Economically Coordinated Job Shop Scheduling and Decision Point Bidding - An Example for Economic Coordination in Manufacturing and Logistics -

Wolfram Conen
XONAR GmbH
Wodanstr. 7, 42555 Velbert, Germany
conen@gmx.de

Abstract

We discuss the application of economic coordination mechanisms to scheduling problems in manufacturing and logistics. We use job shop scheduling as a sample problem domain. We study economically augmented job shop problems (EJSP) which comprise valuation information. We demonstrate how instances of EJSP can be mapped to combinatorial job shop auction problems (CJSAP). A discussion of the tractability of approaches that determine optimal solutions identifies the need for heuristics. We suggest a heuristic that combines the merits of economic coordination with the benefits of closely domain-related solution procedures. We conclude that the application of newly suggested and/or well-known economic coordination mechanisms to job shop and related scheduling problems in manufacturing and logistics seem promising.

Keywords: Scheduling, Economic Coordination, Combinatorial Auctions

1 Introduction

The work presented in the following is based on some of the key assumptions of research in economics and game theory. A general description of the scenario under study could be given as follows:

Scenario 1. *A set of actors decides to collaborate. The actors own a set of resources. The primary objective of the collaboration is to utilize this set of resources as economically efficient as possible. Economic efficiency is measured in terms of money. Each actor has individual preferences for resource allocations. This preference information is (to a large extend) private to the actor—it cannot be assumed that the actors reveal their preferences truthfully without an incentive to do so. Paying tribute to the objective and the restrictions, a coordination mechanism*[1] *is designed to repeatedly determine allocations of the available resources to maximize efficiency. Each actor is free to opt out of the collaboration if he is (repeatedly) dissatisfied with the determined outcomes. Each actor acts rationally (within certain limitations) as an utility maximizer.*

To transfer this scenario to a multi agent setting is straightforward: Each actor or each group of actors can be represented by a (potentially computerized) agent. Furthermore, agents may represent parts of the necessary organizational infrastructure (namely the manifested parts of the coordination mechanism). We primarily use agents as a convenient metaphor for the study of coordination issues. In addition, issues relevant for an implementation of the abstract infrastructure, namely computational tractability and communication complexity, are considered. Due to the conceptual proximity of modeled and implemented multi-agent systems, carrying over some of the theoretical results to the implementation is straightforward.[2]

The scenario emphasizes the relevance of the agents' self-interest. We do not assume that the decision to collaborate with other agents implies that an agent is willing to reveal all his private information or that he is willing to deviate dramatically from his utility-maximizing behavior and turns into an altruist. The reasons for the formation of collaborations can be manifold—most relevant to our analysis is the assumption that the actors decide to become part of the collaboration because they consider this to be the economically most beneficial behavioral alternative, given their expectations and their attitude towards risk. The assumption of rational behavior implies that the agents re-assess this evaluation continually and, consequently, may decide to leave the collaboration. This also explains, why an objective of the design of the coordination mechanism is to give incentives to participate in the mechanism continually.

Also note that this scenario is an abstraction of many real-world scenarios in logistics and manufacturing: a number of units compete for the utilization of resources that are owned by the group of units (or by some larger unit that encloses/owns the considered units). Each unit is, to a certain extent, economically independent in its decisions and has information that is often not only local but also private in the strict sense: it is only communicated truthfully to other agents if an incentive to do so is present. This situation can be found in virtual or extended enterprises (consider, for example, a collaboration of small and medium-size transportation companies with overlapping transportation capacities) or even in comparatively small production environments such as a single workshop (groups of workers, sales people, shop-floor managers, planners influence as actors with self-interest and (conflicting) individual objectives the orga-

[1] To explain the use of the term *coordination*: The mechanism is executed to coordinate the interests of the actors and to determine an outcome that implements a coordination of the behaviors of the agents—both with respect to the resources.

[2] For example, if the implementation of a MAS that was modeled as a distributed system is also distributed, the (theoretical) communication complexity results can easily be turned into approximations for communication latency if the details of the underlying networking infrastructure are worked into the formulas.

nization and enactment of work in the workshop). We chose job shop scheduling as an example for this type of coordination problems.

There might be more natural application areas for value-augmented scheduling/planning problems than job shop scheduling problems (JSP). However, both the inherent computational complexity and the structural simplicity make job shop problems a good target for the demonstration of the applicability of economic coordination mechanisms to value-augmented scheduling/planning or, more general, economic resource allocation problems. Additionally, the work that has been done in areas like holonic manufacturing call for an application of the techniques that are outlined below because it allows us

1. to model jobs, machines, production managers, sales officers, planners etc. uniformly as autonomous, self-interested entities with individual tasks to fulfill and individual goals to accomplish.
2. to study directly the economic impact of decisions across all inter-related units of operations and planning (as long as the interdependencies are also mapped into the problem, which they will be if the actors driving the operations get a handle to influence decisions transparently—giving them useful budgets and allowing for active payments is one possible way, see (Adelsberger, Conen, & Krukis 1995)).
3. to apply a large number of relevant results from game theory and microeconomics to the analysis of operational and structural phenomena in manufacturing/logistics systems and beyond. This includes strategic considerations (with the possible goal to design mechanisms to give incentives to the parties interested in the result to communicate their preferences truthfully), economic efficiency (a goal can be to design coordination mechanisms that compute economically efficient solutions), participation constraints (design the mechanism to make sure that the participants feel as being treated in a fair manner to ensure their continued participation) etc.
4. to design mechanisms that scale, both vertically and horizontally, because they speak a language that is understood at every level of a company, between companies, and between companies and customers: the language of value, expressed in terms of money.

To be specific: The goal of our paper is to demonstrate the applicability of economic coordination mechanisms to economically augmented scheduling problems, and to offer solutions for computing economically efficient and/or straightforwardly computable heuristic solutions that are tailored to the problem domain. We also discuss some of the intricacies of economic coordination such as strategic agent behavior, equilibrium considerations and complexity issues.

1.1 Related Work

Related work is numerous, basically all of microeconomic literature is motivated by and related to resource allocation problems. The study of scheduling problems in an economic context has also a significant tradition in AI literature. For an excellent overview with an emphasis on economics, as well as for interesting and significant recent results related to both areas, see (Parkes 2001) or (Wurman 1999). For an instructive overview of methods and problems in the wider context of self-interested agents, consider (Sandholm 1996), whose further work, for example (Sandholm 2002), also contributes significantly to advances in the area. A relevant example for work related to the economics of scheduling is also the work of Wellman et. al (Walsh & Wellman 1998; Wellman *et al.* 2001; Walsh, Wellman, & Ygge 2000). More closely related to the application of economic principles in manufacturing environments are, for example, the work of Van Dyke Parunak (Parunak 1996) or A.D. Baker (Baker 1996).[3] More recently, this has also been discussed in the context of new modes of manufacturing, for example Holonic Manufacturing, as a promising paradigm for controlling scheduling and planning processes in a distributed environment with local goals, private information and general efficiency objectives (see, for example, (Adelsberger & Conen 2000)).

From an economic perspective, a key issue in the study of resource allocation problems is the design of coordination mechanisms that

- enable the computation of economically efficient solutions. In domains where monetary valuations are available (as we assume below), this amounts to determine allocations maximizing the aggregated welfare of the participating agents. Note that in a setting with private information, this requires that the mechanism gives some incentive to the agents to reveal their preferences truthfully.
- satisfy participation constraints for rational agents. A mechanism is considered to be individually rational if the agents can expect participation to be beneficial. It is also relevant that the agents are satisfied with the outcome of the mechanism to ensure further participation. The question of satisfaction boils down to answer the question to what extend it is possible for the agent to realize his most-preferred solution in a given situation.

In view of related work, the work presented here seems justified as most of the results that have been obtained for general resource allocation problems with an emphasis on economics have not been specifically tailored towards shop floor and related scheduling problems in manufacturing or logistics. Furthermore, recent results in studying combinatorial auctions and exchanges show that not all questions have been answered yet (relevant work is mentioned throughout the paper). Also, economically motivated approaches in scheduling/planning literature are sometimes difficult to analyze with respect to the key issues outlined above or tackle simplified settings.

[3] Though the concept of (Pareto) efficiency has been introduced much earlier into the scheduling literature, see (Wassenhove & Gelders 1980).

1.2 Overview

First, the basic notation and terminology for discussing job shop scheduling problems (JSPs) is introduced. We define the class of economically augmented job shop problems (EJSPs), where job agents have utility for schedules, and related them to typical minsum criteria. We map these problem to combinatorial job shop auction problems (CJSAPs). We introduce prices as a means to design individually rational, economically efficient coordination mechanism and discuss the issues of truthful revelation and the Vickrey principle. We show that both EJSP and CJSAP are (strongly) NP"=hard. In view of this result, we suggest a heuristic, economically-driven coordination mechanism that is build on top of a well-known, domain-specific solution procedure, the Giffler and Thompson algorithm. This decision point bidding approach can be applied to a number of search-based solution procedures. It is an example of an economic coordination mechanisms that is tailored to fit the problem domain. This, in turn, may allow participating agents to fully comprehend and purposefully influence the execution of the mechanism by means of bidding. The paper is concluded with a brief discussion of possibilities to extend and apply the obtained results.

2 From Job Shop Scheduling to Economic Coordination

Scheduling allocates resources over time to enable the execution of tasks. An important subclass of scheduling problems are job shop problems. The tasks are given by a set of jobs. Each job consists of a sequence of operations that have to be performed on specific machines in a given order. We base the following on the constraint optimization version of discrete, deterministic Job Shop Scheduling (JSS) as defined in (Ausiello *et al.* 1999).[4]

Definition 1 (Job Shop Scheduling – Basic Setting). *An instance of the class of job shop scheduling problems (JSP) consists of a set $M = \{1, \ldots, m\}$ of m machines, and a set $J = \{1, \ldots, n\}$ of n jobs, each consisting of a set $O_j = \{o_j^1, \ldots, o_j^{n_j}\}$ of n_j operations. For each such operation, o_j^i, a machine $m_j^i \in M$ and a processing time $p_j^i \in \mathbb{N}$ is given. A ready time, $r_j \in \mathbb{N}_0$, denotes the earliest possible start time for the first operation of each job $j \in J$.*

In a straightforward notation, a job j is given as $(r_j, [(m_j^1), p_j^1), \ldots, (m_j^{n_j}, p_j^{n_j}])$. The following example is used throughout the paper:

Example 1. *Job 1:* $(0, [(2,2), (1,3), (3,2)])$
Job 2: $(0, [(3,1), (2,2), (1,2)])$

In the following, we refer to an arbitrary, but fixed JSS problem P. A *potential schedule* for a set of operations is a mapping from the operations to start times. A potential schedule which assigns start times to all operations in P is called *complete*. A potential schedule is called *feasible* if (1) no first operation starts too early, (2) no sequence constraint is violated, and (3) no overlap in the processing times of assigned operations occurs on any machine. A schedule that is complete and feasible with respect to P is called a *valid* schedule or a *solution* of P (s. Fig. 1 for an example). Formally:

Definition 2 (Valid Schedule). *Given an instance P of JSP. A mapping $s : \cup_{j \in J} O_j \to \mathbb{N}_0$ is a* valid schedule *for P iff (1) $o_j^1 \geq r_j$ for all $j \in J$, (2) $s(o_j^i) + p_j^i \leq s(o_j^{i+1})$ for all $j \in J$ and all $i \in \{1, \ldots, n_j - 1\}$, and (3), for every pair of distinct[5] operations o_j^h, o_k^i, either $m_j^h \neq m_k^i$ or $s(o_j^h) + p_j^h \leq s(o_k^i)$ or $s(o_k^i) + p_k^i \leq s(o_j^h)$.*

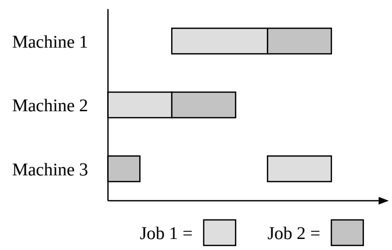

Figure 1: The valid schedule $(0,2,5)_1$, $(0,2,5)_2$ (schedules are given as n_j-ary sequences of start times for each job j) for the example 1.

We may study a problem P relative to a time horizon $[TS, TE]$, $TS, TE \in \mathbb{N}_0$ and write $P^{[TS,TE]}$. The set of all valid schedules restricted to a specific time horizon is denoted with $S^{[TS,TE]}$. If the time horizon is irrelevant, no index is shown. For a given set S, the subset S_j contains all schedules that are restricted to the operation of a given job j. For a given schedule s, the subschedule s_j denotes the elements of s that schedule operations of job j.

We now turn our attention to comparisons of the *quality* of schedules. Often measures that depend on the completion times of the jobs are used, for example, minimize $C^{max} = \max_{j \in J} C_j$, where $C_j = s(o_j^{n_j}) + p_j^{n_j}$ is the time of completion of the last operation $o_j^{n_j}$ of job j in a given (valid) schedule s. Regularly it is also assumed that each job has a due date d_j which allows to use derivations from this due date as a measure.

In a value-oriented business environment, this way of measuring the quality of schedules is only convincing if the measure has a close correspondence to the cashflow pattern that is induced by the imposed solutions (usually, it is tried to find a solution that optimizes the chosen measure or that approximates an optimal solution). With respect to the core measure of the success of management and operations, ie. value, such time (or capacity etc.) related measures can only be considered as being surrogates that may render decisions intransparent because they introduce incomparabilities between manufacturing shops competing for resources and give no handle to tie local decisions to the decisions of upstream and downstream manufacturing/logistics units.

[4]Which, in turn, is the optimization version of the decision problem SS18 as defined by Garey and Johnson in their seminal work (Garey & Johnson 1979) – with one difference: the additional constraint of Garey and Johnson which required that each pair of consecutive operations has to be performed on different machines has been dropped.

[5]That is, $j, k \in J$, $h \in \{1, \ldots, n_j\}$, $i \in \{1, \ldots, n_k\}$ and, if $j = k$, $h \neq i$.

This, however, would be an important prerequisite for an unified analysis of the plans, schedules, decisions and operations of all interrelated business units.

But even within one shop, it is usually not trivial to map the goals that are related to jobs to *one* measure only. Some jobs may be produced directly to customer orders (due date is important), some jobs are produced to fill up the stock (minimizing production cost is important) etc. The resulting multi-criteria problem does not seem to be convincingly solvable without introducing a unifying measure. As the *value* related to strategic, tactic, and operational decisions determines the success of business operations, mapping goals to value and measuring quality with money is a natural choice.

We will now augment job shop problems with value information and view them as economic coordination problems.

2.1 JSPs as Economic Coordination Problems

Job shop scheduling problems can be extended to economic coordination problems—the jobs are identified with agents competing for resources. The goods to be traded are slots of machine time. To augment a JSS problem P economically, we assume that each agent has utility for schedules, that is, a value function $v_j^{'} : S \rightarrow \mathbb{N}_0$ is available for every job agent j which assigns a non-negative value to each possible valid[6] schedule.

Definition 3 (EJSP). *An instance P (possibly restricted to a time horizon) of the class JSS of job shop scheduling problems and a set V of j value functions $v_j^{'} : S \rightarrow \mathbb{N}_0$, one for each job j, defines an instance EP of the class of economically augmented job shop scheduling problems (EJSP).*[7]

The general objective of solving EJSPs is to select a schedule from the set of possible valid schedules that achieves *economic efficiency*:

$$\arg\max_{s \in S} \sum_{j \in J} v_j^{'}(s) \tag{1}$$

As we show below, this class of scheduling problems encompasses a number of traditional job shop scheduling problems. To do so, we first consider a restricted variant of the EJSP, where, for each agent j, the value of a schedule does only depend on the operations in O_j. In the general variant, an agent j may value two schedules differently though both assign the same start times to the operations in O_j. For example, this allows an agent k to express that she prefers if a machine, say 2, is assigned to agent l instead of agent m in the interval $[3, 5]$. This expressiveness is not always required. [8]

[6] It can also be useful to consider incomplete feasible schedules.

[7] Remember that S is the set of all valid schedules that solve P. If a time horizon is given, this set may be empty. If no time horizon is given, the set is countably infinite. We will therefore usually assume that a time horizon is given (without displaying it).

[8] Though, to map a very common objective, namely max completion time, it is required, see the footnote below.

Definition 4 (EJSP without allocative externalities). *Let EP be given as above. EP is an EJSP* without allocative externalities *iff, for any agent j and any pair of schedules $a, b \in S$, such that the restriction a_j of a to operations of j coincides with the restriction b_j of b, $v_j^{'}(a) = v_j^{'}(b)$.*

Now consider a JSS problem $P^{'}$ and a typical minsum criterion (Hoogeveen, Schuurman, & Woeginger 1998), *total job completion time*. We map this $J||\sum C_j$ problem to a restricted EJSP $EP^{'}$ as follows:

First, determine a valid schedule by timetabling the operations of job 1 as early as possible and without slack and proceed with the operations of job 2, starting with o_2^1 at time C_1. Continue this until all operations are scheduled. This produces (with effort linear in the number of operations) a valid schedule $s^{'}$ of length $l = \sum_{j \in J} \sum_{1 \leq i \leq n_j} p_j^i$. This determines the time horizon $[0, l]$ to be considered. Now, a value function $v_j^{'}(\cdot)$ for each agent j can be given that determines, with effort linear in the number of operations, the value of any given input schedule:

Function $v_j^{'}$(In: Valid Schedule s, Out: Value v)
  $C_j \leftarrow s(o_j^{n_j}) + p_j^{n_j}$;
  $v \leftarrow l - C_j$; **return** v;[9]

Proposition 5. *A schedule maximizes the economic efficiency in $EP^{'}$ over all schedules in S if and only if it minimizes the minsum criterion* total completion time[10] *in $P^{'}$ over all schedules in S.*

Proof. Let s be a schedule that maximizes efficiency in $EP^{'}$. Assume that s does not minimize the total completion time in $P^{'}$, that is, there is a schedule $r \in S$ such that $\sum_{j \in J} s(o_j^{n_j}) + p_j^{n_j} > \sum_{j \in J} r(o_j^{n_j}) + p_j^{n_j}$, or, shorter, $\sum_{j \in J} s(o_j^{n_j}) > \sum_{j \in J} r(o_j^{n_j})$. Because s maximizes efficiency in $EP^{'}$, $\sum_{j \in J}(l - s(o_j^{n_j}) + p_j^{n_j}) \geq$

[9] Note that this transformation is polynomial because it makes use of the fact that there is significant structure in the problem. Would we have to enumerate all (or almost all) schedule/value pairs explicitly, the transformation would not be polynomial. However, criteria that would require such an effort are virtually never used in scheduling literature (they would be random in the sense of algorithmic complexity theory, (Li & Vitanyi 1997)).

[10] Similar results can be obtained for restrictions to other minsum criteria (e.g., total tardiness, weighted total completion/tardiness, holding costs, early/tardy penalties). Note also, that EJSP *without allocative externalities* cannot model criteria like $C^{\max}$, because a global optimum for the desired criterion is not obtainable from local considerations. The jobs *cannot* be modeled as being independent in this case. If the value function should reflect the benefit of achieving a minimal $C^{\max}$, they *must* reflect this dependency in their valuation of schedules, or otherwise, self-interest prevents the optimization of the desired criterion. For $C^{\max}$, the jobs should value schedules with an earlier completion time for all jobs higher than, for example, schedules that give them an earlier individual completion time (the time the last *operation* of j finishes) but a later overall completion time (the time the last *job* finishes)—in other words, the value of a schedule does not only depend on the operations in O_j but on all (final) operations (that is, allocative externalities are present).

$\sum_{j\in J}(l - r(o_j^{n_j}) + p_j^{n_j})$, or, written differently, $\sum_{j\in J} l - \sum_{j\in J} s(o_j^{n_j}) - \sum_{j\in J} p_j^{n_j} \geq \sum_{j\in J} l$ - $\sum_{j\in J} r(o_j^{n_j}) - \sum_{j\in J} p_j^{n_j}$ respectively $\sum_{j\in J} r(o_j^{n_j}) \geq \sum_{j\in J} s(o_j^{n_j})$, contradicting the assumption. The other direction follows immediately as well. □

Proposition 6. *EJSPs are NP-hard in the strong sense.*

Proof. As above proposition shows, EJSP can be restricted to $J||\sum C_j$. Furthermore, the transformation is polynomial (see above). From the strong NP-hardness of $J||\sum C_j$ (s. (Lawler *et al.* 1992) or (Hoogeveen, Schuurman, & Woeginger 1998)) the proposition follows. □

To adapt the problem to our economic setting with self-interested, (bounded) rational agents, the following assumption are necessary.

In line with our basic motivation, we assume that the value function is *information private to the agent*, that is, *no (central) institution has access to this information without prior consent of the agent*. In addition, *utility is transferable*[11] between agents (ie, a meaningful currency has to be available to express valuations and to transfer payments).

We also assume that no agent can be forced to act against his will. With these assumptions, we can translate economically augmented job shop scheduling problems to a specific subclass of combinatorial auction problems (CAPs), which are currently studied extensively in AI and microeconomic literature (s. (de Vries & Vohra 2001) for a survey).

2.2 Transforming EJSP to CJSAP

Let EP be in EJSP and let $[TS, TE]$ be a time horizon. An economic coordination problem can now be formulated as follows:

Agents: A set of n job agents, $N = J = \{1, \ldots, n\}$, and an arbitrator, 0. The arbitrator is modeled as a *supplier*. The job agents are modeled as *consumers*.

Goods: The set of goods, Ω, is the set of all machine-specific unit intervals within the time horizon (remember that M is the set of machines), that is:

$$\Omega = \{[z, z+1]_i \mid i \in M \wedge z = TS, \ldots, TE-1\}$$

Any subset $B \subseteq \Omega$ (alternatively writable as $B \in 2^{\Omega}$, the power set of Ω) is called *bundle*.

Before we can define value functions, a function $A : S \rightarrow 2^{\Omega}$ that partitions a schedule s into the covered machine-specific unit intervals is required (note that s is a function, and thus a relation, so that it is appropriate to write $(o, t) \in s$ instead of $t = s(o)$).

$$A(s) = \{[z, z+1]_m \mid (o_j^i, t) \in s \wedge m = m_j^i \wedge z \in \{t, \ldots, t + p_j^i - 1\}$$

Value functions: The value function of consumer j, $v_j : 2^{\Omega} \rightarrow \mathbb{N}_0$ is defined as follows: if at least one $s \in S$ exists such that the schedule s is covered by the unit intervals in B (ie., $A(s) \subseteq B$), the value of B is set to the value (in the underlying EJSP) of the best schedule among the covered schedules, that is, $v_j(B) = v_j^*(B)$ with $v_j^*(B) = \max_{\{s: s\in S \wedge A(s)\subseteq B\}} v_j^{'}(s)$. If no schedule is covered by B, $v_j(B)$ is set to 0.

[11] Technically, the value functions have to be quasilinear in money, compare, for example, (Mas-Colell, Whinston, & Green 1995). Quasilinearity allows to interpret the utility for a good (time slots in our case) as the willingness to pay for it.

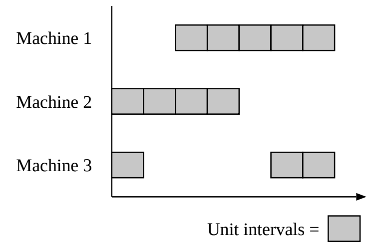

Figure 2: The unit intervals (or *time slots*) of example 1.

Example 2. *Reconsult example 1. The time horizon to be considered is $[0, 9]$. The following consideration leads to the preference relations, which underly the value functions: let j be a job agent and s_1 and s_2 be schedules which are complete with respect to j. If the last operation of agent j in s_1 is completed earlier than s_2, agent j prefers schedule s_1 (written as $s_1 \succ s_2$). If the completion times are equal, the agent is indifferent between s_1 and s_2 (written as $s_1 \sim s_2$).*

For agent 1 the preference relation given below follows. for convenience, a rank is assigned to the equivalence classes of equally preferred schedules.

1: $(0, 2, 5) \succ$
2: $(0, 2, 6)$, $(0, 3, 6)$, $(1, 3, 6) \succ$
3: $(0, 2, 7)$, $(0, 3, 7)$, $(0, 4, 7)$, $(1, 3, 7)$, $(2, 4, 7) \succ$
4: All invalid schedules.

For agent 2, the preference relation is only partially displayed:

1: $(0, 1, 3) \succ$
2: $(0, 1, 4)$, $(0, 2, 4)$, $(1, 2, 4) \succ \ldots$
3: $(0, 1, 5)$, $(0, 2, 5)$, $(0, 3, 5)$, $(1, 2, 5)$, $(1, 3, 5)$, $(2, 3, 5) \succ$
4: $(0, 1, 6), \ldots \succ$ *5:* $(0, 1, 7), \ldots \succ$ *6: All invalid schedules.*

Agent 1 values the earliest possible completion (at 7) with 20 currency units (CU) and agent 2 (at 5) with 16 CU. a delay per time unit reduces the value of the schedule for agent 1 by 3 CU and for agent 2 by 2 CU, resulting in the following functions $v_j^R(\cdot)$, which assign values to ranks:

Agent 1: $v_1^R(1) = 20$, $v_1^R(2) = 17$, $v_1^R(3) = 14$, $v_1^R(4) = 0$.

Agent 2: $v_2^R(1) = 16$, $v_2^R(2) = 14$, $v_2^R(3) = 12$,
$v_2^R(4) = 10$, $v_2^R(5) = 8$, $v_2^R(6) = 0$.

For a given bundle B, its valuation can be determined by mapping B with a function r to the preference rank which corresponds to the best schedule which is covered *by B. For example for the largest bundle, Ω, which covers all schedules that lie within the time horizon, rank 1 is the result of the mapping for both agents, that is $v_1(\Omega) = v_1^r(r(\Omega)) = v_1^r(1) = 20$ respectively $v_2(\Omega) = 16$.*

Proposition 7 (Monotony, Free Disposal). $v_i(A) \leq v_i(B)$ *for all* $A \subseteq B \subseteq \Omega$, $i \in N$.

Proof. This follows immediately from the construction of the value functions: a bundle B that contains at least as many unit intervals as a bundle A covers at least the schedules that A covers, thus its value cannot be smaller than the value of A because the maximum value over all covered schedules in the EJSP instance is at least as large for B as for A. □

This is sometimes also called *Free Disposal* because adding further unit intervals will not reduce value, or, in other words, superfluous goods can be disposed off at no cost.

We also assume that *no budget restrictions* exist: every consumer j is in possession of enough money to be able to pay up to the amount of the valuation for his most preferred bundle. The consumers have no endowment beyond money. All goods belong to the arbitrator. If the consumer does not receive any good, his utility shall be 0 (w.l.o.g).

The objective of the allocation of the goods (from the perspective of the arbitrator) is to maximize the aggregated utility of the consumers. Here, this directly corresponds to maximizing the sum of individual utilities. An allocation[12] $X^* = (X_0^*, \ldots, X_n^*)$ conforms to this objective if and only if the allocation is *efficient*, that is, X^* has to be a maximizer for the following problem

$$\max_X \sum_{j=1}^{n} v_j(X_j) \;\; (X \text{ iterates over all allocations}) \quad (2)$$

Example 3. *For the above example, the following allocations are efficient:*
$X_1 \supseteq \{[0,1]_2, [1,2]_2, [2,3]_1, [3,4]_1, [4,5]_1, [5,6]_3, [6,7]_3\}$
$X_2 \supseteq \{[0,1]_3, [2,3]_2, [3,4]_2, [5,6]_1, [6,7]_1\}$
$X_0 = \Omega \backslash (X_1 \cup X_2)$.

Consecutive unit intervals can be consolidated to bundles in an abbreviating notation:
$X_1 \supseteq \{[0,2]_2, [2,5]_1, [5,7]_3\}$
$X_2 \supseteq \{[0,1]_3, [2,4]_2, [5,7]_1\}$
$X_0 = \Omega \backslash (X_1 \cup X_2)$.

Definition 8 (CJSAP). *Let EP be an instance of EJSP. The sets Ω and N, the value functions $v_j(\cdot)$ of the consumers that are obtained by the above transformation of EP, and the objective* (2) *of the arbitrator define an instance C of the class of* Combinatorial Job Shop Auction Problem, *CJSAP.*

We call an allocation that conforms to the objective of the arbitrator a *solution*.

Proposition 9. *A solution exists for every possible instance of a CJSAP.*

Proof. This follows with a straightforward combinatorial argument immediately from the finiteness of the problem (which we assume throughout the paper). □

[12]That is an $(n+1)$-ary partition of the set of goods, Ω, which assigns to agent j the goods in the bundle X_j (some X_j may be empty, so it is not a partition with non-empty subsets, as it is usually assumed).

Proposition 10. *For any $EP \in$ EJSP, if a valid schedule exists, a solution of the corresponding problem $C \in$ CJSAP can be transformed into an optimal schedule for the original, economically augmented job shop scheduling problem EP (and vice versa).*

Proof. Let $X = (X_0, X_1, \ldots, X_n)$ be a tight solution[13] of C. Construct a schedule s^X from X as follows: for each agent $j \in J$, find the best schedule s_j^X that is covered by X_j by simply timetabling the operations as unit intervals are available (due to the tightness of X, this is possible in cost linear to the number of time slots in X). Combine the schedules s_j^X to the schedule s^X. Note that the construction of C ensures that this construction is always possible if a valid schedule exists. Furthermore, as the allocation X is a partition of Ω, no overlap on a machine can occur. The other direction is even simpler: each reservation for an operation in the optimal schedule can be split into the machine-specific unit intervals. The information in the schedule can be used to assign the unit intervals immediately to the correct part of the allocation (cost linear to the processing time). □

We will now turn our attention to the *problems related to finding an efficient allocation*. They can be outlined as follows: (1) a certain amount of value information is necessary to determine an efficient allocation; (2) once the required information is available, the actual computation has to be performed; (3) an incentive has to be given to the agents to report their part of the required information truthfully, or otherwise, efficiency of the solution cannot be guaranteed; and (4) the agents have to be satisfied with the determined outcome. We neglect issue (1), briefly discuss issue (2) and concentrate on (3) and (4). We will, however, return to the issue of complexity later.

2.3 Determination of Efficient Allocations

If we assume for now that the (complete and true) value functions of all agents are known, an efficient allocation of the unit intervals to the job agents can be computed with one of the well-studied methods of winner determination (compare (Sandholm 2002; Sandholm *et al.* 2001; Fujishima, Leyton-Brown, & Shoham 1999)). To save communication, approaches have been suggested that only partially reveal the value functions of the consumers (compare (Parkes 2001) for indirect, iterative auctions and (Conen & Sandholm 2002b) for a progressive direct mechanism based on (Conen & Sandholm 2001b; 2001a)). Note that the winner determination problem is essentially a set-packing problem (Rothkopf, Pekeč, & Harstad 1998; de Vries & Vohra 2001) and that it is NP-hard. We show below that this also holds for CJSAP.

[13]A tight solution is a solution that does not contain time slots that are not used. Note that a tight solution always exists because the value of a bundle is the value of the best covered schedule, and, in consequence, the bundle that covers just the best schedule, is tight and optimal (that is: unused slots cannot add to the value of a bundle due to the construction of the value functions). Further note that is it always possible to tighten a solution with costs less than exponential.

2.4 Prices

In the above example, two core problems remain: first, because only agent 1 can realize his best alternative, agent 2 will envy him. Second, we have tacitly assumed that the agents report their utility truthfully—but why should they do so under our assumption of preferences being private information? Certainly, agent 2 could expect to benefit from over-exaggerating his valuations. If agent 1 would expect agent 2 to over-exaggerate, he would do the same, and so forth. Without an additional way to ensure satisfaction and to make lying unattractive, we cannot expect to compute allocations that are efficient with respect to the true preferences. The way to go is to introduce prices. Each consumer has a net utility function which reflects the impact of prices (negative transfers in the following definition) on the realizable utility.[14]

Definition 11 (Net utility). *The* net utility function $u_j(\cdot) : 2^\Omega \times \mathbb{Z} \rightarrow \mathbb{Z}$ *for each consumer* j *is defined as* $u_j(x,t) = v_j(x) + t$.

To keep every consumer satisfied with the outcome (consisting of allocation and payment), $u_j(X_j, p) \geq u_i(A, p)$ must hold for every consumer j and every bundle $A \subseteq \Omega$,that is, the net utility of the bundle he receives must be at least as good as it would be for any other bundle at the given prices (or, otherwise, the consumer would prefer to receive the bundle that gives him the best net utility).

This leads immediately to the notion of equilibria. An outcome, that is, an allocation and a payment vector (determined by the prices), is an equilibrium if both sides of the market are satisfied with it (the market is *cleared*). In our case this is true if the above condition holds for all consumers and if the allocation is efficient (to satisfy the supplier). There are different restrictions that one can impose on prices—prices for bundles have to be the sum of prices for individual goods (see (Gul & Stacchetti 1999; Kelso & Crawford 1982), prices are independent for every bundle (see (Wurman & Wellman 2000)), prices are determined for the bundles in the efficient allocation and prices for bundles of these bundles are additive (see (Conen & Sandholm 2002a)). Only the independent pricing of all bundles can guarantee the existence of equilibrium price vectors. However, implementing the determined outcome may require enforcement.[15] For the more natural pricing modes, equilibrium price vectors need not exist.[16] We will therefore not discuss (anonymous) equilibrium pricing in detail and turn our attention to a solution concept for which solutions always exists: Vickrey payments. We will demonstrate in the example below that implementing Vickrey payment-based coordination mechanisms give the participating agents no incentive to lie (a very relevant design objective in a private information setting). Please, consider the continued example below.

[14]Note, that we make the usual assumption of quasi-linearity of valuations.

[15]Each agent is only allowed to buy one bundle—thus, if he is interested in AB and we have $p(AB) = 6$, $p(A) = 2$, $p(B) = 2$, he would want to enter the auction with two identities to buy A and B separately.

[16]This negative result holds also for instances of CJSAP.

Example 4. *Bundles on offer in the efficient allocation:*

$$A = \{[0,2]_2, [2,5]_1, [5,7]_3\}$$
$$B = \{[0,1]_3, [2,4]_2, [5,7]_1\}$$

Demand for these bundles:

	$\emptyset$	A	B	AB
Agent 1	*0*	**20**	*0*	*20*
Agent 2	*0*	*0*	**12**	*16*

Vickrey Payments (the vector of payments also fulfill the equilibrium condition if interpreted as prices for the bundles instead of personalized payments):

	A	B	AB
Agent 1's payment	*4*		
Agent 2's payment		*0*	
Equilibrium prices	*4*	*0*	*4*

The payments ensure that there is no (individual) incentive for the agents to lie, as can be seen as follows. First note that the price each agent has to pay does only depend on the reported utilities of the competing agents—it represents the loss of utility that the other agents experiences due to the participation of the former agent. Now, assume that agent 1 would underbid his valuation with, say, 17. Then he risks that agent 2 (or any other agent) would bid just above 17 but below 20, say 18, and would thus receive the good. As the price he has to pay is independent of his own bid and will equal the bid of the other agent, he could have done better by bidding truthfully (exactly 2 = 20 - 18 instead of 0). If he would have known beforehand what the other agent will bid, say x*, he would not have a reason to underbid either, because there would be no difference in net value for agent 1 in bidding 20 or* $x + \epsilon$ *(as the price will be the bid of the other agent anyway). He would also not overbid, because he risks that he receives the overbidded bundle for a price between his true valuation and his bid and would, thus, realize a loss. If he would overbid in the fully-informed situation, he cannot gain any net value from it either. An analogous reasoning applies to agent 2, thus both agents will not have an incentive to misrepresent their valuation if they act rational (this corresponds to an equilibrium in dominant[17] strategies).*

The general principle invoked here has been mentioned in the example: the payments that an agent has to transfer do not depend on her own bid but captures instead the effect of her participation on the other participating agents.[18] It is intuitively clear that, as soon as there is a dependency for a bundle X_j between the reported utility of agent j and the price j has to pay, j will have an incentive to minimize this price by misrepresenting his utility if possible. To do so, he might start to collect information about the other participating agents (which is not necessary above) to become able to behave *strategically*. This, in turn, may make it impossible for the arbitrator to pick the efficient allocation—all that he could do would be to pick the allocation that is efficient with respect to the *reported* utilities. This brief digression into the issue of *incentive compatibility* may suffice to

[17]Weakly dominant in the case of informed agents.

[18]The principle has been discovered and applied independently by Vickrey (in 1961), Clarke (in 1971) and Groves (in 1973) (see, for example, (Vickrey 1961)).

demonstrate one of the most prevalent problems in environments where the agents have private information: the problem of eliciting their preference truthfully to allow for truly optimal decisions.[19]

Proposition 12 (Existence of Vickrey Outcome). *An outcome consisting of an efficient allocation and a vector of related Vickrey payments exists for every instance of CJSAP.*

Proof. (Sketch) The proof follows from a straightforward combinatorial argument: with finite sets of bundles and agents. A computable solution of the maximization problem (2) (to determine the efficient allocation) and the n (or less) maximization problems following from the initial problem by leaving out, for each non-empty bundle in the allocation, the agent that receives it (to determine the effect of his participation, that is the Vickrey payments), is immediately available from enumerating all possible complete and reduced allocations and picking the optima. □

3 Complexity Issues

Two immediate problems of auction mechanisms that try to solve a CJSAP are (1) that the mechanisms may require communication that is exponential in the number of unit intervals (compare the general result of (Nisan & Segal 2002), which extends to CJSAP) and (2) that solving the actual allocation problem once all required value information is received is NP-hard, as the following corollary demonstrates:

Corollary 13. *CJSAPs are NP-hard in the strong sense.*

Proof. This follows immediately from Proposition 6 and the polynomiality of the transformation given in Proposition 10. □

Note also that CJSAP is a subclass of the combinatorial auction problem CAP which is equivalent to maximum weighted set-packing (compare (de Vries & Vohra 2001)), known to be NP-hard.[20]

Consequently, we cannot expect to obtain optimal solutions for problems of reasonable size with reasonable computational efficiency. Furthermore, the approximability results obtained for maximum weighted set packing (Ausiello *et al.* 1999) are not very encouraging. We also do not expect that CJSAP is an especially well-behaved subclass in this respect due to its proximity to EJSP and, thus, typical job shop scheduling problems. We can now try to modify auction mechanism that determine efficient allocations (like iBundle (Parkes & Ungar 2000) or AkBA (Wurman & Wellman 2000)) to relax the efficiency goal and to restrict the search space heuristically. However, we cannot expect this to be easily justifiable with respect to the properties of the initial problem domain.

It is the nature of any heuristic modification that no guarantee to obtain economically efficient solutions can be given. Furthermore, other desirable properties like incentive-compatibility may get lost. With respect to a specific problem domain, it seems reasonable to postulate that the heuristic reflects the properties of the domain naturally, allowing users to make justifiable decisions when dealing with the heuristic mechanisms. We suggest a mechanisms that is closely related to a straightforward solution procedure for job shop scheduling, namely the algorithm of Giffler and Thompson (GT). It allows the agents (resp. the represented actors) to base their bidding behavior on an understanding of the consequences of decisions that are made along an execution of GT.

4 Decision Point Bidding

Most solution procedures for JSPs (or for other resource allocation problems) can be viewed as a sequence of decisions—for the GT, this is the selection of the next operation to schedule. Now, instead of competing for reservations directly, the agents bid for the right to stipulate decisions.

As GT constructs an active schedule, we will restrict our analysis to problem instances of EJSP that have the property that the individual agents always (weakly) prefer earlier allocation of the final operation (to mimic the behavior of regular measures).

We give a brief re-collection of the GT prior to outlining a coordination mechanism that allows to solve restricted EJSP by following the sequence of decisions that are characteristic for an execution of GT. Assume that an instance EP of EJSP is given.

(1) Initialize the set SO of schedulable operations to $o_j^1, 1 \leq j \leq n$.
(2) As long as SO is not empty do
(3a) Determine the decision set DS by first picking from SO one of the operations with the earliest completion time, say o_x^y (ie. $\arg\min_{o_j^i \in SO} r_j + p_j^i$) and
(3b) insert into DS all operations from SO that require m_x^y and start before the potential completion time $r_x + p_x^y$ of o_x^y
(4) (**Decision Point**) Choose one of the operations from DS, say o_a^b and make reservations corresponding to its related ready and processing times.
(5) Update the ready time information of all jobs that have operations in DS by setting r_j of any such job to $r_a^b + p_a^b$. Remove o_a^b from SO and add its successor (if it exists, ie. if $b < n_j$), o_a^{b+1} to SO. Empty DS.

Now, the basic mechanism is as follows: the arbitrator announces to run a GT and starts to ask the agents for information about ready times and about the processing times of the operations to be scheduled. It iterates the Giffler and Thompson procedure as usual, requesting bids from the agents to be able to determine a decider for each decision point (see above). Each agent who has currently an operation in the decision set is entitled to

[19] For pointers to recent work, see (Bikhchandani *et al.* 2001) or Parkes (Parkes 2001).

[20] If we allow partially ordered sequences of operations and alternative routings, the corresponding variant of EJSP and the resulting CJSAP could be mapped to the general combinatorial auction problem bijectively.

bid. The winner will pay the price of the second highest bid and receives, as a result, the earliest possible reservation of his operation in the decision set (ties in the bids are broken arbitrarily).[21]

Corollary 14. *For any finite EJSP, the mechanism terminates and the resulting allocation corresponds to an active schedule.*

This follows immediately from the construction of the mechanism, which follows the execution pattern of the GT and, thus, inherits its property of computing an active schedule.

Proposition 15. *For the restricted EJSP, the set of possible outcomes of the mechanism contains an efficient allocation.*

Proof. (Sketch) There are $z = \sum_j n_j$ decision points to bid on in each run of the mechanism. The sequence of decisions can be described by a z-ary vector of job agent indices. The set of all sequences that lead to active schedules can be obtained through iterated 0-bidding (no agent bids a non-zero amount of currency units in any iteration). This is due to the non-deterministic tie-breaking and the fact that the mechanism follows the execution path of GT and thus inherits the property of GT to generate all active schedules if all decisions are taken in every possible way. The proposition follows from an adaptation of Theorem 2.1 of (French 1982) (suggested by French in Chap. 10), which states that the class of active schedules contains the optimal schedule for regular measures, and the fact that the restricted EJSP mimics a regular measure. □

Note that each bidding decision is necessarily based on calculations of the expected value of a positive decision (the agent wins the decision point). Each agent knows beforehand that he has to win exactly n_j decision points. He can also calculate the influence of each taken decision on the maximally achievable utility (simply by assuming that he will win all following decision points up to and including an allocation for his last operation, which fixes the reservation times for all remaining operations and, thus, allows us to set a precise upper bound on the realizable utility). His attitude towards risk will influence his bidding decisions, as will his knowledge about the bidding behavior of competing agents in this or in prior runs of the mechanism (if agents participate in multiple allocation problems).[22] A solution concept to study such situations is known as Bayes-Nash Equilibrium, though this is not discussed here. Let us only point out that augmenting solution procedures that show a natural fit to the problem domain may be an interesting opportunity to design economic coordination mechanism that are easily understandable for the participating agents and that give the agents a direct handle to use their understanding to influence the outcome of the mechanism. The decision point approach can be applied to numerous, decision-based solution procedures. Note that obtaining analytical results will not be an easy task, however, advancing this concept looks promising to us due to the similitude of mechanism and domain-specific problem structure.

[21]Instead of the sealed-bid second-price auction suggested here, an open-cry English auction could be chosen, which would convey more information about the competitors to the agents.

[22]The mechanism can be extended to resale to take limited foresight into account.

5 Discussion

This paper has discussed economically augmented job shop problems and their transformation into a specific subclass of combinatorial auction problems. We have briefly discussed the problem of finding efficient solutions and of doing this in a way to ensure that the agents participate and report their valuations truthfully. Motivated by problems of computational efficiency, we have suggested to combine domain-related decision making procedures with bidding. This hybrid approach results in a heuristic that, while remaining problem driven, is applicable in the context of self-interested agents with private information. We consider the study of such application scenarios in the context of scheduling for manufacturing and logistics as important due to the increasing tendencies to (1) restructure companies towards collaborations of (semi-)autonomous units on all levels of granularity and to (2) form (potentially) volatile collaborations between autonomous enterprises. In both cases, monetarian value may turn out to be a key issue for the integration of diverging interests—economically efficient mechanisms that keep the participating agents individually satisfied can contribute to this integration without violating autonomity and information privacy more than necessary.

5.1 Extensions and future work

The mapping of economic job shop problems to combinatorial auction problems can easily be extended to accommodate other classes of job shop problems, for example problems with alternative routings, reservation costs for resources, or sequence-dependent set-up costs.

Further effort is necessary to explore the options and consequences of decision point bidding, with respect to both a theoretical analysis of its effect on efficiency, strategic behavior etc. and a practical or experimental analysis of its applicability. It will also be interesting to apply the basic concept of decision point bidding to other solution procedures for resource allocation problems.

References

Adelsberger, H. H., and Conen, W. 2000. Economic coordination mechanisms for holonic multi agent systems. In *Proc. of HoloMAS Workshop, DEXA 00*.

Adelsberger, H. H.; Conen, W.; and Krukis, R. 1995. Scheduling utilizing market models. In *Proc. of the ICCIM 95*. Singapore: World Scientific.

Ausiello, G.; Crescenzi, P.; Gambosi, G.; Kann, V.; Marchetti-Spaccamela, A.; and Protasi, M. 1999. *Complexity and Approximation*. Springer.

Baker, A. 1996. A case study where agents bid with actual costs to schedule a factory. In Clearwater, S. H., ed., *Market-Based Control*. Singapore: World Scientific. 184–223.

Bikhchandani, S.; de Vries, S.; Schummer, J.; and Vohra, R. V. 2001. Linear programming and Vickrey auctions.

Conen, W., and Sandholm, T. 2001a. Preference elicitation in combinatorial auctions. In *Proc. of the ACM E-commerce conference (ACM-EC01)*. Tampa, FL: ACM. Short Paper.

Conen, W., and Sandholm, T. 2001b. Minimal preference elicitation in combinatorial auctions. In *IJCAI-2001 WS on Economic Agents, Models, and Mechanisms*, 71–80.

Conen, W., and Sandholm, T. 2002a. Coherent pricing of efficient allocations in combinatorial economies. In *Prooceddings of the AAAI-02 Workshop on Game theoretic and Decisiosn Theoretic Agents (GTDT-02)*.

Conen, W., and Sandholm, T. 2002b. Partial-revelation VCG mechanism for combinatorial auctions. In *Proceddings of the National Conference on Artificial Intelligence (AAAI)*.

de Vries, S., and Vohra, R. 2001. Combinatorial auctions: A survey. Draft of 25. Oktober 2001.

French, S. 1982. *Sequencing and Scheduling*. Ellis Horwood Limited.

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

Garey, M. R., and Johnson, D. S. 1979. *Computers and Intractability*. W. H. Freeman and Company.

Gul, F., and Stacchetti, E. 1999. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory* 87:95–124.

Hoogeveen, H.; Schuurman, P.; and Woeginger, G. 1998. Non-approximability results for scheduling problems with minsum criteria. In Bixby, R. E.; Boyd, E. A.; and Rios-Mecado, R. Z., eds., *Proceedings of the 6th Conference on Integer Programming and Combinatorial Optimization,*, number 1412 in LNCS, 353–366. Springer.

Kelso, A. S., and Crawford, V. 1982. Job matching, coalition formation, and gross substitues. *Econometrica* 50:1483–1504.

Lawler, E.; Lenstra, J.; Kan, A. R.; ; and Shmoys, D. 1992. *Sequencing and Scheduling: Algorithms and Complexity*, volume 4 of *Handbooks in Operations Research and Management Science*. North-Holland.

Li, M., and Vitanyi, P. 1997. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, 2nd edition edition.

Mas-Colell, A.; Whinston, M.; and Green, J. R. 1995. *Microeconomic Theory*. Oxford University Press.

Nisan, N., and Segal, I. 2002. The communication complexity of efficient allocation problems. Working paper.

Parkes, D., and Ungar, L. 2000. Iterative combinatorial auctions: Theory and practice. In *Proc. of the AAAI-00*.

Parkes, D. C. 2001. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. Ph.D. Dissertation, Computer and Information Sciences, University of Pennsylvania.

Parunak, H. V. D. 1996. Applications of distributed artificial intelligence in industry. In O'Hare, G., and Jennings, N., eds., *Foundations of Distributed Artificial Intelligence*. New York: Wiley. 139–164.

Rothkopf, M. H.; Pekeč, A.; and Harstad, R. M. 1998. Computationally manageable combinatorial auctions. *Management Science* 44(8):1131–1147. Early version: Rutgers Center for Operations Research technical report 13-95.

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

Sandholm, T. 1996. *Negotiation among Self-Interested Computationally Limited Agents*. Ph.D. Dissertation, University of Massachusetts, Amherst. Available at http:// www.cs.cmu.edu/ ˜sandholm/ dissertation.ps.

Sandholm, T. 2002. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence* 135:1–54.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16:8–37.

Walsh, W., and Wellman, M. 1998. A market protocol for decentralized task allocation. In *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS)*, 325–332.

Walsh, W.; Wellman, M.; and Ygge, F. 2000. Combinatorial auctions for supply chain formation. In *ACM Conference on Electronic Commerce 2000: 260-269*.

Wassenhove, L. V., and Gelders, L. 1980. Solving a bicriterion scheduling problem. *Eur.J. Opl. Res.* 4:42–48.

Wellman, M.; Walsh, W.; Wurman, P.; and MacKie-Mason, J. 2001. Auction protocols for decentralized scheduling. *Games and Economic Behavior* 35:271–303.

Wurman, P. R., and Wellman, M. P. 2000. AkBA: A progressive, anonymous-price combinatorial auction. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, 21–29.

Wurman, P. R. 1999. *Market Structure and Multidimensional Design Auction for Computational Economies*. Ph.D. Dissertation, Computer Science and Engineering, University of Michigan.