INFORMS TutORials in Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Cooperative and Noncooperative Games for Capacity Planning and Scheduling

Nicholas G. HallZhixin Liu

TUTORIALS IN OPERATIONS RESEARCH
INFORMS 2008

 | ISBN 978-1-877640-23-0
DOI 10.1287/educ.1080.0041

Cooperative and Noncooperative Games for Capacity Planning and Scheduling

Nicholas G. Hall
Department of Management Sciences and 3D Laboratory, Fisher College of Business, The Ohio State University, Columbus, Ohio 43210, hall_33@fisher.osu.edu

Zhixin Liu
Department of Management Studies, School of Management, University of Michigan–Dearborn, Dearborn, Michigan 48126, zhixin@umd.umich.edu

Abstract This tutorial provides an overview of cooperative and noncooperative games for capacity planning and scheduling decisions. For both cooperative and noncooperative games, we present some basic definitions and concepts, review the recent literature, and discuss our own current research. The discussion of the literature of cooperative games includes sequencing games, models of outsourcing operations, and economic lot-sizing and schedule planning games. Our research considers a combined capacity planning and scheduling problem where a supplier allocates capacity to distributors, who may then cooperate to share their capacity allocations and resubmit revised orders. The discussion of the literature of noncooperative games includes capacity allocation mechanisms based on order sizes and sales, incentives for truth telling, models of decentralized versus centralized decision making, and auction mechanisms for capacity allocation. Our research develops an auction mechanism to allocate capacity to various agents with jobs that can be scheduled profitably; depending upon the choice of market good in the auction, there may be some flexibility in converting the allocated capacity into a feasible schedule. We also provide a discussion of future research directions, where we identify several issues that can be studied using new cooperative and noncooperative games.

Keywords capacity planning and scheduling; cooperative game; noncooperative game; auction

1. Introduction

We first define the scope of this tutorial. We consider *cooperative games*, where several players have the opportunity to form a coalition for mutual benefit, provided that they cannot improve their profit or cost by defecting from the coalition. We also consider *noncooperative games*, where the players make self-interested decisions that may nonetheless converge to an equilibrium solution. Within noncooperative games, we consider various market mechanisms including auctions. Our discussion of *capacity planning* includes (a) setting total capacity levels (Cachon and Lariviere [5]), and (b) given a total available capacity, allocating it between different product lines or retailers (Cachon and Lariviere [6]). Our discussion of *scheduling* includes the allocation of scarce resources to jobs over time in various planning environments (Pinedo [41]). The sequencing of jobs is an important special case. The second "and" that appears in the title should be interpreted as "and/or" throughout the tutorial.

Most traditional research on capacity planning and scheduling issues focuses on tactical decision making by single agents using optimization methods. This approach relies on the assumption that the outcome of a particular decision is independent of the decisions of other agents. However, as a result of the recent research focus on *supply chain management*, an alternative perspective is becoming more common. Specifically, many recent research papers recognize the interconnectedness of the decisions of multiple agents within supply chains. These agents are often independently owned and motivated companies. The fact that the outcomes from the agents' decisions depend partly on the decisions of other independent agents makes game theory a natural approach to modelling those decisions. In practice, the agents may behave either cooperatively or noncooperatively; therefore it is important to consider both cases, as in this tutorial.

The recent literature contains two comprehensive general reviews, by Cachon and Netessine [7] for both cooperative and noncooperative supply chain games and by Nagarajan and Sošić [36] for cooperative supply chain games. However, our focus is specifically on capacity planning and scheduling games.

We now provide an overview of the tutorial. Section 2 provides basic definitions and concepts for the capacity planning and scheduling problems we consider. Section 3 discusses applications of cooperative games. In this section, the work discussed focuses on the existence or nonexistence of fair allocations of savings and/or capacity that ensure the cooperation of all the players. An important class of cooperative games is *sequencing games*, which require the fair allocation of savings from sequence changes. We present some results of our own research about a supply chain where several distributors form a cooperative game to share their capacity allocations. Section 4 discusses applications of noncooperative games. For example, a supply chain may contain a supplier and several retailers, each of which is privately motivated and perhaps even reluctant to share information with the others. Among the relevant questions is whether, if the decisions of the supply chain converge to a Nash equilibrium, this solution is also globally optimal for the supply chain. It is also valuable to identify mechanisms for capacity planning that induce the players to provide their information truthfully. We present some results of our own research about an ascending auction where agents buy capacity to schedule their jobs. Section 5 presents a list of potential future research directions. Finally, §6 provides a summary and some concluding remarks.

2. Capacity Planning and Scheduling

Due to the unpredictability of customer demand, it is common for suppliers to lack sufficient capacity to meet all the orders they receive. Hence, capacity planning problems arise. Such problems frequently need to be solved in industries with high fashion content, rapid technological development, or occasional demand surges (Fisher [19]). Iyer et al. [26] identify several practical examples in the fashion goods, telecommunications, and electricity supply industries. Durango-Cohen and Yano [16] observe that capital-intensive industries have to deal with both demand variability and the inability to change their capacity quickly to match demand, which often results in capacity shortfalls.

We consider two problems related to capacity planning. The first problem relates to the supplier's choice of overall capacity. This choice is essentially a trade-off between investment cost and the ability to meet orders from retailers. The second problem assumes that the total capacity is given and considers how to allocate it among the retailers.

Consider a supplier and $n \geq 2$ retailers. The retailers place orders with sizes $m_1, \ldots, m_n$, where we assume $m_1 \geq \cdots \geq m_n$. The supplier's available capacity is denoted by T. If $\sum_{i=1}^{n} m_i > T$, then the supplier needs to develop a *capacity allocation mechanism* to allocate capacity to the retailers. Let x_i denote the amount of capacity allocated to retailer i. Some capacity allocation mechanisms that are discussed below include the following:

Proportional allocation:

$$x_i = \min\left\{m_i, m_i K \Big/ \sum_{j=1}^{n} m_j\right\}, \quad i = 1, \ldots, n. \tag{1}$$

Linear allocation:

$$x_i = \begin{cases} m_i - \max\left\{0, \sum_{j=1}^{\tilde{n}} m_j - T\right\} \Big/ \tilde{n}, & \text{if } i \le \tilde{n} \\ 0, & \text{otherwise,} \end{cases} \tag{2}$$

where $\tilde{n}$ is the largest integer no greater than n such that $m_{\tilde{n}} - \max\{0, \sum_{j=1}^{\tilde{n}} m_j - T\}/\tilde{n} \ge 0$.

Relaxed linear allocation:

$$x_i = \min\left\{m_i, m_i - \frac{1}{n}\left(\sum_{j=1}^{n} m_j - T\right)\right\}, \quad i = 1, \ldots, n. \tag{3}$$

Uniform allocation:

$$x_i = \begin{cases} \frac{1}{\hat{n}}\left(T - \sum_{j=\hat{n}+1}^{n} m_j\right), & \text{if } i \le \hat{n} \\ m_i, & \text{otherwise,} \end{cases} \tag{4}$$

where $\hat{n}$ is the largest integer no greater than n such that $x_{\hat{n}} \le m_{\hat{n}}$.

Lexicographic allocation:

$$x_i = \min\left\{m_i, \left(T - \sum_{j=1}^{i-1} x_j\right)\right\}, \quad i = 1, \ldots, n. \tag{5}$$

Capacity allocation at the manufacturing level is generally defined as scheduling. Scheduling is the allocation of scarce resources to jobs over time to optimize one or more given objectives. A wide variety of scheduling problems have been studied. For example, Brucker [2] and Pinedo [41] discuss many applications of scheduling problems and algorithms for solving them. The research discussed here includes two types of scheduling environments. The first environment is a single machine and the second is a job shop. Single machine problems are frequently studied because in this environment it is easiest to obtain intuition about job priorities and to solve scheduling problems. See Johnson [28], Jackson [27], and Smith [53] for early discussions of single machine scheduling problems. A job shop environment includes multiple processing machines and jobs that follow different paths through them, and generically describes a traditional factory floor. See Muth and Thompson [35] for an early discussion of the job shop problem.

A regular scheduling cost is a nondecreasing function of job completion time. Let p_j denote the processing time of a job, and w_j its weight or value. In some of the problems considered, a job j also has a due date d_j. In a schedule σ, we let $C_j(\sigma)$ denote the completion time of job j, where the argument is omitted if the schedule is clear from context. An arbitrary regular scheduling cost of job j with completion time C_j is denoted by $f_j(C_j)$. The various works to be discussed consider the minimization of the following objectives:

$\sum w_j C_j$, the total weighted completion time;
$\sum w_j \max\{C_j - d_j, 0\}$, the total weighted tardiness; and
$\sum f_j(C_j)$, the total scheduling cost.

Pinedo [41] discusses many practical applications where it is necessary to schedule jobs *nonpreemptively*, i.e., from start through completion without interruption. Here we discuss both nonpreemptive and preemptive scheduling modes. In the latter case, the scheduler can interrupt the processing of a job and resume it later at the same point.

3. Cooperative Games

Section 3.1 describes some basic definitions and concepts of cooperative games. In §3.2, we review the literature of cooperative game applications to capacity planning and scheduling problems. In §3.3, we discuss our recent research study, which applies cooperative game theory to a capacity allocation and scheduling problem.

3.1. Basic Concepts

A cooperative game (Peleg and Sudhölter [40]) in characteristic function form is an ordered pair (N, v), where N is a finite set of players and v: $2^N \to \mathcal{R}_+$ is a function that assigns a nonnegative real number to every subset $S \subseteq N$, where $v(\varnothing) = 0$. The function $v(S)$ denotes the profit that subset S can achieve as a separate *coalition*. A coalition consisting of all the players is the *grand coalition*. We let $v(N)$ denote the profit of the grand coalition. In a particular instance, if there exists a division of $v(N)$ such that no subset of players can improve its total payoff by departing from the coalition, then that division is *fair*. The set of all fair payoff divisions is the *core* of a game. Formally, the core of (N, v) is denoted by $C(v)$ and defined by

$$C(v) = \left\{ x \in \mathcal{R}_+^n \,\middle|\, \sum_{i \in N} x_i = v(N); \sum_{i \in S} x_i \geq v(S),\ \forall S \subseteq N \right\}. \tag{6}$$

A game with a nonempty core for all instances is *balanced*. A class of games that is balanced is *convex games* (Shapley [47]), which satisfy

$$v(S \cup \{i\}) - v(S) \leq v(T \cup \{i\}) - v(T), \quad i \in N, S \subset T \subset N \backslash \{i\}. \tag{7}$$

More generally, the core of a game can be empty. Intuitively, this means that there does not exist a payoff division that can prevent some coalition being more profitable by defecting from the grand coalition.

Let $\pi(1), \ldots, \pi(N)$ denote a sequence of the N players, where $\pi \in \Pi$, and Π denotes all possible sequences. Also, let

$$\psi_i^\pi(v) = v(\{\pi(1), \ldots, \pi(n)\}) - v(\{\pi(1), \ldots, \pi(i-1)\} \cup \{\pi(i+1), \ldots, \pi(n)\}), \quad i \in N \tag{8}$$

denote the marginal value of adding player $\pi(i)$ to a coalition consisting of players sequenced as $\pi(1), \ldots, \pi(i-1), \pi(i+1), \ldots, \pi(n)$. Then, for the scheduling problems we consider, the *Shapley value* $\Phi(v)$ (Shapley [46]) is given by

$$\Phi_i(v) = \frac{1}{n!} \sum_{\pi \in \Pi} \psi_i^\pi, \quad i \in N. \tag{9}$$

Intuitively, the Shapley value is the marginal value added by a player when he joins a coalition in a specific sequence, averaged over all possible joining sequences.

The *maximal payoff*, $M_i(v)$, of player i is given by

$$M_i(v) = v(N) - v(N \backslash \{i\}), \quad i \in N. \tag{10}$$

The maximal payoff is an upper bound on what player i can obtain; clearly, if he asks for more, the remainder of the coalition can do better without him. Then the τ *value* (Tijs [54]) of player i, $\tau_i(v)$, is given by

$$\tau_i(v) = \max_{S \ni i} \left\{ v(S) - \sum_{j \in S \backslash \{i\}} M_j(v) \right\}, \quad i \in N. \tag{11}$$

The maximum in (11) represents player i's choice of coalition S to join.

3.2. Literature Review

We first discuss the sequencing games studied by Curiel et al. [13, 14], Maniquet [33], and Slikker [52]. Then we introduce the work of Cai and Vairaktarakis [8], Aydinliyim and Vairaktarakis [1], and Vairaktarakis and Aydinliyim [55], who study cooperation between manufacturers who outsource production from a common third party. Also, we describe the analysis of cooperation among the retailers during the ordering procedure in an *economic lot-sizing game* studied by Chen and Zhang [9]. Finally, we review the work of Schulz and Uhan [44], who examine cooperative games with supermodular costs. Such games include a type of *schedule planning game*.

Curiel et al. [13] define a class of cooperative games known as sequencing games. Suppose that a sequence of n jobs, without idle time, is given. In practice, this is typically the order in which the jobs arrived. However, the cost of this sequence is not necessarily minimum. Therefore, a pairwise interchange of two adjacent jobs may result in cost savings. A sequencing game considers the various possible distributions of such savings. The paper shows that sequencing games are convex, and therefore from Shapley [47] they have a nonempty core. More specifically, it is shown that distributions that result from the *equal gain splitting rule* are in the core. Whenever an adjacent pairwise interchange creates savings, this rule divides the savings equally between the two adjacent jobs. This rule is also the unique rule that satisfies three relevant properties. Because a sequencing game is convex, it is known that the Shapley value distribution is in the core, and the paper provides a closed-form expression for this distribution. The Shapley value distribution equally divides the gain that two players can make, between themselves and the players sequenced between them. Similar results are obtained for the τ-value of the game (Tijs [54]). The τ-value divides the cost savings in proportion to the maximal payoffs, as defined by (10), of the players.

Curiel et al. [14] extend the work of Curiel et al. [13] to consider a more general problem, where cost is measured by a weakly monotonic function of the completion times of the jobs. They describe a closed-form solution for a particular distribution and show that it is a core solution, for any instance. Under this distribution, the payoff to player i is the average of (a) the marginal value of i to the coalition of players that precede i in the initial distribution and (b) the marginal value of i to the coalition of players that follow i in the initial sequence. Slikker [52] also generalizes the work of Curiel et al. [13], but in a different direction. Whereas that earlier work allows only interchanges between jobs that are contiguous in the given schedule, here more general interchanges are allowed. Specifically, coalition $S \subset N$ is allowed to decrease (but not increase) the starting times of jobs in $N \backslash S$, and those jobs are allowed to change positions in the sequence. The paper shows that this game is balanced. An updated survey of the literature of sequencing games is provided by Curiel et al. [12].

Maniquet [33] discusses a problem of sequencing agents who require a service. Their sequence is given by their arrival order. However, efficiency requires that they be served in nondecreasing order of waiting cost. Although this problem is studied in a queueing context, it is mathematically equivalent to minimizing the total completion time in nonpreemptive single machine scheduling, or problem $1||\sum C_j$. The paper considers two alternative cooperative game formulations of the problem. In the first formulation, the value of a coalition is the total waiting cost that its members would incur if they collectively arrived first, whereas, in the second formulation, the coalition's value is found by assuming that they collectively arrive last. Because the value of a coalition should not depend on the position in which the members arrive, it is intuitive that these two values should be equal. The paper shows that, if the payoff assigned to each agent is defined by the Shapley value, then these two coalition values are equal. This result justifies the use of the Shapley value. The paper further presents an axiomatic justification for the use of the Shapley value.

Aydinliyim and Vairaktarakis [1] consider a set of manufacturers that outsource their operations to a single third party, for example, a large-scale electronics contract supplier. Each time window of the third-party capacity can be booked by the manufacturers at a known cost, in the order in which they arrive. Each manufacturer's total cost is this booking cost plus the weighted flow time of its jobs; minimizing this total cost is shown to be unary *NP*-hard (see Garey and Johnson [21] for related definitions). However, if a coalition of the manufacturers agrees, the third party reschedules their jobs to minimize their total cost. Moreover, if the manufacturers agree to form the grand coalition, a fraction of the booking cost is returned to all of them by the third party as an incentive, because it benefits from this situation. This problem is modelled as a cooperative game that is superadditive. The paper describes a closed-form expression for a particular savings allocation and shows that it is a core solution of the game. Computational experiments show that the total cost can be reduced by an average of 32%. The payoff division uses 40%–55% of those savings to compensate the approximately one-third of the manufacturers whose costs increase as a result of coordination. Cai and Vairaktarakis [8] provide a similar analysis for a group of manufacturers who outsource their finishing operations to a single third party. The third party prices capacity at two levels—a regular price and a higher peak price. If capacity is fully booked, then overtime can be purchased at a still higher price. Manufacturers book capacity to minimize their booking and tardiness costs. After bookings have been completed, the third party computes the savings if all manufacturers cooperate to form the grand coalition.

Vairaktarakis and Aydinliyim [55] consider a group of manufacturers who individually subcontract part of their workloads to a third party, with the objective of minimizing their total job completion time. The resulting solution is in general not optimal. The third party's rescheduling problem is considered under traditional nonpreemptive and preemptive job disciplines, as well as an overlapping situation where part of a preempted job can be processed simultaneously by the manufacturer and the third party. In the first case, the problem is binary *NP*-hard. However, in the last two cases, closed-form solutions are found for the optimal amount of subcontracting. In each case, the third party's solution provides substantial savings relative to the Nash equilibrium or first-come first-served solutions. The distribution of those savings is modelled as a cooperative game. The paper shows that this game is convex, describes a closed-form expression for a particular savings allocation, and shows that it is a core solution of the game. A computational study shows that decentralized solutions tend to underutilize third-party capacity. Moreover, relative to Nash equilibrium (respectively, first-come first-served) schedules, the savings scheme allocates 70% (respectively, 50%) of the savings to the 53% (respectively, 42%) of the players who lose by cooperation. Therefore, the incentive for coordination is stronger relative to the Nash equilibrium than to the first-come first-served solution.

Chen and Zhang [9] study an economic lot-sizing game with general concave ordering costs. Multiple retailers form a coalition by placing joint orders with a single supplier to reduce their total ordering cost. Although this problem is uncapacitated, the coordination of orders that the authors model is also useful in finite capacity environments. When both the inventory holding cost and the backlogging cost are linear functions, it is shown that the core of this game is nonempty. The analysis begins with an integer programming model of the game. The linear relaxation of this model has zero integrality gap, which makes it possible to analyze the game by using linear programming duality, following the work of Owen [39]. It is shown that, by contrast with core solutions that are found using duality for other games, not every optimal dual solution defines a core allocation. Nevertheless, the addition of specific inequalities to the linear relaxation generates a dual solution that is in the core. Consequently, a core allocation can be computed in polynomial time.

Schulz and Uhan [44] study cooperative games with supermodular costs, which include a schedule planning game with total weighted completion time objective as a special case. In such games, as a coalition grows, the cost of adding a particular agent increases. This makes it hard to build a grand coalition, and in general, such games do not have a core solution. However, a coalition of all the players can be sustained by penalizing a coalition for defecting. The minimum such penalty that is required to ensure the existence of a core solution is the *least core value*. The paper considers a schedule planning game, where the cost incurred by any coalition S is the minimum weighted sum of completion times of the jobs of S if they are all scheduled on the same machine. In this context, the least core value has a practical interpretation: the amount charged to any coalition for opening its machine. The paper shows that the computation of the least core value is binary *NP*-hard. It is also shown that a ρ-approximation algorithm for the $\bar{x}$-maximally violated constraint problem, where $\bar{x}$ is a specific cost allocation, provides a ρ-approximation for computing the least core value of schedule planning games. This result is used to develop a fully polynomial time approximation scheme (see Schuurman and Woeginger [45] for definitions) for computing the least core value of schedule planning games.

3.3. Our Research

For details of the results in this section, we refer the reader to Hall and Liu [22]. Consider a supplier who manufactures products in response to orders from several distributors. Let $K = \{1, \ldots, k\}$ denote the k products produced by the supplier, and $N = \{1, \ldots, n\}$ denote the n distributors. Let T denote the supplier's total production capacity, in time units. An order from distributor j for product i is denoted by ν_{ij}. Let p_{ij} denote the integer number of units of capacity required by order ν_{ij}. Without loss of generality, we assume that $p_{ij} > 0$ and that a unit of capacity is not divisible for scheduling purposes by the supplier. Let $P_j = \sum_{i \in K} p_{ij}$ and $P_{\max} = \max_j \{P_j\}$. The scheduling cost is the total weighted completion time, $\sum_{\nu_{ij} \in S} w_{ij} C_{ij}(\sigma)$, where w_{ij} represents either a holding cost per unit time or the value already added to order ν_{ij}. Suppose that each order ν_{ij} generates a revenue u_{ij} for the manufacturer and a profit v_{ij} for distributor j. Then, the profit per unit of manufacturing capacity from order ν_{ij} is $\bar{v}_{ij} = v_{ij}/p_{ij}$ for distributor j.

Traditional capacity allocation mechanisms are equitable between the distributors, in that the allocation is nondecreasing with order size. However, traditional mechanisms such as (1)–(5) are based solely on order sizes. Because they do not consider revenue or scheduling cost, the resulting capacity allocations may generate small revenues and large costs for the supplier. To resolve this problem, we propose the following algorithm, which integrates both scheduling information and equity considerations into the supplier's capacity allocation decisions.

Proportional Allocation Algorithm (PAAr). Given $v_{ij}, u_{ij}, p_{ij}, w_{ij}, T$, and a measure $r \geq 1$ of the relative bargaining power of the distributors.

1. Find a schedule with integer capacity requirements that maximizes the total net profit of the supplier.
2. Allocate capacity $X_j = \min\{P_j, P_j T / \sum_{i \in N} P_i\}$ to distributor j, for $j = 1, \ldots, n$.
3. Allocate a set of resubmittable orders with total capacity requirement $\widehat{P}_j \geq r X_j$ to distributor j, for $j = 1, \ldots, n$. Include in the set of orders all those that appear in the supplier's optimal schedule in Step 1. If the total capacity of those orders is less than $r X_j$, then select other orders in nonincreasing order of revenue per unit of capacity, allowing partial orders where applicable, until $\widehat{P}_j \geq r X_j$.

The implementation of the PAAr allocation mechanism requires the supplier to solve an optimization problem to maximize its total net profit. We have the following results.

Theorem 1. *For the manufacturer's scheduling problem in Step* 1 *of PAAr, if integer-size partial orders are allowed, then an optimal schedule can be found in polynomial time. If partial orders are not allowed, then it is binary* NP*-hard to find an optimal schedule, and an optimal schedule can be found in pseudopolynomial time.*

Once the distributors receive capacity allocations from the supplier, they can either resubmit their orders individually or by sharing their capacity allocations cooperatively. The latter case defines a class of *capacity sharing games*. First, we allow partial orders. That is, we consider a situation where the distributors can resubmit an order that is no larger than the corresponding original size and has an integer capacity requirement. The capacity sharing and order revision problem of each distributor or coalition of distributors can be modelled as a mathematical program. Let p_{ij} denote the capacity requirement of resubmittable order ν_{ij}. If a coalition contains a set N' of distributors, each with allocated capacity X_j, then this mathematical program is as follows:

$$\max \sum_{i\in M}\sum_{j\in N'} \bar{v}_{ij} x_{ij} \tag{12}$$

$$\text{s.t. } x_{ij} \le p_{ij}, \quad i\in M,\ j\in N' \tag{13}$$

$$\sum_{i\in K}\sum_{j\in N'} x_{ij} \le \sum_{j\in N'} X_j \tag{14}$$

$$x_{ij} \ge 0 \text{ and integer}, \quad i\in M,\ j\in N'. \tag{15}$$

The objective function (12) maximizes the total profit of the coalition. Constraints (13) ensure that each revised order is no larger than the corresponding original size. Constraint (14) ensures that the total capacity requirement of the revised orders does not exceed the total allocated capacity. It can be shown that the constraint matrix of the above mathematical program is totally unimodular. Therefore, the linear program (LP) defined by replacing (15) by $x_{ij}\ge 0$, for $i\in k$, $j\in N'$, has an integer solution. We develop an efficient algorithm to find a core member of this *linear programming game* (Owen [39]).

Theorem 2. *A core member of the distributors' LP game can be found in $O(kn)$ time.*

For some special classes of LP games, the *Owen set* of core solutions (Owen [39]) covers the entire core (Shapley and Shubik [48], Kalai and Zemel [29], Samet and Zemel [43]). However, the following result shows that this is not the case for the distributors' LP game.

Theorem 3. *In the distributors' LP game, there exist core members that are not in the Owen set.*

Indeed, core solutions that are not in the Owen set may be more reasonable than those that are. In this case, it is likely that the distributors will propose a different payoff division. To evaluate whether this payoff division is in the core, we need to answer the *core membership test*: given an instance of the distributors' game and a payoff vector, is that vector in the core of the instance? In a general LP game, the core membership test is unary co-*NP*-complete (Fang et al. [18]).

The next two results are interesting in two respects. First, in most cases in the literature where the core membership test is not unary *NP*-complete, the Owen set covers the entire core. However, we show that, although the Owen set is a proper subset of the core (Theorem 3), the core membership test in the distributors' LP game is pseudopolynomially solvable. Second, this is an example within cooperative games of a core membership test that is binary co-*NP*-complete, but solvable in pseudopolynomial time.

Let $\widehat{P}_{\max} = \max_j\{\widehat{P}_j\}$, and let $\bar{X} = \sum_{j\in N} X_j \le T$.

Theorem 4. *For a given instance of the distributors' LP game, whether a given payoff vector is in the core of the instance can be determined in $O(n\max\{k,\bar{X}\}\min\{\bar{X},\widehat{P}_{\max}\})$ time.*

The following result shows that we cannot expect to find a polynomial time algorithm for the core membership test in the distributors' LP game, unless $P = NP$.

Theorem 5. *The core membership test in the distributors' LP game is binary co-NP-complete, even for a single product.*

We now consider the situation where partial orders are not allowed. Thus, an order must be accepted in full or rejected. This problem is modeled as a *knapsack game* because each coalition $N' \subseteq N$ faces a knapsack problem with knapsack capacity $\sum_{j \in N'} X_j$, where each item has a size $\hat{p}_{ij}$ and a value $\hat{v}_{ij} = \hat{p}_{ij}\bar{v}_{ij}$. Unlike the LP game, the knapsack game is not balanced. Therefore, we need to solve the *core emptiness test*: determine whether an instance of the knapsack game has a core member or not. Because in practice monetary payoffs are not fractional, we develop a dynamic programming algorithm to test whether an instance of the knapsack game has an integer core member or not. Let V denote the maximum profit of the grand coalition.

Theorem 6. *It is binary* NP*-hard to decide whether an instance of the knapsack game has an integer core member, even for a single product. An integer core member of an instance of the knapsack game, if one exists, can be found in* $O(n\bar{X}\min\{\bar{X}, \widehat{P}_{\max}\}V^2)$ *time.*

Theorem 6 is an example of a core emptiness test that is binary *NP*-hard but solvable in pseudopolynomial time. Moreover, for the core membership test in the knapsack game, we have the following result.

Theorem 7. *It is binary* NP*-hard to decide whether a given payoff vector is in the core of an instance of the knapsack game, even for a single product. Whether a given payoff vector is in the core of the instance can be determined in* $O(n\max\{m, \bar{X}\}\min\{\bar{X}, \widehat{P}_{\max}\})$ *time.*

To evaluate the benefit of coordination by the distributors under various parameter settings, we conduct a computational study of the effect of four parameters. We consider five capacity allocation mechanisms: proportional (1), linear (2), and PAAr with $r \in \{1.00, 1.15, 1.33\}$. The number of products, k, is set to 5, 10, 20, 50, or 100. Order revenues for the suppliers are either independent, positively correlated, or negatively correlated with the distributors' profits. Finally, partial orders are either allowed or not. In each case, we consider the effect of parameter changes on the ratio of the distributors' total profit with coordination over that without coordination. Our main findings are as follows.

First, coordination between the distributors is more valuable when the supplier uses PAAr, compared with using the proportional and linear mechanisms, especially for small r values. By implicitly forcing the distributors to follow the supplier's schedule, PAAr potentially hurts the distributors' profit, especially when the supplier's and distributors' priorities conflict. However, coordination by the distributors mitigates this effect.

Second, coordination generates more profit for the distributors when the number of products is relatively small, for example, 7.08% more profit for five products compared to 0.42% more profit for 100 products. As k decreases, each distributor has fewer orders to revise within its allocated capacity, and thus may utilize its allocated capacity less effectively. Consequently, capacity sharing by the distributors becomes more valuable.

Third, slightly more profit is achieved by the distributors' coordination when manufacturer's revenues and distributors' profits are negatively correlated, specifically 2.74% additional profit when compared with 2.58% for independent and 2.41% for positively correlated data. This means that failure to coordinate hurts the distributors more when the priorities of the supplier and the distributors conflict.

Fourth, coordination is more valuable when partial orders are not allowed, specifically 2.99% more profit compared with 2.17% additional profit with partial orders. When partial orders are not allowed, each distributor faces a zero-one knapsack problem with a small number of items, hence the allocated capacity may be used ineffectively.

4. Noncooperative Games

Section 4.1 describes some basic definitions and concepts of noncooperative games, including auctions. In §4.2, we review the literature of cooperative game applications to capacity planning and scheduling problems. In §4.3, we discuss our recent research study, which develops an ascending price auction for a combined capacity allocation and scheduling problem.

4.1. Basic Concepts

A noncooperative game (Cachon and Netessine [7]) in normal form consists of players indexed by $i = 1, \ldots, n$, strategies denoted by $x_1, \ldots, x_n$ available to each player, and payoffs $\pi_i(x_1, \ldots, x_n)$, $i = 1, \ldots, n$, received by each player. If each player chooses a strategy and no player can benefit by changing its strategy while the other players keep theirs unchanged, then the current set of strategy choices and the corresponding payoffs constitute a *Nash equilibrium* (Nash [37]). If each player has a strategy that maximizes its payoff regardless of the strategy choices of the other players, then the current set of strategy choices and the corresponding payoffs constitute a *dominant strategy equilibrium*.

The *Pareto frontier* is the set of strategies such that each player cannot improve its payoff unless some other player's payoff is reduced. Given a capacity allocation mechanism from the supplier, the retailers compete for capacity and face a *capacity allocation game*. A *Pareto allocation mechanism* is a mechanism that maximizes the sum of retailer profits, assuming that all the retailers truthfully submit their optimal orders. Under a *truth inducing mechanism*, each retailer who truthfully submits his optimal order quantity is at a Nash equilibrium.

An *auction mechanism* (McAfee and McMillan [34]) is a set of rules for formulating an allocation of scarce resources based on received bids. In an *ascending auction*, agents submit successively higher bids. At each round, all agents are allowed to submit bids simultaneously, and the seller determines which bids are admitted. If at any round no agent submits a bid, then the auction reaches closure and stops. Let $\varepsilon(\cdot)$ denote a weakly concave nondecreasing bid increment function, which defines a bid's minimum increment over the current price for the bid to be admissible.

An ascending auction mechanism is defined as follows for each market good. At any time, the *bid price* for market good j, denoted by β_j, is the highest bid it has received so far. If the market good has received no bids, then its bid price is undefined. The *ask price* of a market good j equals $\alpha_j = \beta_j + \varepsilon(\beta_j)$ if β_j is defined, and equals the reserve value a_j if β_j is not defined. Given a set $\alpha = \{\alpha_1, \ldots, \alpha_n\}$ of ask prices, each agent solves its *bid determination problem*. This problem requires determining which market goods to bid for and how much to bid, to maximize the agent's profit if it wins. A bid that is at least equal to the ask price of the market good is admissible; other bids are rejected. Agents cannot withdraw bids, but an agent can replace its bid with a new bid at later rounds. Given a set of admissible bids, including those that are already admitted, the seller solves its *winner determination problem*. This problem requires labelling each admissible bid as winning or losing so as to maximize the seller's revenue. This mechanism guarantees that prices increase monotonically and that the bidding process terminates. Finally, a *combinatorial auction* is an auction in which bidders can place bids on combinations of items, or bundles, rather than only on individual items.

4.2. Literature Review

First, we discuss the work of Cachon and Lariviere [4–6], who analyze incentives, equilibrium, and supply chain performance under a variety of capacity planning and allocation mechanisms based on order sizes or past sales. Next, we discuss research by Fang and Whinston [17] and Ganesh et al. [20], who study capacity allocation using option contracts and pricing. Also, we describe capacity planning and allocation models with applications to semiconductor manufacturing (Mallik and Harker [32], Karabuk and Wu [30]) and to network communications (Niyato and Hossain [38]). Finally, we review the studies of noncooperative decentralized scheduling problems by Kutanoglu and Wu [31], Wellman et al. [57], Reeves et al. [42], and Bukchin and Hanany [3].

Cachon and Lariviere [5] consider a two-stage supply chain consisting of a supplier and several retailers who make decisions in a single period. The supplier can choose between two classes of capacity allocation mechanisms: those that induce truth telling by the retailers, and those that are subject to strategic ordering. Considering both classes of mechanisms, the supplier announces the combination of mechanism and capacity that maximizes expected profit. The calculation of expected profit is based on order submissions by the retailers that are filled using the announced mechanism. Various capacity allocation mechanisms are studied in the paper. The Pareto allocation mechanism does not induce truth telling. However, both uniform (4) and lexicographic (5) allocation mechanisms induce truth telling. The relaxed linear allocation mechanism (3) maximizes each retailer's profits in expectation, assuming that other retailers follow the same strategy. A computational study indicates that the lexicographic and especially the uniform mechanism are effective when capacity is cheap. The relaxed linear mechanism gives excellent performance that is robust across different parameter settings. It is further shown that a manipulable mechanism may lead the supplier to choose a higher level of capacity than under a truth inducing mechanism. Overall, truth telling helps with the allocation of capacity among the retailers but may distort the total capacity level.

Cachon and Lariviere [6] compare proportional (1), linear (2), and uniform (4) capacity allocation mechanisms in a supply chain consisting of a supplier and two retailers. The proportional and linear mechanisms are subject to order inflation, whereas the uniform mechanism is truth inducing. The paper shows that ordering behavior is complex under either proportional or linear allocation. For one thing, there may exist no pure strategy Nash equilibrium; based on the results of a computational study, this outcome seems to occur frequently. Another possibility is that there are multiple Nash equilibria. Under linear allocation these can all be found, whereas under proportional allocation a numerical search procedure may not find all of them. With uniform allocation, however, every retailer that chooses its optimal order size is always at a unique Nash equilibrium. In instances where the proportional or linear mechanism has an equilibrium, supply chain profits are on average higher than with the uniform mechanism. This occurs because order inflation increases the supplier's profits. However, in cases where the proportional or linear allocation does not have an equilibrium, order inflation is likely. The resulting randomness in capacity allocation often leads to poor inventory allocation among the retailers. Consequently, even though the supplier's profit may increase, this may be more than offset by reduced profit at the retailers. The net effect on supply chain profit tends to be positive when the wholesale price is high, and strongly negative—thus, recommending uniform allocation as an alternative—when it is low.

Cachon and Lariviere [4] consider a supplier that sells to two retailers over a two-period planning horizon. Demand in each period has two possible states, high and low. The supplier chooses a wholesale price and a capacity level. The retailers choose their order quantities and, subject to demand, their sales quantities. The paper compares the performance of a fixed capacity allocation policy such as (1)–(5) based on order size, and a turn-and-earn allocation policy, which is based on sales at the previous period. By deriving unique Nash equilibrium solutions for all the players, the authors demonstrate that the turn-and-earn policy may benefit the supplier at the expense of the retailers and the overall supply chain. This occurs because the turn-and-earn policy may induce the retailers to sell more than is optimally profitable, to increase their future capacity allocations. Problematically, because both retailers do this, no such increase actually occurs. However, the supplier benefits from higher capacity utilization. On the other hand, the supplier may change its wholesale price and capacity level to account for the retailers' increased sales, which can compensate the retailers' and overall supply chain profits. When capacity is mildly tight, this compensation is sufficient and overall supply chain performance improves. However, when capacity is extremely tight, the retailers' competition wastes more profit than the supplier gains, and overall supply chain performance degrades.

Fang and Whinston [17] consider a supply chain in which a monopolistic supplier sells to two risk neutral customers. The paper designs an option contract that guarantees either delivery by the supplier or compensation for failure to deliver. From the supplier's perspective, each customer's marginal utility for the good is assumed to be either "high" with probability λ or "low" with probability $(1-\lambda)$. Each customer receives total utility equal to the utility of the good supplied, less the payment to the supplier. The supplier's profit equals the total payments received from the customers, less the cost of capital investment. The model consists of two phases. In the contracting phase, the supplier announces the option price and strike price; then each customer chooses how many options to purchase. In the consumption phase, the supplier decides on the level of capacity investment, then customer demands are realized, then customers decide how many options to exercise, and finally, the supplier satisfies first the options and then the remaining demands. The option contract achieves three things. First, high-utility customers pay higher marginal prices. Second, high-utility customers receive first priority for available product, which results in effective allocations. Third, the number of options purchased informs the supplier about the number of customers of different types, because only high-utility customers will purchase it. Consequently, the supply chain with the option contract achieves the same expected total return as where the supplier knows the number of customers of each type before investing in capacity.

Ganesh et al. [20] develop a congestion pricing mechanism for allocating bandwidth in communication networks. There are n competing users. At each time period, each user i decides a quantity x_i of data packets to transmit onto the communication link. More formally, user i solves the problem

$$\max_{x_i \geq 0} \pi_i(x) = u_i(x_i) - x_i p(x), \tag{16}$$

where $\pi_i(x)$ is the single-stage profit of user i as a function of the decisions x of all the users, $u_i(x_i)$ is the utility from scheduling x_i packets, and $p(x)$ is the unit price for bandwidth. A traditional equilibrium model assumes that the price function and all utility functions are known. However, this is typically an unrealistic assumption in modern communication networks. Therefore, beliefs about the price are modelled using the expectation of a probability distribution that is derived from historic prices. Under standard regularity assumptions about price and utility, it is shown that there exists a unit price that is *self-consistent*, in that if all users predict that price and select their transmission rates accordingly, then the resulting price coincides with the prediction. Also, the unit price approximately converges to a Nash equilibrium for a large number of players. Moreover, this equilibrium price equates the price and the marginal utility of bandwidth for each player, which results in a capacity allocation that maximizes total welfare.

Mallik and Harker [32] consider the problem of allocating capacity to products at a major U.S. semiconductor manufacturer. A central coordinator is responsible for the allocation decisions. Problematically, product managers have an incentive to inflate their demand forecasts to gain greater capacity allocation. Moreover, manufacturing managers have an incentive to deflate their forecasts of available capacity to allow for production uncertainties. The paper develops a combined capacity allocation and incentive payment scheme that induces truth telling. The model proceeds as follows. First, the managers learn their demand or capacity distributions privately; then the coordinator announces the allocation mechanism and the payment scheme; then the managers submit their demand and capacity forecasts; then the coordinator allocates capacity to products; then production takes place and capacity is realized; and finally, demand is realized and managers are rewarded. Both product managers and manufacturing managers act to maximize their expected utility. The paper proposes a modified lexicographic mechanism, which is a variant of (5). Under this mechanism, each product line receives the minimum of the so-far unallocated capacity and the product manager's optimal newsvendor quantity. This scheme by itself is enough to induce truth telling by the product managers. However, a payment scheme is additionally required to induce truth telling by the manufacturing managers. This combination of capacity allocation and payment scheme achieves an optimal allocation as a dominant equilibrium.

Karabuk and Wu [30] also study incentive issues in capacity allocation for semiconductor manufacturing. As in Mallik and Harker [32], product managers have an incentive to inflate their demand estimates to increase their capacity allocations. The paper resolves this issue by recommending a specific bonus payment to the product managers. Each product manager receives an initial capacity allocation, based on strategic capacity planning; however, this allocation may change at tactical decision points. The capacity allocation decision is modelled as a noncooperative game, with the following sequence of events. At the start of a fiscal period, strategic capacity allocation decisions are made and form the basis for operational decisions; at the next tactical planning point, the product managers report mean demand and the coordinator announces an allocation mechanism and a bonus function; then the product manager observes the true demand and announces a mean demand; and finally, capacity is allocated by applying the allocation mechanism to the mean demand. The bonus paid to the product managers effectively neutralizes their incentive to inflate demand estimates. The main conclusion of the paper is that keeping a fraction of the available capacity in hand is both necessary and sufficient for making the appropriate bonus payments to the product managers. A case study indicates that this fraction need only be about 10%.

Niyato and Hossain [38] develop noncooperative games for three problems that arise in fourth generation (4G) network communications. First, they consider the problem of making long-term bandwidth allocations to a service area from the different access networks that are available in that area. This problem is formulated as a noncooperative game, where the players are the wireless metropolitan area network and the cellular network for a particular area. The Nash equilibrium is found by using best response functions that maximize the players' utility and then solving the resulting equations. Moreover, the paper shows that the Nash equilibrium solution maximizes the total utility of the network. Second, using the bandwidth allocation, it is necessary to find capacity reservation thresholds for high-priority connections that satisfy connection-level requirements for quality of service. This problem is formulated as a *bargaining game*, where the players are defined as an incoming new connection, a vertical handoff connection, and a horizontal handoff connection. The Pareto optimal solutions of the game provide several candidate strategies, one of which is an equilibrium. Third, it is necessary to decide the amount of short-term bandwidth to allocate to an arriving connection. This problem is formulated as a noncooperative game where the players are the networks available, and they make decisions about how much bandwidth to offer to an incoming connection. The Nash equilibrium solution, which is obtained by using best response functions, is unique. This solution can be found using an iterative search algorithm that in practice converges quickly enough for online use.

For scheduling problems, when either centralized or cooperative decentralized decision making is possible, it often yields higher system values than noncooperative decentralized approaches. However, there are many situations where agents have competing interests and privately held information about their job requirements and values. Here the agents may not communicate their information reliably, since the use of the information may create incentives for the agents to misrepresent their information (Clearwater [11]). These difficulties associated with centralized or cooperative decentralized decision making motivate the following studies of competitive decentralized scheduling.

Bukchin and Hanany [3] analyze a dispatching and sequencing model in which jobs are owned by different departments of a company. Each department has a set of jobs and needs to decide whether to process each job using an in-house resource with limited capacity or using a less efficient subcontractor with unlimited capacity. The objective is to minimize the total completion time of the jobs. A centralized schedule minimizes the total completion time of jobs across all the departments. Assuming that all job processing times are publicly known, a decentralized Nash equilibrium schedule is a schedule in which each department cannot reduce its total scheduling cost given the schedule of all the other jobs. It is shown that a centralized schedule may not be a decentralized Nash equilibrium schedule. As a result, departments have an incentive to dispatch their jobs suboptimally to either resource for processing. Also, the decentralized Nash equilibrium schedule may not be unique and typically is intractable to find. The authors present techniques for finding lower bounds on the decentralized Nash equilibrium cost. A computational study is conducted to investigate the ratio between the Nash equilibrium cost and the cost of the centralized optimal solution. For the instances tested, this ratio ranges from 1.00 to 1.35. Finally, the authors design a scheduling-based coordination mechanism, which guarantees that at equilibrium a centralized optimal solution is found.

Kutanoglu and Wu [31] design a combinatorial auction mechanism for allocating capacity to jobs in a job shop scheduling problem. They establish a connection between subgradient search in Lagrangian relaxation and the progress of their combinatorial auction. They compare two different approaches to price adjustment, depending on whether the auctioneer makes more aggressive price changes earlier in the auction to assess demand (adaptive) or keeps the size of price changes constant (nonadaptive). They also compare two different payment functions, depending on whether price discrimination is used (augmented) or is not used (regular). The problem with the regular payment function is that there is usually no optimal set of prices that supports an optimal resource allocation. A computational study suggests that the adaptive price changes auction converges more quickly than the nonadaptive auction. Also, the auction with regular payment functions converges more quickly than the one with augmented payment functions.

Wellman et al. [57] consider a variety of issues in decentralized scheduling. Several agents have a single job that can be processed preemptively on a common facility. The agents wish to maximize their profit, which is the value of their scheduled jobs, less their scheduling cost and cost of purchasing capacity. The facility owner wishes to maximize its total revenue, which is its revenue from selling capacity plus its reserve value from unsold time slots. The system value is the value of the scheduled jobs plus the reserve value of the unsold time slots, less the scheduling cost. For the general problem with arbitrary processing times, the paper describes ascending auction mechanisms with two alternative choices for market goods. First, if *time slots* are chosen as market goods, then an equilibrium solution is globally optimal. However, it is possible that no equilibrium solution exists, due to the existence of complementarities; i.e., the value of holding a time slot is increased if another time slot is also held. Second, the paper discusses the use of *time slot bundles* as market goods. A time slot bundle (p_i, y) consists of p_i time slots, the last of which is at time y. Using this market good, the system value of an equilibrium solution is not always bounded in a practically useful way. For a special case of the problem with unit processing times, i.e., $p_j = 1$ for $j \in N$, it is possible to establish bounds on the difference between the optimal and auction closing prices, as well as bounds on the suboptimality of the system value. A direct revelation mechanism in the form of a generalized Vickrey auction (Vickrey [56]) is also discussed. The problem with this approach lies in solving a large combinatorial problem once all the information has been revealed. Some other results in Wellman et al. [57] are discussed in the context of our research in §4.3. In a similar capacity allocation and scheduling environment, Reeves et al. [42] illustrate the difficulty of making strategy choices in a simple ascending auction game. They show that straightforward bidding policies and their variants are not approximately optimal. Moreover, analytic methods are incapable of identifying appropriate strategies within the very large solution space. Shen et al. [51] provide a survey of the literature of agent-based allocation mechanisms for intelligent manufacturing systems.

4.3. Our Research

For details of the results in this section, we refer the reader to Hall and Liu [23]. We consider the following capacity allocation and scheduling problem. There are n competing agents, $\mathcal{A} = \{1, \ldots, n\}$. Each agent j has a single job j with a processing time p_j. This job can be processed nonpreemptively on a common facility. The facility can process at most one job at a time. We assume that all the jobs are available at the start of the planning horizon. From the facility owner's perspective, assuming that each agent has only a single job is not restrictive.

Each job j generates a revenue v_j for agent j if processed, and incurs a scheduling cost $f_j(C_j)$ if completed at time C_j. We consider a general cost function $f_j(C_j)$, which is a nondecreasing function of C_j, for $j = 1, \ldots, n$. Unless otherwise specified, all of our results hold for this function. Assume that the price for purchasing capacity to process job j is β_j. Job j generates a profit $v_j - f_j(C_j) - \beta_j$ for agent j if it completes processing at time C_j, and generates zero profit for agent j if it is not processed. The objective of each agent is to maximize its profit.

Let e represent the facility owner that sells capacity. The facility owner has a set of consecutive time slots, $\mathcal{T} = \{1, \ldots, T\}$, to allocate to agents. Each time slot t that defines a time interval $[t-1, t]$ is a resource that corresponds to a unit of processing time in a schedule. The facility owner sets a reserve value q_t for time slot t. The objective of the facility owner is to maximize its revenue from selling capacity plus its reserve value from holding unallocated time slots.

We design an auction mechanism with three alternative market goods. First, a *fixed time block*, $(p, \bar{u})$, $p \le \bar{u}$, is a set of p consecutive time slots from $\bar{u} - p + 1$ to $\bar{u}$. Even if they are not formally defined as such, fixed time blocks are used as market goods in the literature (Shaw [49, 50]). A fixed time block $(p, \bar{u})$ is feasibly scheduled if it is allocated at least the p consecutive time slots $\bar{u} - p + 1, \ldots, \bar{u}$. Second, a time slot t is a fixed time block $(1, t)$, as discussed by Kutanoglu and Wu [31] and Wellman et al. [57]. Although time slots are the most natural market goods, their use may result in allocations that cannot feasibly process nonpreemptive jobs. Finally, we introduce a new market good which is a *flexible time block*, (p, u), $p \le u$, consisting of a set of any p consecutive time slots, with the last one being no later than u. A flexible time block (p, u) is feasibly scheduled if it is allocated no less than p consecutive time slots, with the last one being no later than u. Let $\mathcal{G}_1$ and $\mathcal{G}_2$ denote the set of fixed and flexible time blocks, respectively. Agents simultaneously submit bids $\mathcal{B} = \{B_1, \ldots, B_n\}$ at each round. A bid is a tuple $B_j = \langle g_j, \beta_j \rangle$, where g_j is a market good, and β_j is a bid price.

A *solution* is a mapping $F\colon \mathcal{T} \to \mathcal{A} \cup e$, indicating which agent, if any, receives each time slot. Let $F_j = \{t \mid F(t) = j\}$ denote the set of time slots allocated to agent j, and $F_e = \{t \mid F(t) = e\}$ the set of unallocated resources in F. We assume that each agent schedules its job as early as possible within its allocated capacity. Thus, each solution uniquely defines a schedule. The *system value* $V(F)$ of a solution F is the total profit of the n agents and the facility owner. A *globally optimal solution* maximizes the system value, under the condition that all the agents' information is known.

In defining our auction mechanism, we need to consider the solvability of the bid determination problem. For both the fixed and flexible time block auctions, this problem can be solved by a straightforward enumeration procedure in $O(T)$ time.

We also consider the solvability of the winner determination problem. Suppose there are n bids $\bar{u}_1 \le \cdots \le \bar{u}_n$ (respectively, $u_1 \le \cdots \le u_n$) for fixed (respectively, flexible) time blocks. We have the following result.

Theorem 8. *For the winner determination problem with fixed time blocks as market goods, a schedule with maximum total revenue for the facility owner can be found in* $O(\max\{n \log n, T, n\bar{u}_n\})$ *time.*

By contrast, for the winner determination problem with flexible time blocks as market goods, we have a negative result.

Theorem 9. *The recognition version of the winner determination problem with flexible time blocks as market goods and with arbitrary reserve values is unary* NP*-complete.*

However, two natural special cases of this problem can be solved efficiently, as shown in the next two results.

Theorem 10. *For the winner determination problem with flexible time blocks as market goods and with nonincreasing reserve values, a schedule with maximum total revenue for the facility owner can be found in* $O(\max\{n\log n, T, nu_n\})$ *time.*

Theorem 11. *For the winner determination problem with flexible time blocks as market goods, if there exists an optimal solution without unallocated time slots, then an optimal solution can be found in* $O(\max\{n\log n, T, nu_n\})$ *time.*

Using the citation from Theorems 10 and 11, we design a heuristic for the unary *NP*-hard general winner determination problem. Our computational results show that this algorithm routinely delivers close to optimal solutions, with an average error of less than 3.2%. Thus, despite the result in Theorem 9, the winner determination problem with flexible time blocks can be solved effectively in practice.

Next, we address the properties of equilibrium solutions in the fixed and flexible time block auctions. The definition of an equilibrium solution is that each buying agent and the facility owner obtains an allocation that maximizes its profit, given the current prices.

Theorem 12. *There exists an equilibrium solution for any instance of the resource auction problem with fixed or flexible time blocks as market goods. In general, the equilibrium solution is not unique.*

Even though an equilibrium solution always exists with time blocks as market goods, the globally optimal solution typically is not provided by an equilibrium solution, as shown by the next two results.

Theorem 13. *With time blocks as market goods, the scheduling cost* (*respectively, system value*) *of a solution at equilibrium can be arbitrarily large* (*respectively, small*) *compared with that of an optimal solution, even when all reserve values are* 0*, the facility owner has sufficient time slots that all jobs can be processed, and the scheduling cost of each job is the weighted completion time.*

Theorem 14. *Even with nonincreasing reserve values for time slots, with time blocks as market goods, there may not exist an equilibrium solution that is globally optimal.*

Theorem 14 contradicts Theorem 7.2 of Wellman et al. [57]. However, we provide the following positive result under stronger conditions.

Theorem 15. *For auctions with fixed or flexible time blocks as market goods, if* $p_1 = \cdots = p_n$ *and* $f_i(p_i) = f_i(p_i+1) = \cdots = f_i(T)$ *for* $i = 1, \ldots, n$*, then each globally optimal schedule is supported by an equilibrium solution.*

The conditions of Theorem 15 specify that all jobs have the same processing time and each job has a scheduling cost that is independent of its completion time. However, the revenues of the jobs and the reserve values of the time slots are arbitrary.

We also conduct a computational study on randomly generated instances of the problem to investigate the effectiveness and efficiency of ascending auctions with the three market goods considered. Our data set consists of 4,860 problem instances, which are randomly generated following the guidelines established by Hall and Posner [24]. We vary the number of agents, the way in which job weights are generated, the ratio of the available capacity to the total requested capacity, the way in which reserve values of time slots are generated, and the bid increment function.

We compare system values generated by the auctions of the three market goods: time slots, fixed time blocks, and flexible time blocks. We use the system value of a schedule generated by a centralized scheduling method as a benchmark. On average, the auctions of time slots, fixed time blocks, and flexible time blocks generate 74.21%, 91.25%, and 94.33% of the benchmark value, respectively. Thus, flexible time block auctions generate 27.11% and 3.38% more system value than time slot and fixed time block auctions, respectively.

Different parameters affect the system value as follows. First, as the number of buying agents increases, time slot auctions perform worse, fixed time block auctions perform slightly better, and flexible time block auctions perform much better. Second, time slot auctions perform worst and flexible time block auctions perform best when the weight of a job is proportional to its processing time. Third, as the available scheduling resources increase, all three market goods generate system value that is closer to optimal. Finally, as the bid increment function increases, time slot auctions perform worse, whereas both time block auctions show stable performance. This last observation suggests that auctions with time blocks as market goods are insensitive to different bidding policies employed by the buying agents.

Next, we summarize the value generated from the three auctions for the buying agents and the facility owner. For time slot auctions, the total net profit of the buying agents is negative in 2,821 out of 4,860 instances, or 58.05%. This is because a buying agent may win some time slots but be unable to process its job. However, in time block auctions, an agent's profit is never negative. On average, flexible time block auctions generate 20.03% more profit for the buying agents than fixed time block auctions. However, the fixed and flexible time block auctions generate 87.21% and 91.39% of the revenue from time slot auctions for the facility owner, respectively. The last result suggests that the facility owner may prefer time slots as market goods.

However, because the buying agents may wish to avoid negative profits when bidding in time slot auctions, they may withdraw or bid less. Our results show that even if only 6% of agents withdraw, flexible time block auctions generate more revenue for the facility owner than time slot auctions, even without any reduction in bid prices. Alternatively, if the average reduction in bid prices at closure is over 25%, then the facility owner achieves more revenue in the flexible time block auction than the time slot auction, even if no agents withdraw.

We compare the number of rounds needed to reach closure in the three market good auctions. On average over all of the instances considered, the time slot, fixed, and flexible time block auctions require 167.8, 78.9, and 80.4 rounds, respectively. All the three market goods require more rounds to reach closure when the number of bidding agents is larger, more scheduling resources are available, or the bid increment function is smaller.

Our computational results enable us to compare the performance of auctions with the three market goods, as follows. Using time slots as market goods is not recommended, because in many instances this provides low system value and negative profit for the buying agents and is inefficient at reaching closure. Moreover, the facility owner's choice of market good may be influenced toward time blocks by likely reductions in the number of bidders and their bid prices when using time slots. Comparing the two time block auctions, they require a similar number of rounds to reach closure. However, auctions with flexible time blocks significantly outperform those with fixed time blocks with respect to the facility owner's value, the buying agents' value, and the system value. Therefore, we recommend flexible time blocks as the most effective market good for auctions of scheduling capacity.

5. Future Research

For capacity planning and scheduling with cooperative players, we propose the following interesting directions for further research.

1. Existing discussions of capacity planning assume that the cost of capacity is incurred solely by the suppliers, and that the retailers submit orders after the capacity level is determined. In such two-step capacity planning, suppliers have limited access to budget and future demand information and may build capacity at inappropriate levels. Cooperation, and even cost sharing, between the suppliers and the retailers at the capacity planning stage may overcome such difficulties. Cooperative game models of this decision process deserve detailed investigation. See Cachon and Lariviere [5].

2. Capacity sharing games need to be studied more thoroughly. Most of the existing literature assumes that the retailers are not in direct competition with each other, and allocated capacity is not transferrable between them. However, secondary markets and side payments are common in business, especially when available resources are scarce. Capacity sharing games, involving both multiple retailers and multiple suppliers, deserve careful exploration. See Hall and Liu [22], Aydinliyim and Vairaktarakis [1], and Vairaktarakis and Aydinliyim [55].

3. Cooperative schedule planning games have been investigated to only a limited extent. Such games arise when various scheduling resources are owned by different agents. An initial step toward the investigation of schedule planning games is to find Pareto optimal schedules with limited numbers of agents. See Schulz and Uhan [44].

4. The typical starting assumption of a predetermined sequence in sequencing games is restrictive, in that only limited cooperation between the players is available. It is predictable that the study of more general cases where agents have greater flexibility to cooperate is challenging. But it would be interesting to generalize sequencing games to consider more general scheduling environments and find implementable cooperative mechanisms. See Curiel et al. [13] and Maniquet [33].

5. A typical assumption in cooperative games is full information sharing. However, in supply chains, players may not release all their relevant information. It would be interesting to study cooperation with limited information sharing. Furthermore, it should be possible to design mechanisms and incentives to encourage the players to release full and true information. See Vickrey [56] and Cachon and Lariviere [5].

6. Several papers have studied the integration of capacity planning and scheduling in a supply chain, generally described as *supply chain scheduling*. However, these works do not study incentives, stability of cooperation, and strategic behavior. For example, it would be useful to design truth inducing mechanisms that benefit all the participants in supply chain scheduling. See Hall and Potts [25] and Chen and Hall [10].

7. Beyond studying specific games in a given capacity allocation or scheduling environment, it is important to consider the general structure supporting such games. General theoretical structures may play an important role in successfully solving a specific game. Recently, for example, linear programming games and duality have been applied to solve several important games in supply chain cooperation. See Schulz and Uhan [44] and Chen and Zhang [9].

8. A key issue within cooperative games is how to find a division of the total payoff that is fair to all the players. All the studies discussed in this tutorial use one of the following two approaches to find fair payoff divisions: (a) find a closed-form expression and prove that it is fair, or (b) provide a solution approach and prove that it finds a fair division if one exists. However, other important questions, such as testing whether a given instance has a fair payoff division, and testing whether a given payoff division is fair, have been studied to only a limited extent. See Deng and Papadimitriou [15] and Hall and Liu [22].

9. Most of the literature addresses fairness through considering payoff divisions in the core. However, many other payoff divisions are also fair, including the Shapley value, nucleolus, and others. Moreover, desirable properties of fair solutions, such as population monotonicity, should also be considered. Therefore, it is valuable to extend the current research to consider more general measures of fairness and to distinguish between fair solutions based on secondary criteria. See Deng and Papadimitriou [15] and Zhang [58].

10. One important problem related to fair payoff division is to find a weakly fair payoff division if a fair payoff division does not exist. This is motivated by the observation that many games are naturally unbalanced. In such cases, it is important to find payoff divisions that are fair under some restrictions, such as ϵ-core, least core, and others. See Schulz and Uhan [44].

For capacity planning and scheduling with noncooperative players, we propose the following interesting directions for further research.

1. Noncooperative capacity planning games involving multiple suppliers should be studied. Most current research focuses on the capacity planning problem faced by a single supplier with uncertain demand. If multiple suppliers compete for demand, then both the solution approaches and the managerial insights are expected to be different and more valuable. See Cachon and Lariviere [6].

2. Competition between the retailers should be studied in the context of capacity allocation games. Because retailers compete for capacity from a common supplier, it is natural for them also to compete with each other for customer demand. Interesting results are expected to be obtained from the analysis of such games. See Cachon and Lariviere [6].

3. Capacity allocation with discrimination, either in prices or in capacity allocation mechanisms, or in both, is an interesting topic for future study. The supplier can seek information from the retailers besides order quantities and past sales and then use the information obtained to facilitate capacity allocation decisions by the individual retailers. See Cachon and Lariviere [4] and Fang and Whinston [17].

4. Existing capacity allocation games involve orders that are either not time sensitive or are delivered simultaneously by the supplier. It is also interesting to consider a generalization where retailers place orders that compete for both production capacity and order delivery time. See Ganesh et al. [20].

5. A common assumption of capacity allocation games is that retailers do not hold inventory. Multiechelon capacity allocation games with inventory-holding retailers model additional practical issues and therefore deserve investigation. See Cachon and Lariviere [4] and Chen and Zhang [9].

6. Scheduling issues are largely omitted in the current literature of noncooperative games for capacity allocation. However, it benefits suppliers, especially those who offer multiple products, to consider scheduling costs while making capacity allocation decisions. Opportunities exist for modelling scheduling costs in various ways. See Hall and Liu [23].

7. Combinatorial auctions are naturally useful for the allocation of discrete capacity and other scheduling resources. In combinatorial auctions, typically the bid determination and winner determination problems are intractable, even without considering scheduling cost. Therefore, research is needed to study the optimization of the bid determination and winner determination problems. See Kutanoglu and Wu [31] and Hall and Liu [23].

8. It is interesting to investigate heuristic-based equilibrium for scheduling jobs owned by noncooperative players. The literature recognizes the intractability of noncooperative scheduling games, especially where each player seeks an optimal schedule for itself based on the strategies of others. If we focus on solutions where players use heuristic rules rather than optimal algorithms, then the problem is more tractable and additional interesting results may be obtainable. See Wellman et al. [57] and Hall and Liu [23].

9. It is important to conduct further studies of auctions of scheduling resources. This type of auction has many emerging applications. Despite the many applications and simulation studies in the literature, few practically effective auction mechanisms have been developed. Also, bidders' strategic behavior has been analyzed to only a limited extent. Results from combinatorial auctions are useful tools for the study of auctions of scheduling resources. See Kutanoglu and Wu [31] and Wellman et al. [57].

10. Due to the intractability of the bid determination and winner determination problems in auctions, it is difficult to predict the bidding strategies and policies of bidders. Therefore,

it is important to design robust auction mechanisms that are insensitive to players' bidding strategies and policies. See Wellman et al. [57] and Reeves et al. [42].

6. Conclusions

In this tutorial, we examine both cooperative and noncooperative games for capacity planning and scheduling problems. In both types of games, we review the recent literature, discuss our own current research, and provide guidelines for future research. The discussion of the literature of cooperative games includes models of sequencing games, outsourcing operations, and economic lot-sizing and schedule-planning games. Our research considers a combined capacity allocation and scheduling problem where a supplier allocates capacity to distributors, who may then cooperate to share their capacity allocations and resubmit their orders. The discussion of the literature of noncooperative games includes capacity allocation mechanisms based on order sizes and past sales, incentives for truth telling, option contracts, scheduling under decentralized versus centralized decision making, and auctions. Our research develops an auction mechanism to allocate capacity to various agents with jobs that can be scheduled profitably using that capacity.

We believe that even though game theory has been widely applied to the study of capacity planning and scheduling, there is considerable potential for future research. Although centralized decision making can potentially improve supply chain performance to a significant degree, such approaches are often difficult to implement when supply chain members have different incentives and look for strategies to benefit themselves. Therefore, game theory models provide a valuable approach to studying incentives, stability of cooperation, and mechanisms for decentralized decision making. We hope that our work will stimulate further research in these exciting directions.

References

[1] T. Aydinliyim and G. L. Vairaktarakis. Coordination of outsourcing operations. *Manufacturing & Service Operations Management*, Forthcoming, 2008.

[2] P. Brucker. *Scheduling Algorithms.* Springer, Berlin, 1998.

[3] Y. Bukchin and E. Hanany. Decentralization cost in scheduling: A game-theoretic approach. *Manufacturing & Service Operations Management* 9:263–275, 2007.

[4] G. P. Cachon and M. A. Lariviere. Capacity allocation using past sales: When to turn-and-earn. *Management Science* 45:685–703, 1999.

[5] G. P. Cachon and M. A. Lariviere. Capacity choice and allocation: Strategic behavior and supply chain performance. *Management Science* 45:1091–1108, 1999.

[6] G. P. Cachon and M. A. Lariviere. An equilibrium analysis of linear, proportional and uniform allocation of scarce capacity. *IIE Transactions* 31:835–849, 1999.

[7] G. P. Cachon and S. Netessine. Game theoretic applications in supply chain analysis. D. Simchi-Levi, S. D. Wu, Z.-J. Shen, eds. *Supply Chain Analysis in the eBusiness Era.* Kluwer, Amsterdam, 2004.

[8] X. Cai and G. L. Vairaktarakis. Cooperative strategies for manufacturing planning with negotiable third-party capacity. Technical Report TM-820, Weatherhead School of Management, Case Western Reserve University, Cleveland.

[9] X. Chen and J. Zhang. Duality approaches to economic lot-sizing games. Working paper, Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana-Champaign, Urbana, 2007.

[10] Z.-L. Chen and N. G. Hall. Supply chain scheduling: Conflict and cooperation in assembly systems. *Operations Research* 55:1072–1089, 2007.

[11] S. H. Clearwater, ed. *Market-Based Control: A Paradigm for Distributed Resource Allocation.* World Scientific, Hackensack, NJ, 1996.

[12] I. Curiel, H. Hamers, and F. Klijn. Sequencing games: A survey. P. Borm, H. Peters, eds. *Chapters in Game Theory: In Honor of Stef Tijs.* Kluwer Academic Publishers, Boston, 27–50, 2002.

[13] I. Curiel, G. Pederzoli, and S. Tijs. Sequencing games. *European Journal of Operational Research* 40:344–351, 1989.

[14] I. Curiel, J. Potters, R. Prasad, S. Tijs, and B. Veltman. Sequencing and cooperation. *Operations Research* 42:566–568, 1994.

[15] X. Deng and C. H. Papadimitriou. On the complexity of cooperative solution concepts. *Mathematics of Operations Research* 19:257–266, 1994.

[16] E. J. Durango-Cohen and C. A. Yano. Supplier commitment and production decisions under a forecast-commitment contract. *Management Science* 52:54–67, 2006.

[17] F. Fang and A. Whinston. Option contracts and capacity management: Enabling price discrimination under demand uncertainty. *Production and Operations Management* 16:125–137, 2007.

[18] Q. Fang, S. Zhu, M. Cai, and X. Deng. On computational complexity of membership test in flow games and linear production games. *International Journal of Game Theory* 31:39–45, 2002.

[19] M. L. Fisher. What is the right supply chain for your product? *Harvard Business Review* 75(2):105–116, 1997.

[20] A. Ganesh, K. Laevens, and R. Steinberg. Congestion pricing and noncooperative games in communication networks. *Operations Research* 55:430–438, 2007.

[21] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, San Francisco, 1979.

[22] N. G. Hall and Z. Liu. Capacity allocation and scheduling in supply chains. Working paper, Fisher College of Business, The Ohio State University, Columbus, 2007.

[23] N. G. Hall and Z. Liu. Auctions for competitive capacity allocation and scheduling. Working paper, Fisher College of Business, The Ohio State University, Columbus, 2008.

[24] N. G. Hall and M. E. Posner. Generating experimental data for computational testing with machine scheduling applications. *Operations Research* 49:854–865, 2001.

[25] N. G. Hall and C. N. Potts. Supply chain scheduling: Batching and delivery. *Operations Research* 51:566–584, 2003.

[26] A. V. Iyer, V. Deshpande, and Z. P. Wu. A postponement model for demand management. *Management Science* 49:983–1002, 2003.

[27] J. R. Jackson. Scheduling a production line to minimize maximum tardiness. Research Report 43, Management Science Research Project, University of California, Los Angeles, 1955.

[28] S. M. Johnson. Optimal two- and three-stage production schedule with set up time included. *Naval Research Logistics Quarterly* 1:61–68, 1954.

[29] E. Kalai and E. Zemel. Generalized network problems yielding totally balanced games. *Operations Research* 30:998–1008, 1982.

[30] S. Karabuk and S. D. Wu. Incentive schemes for semiconductor capacity allocation: A game theoretic analysis. *Production and Operations Management* 14:175–188, 2005.

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

[32] S. Mallik and P. T. Harker. Coordinating supply chains with competition: Capacity allocation in semiconductor manufacturing. *European Journal of Operational Research* 159:330–347, 2004.

[33] F. Maniquet. A characterization of the Shapley value in queueing problems. *Journal of Economic Theory* 109:90–103, 2003.

[34] R. P. McAfee and J. McMillan. Auctions and bidding. *Journal of Economic Literature* 25:699–738, 1987.

[35] E. J. Muth and G. L. Thompson, eds. *Industrial Scheduling*. Prentice Hall, Englewood Cliffs, NJ, 1963.

[36] M. Nagarajan and G. Sošić. Game-theoretic analysis of cooperation among supply chain agents: Review and extensions. *European Journal of Operational Research* 187:719–745, 2008.

[37] J. Nash. Non-cooperative games. *The Annals of Mathematics* 54:286–295, 1951.

[38] D. Niyato and E. Hossain. A noncooperative game-theoretic framework for radio resource management in 4G heterogeneous wireless access networks. *IEEE Transactions on Mobile Computing* 7:332–345, 2008.

[39] G. Owen. On the core of linear production games. *Mathematical Programming* 9:358–370, 1975.

[40] B. Peleg and P. Sudhölter. *Introduction to the Theory of Cooperative Games*. Kluwer, Boston, 2003.

[41] M. Pinedo. *Scheduling, Theory, Algorithms and Systems*, 2nd ed. Prentice Hall, Englewood Cliffs, NJ, 2002.

[42] D. M. Reeves, M. P. Wellman, J. K. MacKie-Mason, and A. Osepayshvili. Exploring bidding strategies for market-based scheduling. *Decision Support Systems* 39:67–85, 2005.

[43] D. Samet and E. Zemel. On the core and dual set of linear programming games. *Mathematics of Operations Research* 9:309–316, 1984.

[44] A. S. Schulz and N. A. Uhan. Encouraging cooperation in sharing supermodular costs. Working paper, Sloan School of Management, MIT, Cambridge, MA, 2007.

[45] P. Schuurman and G. Woeginger. Approximation schemes—A tutorial. R. H. Möhring, C. N. Potts, A. S. Schulz, G. J. Woeginger, L. A. Wolsey, eds. *Lectures on Scheduling*, 2008.

[46] L. S. Shapley. A value for n-person games. *Annals of Mathematics Study* 28:307–317, 1953.

[47] L. S. Shapley. Cores of convex games. *International Journal of Game Theory* 1:11–26, 1971.

[48] L. S. Shapley and M. Shubik. The assignment game, 1: The core. *International Journal of Game Theory* 1:111–130, 1972.

[49] M. J. Shaw. A distributed scheduling method for computer integrated manufacturing: The use of local area networks in cellular-systems. *International Journal of Production Research* 25:1285–1303, 1987.

[50] M. J. Shaw. Dynamic scheduling in cellular manufacturing systems: A framework for networked decision-making. *Journal of Manufacturing Systems* 7:83–94, 1988.

[51] W. Shen, Q. Hao, H.-J. Yoon, D. H. Norrie. Applications of agent-based systems in intelligent manufacturing: An updated review. *Advanced Engineering Informatics* 20:415–431, 2006.

[52] M. Slikker. Relaxed sequencing games have a nonempty core. *Naval Research Logistics* 53:235–242, 2006.

[53] W. E. Smith. Various optimizers for single-stage production. *Naval Research Logistics Quarterly* 3:59–66, 1956.

[54] S. H. Tijs. Bounds for the core and the τ-value. O. Moeschlin, D. Pallaschke, eds. *Game Theory and Mathematical Economics*. North-Holland, Amsterdam, 122–132, 1981.

[55] G. L. Vairaktarakis and T. Aydinliyim. Centralization vs. competition in subcontracting operations. Working paper, Department of Operations, Case Western Reserve University, Cleveland, 2008.

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

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

[58] J. Zhang. Joint replenishment game and maximizing an H-Schur concave function over a polymatroid. Working paper, Stern School of Business, New York University, New York, 2008.