Auction Protocols for Decentralized Scheduling*

Michael P. Wellman, William E. Walsh
Computer Science and Engineering, University of Michigan, Ann Arbor, MI 48109-2110 USA
(wellman, wew)@umich.edu

Peter R. Wurman
Department of Computer Science, North Carolina State University, Raleigh, NC 27695 USA
wurman@csc.ncsu.edu

Jeffrey K. MacKie-Mason
School of Information, University of Michigan, Ann Arbor, MI 48109-1092 USA
jmm@umich.edu

August 15, 2001

Abstract

Decentralized scheduling is the problem of allocating resources to possible uses over time, with alternative uses represented by autonomous agents. Market mechanisms use prices derived through bidding protocols to determine schedules. We investigate the existence of equilibrium prices for some simple kinds of scheduling problem, the quality of equilibrium solutions, and the behavior of an ascending auction protocol. To remedy the potential nonexistence of price equilibria due to complementarities in preference, we introduce additional markets in combinations of goods. We also consider direct revelation mechanisms, and compare to the market approach.

*Abstracted version of an article appearing in *Games and Economic Behavior* **35**, 2001. This in turn is a revised and extended version of "Some economics of market-based distributed scheduling", presented at the *Eighteenth International Conference on Distributed Computing Systems*, Amsterdam, May 1998.

1 Introduction

Allocating resources with and for distributed computing systems presents particular challenges attributable to the decentralized nature of the computation. Consider, for instance, the problem of scheduling network access for users on the Internet. In such an environment, system modules (user programs) represent independent entities (users) with conflicting and competing scheduling requirements, who may possess localized information relevant to their needs (such as the value they place on a particular schedule). To recognize this independence, we treat the modules as *agents*, ascribing each of them autonomy to decide how to deploy resources under their control in service of their interests. We assume that the agents communicate via messages in which they may convey some of their private information.

The challenges for a decentralized solution to

the scheduling problem then include: How do we manage message passing, reach closure, and determine the final schedule? Further, since the value of alternatives depends on the information held privately by the agents, how do we elicit messages that contain the information needed to formulate a desirable schedule?

The first problem is fundamental in distributed computing systems, due to the asynchrony of communication. A distributed system to solve a scheduling problem based on message-passing needs to specify which messages are admissible (have a well-formed syntax), when they may be sent, and when closure (if ever) will be reached and a schedule formulated.

The second problem is the subject of the theory of mechanism design. Given agents' private information about resources and preferences, and some social welfare critera, some schedules can be considered more desirable than others. Then the problem is to design a *mechanism*: to choose rules for formulating a schedule based on received messages, and possibly for exchanging other resources (e.g., money), that will induce the agents to reveal the private information needed to determine the socially more desirable schedules.

Within this setting, the quality of a decentralized scheduling method can be assessed according to how well it exhibits the following properties:

- Self-interested agents can make effective decisions with local (private) information, without knowing the private information and strategies of other agents.
- The method requires minimal communication overhead.
- The method reaches closure in reasonable time and at reasonable computational expense.
- Solutions do not waste resources. If there is some way to make some agent(s) better off without harming others, it should be done. A solution that cannot be improved in this way is called *Pareto optimal*. (As suggested above, it might sometimes be appropriate to adopt some stronger optimality criteria, based on a judgment about social value of the various agents.)

Straightforward distributed scheduling policies—such as first-come first-served, shortest-job-first, priority-first, and combinations thereof—do not generally possess these properties. For example, queue-position schemes are insensitive to relative value based on the *substance* of the task being performed. On the other hand, priority-based schemes beg the question of how to set priorities so that desirable results follow. If self-interested agents are free to set their own priorities, then without some incentive to the contrary, they will specify maximum priority for whatever they are interested in.

Citing such limitations, several have proposed that distributed resource allocation problems be solved via market mechanisms [8], an approach we have called *market-oriented programming* (MOP) [37]. Assuming that a scheduling problem must be decentralized, markets can provide several advantages:

- Markets are naturally decentralized. Agents make their own decisions about how to bid based on the prices and their own relative valuations of the goods.
- Communication is limited to the exchange of *bids* and *prices* between agents and the

market mechanism. In particular settings, it can be shown that price systems minimize the dimensionality of messages required to determine Pareto optimal allocations [14].

- Since agents must back their representations with exchange offers, some mechanisms can elicit the information necessary to achieve Pareto and global optima (or come within some tolerance of optimal) in some well-characterized situations.

Of course, all of these benefits do not automatically accrue as a result of setting up a market-like environment. Although classic welfare theorems [19] guarantee strong performance for some market mechanisms, these results are formally restricted to rather special environments. Scheduling problems often exhibit complementarities and nonconvexities, which violate the ideal conditions for the welfare theorems or for particular market protocols.

Prior work applying market-inspired mechanisms to scheduling [2, 13, 18, 34, 35] and other distributed resource allocation problems [8, 16, 32, 43] has produced promising empirical results. Understanding the scope of these methods, and developing a general design methodology for computational markets, however, requires an analytical characterization of their properties. In our own MOP work, we have adopted the framework of general equilibrium theory [19], and have found that our computational markets behave predictably when conditions of the theory are met [23, 37, 39]. We have also applied the approach to discrete optimization problems—where the conditions guaranteeing desirable outcomes are not satisfied—and have found (not surprisingly) that the methods sometimes work, and other times break down [36, 38].

Since scheduling problems very often involve discrete (indivisible) resource units, we have undertaken to analyze directly the behavior of computational market mechanisms for such problems. We start by defining a general class of discrete allocation problems, and characterizing some distinctions particularly meaningful in the scheduling domain. We show how some recent results in economic theory apply to the scheduling problem, and report our own extensions and analysis.

In the next section, we motivate the work with a concrete example of a simple factory scheduling problem. In those following, we sketch a model of a general version of the problem, and relate some equilibrium and optimality properties. We then describe and analyze an ascending auction protocol for the scheduling problem. To address limitations of the basic market formulation, we present an extended combinatorial market, and a direct revelation mechanism.

2 A Factory Scheduling Economy

Consider a factory with an unscheduled day shift. There are eight one-hour time slots, labeled 9:00 to 16:00 according to their respective end times. Slots can be allocated for the production of customer orders. The factory has a *reserve price* for each time slot, representing the minimum price that the factory is willing to accept in exchange for that time slot.

Assume each customer agent has one job it wants completed. An agent's job is defined by its duration (length), its deadline, and the value (expressed in dollars) the agent places on the job. An agent is willing to spend up to this value (though of course it prefers to pay less) to com-

Figure 1: A factory scheduling economy. Lines connecting the agents to time slots represent one feasible allocation.

plete its job. To do so, the agent must acquire a number of slots no less than the length (not necessarily contiguous), no later than the deadline. The agent gets no value if its job cannot be completed before its deadline. The global value of a solution is the sum of values of the agents holding the goods, which is the sum of the reserve price for each time slot that was not sold, plus the value associated with each customer agent that meets its job deadline.

Example 1 *The agents are shown in Figure 1. Since the sum of lengths exceeds available factory time, it is not possible for all of the agents to produce their orders. The allocation depicted in Figure 1 represents a global optimum.*

Given an assignment of prices to goods, we can define an agent's optimal choice as a set of slots that complete the job at the minimum cost, or the empty set if the the job could not be completed for less than its value. The reader can verify that at the prices shown in Figure 1, each agent makes a locally optimal choice in the globally optimal allocation.

3 Modeling the Scheduling Economy

We define a general discrete resource allocation problem[1] in terms of agents, goods, and prices for each good.

We assume that agents' valuations can be measured in terms of a common numeraire, which for convenience can be taken to be "money". Therefore, we can directly compare the utility of different agents, and meaningfully treat the sum as a measure of global value. An agent's utility equals the value for the goods it holds, plus its quantity of money.

A *solution* is a mapping indicating which agent, if any, gets each good. Goods not allocated are valued by the seller at their *reserve values*. Different time slots could potentially have different reserve values; for instance, a factory may have a higher reserve price for evening hours to cover overtime expenses.

The *global value* of a solution is the sum of the agent values achieved and the reserve value of goods not used by agents.[2] A solution is *optimal* if no other solution has higher value.

In the remainder of this article, we present market schemes and auction protocols for this very general resource allocation problem. However, the theoretical results and examples we present focus on particular subclasses of scheduling problems in which each agent has one job to complete. For these problems, we associate each agent with a job length and some number of deadlines and corresponding values. The value of a particular allocation to an agent is that of the earliest deadline for which the agent has a sufficient number of slots to complete its job. Note that although the domain of the agent's utility function comprises an exponential number of possible time-slot bundles, for the scheduling problem this value function can be encoded

[1] Complete formal definitions are provided in the full paper. Here we provide conceptual definitions and informal statements of theorems, omitting mathematical notation and proofs.

[2] Because all agents have utility that is linear in money, the total value obtained from money is constant and hence can be ignored.

compactly in terms of a linear number of deadlines and values.

If every agent's job length is one, we call the scheduling problem *single unit*. Problems violating this constraint are *multiple unit*. If each agent has a single deadline, we call the problem *fixed deadline*, otherwise it is *variable deadline*.

4 Price Equilibria

Definition 1 (price equilibrium) *A solution is in* equilibrium *at a given set of prices iff*

1. *All agents are allocated goods maximizing their surplus at these prices.*
2. *All goods are priced at least at their reserve value.*
3. *All unallocated goods are priced at exactly their reserve value.*

Intuitively, this definition states that in equilibrium, each agent (including the seller) gets an allocation that maximizes its utility given the current prices. Equilibria sometimes exist, and are generally not unique. Consider Example 1. The solution shown, with only agent 3 receiving no goods, is in equilibrium at the set of prices suggested, with slots 9:00, 10:00, and 11:00 each having a price of $6.25, and all other slots having a price of $3.25. The same solution is also in equilibrium with respective prices of $6.50 and $3.35, and many other combinations. The equilibrium solution has value $40.50, which is optimal. Indeed it had to be, as demonstrated by the following result.

Theorem 1 *For the general discrete resource allocation problem, if a solution is in equilibrium at some set of prices, than that solution is optimal.*

Name	Job Length	Deadline	Value
Agent 1	2	2	$3
Agent 2	1	2	$2

Table 1: A problem with no equilibrium. Adapted from a demonstration [21] that price equilibria may not exist in the FCC market for radio spectrum.

This result confirms the usual consequence of competitive equilibrium: that no further gains from trade are possible and so the result is Pareto optimal. Since we assume that agent values are expressible in price units, Pareto optimality corresponds to global optimality.

Example 2 *There are two agents as described in Table 1, and the reserve price of each good is zero.*

The optimal solution, where agent 1 gets both slots, is not in equilibrium at any prices, and indeed no equilibrium exists in this case. Any equilibrium must price both goods at $2 or more, otherwise agent 2 would demand one of the goods. But if these inequalities hold then agent 1 would not demand the two time slots it requires.

In this example, the nonexistence of equilibrium prices is due to *complementarities* in agent preferences. Agent 1 considers the two time slots complementary in that it values one iff it has the other. Complementarities cannot arise in the single-unit scheduling problem.

Theorem 2 *Any optimal solution to the single-unit scheduling problem (fixed or variable deadline) is supported by a price equilibrium.*

Together, Theorems 1 and 2 establish that a solution to the single-unit scheduling problem is

optimal iff it is supported by a price equilibrium. Example 2 demonstrates that relaxing the single-unit restriction immediately leads to the possibility that an equilibrium will not exist. This holds generally, as Milgrom's result that a single complementarity is sufficient to prevent a price equilibrium [22] applies to our scheduling problem.

In addition to the single-unit restriction of Theorem 2, we can identify a few other conditions that guarantee the existence of equilibrium [5, 15]. One example would be if agents had additive preferences for completing multiple single-unit jobs. Unfortunately, other known conditions do not seem to be immediately expressible in scheduling terms.

Finally, note that earlier slots must be at least as valuable as later ones.

Theorem 3 *In an equilibrium for the scheduling economy, prices for differently allocated time slots must be nonincreasing with their time indices.*

5 Auction Protocols

We use the term *protocol* to refer to a *mechanism*, along with agent *bidding policies*. The mechanisms we consider are generically called *auctions*. McAfee and McMillan provide the following definition [20]:

> An auction is a market institution with an explicit set of rules determining resource allocation and prices on the basis of bids from the market participants.

This definition includes the well known English open-outcry and first-price sealed bid auctions—commonly used to sell art and to award procurement contracts, respectively—as well as a broad range of other mechanisms, including fixed pricing, Dutch auctions, Vickrey auctions, commodities markets, and the ascending, combinatorial, and Generalized Vickrey auction schemes described in Sections 6 through 8.

In order to place some structure on the space of possibilities, and also to provide a common interface to agents, we define a somewhat restricted, but still very general auction mechanism.

1. Agents send bids to the mechanism to indicate their willingness to exchange goods.
2. The auction may post *price quotes* to provide summarized information about the status of the price-determination process.

   Steps 1 and 2 may be iterated.
3. The auction determines an allocation and notifies the agents as to who purchases what from whom at what price.

The above sequence may be performed once or repeated any number of times.

Auctions can be differentiated across many parameters including, but not limited to, those concerning: matching algorithm, price determination algorithm, event timing, bid restrictions, and intermediate price revelation [24, 29, 42]. One of the most important distinctions is whether an individual auction allocates a single resource, or several at once. The latter type, called *combinatorial auctions*, accept bids referring to combinations of basic goods.

In order to predict auction outcomes, we must consider the agents' presumed bidding policies, which in turn we might base on some model of their beliefs and preferences. In some auction contexts we are able to determine analytically that a particular bidding policy is part of

a Bayesian-Nash equilibrium, or even the dominant strategy. In other settings we rely on experimentation and rules of thumb based on economic principles to determine reasonable bidding policies.

The auction mechanisms we discuss are decentralized in the sense that each agent calculates its own bidding strategy, based on local information. Single-resource auctions, as in the ascending auction protocol, are further distributed in that allocations for each good can be computed separately.

6 Ascending Auction

We define the ascending auction protocol for the general discrete resource allocation problem. Separate auctions determine prices for each of the goods. Agents submit successively higher bids to the auctions, and auctions immediately report price quotes to all interested agents upon receiving a bid. We allow that agents and auctions operate asynchronously, that is, we impose no bound on the relative times that agents take to compute and send bids, or auctions to compute and send price quotes. Nor do we assume any ordering on bid actions, other than explicitly stated. (As a consequence of this asynchrony and flexible ordering, the protocol is nondeterministic.) When the bidding stops (i.e., the protocol reaches *quiescence* [40]), each auction allocates its respective good to the highest bidder at the price the agent bid, or the good is retained by the seller if there are no bids.

6.1 Bidding Rules

At any point in time, the *bid price* in an auction is the highest price offered thus far. If an auction has received no bids, its bid price is undefined. An auction's *ask price* is its bid price plus an increment ϵ. If the bid price is undefined, the ask price is the reserve.

The ascending auction rejects any bid less than its ask price. Agents are not allowed to withdraw bids. An agent may replace its bid with another, but the new bid must be at least the current ask price. These rules guarantee that prices do not decrease and that the bidding process terminates.

6.2 Agent Bidding Policies

When an agent enters the market, it bids the ask prices for the set of goods that maximizes its surplus, based on the current ask prices (breaking ties arbitrarily). As other agents continue, the agent may be outbid for some goods. When this occurs, the agent bids the ask price on the set of goods that maximizes its surplus, assuming that it can obtain the goods it is currently winning at their bid prices. For the single-unit scheduling problem, whenever an agent is not already winning a bid, it simply bids the ask price for the single good that maximizes its surplus at the ask prices. If no good would provide it with a positive surplus, then the agent "drops out" of the auction. For the general multi-unit variable-deadline problem, the surplus-maximizing bids can be computed straighforwardly.

This bidding strategy is quite simple, involving no anticipation of other agents' strategies. For the single-unit problem, such anticipation is unnecessary, as the agent would not wish to change its bid even after observing what the other agents did. This is called the *no regret* property [5], and means that from the agent's perspective, no bidding policy would have been a better response to the other agents' bids. The no-regret property does *not* hold, however,

Name	Job Length	Deadline	Value
Agent 1	2	2	\$20
Agent 2	2	3	\$8
Agent 3	1	3	\$2

Table 2: A multiple-unit problem (Example 3).

for the ascending auction in the multiple-unit scheduling problem, regardless of the bidding strategy [5]. In general, an agent might perform better, for example, through accurate prediction of the other agents' behavior. In the absence of a basis for prediction, however, the simple strategy proposed may indeed be reasonable.

6.3 Analysis of the Ascending Auction

Under the ascending auction rules, when the protocol reaches quiescence, the price equals the bid price if defined, otherwise it equals the reserve.

It is possible that the ascending auction can determine prices that differ from an equilibrium of a multiple-unit scheduling economy by arbitrarily large amounts.

Example 3 *The bid increment is* $\epsilon = \$1$ *and the reserve prices are zero. The agents are described in Table 2.*

Although there are many equilibrium price sets (one of which is \$8, \$8, and \$1 for the respective goods), the ascending auction may not find an equilibrium. Agent 2 could bid up good 3 until its ask price exceeds \$2, while it and agent 1 both bid up the prices on 1 and 2. The reader can verify that any equilibrium must have agent 3 winning good 3 at a price no greater than \$2.

In the multiple-unit scheduling problem, the ascending auction can produce allocations that are arbitrarily far from optimal.

Name	Job Length	Deadline	Value
Agent 1	1	1	\$3
Agent 2	2	2	\$11

Table 3: A multiple-unit problem (Example 4).

Example 4 *There are two agents as shown in Table 3. Reserve prices are \$1 and \$9, respectively, and the bid increment is* $\epsilon = \$1$.

If agent 2 places its bids first, it will bid \$1 for 1 and \$9 for 2. Agent 1 will then bid \$2 for 1. The bidding will stop with good 1 allocated to agent 1 and good 2 allocated to agent 2. This solution has a value of \$3 yet the optimal solution, with 2 unallocated, has a value of \$12. It is easy to see—by increasing the reserve for good 2 and agent 2's valuation by the same amount—that the ascending auction can produce a solution that is arbitrarily far from optimal.

If we restrict each agent's requirement to a single time slice, then by Theorem 2 an equilibrium exists. However, the ascending auction protocol is not guaranteed to reach an equilibrium even with this restriction.

Example 5 *The bid increment is* $\epsilon = \$1$. *The respective reserve prices are \$4, \$3, and \$3. The agents are described in Table 4.*

It is possible that agent 2 may bid first, for 2. Then the ask price for good 2 is \$4. Agent 1 will then bid \$4 for either 1 or 2. If it bids for 1 then the bidding will stop and agent 1 will win 1 for \$4 and agent 2 will win 2 for \$3. But since the price for good 2 is less than that for good 1, agent 1 would maximize its surplus by demanding 2 at the final prices.

Name	Job Length	Deadline	Value
Agent 1	1	2	\$6
Agent 2	1	3	\$7

Table 4: A single-unit problem (Example 5).

It is not hard to see that the potential failure to reach equilibrium can be demonstrated for any positive value of ϵ, no matter how small. Nevertheless, unlike the multiple-unit problem, we can bound the distance from the equilibrium price vector.

Theorem 4 *For the variable-deadline, single-unit scheduling problem, the final price of any good determined by ascending auction protocol will differ from the unique minimum equilibrium prices by at most ϵ times the number of agents (or the number of slots, if there are more agents than slots).*

Consider again Example 5. The solution exhibited has a value of \$16. If agent 1 had received good 2 and agent 2 had received good 3 then the value of the solution would be \$17, which is optimal. However, the solution can be suboptimal by only a bounded amount, as established by Bertsekas [4] in a more general setting.

Theorem 5 *The ascending auction protocol produces a solution to the variable-deadline, single-unit scheduling problem that is suboptimal by at most an amount linear in ϵ and quadratic in the number of agents or slots.*

Computing the clearing and price quotes is trivial in the ascending auction. Communication costs dominate the run time, which can therefore be measured in terms of the bids required. Because bids increase by a fixed increment, the number of iterations is inversely proportional to ϵ. Hence, in choosing the value for ϵ, we trade off solution value for communication efficiency.

We have shown that the simple bidding policy is reasonable for individual agents, and produces allocations with desirable system properties in the single-unit problem. The results do not provide strong support for this simple policy in the multiple-unit problem. Other strategies, such as jump bidding—where an agent bids in large increments for sets of goods to signal its willingness to aggressively pursue that set—may provide potential advantages to individuals or the system. However, it is an open question as to whether there exists a policy for the ascending auction (or any complete protocol) that always finds (within some tolerance) an equilibrium when it exists.

6.4 Incremental Auction Closing

In the basic version of the ascending auction mechanism, we close the auctions simultaneously, once the bidding process reaches quiescence. In a variant, we close one or a few at a time, reopening the bidding process after each close. Once an auction closes, the commitment of the winning bidder to buy the good is finalized, and the price paid constitutes a sunk cost. This may cause the bidder to reassess its decisions about other goods, and bid in auctions out of which it had previously dropped.

Example 6 *Reconsider Example 4, with agents described by Table 3, and reserves of \$1 and \$9, respectively. As pointed out above, the ascending auction may reach a solution with good 1 allocated to agent 1 for \$2 and good 2 allocated to agent 2 for \$9. This solution has value \$3.*

If we close the auction for good 2, then agent 2

treats its payment as sunk, and so now values good 1 at \$11. Therefore, with bidding on good 1 reopened, it will clearly outbid agent 1. If we instead close good 1 and reopen bidding on good 2, nothing changes.

In this example, the resulting allocation has value \$11—still suboptimal, but an improvement over the original allocation. Indeed, it can be shown that incrementally reopening bidding on the last good can only improve solution value. Indeed, in broad classes of situations where the ascending protocol is known to work well, we can show that incremental closing will not degrade quality. In general, however, reopening bidding after some auctions close can have positive or negative effects. Which is the case depends critically on the order that auctions close. Unfortunately, in general one cannot tell which order will be advantageous, without knowing the agents' private information.

7 Combinatorial Markets

The ascending auction performs well for single-unit allocation problems. Even when equilibria exist for a multiple-unit problem, the ascending auction may not find one, as shown by Example 4. Further, as Example 2 demonstrates, many scheduling problems cannot support allocations with any price equilibrium.

In light of these limitations, several have proposed *combinatorial auction* mechanisms, where agents submit bids for combinations of goods [3, 6, 17, 28, 30, 41]. Such mechanisms operate in a variety of ways, typically calculating allocations and prices as a function of bids for all the combinations. Prices may refer to individual goods, or to entire bundles.

One of the drawbacks of combinatorial auctions is their potential computational complexity. With n goods, there are 2^n combinations, which can entail complex calculations for both the agents and the mechanism. Typical formulations of the mechanism's computational problem are NP-Complete, but may admit heuristic search procedures effective in practice in certain environments [1, 10, 31]. As Rothkopf et al. [30] point out, moreover, restricting the set of allowed combinations can preserve computational tractability.

We pursue a strategy similar in spirit to the restriction approach, presenting a reformulation of the problem that extends the price system in a controlled way, without admitting an exponential number of markets. The reformulation expands the class of problems solved by price equilibria, and suggests corresponding auction protocols for determining these prices.

7.1 Problem Formulation

In the revised combinatorial formulation, rather than impose a price system over the basic goods, however, we introduce an expanded set of *market goods*. A market good is a pair, (y, z), denoting a bundle of y time slots no later than time z. The price system for this formulation assigns prices to all (y, z) combinations (a quadratic number of prices). We generally require that prices be *monotone* in the number of time slots, y, for any given upper bound z.

A solution is defined as in the original formulation: a mapping from basic goods to agents. A *market allocation* is an assignment of market goods to agents. We say that a market allocation *is consistent with* a solution if that solution gives each agent what it is promised.

Note that although a scheduling agent obtain-

ing a market good cannot be sure exactly *which* time slots it will receive, its utility is completely determined by whether it obtains enough time slots to finish its job, and if so, by what deadline.

Definition 2 (temporally consistent prices) *A monotone price function is* temporally consistent *if*

1. *Increasing the slot bound (z) does not increase the price.*
2. *It is not possible to break a market good into smaller pieces for a lower total price.*

7.2 Equilibrium and Efficiency

Definition 3 (combinatorial price equilibrium) *A market allocation is in* equilibrium *at a given set of prices iff*

1. *All agents are allocated market goods maximizing their surplus at these prices.*
2. *The prices of all market goods is at least the sum of minimal reserve prices for constituent basic goods.*
3. *There is some solution consistent with the allocation such that the prices for unallocated basic goods is as small as possible.*

The conditions relating market prices to reserve prices are complicated by the indeterminate relationship between market and basic goods. We require that the price of a market good be at least the minimum consistent reserve price, else the sellers would not part with the constituent basic goods. For market goods actually allocated, we require the price to (weakly) exceed the reserves of basic goods comprising it in a consistent solution. And when a market good could be satisfied by basic goods unallocated in this solution, the reserves of those goods define an upper bound on its price.

Example 7 *Reconsider Example 2, with parameters illustrated by Table 1, and zero reserve prices. Although no price equilibrium exists for the original formulation, we can support the optimal solution with a combinatorial price equilibrium. Consider prices for 1 slot (by either 1 or 2) of 2.1, and a price for both slots of 2.9. The allocation giving both slots to agent 1 satisfies the combinatorial equilibrium conditions at these prices. Note that no combinatorial equilibrium can support any other allocation.*

Unlike in the basic configuration, however, combinatorial price equilibria are not necessarily efficient.

Example 8 *Consider an extension of the previous example, described by Table 5 (with zero reserve prices). The problem has a basic equilibrium, with the price of each good at 1.6, and agents 2 and 3 each getting one of the slots. This optimal solution is also supported by the combinatorial equilibrium with prices 1.6 for one slot, and 3.2 for both. However, the nonoptimal solution where agent 1 gets both slots is also in equilibrium, at prices 2.1 for one slot, and 2.9 for both.*

Moreover, the degree of suboptimality is not usefully bounded—even without reserve prices. We can extend Example 8 to obtain n-agent problems where equilibrium solutions are a factor of $n-1$ worse than optimal. On the positive side, we have the following results.

Theorem 6 *Optimal solutions supported by price equilibria in the original formulation are*

Name	Job Length	Deadline	Value
Agent 1	2	2	\$3
Agent 2	1	2	\$2
Agent 3	1	2	\$2

Table 5: A problem with both optimal and sub-optimal combinatorial equilibria.

retained (albeit not uniquely) in the combinatorial formulation.

Theorem 7 *For any scheduling problem with monotone reserve prices, there exists an equilibrium supporting an optimal allocation.*

7.3 Combinatorial Auction Protocols

In future work, we intend to define and analyze combinatorial protocols analogous to the ascending auction. A straightforward implementation of this protocol is not well-defined for the combinatorial case, as basic goods may be assigned to various market goods. Accordingly, we must define allocation and price quote policies as a function of combinations of bids for alternative goods, not just individual market goods. Such protocols have been proposed for the case of general combinatorial bidding by Parkes [27] and Wurman [41]. We also plan to investigate adaptations to this scheduling problem of other mechanisms proposed for combinatorial settings [3, 6, 9, 28].

Since combinatorial auctions can support sub-optimal equilibrium solutions, it can be disadvantageous to open combinatorial markets when equilibria exist in basic goods. A natural approach would be to start with markets in basic goods, and open combinatorial auctions only if the protocol does not reach equilibrium. We can apply this incrementally, progressively increasing the maximum number of slots per market good until an equilibrium is reached. Of course, this presumes we have a way to detect equilibrium states, or at least some indication of whether opening additional combination markets will be beneficial.

8 Generalized Vickrey Auction

The preceding analysis characterizes the performance of multiple ascending single-good auctions for the scheduling problem, and the prospects for combinatorial auctions. Neither is guaranteed to produce optimal solutions to all scheduling problems. Another mechanism, the Generalized Vickrey Auction (GVA) [33], does find efficient schedules for all of our problems. Although the main results are not new, we briefly present the GVA because it takes an important place in our spectrum of mechanisms for scheduling.

The GVA is a direct revelation mechanism (DRM), and thus is not a price system. Rather, it computes overall payments for agents' allocations that sometimes, but not always, translate into meaningful prices for individual goods. If agents play Bayesian-Nash or dominant strategies, any desirable choice function that can be implemented by a mechanism can be implemented by a DRM, so the appeal of this type of mechanism is quite powerful.[3] The GVA can implement optimal solutions for multi-unit

[3] Specifically, the GVA is a DRM relying on dominant strategies in the class of Groves [12] and Clarke [7] mechanisms. Green and Laffont have shown under rather general conditions that the only efficient social choice functions that are implementable in dominant strategies are those that are implementable by Groves-Clarke mechanisms [11].

scheduling problems, as well as extended versions of the problem involving, for example, multiple jobs or externalities (i.e., values for one agent that depend on the allocations obtained by other agents).

8.1 Bidding Rules

As in any DRM, each agent announces its alleged utility function–in the case of our scheduling problem, its job length and deadline-value correspondence. The agent is not constrained to be truthful by the rules per se (but will have incentives to do so as we see below). The auction knows the reserve values. After receiving the bids, the GVA returns an allocation, and a vector of positive or negative payments to be made to the agents.

8.2 Allocation Rules and Optimality

Given the agents' reports, the auction mechanism computes the optimal solution, under the assumption that agents reported their true values. It then computes payments to agents. Specifically, each agent receives the value of the computed solution to all of the *other* agents (based on reported values, and including reserves of the seller), minus the value of the best solution that would be obtainable if this agent did not exist.

Given this allocation rule, truthful bidding of the utility function is a dominant strategy [33]. The GVA computes the optimal allocation based on the bids, and since all bids are truthful, the allocation is globally optimal.

Example 9 *Consider the setup of Example 8 (Table 5). If the agents truthfully report their value functions, the auction mechanism finds an optimal solution: agent 2 gets the first slot, and agent 3 the second. Agent 1 pays nothing, agents 2 and 3 pay \$1, and the mechanism has a net revenue of \$2.*

Recall that on this same example, the single-good ascending auction protocol does not guarantee convergence, and the combinatorial scheme admits inefficient equilibria.

8.3 Limitations on the GVA

A mechanism is *individually rational* if no agent can be worse off from participating in the auction than had it declined to participate. A mechanism is *budget balanced* if the net payment over all agents is nonnegative. Generally, these, along with optimality, are the properties we desire when agents play their equilibrium strategies in a mechanism. However, Groves-Clarke mechanisms are not always guaranteed to be budget-balanced: they may require an outside injection of resources (subsidy). Alternatively, even when a social choice function can be implemented with a budget-balanced Groves-Clarke mechanism, we cannot guarantee that rational agents will agree ex post to participate in the allocation.

For our scheduling problem, we can show that it *is* possible to obtain all three desirable properties if the mechanism designer knows the reserve value for each time slot. The payment function transfers to each agent the net value increment to all others that results from that agent's participation in the auction. An agent's only effect on others is that it may get time slices that others desire, so its participation always makes other agents weakly worse off. Thus, agents always pay nonnegative amounts, and the auction mechanism runs a surplus. In this situation the GVA is essentially an ideal mechanism for the

scheduling problem, if the computational cost is not too high.

Theorem 8 *For the GVA as described above, the individual rationality constraint is satisfied and the net monetary payments to the auction mechanism are nonnegative.*

If instead the reserve prices are the private information of seller agents, then the mechanism needs to elicit this information to satisfy all of the desirable properties. Myerson and Satterthwaite [25] proved that no mechanism can obtain more than two out of the three desired properties for bilateral exchange problems.

Example 10 (Bilateral exchange) *Suppose there is one buyer, who has a single-unit job with deadline 1 and value v. Let the seller be an agent, with reserve value q. Suppose $v > q$. The GVA would induce truthful reporting of v and q, give the good to the buyer, require the buyer to pay q, and pay v to the seller. Although the mechanism is individually rational and would produce the optimal allocation, the auction would run a deficit of $v - q$.*

In any case, it is easy to show that for all scheduling problems in our class, it is possible to achieve any two out of three of the desirable properties.

8.4 GVA Computation

For a general problem, the heart of the GVA allocation mechanism requires the auction to solve a possibly complex (e.g., nonlinear, nonconvex, integer-constrained) optimization problem multiple times. As a baseline for computational efficiency, we note that Neapolitan and Naimipour [26] show that a simple centralized greedy algorithm solves the single-unit, fixed-deadline scheduling problem optimally, in time $\Theta(m \lg m)$. The GVA mechanism must solve multiple optimization problems to process the bids: one to determine the optimal allocation, and one for each agent with its bid removed to determine the agent payments. For a single-unit, fixed-deadline problem we can use the centralized algorithm for each optimization, with a total runtime of $\Theta(m^2 \lg m)$. Thus, inducing preference revelation (and thereby obtaining full optimality) via straightforward implementation of the GVA raises computational cost by a factor of m.

If we remove the single-unit restriction, then any centralized algorithm that can solve the scheduling problem optimally can solve the Integer Knapsack problem. Hence the multiple-unit scheduling problem is NP-Complete. By the preceding argument, distributing the multiple-unit problem via the GVA contributes a factor of m to the computation.

9 Discussion

We have presented two auction mechanisms—ascending single-good markets and the GVA—that can compute optimal or near-optimal solutions to the single-unit distributed scheduling problem in a computationally efficient manner. The multiple-unit problem is significantly more difficult and entails a sharper tradeoff among solution quality, computational efficiency, and the degree to which the mechanism is decentralized. The computation performed by the ascending auction is trivial, and can be distributed by goods. However, we cannot guarantee the quality of solutions produced by this mechanism for the multiple-unit problem. Combinatorial auc-

tions support equilibria in cases where single-good markets do not, but may also admit sub-optimal solutions. It remains to be seen whether we can design mechanisms for combinatorial auctions that produce desirable outcomes for plausible agent behavior. The GVA always finds the optimal solution and implements it in dominant strategies, but must in general solve multiple combinatorial problems, and may require a subsidy when seller reserves are not known.

The three categories of mechanisms investigated here can be viewed on a spectrum,

single-good $\leftrightarrow$ combinatorial $\leftrightarrow$ direct revelation

where the mechanism's *scope of concern* increases as we move to the right. Toward the left, the overall mechanism decomposes into sub-mechanisms, where each sub-mechanism has a limited scope (i.e., subsets of the resources, ultimately singletons). For large-scale systems, we suspect that this decomposition is essential, as no single designer will even be aware of all of the resources of interest to some of the agents. Even when we imagine that all concerns are covered (as for direct mechanisms), the very use of monetary payments suggests that there exist some other concerns not included (else what use is money?), but assumed separable. Thus, we suspect that mechanisms operating at all points of the spectrum will play a role in computational markets for complex allocation problems.

We view this work as a first step in developing a broad framework for using markets to solve distributed scheduling problems. In order to move forward we must identify broader classes of scheduling problems and design associated mechanisms for which we can effectively predict agent behavior and analyze resultant protocols. We do not expect to find a single mechanism that reaches an optimal equilibrium in all situations where such equilibria exist. Instead we aim to produce a suite of mechanisms that collectively cover a broad range of problems. Ideally, we would like to be able to choose a mechanism for a given problem and know that it will reach an optimal solution when one would be supportable, or else perform acceptably in some other respect when this is not possible.

Acknowledgments

Thanks to Wolfram Conen, Fredrik Ygge, David Pennock, Terence Kelly, and anonymous reviewers for constructive comments. This work was supported by DARPA grant F30602-97-1-0228, and National Science Foundation grant IIS-9988715.

References

[1] Arne Andersson, Mattias Tenhunen, and Fredrik Ygge. Integer programming for combinatorial auction winner determination. In *Fourth International Conference on Multiagent Systems*, Boston, 2000.

[2] Albert D. Baker. Metaphor or reality: A case study where agents bid with actual costs to schedule a factory. In Clearwater [8].

[3] Jeffrey S. Banks, John O. Ledyard, and David P. Porter. Allocating uncertain and unresponsive resources: An experimental approach. *RAND Journal of Economics*, 20:1–25, 1989.

[4] Dimitri P. Bertsekas. Auction algorithms for network flow problems: A tutorial intro-

duction. *Computational Optimization and Applications*, 1:7–66, 1992.

[5] Sushil Bikhchandani and John W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74:385–413, 1997.

[6] Paul J. Brewer and Charles R. Plott. A binary conflict ascending price (BICAP) mechanism for the decentralized allocation of right to use railroad tracks. *International Journal of Industrial Organization*, 14:857–886, 1996.

[7] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[8] Scott Clearwater, editor. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, 1995.

[9] Christine DeMartini, Anthony M. Kwasnica, John O. Ledyard, and David Porter. A new and improved design for multi-object iterative auctions. Technical report, California Institute of Technology, November 1998.

[10] Yuzo Fujishima, Kevin Leyton-Brown, and Yoav Shoham. Taming the computational complexity of combinatorial auctions. In *Sixteenth International Joint Conference on Artificial Intelligence*, pages 548–553, Stockholm, 1999.

[11] J. R. Green and J.-J. Laffont. *Incentives in Public Decision Making*. North-Holland, 1979.

[12] Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[13] Patrick T. Harker and Lyle H. Ungar. A market-based approach to workflow automation. Presented at the NSF Workshop on Workflow and Process Automation in Information Systems (Athens, GA, May 8-10). Available at `http://www.cis.upenn.edu/~ungar/NSF.html`, 1996.

[14] J. S. Jordan. The competitive allocation process is informationally efficient uniquely. *Journal of Economic Theory*, 28:1–18, 1982.

[15] Alexander S. Kelso and Vincent P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50:1483–1504, 1982.

[16] James F. Kurose and Rahul Simha. A microeconomic approach to optimal resource allocation in distributed computer systems. *IEEE Transactions on Computers*, 38:705–717, 1989.

[17] Erhan Kutanoglu and S. David Wu. On combinatorial auction and Lagrangean relaxation for distributed resource scheduling. *IIE Transactions*, 31:813–826, 1999.

[18] T. W. Malone, R. E. Fikes, K. R. Grant, and M. T. Howard. Enterprise: A market-like task scheduler for distributed computing environments. In B. A. Huberman, editor, *The Ecology of Computation*, pages 177–205. North-Holland, 1988.

[19] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.

[20] R. Preston McAfee and John McMillan. Auctions and bidding. *Journal of Economic Literature*, 25:699–738, 1987.

[21] R. Preston McAfee and John McMillan. Analyzing the airwaves auction. *Journal of Economic Perspectives*, 10(1):159–175, 1996.

[22] Paul Milgrom. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108:245–272, 2000.

[23] Tracy Mullen and Michael P. Wellman. A simple computational market for network information services. In *First International Conference on Multiagent Systems*, pages 283–289, San Francisco, CA, 1995.

[24] Tracy Mullen and Michael P. Wellman. Market-based negotiation for digital library services. In *Second USENIX Workshop on Electronic Commerce*, pages 259–269, Oakland, CA, 1996.

[25] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.

[26] Richard E. Neapolitan and Kumarss Naimipour. *Foundations of Algorithms*. D. C. Heath and Company, Lexington, MA, 1996.

[27] David C. Parkes. *i*bundle: An efficient ascending price bundle auction. In *ACM Conference on Electronic Commerce*, pages 148–157, Denver, 1999.

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

[29] Juan A. Rodríguez-Aguilar, Francisco J. Martín, Pablo Noriega, Pere Garcia, and Carles Sierra. Competitive scenarios for heterogeneous trading agents. In *Second International Conference on Autonomous Agents*, pages 293–300, Minneapolis, 1998.

[30] Michael H. Rothkopf, Aleksander Pekeč, and Ronald M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44:1131–1147, 1998.

[31] Tuomas Sandholm. An algorithm for optimal winner determination in combinatorial auctions. In *Sixteenth International Joint Conference on Artificial Intelligence*, pages 542–547, Stockholm, 1999.

[32] Michael Stonebraker, Paul M. Aoki, Witold Litwin, Avi Pfeffer, Adam Sah, Jeff Sidell, Carl Staelin, and Andrew Yu. Mariposa: A wide-area distributed database system. *VLDB Journal*, 5:48–63, 1996.

[33] Hal R. Varian and Jeffrey K. MacKie-Mason. Generalized Vickrey auctions. Technical report, Department of Economics, University of Michigan, June 1994.

[34] Carl A. Waldspurger, Tad Hogg, Bernardo A. Huberman, Jeffrey O. Kephart, and Scott Stornetta. Spawn: A distributed computational economy. *IEEE Transactions on Software Engineering*, 18:103–117, 1992.

[35] Carl A. Waldspurger and William E. Weihl. Lottery scheduling: Flexible proportional-share resource management. In *Proceedings of the First Symposium on Operating System Design and Implementation (OSDI)*, pages 1–11, 1994.

[36] William E. Walsh and Michael P. Wellman. Efficiency and equilibrium in task allocation economies with hierarchical dependencies. In *Sixteenth International Joint Conference on Artificial Intelligence*, pages 520–526, Stockholm, 1999.

[37] Michael P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993.

[38] Michael P. Wellman. A computational market model for distributed configuration design. *AI EDAM*, 9:125–133, 1995.

[39] Michael P. Wellman. Market-oriented programming: Some early lessons. In Clearwater [8].

[40] Michael P. Wellman and William E. Walsh. Distributed quiescence detection in multiagent negotiation. In *Fourth International Conference on Multiagent Systems*, pages 317–324, Boston, 2000.

[41] Peter R. Wurman. *Market Structure and Multidimensional Auction Design for Computational Economies*. PhD thesis, University of Michigan, 1999.

[42] Peter R. Wurman, Michael P. Wellman, and William E. Walsh. A parametrization of the auction design space. *Games and Economic Behavior*, 35:304–338, 2001.

[43] Fredrik Ygge and Hans Akkermans. Decentralized markets versus central control: A comparative study. *Journal of Artificial Intelligence Research*, 11:301–333, 1999.