow safe exchange.

item name	number of units to exchange	$\frac{\Delta v_s}{\Delta x_i}$	$\frac{\Delta v_d}{\Delta x_i}$
MSFT	60	\$90.81	\$98
HWP	30	\$102.63	\$109.9
DELL	10	\$41.81	\$42.5
EBAY	100	\$120.25	\$127.25
PG	40	\$85.1	\$88.44

Table 1: *Input form for independent items, independent units.*

Based on this input, *eExchangeHouse* determines an optimal safe exchange plan, which is shown in Figure 4.

4.2 Example of exchanging dependent units of dependent items

Consider a software system consisting of two parts, a main package and a plug-in. The former takes four floppy disks (units) whereas the latter takes only one. They depend on each other in the sense of the cost function of the seller (v_s) and the value function of the demander (v_d), as the user has expressed in Table 2.¹¹

plug-in	package	v_s	v_d
0	0	0	0
0	1	5	3
0	2	9	6
0	3	13	10
0	4	17	14
1	0	10	6
1	1	13	10
1	2	16	15
1	3	18	20
1	4	19	26

Table 2: *Input table for exchanging dependent units of dependent items.*

We see in this example that—without a plug-in—the cost of one disk of the package decreases (5, 4, 4, 4) depending on how many disks have been delivered so far. Also, the value of one disk increases (3, 3, 4, 4) depending on how many disks have been delivered so far. Therefore, the units of the item ”package” are dependent. Moreover, once the plug-in is delivered, the cost and value of one disk of the package change (respectively (4, 4, 3, 2) and (4, 4, 5, 5)). Therefore, the item ”package” depends on the item ”plug-in”. (Similarly, the plug-in depends on the package.)

¹¹In settings where the items are dependent, but the units within each item are independent, the amount of input required from the user can be reduced [31].

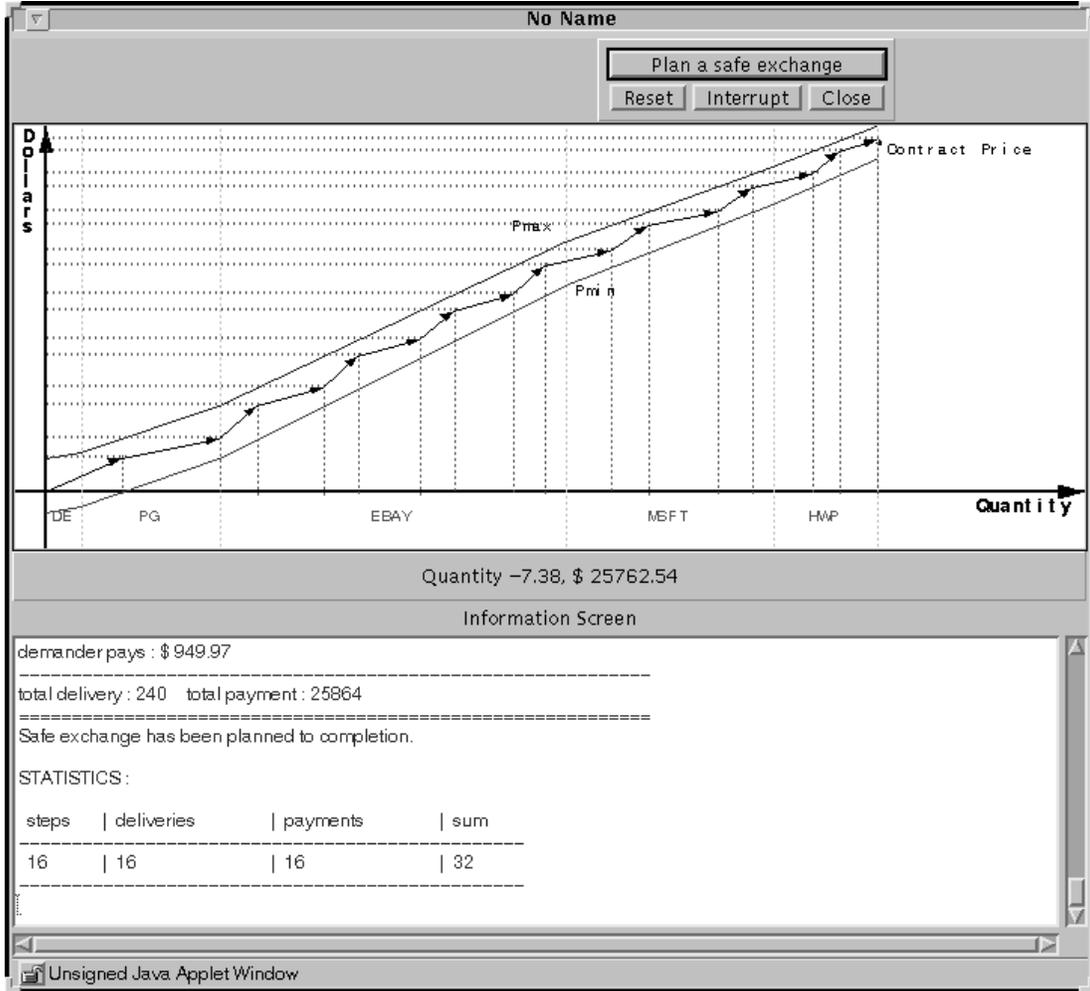


Figure 4: Output from the planner. This optimal plan is 16 steps long. Each step of the exchange is described in detail in the bottom window (only the end of the safe exchange plan, with summary statistics, is visible).

Let the contract price be \$22, and let each agent be willing to forego \$2 to avoid negative reputation effects associated with defecting.

Given this input, and the objective of minimizing the number of exchange steps (where each step could include both a delivery and a payment), *eExchangeHouse* returns the following optimal safe exchange plan. In the first step, the supplier delivers 2 units of item “package” and the demander pays \$5. In the second step, the supplier delivers 1 unit of item “package” and the demander pays \$9. In the third step, the supplier delivers 1 unit of item “package” and the demander pays \$4. In the fourth step, the supplier does not deliver, but the demander pays \$4. In the fifth step, the supplier delivers 1 unit of item “plug-in” and the demander does not pay anything. So, the exchange is completed safely in five steps, containing four deliveries and four payments.

In a recent theoretical paper, we develop a unified model of safe exchange which captures the known safe exchange mechanisms such as cryptographic signatures, cryptographic coin ripping, and our chunking technique (as well as new safe exchange mechanisms). In that general framework we uncover the conditions under which safe exchange is inherently possible or impossible—not just using our chunking mechanism, but using *any* mechanism [40].

5 Conclusions

The *eMediator* electronic commerce server prototype exemplifies several new features that can facilitate more efficient ecommerce in the future. Each one of the three components exhibits an interesting interplay between algorithms and game-theoretic incentive engineering.

eAuctionHouse, the configurable auction server, includes a variety of generalized combinatorial auctions and exchanges, pricing schemes, bidding languages, mobile agents, and user support for choosing an auction type. We introduced two new logical bidding languages for combinatorial markets: the XOR bidding language and the OR-of-XORs bidding language. Unlike the traditional OR bidding language, they are fully expressive. They therefore enable the use of the Clarke-Groves pricing mechanism for motivating the bidders to bid truthfully. *eAuctionHouse* also supports supply/demand curve bidding for single-item multi-unit markets.

eCommitter, the leveled commitment contract optimizer, determines the optimal contract price and decommitting penalties for a variety of leveled commitment contracting mechanisms, taking into account that rational agents will decommit strategically. It also determines the optimal decommitting strategies for any given leveled commitment contract.

eExchangeHouse, the safe exchange planner, enables unenforced anonymous exchanges by dividing the exchange into chunks and sequencing those chunks to be delivered safely in alternation between the buyer and the seller.

Each component can be used separately, but they can also be used together. For example, if a deal is reached using *eAuctionHouse* or *eCommitter*, then *eExchangeHouse* could be used to carry out the deal. As another example, the contract determined using *eCommitter* could be auctioned using *eAuctionHouse*.

In the future we are planning to add to *eMediator* nonmanipulable reputation maintenance algorithms, product evaluation methods, a nonmanipulable voting server, and coalition formation support.

6 Acknowledgments

I thank Qianbo Huai for programming most of the auction house component. I thank Qianbo Huai, Alan Huffman, and Pradeep Gore for implementing the mobile agent component. I thank Sandeep Sikka, Samphel Norden, and Yunhong Zhou for their contributions to analyzing and programming the leveled commitment contract optimizer. I thank Vincent Ferrandon for programming the safe exchange planner.

References

- [1] L. M. Ausubel and P. C. Cramton. Demand reduction and inefficiency in multi-unit auctions. Technical Report 96-07, University of Maryland, Department of Economics, 1996.
- [2] C. Boutilier, M. Goldszmidt, and B. Sabata. Sequential auctions for the allocation of resources with complementarities. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 527–534, Stockholm, Sweden, Aug. 1999.
- [3] B. Chandra and M. M. Halldórsson. Greedy local search and weighted set packing approximation. In *10th Annual SIAM-ACM Symposium on Discrete Algorithms (SODA)*, pages 169–176, Jan. 1999.
- [4] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [5] Y. Fujishima, K. Leyton-Brown, and Y. Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 548–553, Stockholm, Sweden, Aug. 1999.

- [6] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [7] M. M. Halldórsson. Approximations of independent sets in graphs. In K. Jansen and J. Rolim, editors, *The First International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 1–14, Aalborg, Denmark, July 1998. Springer LNCS 1444.
- [8] M. M. Halldórsson and H. C. Lau. Low-degree graph partitioning via local search with applications to constraint satisfaction, max cut, and 3-coloring. *Journal of Graph Algorithms and Applications*, 1(3):1–13, 1997.
- [9] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182:105–142, 1999.
- [10] D. S. Hochbaum. Efficient bounds for the stable set, vertex cover, and set packing problems. *Discrete Applied Mathematics*, 6:243–254, 1983.
- [11] R. M. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, NY, 1972.
- [12] D. Lehmann, L. I. O’Callaghan, and Y. Shoham. Truth revelation in rapid, approximately efficient combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 96–102, Denver, CO, Nov. 1999.
- [13] K. Leyton-Brown, M. Tennenholtz, and Y. Shoham. An algorithm for multi-unit combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Austin, TX, Aug. 2000.
- [14] W. A. Lupien and J. T. Rickard. Crossing network utilizing optimal mutual satisfaction density profile. US Patent 5,689,652, granted Nov. 18 to Optimark Technologies, 1997.
- [15] J. K. MacKie-Mason and H. R. Varian. Generalized Vickrey auctions. Technical report, University of Michigan, 1995.
- [16] R. P. McAfee and J. McMillan. Analyzing the airwaves auction. *Journal of Economic Perspectives*, 10(1):159–175, 1996.
- [17] J. McMillan. Selling spectrum rights. *Journal of Economic Perspectives*, 8(3):145–162, 1994.
- [18] National Consumers League. New NCL survey shows consumers are both excited and confused about shopping online, 1999. www.natconsumersleague.org/BeEWisep.html, Oct. 20. Survey conducted by Opinion Research Corporation.
- [19] N. Nisan. Bidding and allocation in combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 1–12, Minneapolis, MN, 2000.
- [20] D. C. Parkes and L. Ungar. Iterative combinatorial auctions: Theory and practice. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 74–81, Austin, TX, Aug. 2000.
- [21] M. Penn and M. Tennenholtz. Constrained multi-object auctions and b -matching. *Information Processing Letters*, 75(1–2):29–34, July 2000.
- [22] H. Raiffa. *The Art and Science of Negotiation*. Harvard Univ. Press, Cambridge, Mass., 1982.
- [23] E. Rasmusen. *Games and Information*. Basil Blackwell, 1989.
- [24] S. J. Rassenti, V. L. Smith, and R. L. Bulfin. A combinatorial auction mechanism for airport time slot allocation. *Bell J. of Economics*, 13:402–417, 1982.

- [25] J. A. Rodriguez-Aguilar, P. Noriega, C. Sierra, and J. Padget. FM96.5: A Java-based electronic auction house. In *In Proceedings of the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'97)*, 1997.
- [26] M. H. Rothkopf, A. Pekeć, and R. M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.
- [27] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 256–262, Washington, D.C., July 1993.
- [28] T. Sandholm. *Negotiation among Self-Interested Computationally Limited Agents*. PhD thesis, University of Massachusetts, Amherst, 1996. Available at <http://www.cs.cmu.edu/~sandholm/dissertation.ps>.
- [29] T. Sandholm. Issues in computational Vickrey auctions. *International Journal of Electronic Commerce*, 4(3):107–129, 2000. Special Issue on Applying Intelligent Agents for Electronic Commerce. A short, early version appeared at the Second International Conference on Multi-Agent Systems (ICMAS), pages 299–306, 1996.
- [30] T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, Jan. 2002. First appeared as an invited talk at the First International Conference on Information and Computation Economics, 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.
- [31] T. Sandholm and V. Ferrandon. Safe exchange planner. In *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS)*, pages 255–262, Boston, MA, July 2000.
- [32] T. Sandholm and V. R. Lesser. Equilibrium analysis of the possibilities of unenforced exchange in multiagent systems. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 694–701, Montreal, Canada, Aug. 1995.
- [33] T. Sandholm and V. R. Lesser. Leveled commitment contracts and strategic breach. *Games and Economic Behavior*, 35:212–270, 2001. Special issue on AI and Economics. Early versions appeared as *Advantages of a Leveled Commitment Contracting Protocol* in the proceedings of the National Conference on Artificial Intelligence (AAAI), pp. 126–133, Portland, OR, 1996, and as a University of Massachusetts at Amherst, Computer Science Department technical report 95-72.
- [34] T. Sandholm, S. Sikka, and S. Norden. Algorithms for optimizing leveled commitment contracts. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 535–540, Stockholm, Sweden, 1999. Extended version: Washington University, Department of Computer Science technical report WUCS-99-04.
- [35] T. Sandholm and S. Suri. Improved algorithms for optimal winner determination in combinatorial auctions and generalizations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 90–97, Austin, TX, 2000.
- [36] T. Sandholm and S. Suri. Market clearability. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1145–1151, Seattle, WA, 2001.
- [37] T. Sandholm and S. Suri. Optimal clearing of supply/demand curves. In *AAAI-02 workshop on Agent-Based Technologies for B2B Electronic Commerce*, Edmonton, Canada, 2002.

- [38] T. Sandholm, S. Suri, A. Gilpin, and D. Levine. CABOB: A fast optimal algorithm for combinatorial auctions. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1102–1108, Seattle, WA, 2001.
- [39] T. Sandholm, S. Suri, A. Gilpin, and D. Levine. Winner determination in combinatorial auction generalizations. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Bologna, Italy, July 2002. Early version appeared at the AGENTS-01 Workshop on Agent-Based Approaches to B2B, pp. 35–41, Montreal, Canada, May 2001.
- [40] T. Sandholm and X. Wang. (Im)possibility of safe exchange mechanism design. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Edmonton, Canada, 2002.
- [41] T. Sandholm and Y. Zhou. Surplus equivalence of leveled commitment contracts. *Artificial Intelligence*, 2002. To appear. Early versions appeared in ICMAS-00 and in the AAAI-99 Workshop on Negotiation: Settling conflicts and identifying opportunities.
- [42] M. Temenholtz. Some tractable combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Austin, TX, Aug. 2000.
- [43] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
- [44] E. Wolfstetter. Auctions: An introduction. Technical report, Humboldt University of Berlin, 1994.
- [45] P. R. Wurman, M. P. Wellman, and W. E. Walsh. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. In *Proceedings of the Second International Conference on Autonomous Agents (AGENTS)*, pages 301–308, Minneapolis/St. Paul, MN, May 1998.