Economically-Augmented Job Shop Scheduling

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

Abstract

We present economically augmented job shop scheduling (EJSP) as an example of a coordination problem among self-interested agents with private information. We discuss its significance in modern organizational supply-web structures, analyze its complexity and present a specific type of combinatorial auctions as a solution mechanism. We relate EJSP to results from the area of economic mechanism design and especially emphasize the need for solution mechanisms, which give the agents no incentive to lie. This requirement is a significant extension to the side constraints that are usually considered in the scheduling literature.

Keywords: scheduling and economics, distributed and multi-agent planning and scheduling, decision-theoretic planning and scheduling

"Cooperation in the economic tradition is mutual assistance between egoists."

Hervé Moulin in (Moulin 1995), p. 5

Introduction

In recent years, two related trends influenced the development of organizational structures in and between production enterprises. Within enterprises, an emphasis on local competencies and decentralized decision making promoted the creation of (semi-)autonomous substructures. Between enterprises, this was complemented by the formation of innovatively designed cooperations of otherwise economically indepent enterprises. Both tendencies were driven by the conviction that a network of "units" (to be called *economic cooperation* below), each of which concentrates on its core competency and is equipped with enough authority for locally autonomous and economically sensible decision making, is an answer to the challenges caused by an increasingly complex environment with increasingly agile competitors.

This was accompanied by a shift from the typical centralized decision making in monolithic, hierachically organized structures to a decentralized decision mode with its need for the explicit *coordination of objectives and activities*. While actors and groups of actors have often been (ignorantly) treated as well-functioning, self-motivating entities in the former paradigm, the (partial) economic autonomity of the actors in a collaboration required to put an emphasis on their motivation (or, to speak in economic terms, their *incentive*) to collaborate properly.

As usual, the cooperation is facing a number of problems (producing goods, designing new products, acquiring human and financial resources etc.). The cooperation is driven by the individual objectives of the participants. We make the usual assumption in an economic environment with rational agents that the *individual objective is to maxmize utility*. We also assume that the participants share some objectives (if this would not be the case, there would not be a point in speaking of a cooperation). We will concentrate on types of cooperations whose objective is *to maximize economic efficiency of the cooperation as a whole*. Note that this makes sense from the perspective of the participants that will negotiate their overall goal before the precise details of day-to-day business will be known: if the *cooperation as a whole* strives for economic efficiency, the amount of money that will be earned (or saved) will be maximized – and, in consequence, more money (or less costs) will be distributed to the participants than it would be possible with other objectives. Of course, if judged from the perspective of an individual actor, other objectives may be more beneficial (e.g., "let me do every interesting task and let me keep all the benefits"), but, as we discuss cooperations, it seems sensible to use efficiency as a least common denominator – the actual mode of distributing earnings or costs can still be discussed, but if more money than necessary has been spent or less money than possible has been earned, the amount to distribute will inevitably be smaller than necessary.

In total, the cooperation will search for a mechanism that implements a solution to the problem at hand which satisfies the individual objectives "utility maximisation" and the common objective "efficiency". The solution will coordinate the activites of the actors and their individual objectives. We will therefore call the mechanism an (economic) coordination mechanism. To be able to guarantee efficiency, it is necessary that the actors report the required information about their objectives, their utilities etc. truthfully. Therefore, the mechanism should give the right incentives to do so.[1]

[1]Note that paying respect to the assumption of utility maximizing behaviour, we cannot assume that the actors would report information truthfully without an economic incentive to do so. Al-

A significant and important subset of the problems that economic cooperations are facing are resource allocation problems.

While numerous mechanisms for the general case (Parkes & Ungar 2000; Wurman & Wellman 2000; Conen & Sandholm 2002c; 2002b) and for certain limited subcases (Ausubel & Milgrom 2002; Lehmann, Lehmann, & Nisan 2001; Tennenholtz 2000) have been investigated recently, only very few studies have been carried out to understand and solve specificially structured problems in relevant application domains.

An important subclass of resource allocation problems, which frequently occurs in manufacturing and logistics, are of the shop scheduling type. Its occurence is not confined to single production shops. It can also be used to model a large set of structurally equivalent or similar problems in fine- and medium-grain manufacturing planning, in supply-web logistics, resource-bounded project planning, task assignment and so forth.

In the following, we will introduce economically augmented job shop scheduling problems (Sec.). We will discuss some complexity issues, cast the problems as a specific type of combinatorial auction problem, discuss complexity again (Sec.) and finally relate the resulting economic coordination problem to results from economic mechanism design.

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 (Wellman *et al.* 2001; Walsh, Wellman, & Ygge 2000). Especially (Wellman *et al.* 2001) tackles "factory" scheduling problem.[2] 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 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 compare (Vohra 2001)). 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.

From Shop Scheduling to Economic Coordination

Scheduling allocates resources over time to enable the execution of tasks. We study an important subclass of scheduling problems, namely 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 (JSS) 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*

though, from the assumption of (partial) autonomity, we cannot assume that we will be able generally to enforce or verify truthfulness (neglecting this is one of the most significant weaknesses of old-fashioned decision-making mechanisms, see below)

[2]Restricted to the single machine case. It is, however, straightforward to extend a number of their results, especially those on the (non-)existence of equilibria, to various shop scheduling settings. The paper gives also an excellent introduction to the microeconomic notion of equilbria.

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

[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.

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 $j \in J$.

The following example with two agents and three machines is used throughout the paper:

Example 1. *Job 1: Ready time 0, Operation 1 on machine 2 requires 2 time units. Operation 2 on machine 1 requires 3 time units, operation 3 on machine 3 requires 2 time units.*

Job 2: Ready time 0, Operation 1 on machine 3 requires 1 time unit. Operation 2 on machine 2 requires 2 time units, operation 3 on machine 1 requires 2 time units.

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) $s(o_j^1) \geq r_j$ 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 $s(o_k^i) + p_k^i \leq s(o_j^h)$.*

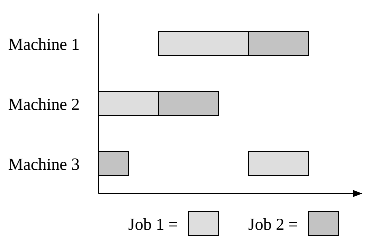

Figure 1: The valid schedule $(0, 2, 5)_1$, $(0, 2, 5)_2$ (schedules are given as n_j-ary sequences of start times for each job j) for example 1.

We frequently 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 sub-schedule 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.

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

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 cash-flow 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. 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.

JSS as a Economic Coordination Problem

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 $v_j^{'} : S \to \mathbb{N}_0$ 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 $v_j^{'} : S \to \mathbb{N}_0$, 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*:

[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).

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

We show below that 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]

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 (the mapping will be denote with $g(\cdot)$ below):

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 $l = \sum_{j \in J} \sum_{1 \le i \le n_j} p_j^i$. 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. *Let $P^{'}$ be an arbitrary, but fixed instance of $J||\sum C_j$ and set $EP^{'} = g(P^{'})$. A schedule s minimized the criterion $\sum C_j$ in $P^{'}$ if and only if it maximizes the economic efficiency in $EP^{'}$.* [10]

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

Proposition 6. *EJSP is (strongly) NP-hard.*

Proof. First note that if we formulate EJSP as a decision problem that asks, if there exists a schedule with utility above a certain bound, we can check if a certain schedule is a "Yes" instance in time polynomial to the size of the input easily. Thus, the decision problem variant of EJSP is in NP. As the above proposition 5 shows, $J||\sum C_j$ can be reduced to EJSP. Furthermore, the transformation g is polynomial (see above, which also applies to the decision problem variants). 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).

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

[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.

[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).

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

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:

$$\Omega = \{[z, z+1]_i \mid i \in M,\ z \in \mathbb{N}_0,\ z = TS, \ldots, TE-1\}$$

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

Value functions: Before we can define value functions, a function $A(\cdot)$ 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 $t = s(o)$). This function $A : S \to 2^\Omega$ is defined as

$$A(s) = \{[z, z+1]_m \mid (o^i_j, t) \in s,\ m = m^i_j,\ z = t, \ldots, t + p^i_j - 1\}.$$

Furthermore, let B be a bundle and s a schedule. B covers s iff $A(s) \subseteq B$. B *precisely covers* s iff $A(s) = B$. We will also say that s covers B precisely.

In an allocation, every consumer will obtain a possibly empty subset of the unit intervals, that is a bundle. The value of a bundle for the consumer is equal to the value of the best schedule that is covered by the bundle. The value function of consumer j, $v_j : 2^\Omega \to \mathbb{N}_0$ can now be defined as follows:

$$u_j(B) = \begin{cases} 0 & \text{if no } s \in S_{r[j]} \text{ exists, such that } A(s) \subseteq B \\ \max_{\{s: s \in S_{r[j]},\ A(s) \subseteq B\}} u'_j(s) & \text{otherwise.} \end{cases}$$

Here, $S_{r[j]}$ is the set of all schedules such that their restriction to the operations of agent j is feasible.

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: $(0, 1, 4)$, $(0, 2, 4)$, $(1, 2, 4) \succ \ldots$ *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 $v^R_j(\cdot)$, which assign values to ranks:

Agent 1: $v^R_1(1) = 20$, $v^R_1(2) = 17$, $v^R_1(3) = 14$, $v^R_1(4) = 0$.

Agent 2: $v^R_2(1) = 16$, $v^R_2(2) = 14$, $v^R_2(3) = 12$,
$v^R_2(4) = 10$, $v^R_2(5) = 8$, $v^R_2(6) = 0$.

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* $v_1(\Omega) = v^r_1(r(\Omega)) = v^r_1(1) = 20$ *respectively* $v_2(\Omega) = 16$.

Proposition 7 (Monotony).
$v_i(A) \le v_i(B)$ *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 formalizes the *Free Disposal* condition, 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

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

Example 3. *For the above example, the following assignments lead to efficient allocations (consecutive unit intervals are consolidated into bundles):* $X_1 \supseteq \{[0,2]_2, [2,5]_1, [5,7]_3\}$,
$X_2 \supseteq \{[0,1]_3, [2,4]_2, [5,7]_1\}$, $X_0 = \Omega \backslash (X_1 \cup X_2)$.

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 Problems, *CJSAP.*

[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).

We call an allocation that conforms to the objective of the arbitrator a *solution*. A solution exists for every possible instance of a CJSAP. This follows with a straightforward combinatorial argument immediately from the finiteness of the problem (which we assume throughout the paper).

Proposition 9. *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.

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 2002c; 2002b) for progressive direct mechanisms based on (Conen & Sandholm 2001)). 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.

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.

Definition 10 (Net utility). *The* net utility function $u_j(\cdot) : 2^\Omega \times \mathbb{Z} \to \mathbb{Z}$ *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.[14] For the more natural pricing modes, equilibrium price vectors need not exist. This negative result holds also for instances of CJSAP. We will therefore not discuss (anonymous) equilibrium pricing in detail and turn our attention to a solution concept for which solutions always

[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.

[14] 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.

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.

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 under-bid 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 under-bid 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 over-bid, because he risks that he receives the over-bidded bundle for a price between his true valuation and his bid and would, thus, realize a loss. If he would over-bid 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[15] *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.[16] 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 whenever 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 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 (for pointers to recent work, see (Bikhchandani *et al.* 2001; Parkes 2001)).

Proposition 11 (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. The proof follows from a straightforward combinatorial argument for 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. □

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 12. *CJSAP is (strongly) NP-hard.*

Proof. The decision problem variant of CJSAP is in NP. The proposition follows immediately from Proposition 6 and the polynomiality of the transformation given in Proposition 9. □

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.[17]

Consequently, we cannot expect to obtain optimal solutions for every problem 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

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

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

[17] 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.

its proximity to EJSP and, thus, typical job shop scheduling problems.

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. The presented approach to mechanism design 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 autonomy and information privacy more than necessary.

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, or of other open and flow shop scheduling problems.

Further effort is necessary to develop a useful set of benchmark problems and to study the specific structure of real-world problem scenarios. The impact of the utilization of heuristic modifications of standard combinatorial auction procedures on the desired properties efficiency and incentive-compatibility has to be studied in these specific scenarios.

On the other hand, the detailed knowledge about various subclasses of shop scheduling problems (with specific one or multi-dimensional optimization criteria) can have an impact on the analysis of structural phenomena in the corresponding combinatorial auction problems that may help to develop improved mechanisms and algorithms for specific problems.

References

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

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

Ausubel, L. M., and Milgrom, P. 2002. Ascending auction with package bidding. Technical report, University of Maryland, Department of Economics. April 25, 2002.

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. 2001. Preference elicitation in combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, 256–259.

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

Conen, W., and Sandholm, T. 2002b. Differential-revelation vcg mechanisms for combinatorial auctions. In *Proceedings of the AAMAS Workshop on Agent Mediated Electronic Commerce IV*, LNAI 2531. Springer.

Conen, W., and Sandholm, T. 2002c. 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.

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 substitutes. *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.

Lehmann, B.; Lehmann, D.; and Nisan, N. 2001. Combinatorial auctions with decreasing marginal utilities. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, 18–28.

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

Moulin, H. 1995. *Cooperative Microeconomics*. Princeton 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 *Proceddings of the National Conference on Artificial Intelligence (AAAI)*, 74–81.

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. Eine frühere Version erschien als Technical Report RRR 13-95 des Rutgers Center for Operations Research, Rutgers University, 1995.

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. Verfügbar unter `http://www.cs.cmu.edu/~sandholm/dissertation.ps` (Letzter Zugriff 08-08-2002).

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

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

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

Vohra, R. 2001. Research problems in combinatorial auctions. In *Online-Proceedings of the DIMACS Workshop on Computational Issues in Game Theory and Mechanism Design, 31.10.–2.11.2001*. `http://dimacs.rutgers.edu/Workshops/gametheory/vohra-paper.pdf`, (letzter Zugriff 08-08-2002).

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.