Supply Chain Coordination by Means of Automated Negotiations

Andreas Fink
University of Hamburg, Germany
fink@econ.uni-hamburg.de

Abstract

The coordination within supply chains depends on appropriate forms of distributed decision making. Considering joint decisions as formal contracts, the coordination problem may be regarded as a search process in a corresponding contract space. Automated negotiations, with firms or decision making units represented as software agents, can provide an effective mechanism to determine mutually beneficial contracts. The generic negotiation approach examined in this paper is based on a formal specification of contracts that represent bilateral collaborations between two firms (agents) which aim for the coordination of their production sequences. Taking into account asymmetric information and opportunistic behavior, a mediator supports the negotiation process. This mediator repeatedly generates new candidate contracts, which are accepted or rejected by the agents according to particular strategies. We define an explicit mechanism for implementing a cooperative acceptance criterion, whereby agents conditionally agree on utility deteriorations according to a probabilistic criterion similar to that of simulated annealing. The proposed design enables the definition of negotiation rules to be verified by the mediator, forcing both agents to behave in a cooperative manner. The negotiation approach is validated for different supply chain sequencing scenarios. In spite of the simplicity and generality of the negotiation mechanism, the experimental results are very promising. Thus automated negotiations may constitute an effective means for coordinating decisions within supply chains.

1. Introduction

The increasing use of information technology to support collaborative planning within supply chains affects related coordination problems and the means for distributed decision making. In this context, Malone et al. conjecture the growing use of market coordination mechanisms relative to hierarchies ("electronic market hypothesis" [23] [24]). For example, the application of multilateral auctions may be reasonable under specific assumptions. However, long-term relationships between firms collaborating in a supply chain may necessitate other forms of coordination. With markets and hierarchies as extremes of the coordination spectrum, Clemons et al. conjecture the growing use of hybrid forms of coordination, with a firm typically having only a small number of business partners ("move to the middle hypothesis" [4]). But even when assuming long-term relationships, selfish actions of independent decision making units must be taken into account.[1] That is, in general one cannot take benevolent collaborative planning for granted but must design sensible mechanisms that contribute to achieving cooperative behavior. To date, research has focused mainly on devising incentive policies (such as price discounts or penalty costs, which are set up in basic agreements) that direct the behavior of decision making units in such a way that the performance of some hypothetical centralized planning is approached.[2] However, in relation to the amount of supply chain management research on strategic questions and stochastic models, there is a lack of investigations that consider deterministic models of operations management [11]. In the context of this observation, we deal with the short-term coordination of production schedules between two firms within the structure of a two-stage supply chain, focusing on the problem of matching the production schedules (sequences). This key problem is to be solved repeatedly subject to actual data such as external demand and capacity constraints.

[1] In spite of this, a large proportion of the research literature on solving decision problems within supply chain management is set up under the paradigm of centralized planning in a vertically integrated firm (in the sense of one decision making unit). Surveys on the integrated analysis and planning in supply chains are provided in, e.g., [8] [35] [36] [39].

[2] For example, there are various articles that consider joint decision making and corresponding contracts in two-stage supply chains facing static and/or stochastic demand [19] [26]. General surveys on supply chain contracting can be found in, e.g., [3] [40].

Taking into account asymmetric information and opportunistic behavior of both firms[3], automated negotiations may provide an effective means for achieving collaborative planning. We assume that both firms have already gained agreement about the set of potential collaborations. Achieving such an understanding as well as reaching an agreement about the negotiation procedure may itself be regarded as a (meta-)negotiation, which is not subject of this paper. Yet such requirements may make the use of complex automated negotiation approaches impractical for ad-hoc collaborations. However, long-term relationships promote the design and application of specific coordination schemes if those provide a fair as well as effective mechanism for collaborative planning. Thus appropriate negotiation procedures may foster both maintaining and building long-term collaborations between self-interested firms.

The representation of collaborations by formal contracts results in a corresponding set of potential contracts. This contract space generally constitutes a complex search space, which may even be intractable when aiming for a single objective optimum in the context of centralized planning. If one of the firms dominates the collaboration and has the power to enforce its decisions, the problem may decompose in two successive decision problems. However, typically one can envisage win-win opportunities due to collaborative planning.[4] Thus the problem is about determining some mutually beneficial contract. Automated negotiations may provide an effective coordination mechanism for such problem settings. Designing adequate negotiation procedures depends on the assumptions and the criteria for the quality of resulting contracts. We assume that the involved firms opportunistically pursue maximization of individual utility. That is, the quality of a contract is assessed by each firm on the basis of the deviation from a hypothetical individually optimal contract, respectively. The utility functions represent private information for each party. To ensure the acceptance of the rules of the negotiation procedure by both firms involved, some kind of fairness should be provided. In particular, no firm should be able to profit from untruthful utterances. Furthermore, the joint gain should be shared out among the firms with regard to some sensible normative judgment on fairness. Taking into account the questionable interpersonal comparison of utility, the assessment of the overall quality of the outcome of negotiations is a problem in itself.

As information and communication technology enables the automation of formal negotiation procedures, automated negotiations have recently received considerable attention in information systems research. In particular, automation enables devising and implementing new negotiation schemes which would cause prohibitive costs when executed manually. The bilateral negotiation approach that is examined in this paper is based on novel ideas of Klein et al. [20].

In the following section, we discuss specific supply chain sequencing scenarios, which are later used for illustration and computational experiments. In Section 3, we describe a general mechanism for negotiating about complex contracts, which may be applied to any kind of collaborative planning on the basis of formal contract spaces. Experimental results are presented in Section 4, where we aim for prescriptive insights with regard to achieving high quality negotiation outcomes. In Section 5, we argue about actually obtaining such results, taking into account strategic behavior of the decision making units from a game theory perspective. Finally, we summarize our key findings and pose open research questions.

2. Problem description

Supply chain management is about the coordination of logistics processes between different facilities within a firm as well as encompassing more than one firm. In recent years, supply chain management has received considerable attention both in practice and in research. In this paper, we consider the short-term coordination of production decisions in two-stage one-to-one supply chains. Obviously, most real-world scenarios involve more intricate characteristics such as multilateral relationships and contract spaces of arbitrary complexity. However, in this paper we aim to study the principal potential of the proposed negotiation procedure for a straightforward application as a primary step. The generalization of resulting insights to more realistic scenarios is discussed in the final section.

We assume that firm 1 manufactures parts that constitute input to the production processes of firm 2 in the context of a tight coupling of both production processes (in the sense of just-in-time (JIT)). We focus on the problem of determining a common production sequence, while disregarding features such as batch deliveries as well as the actual production system within a

[3] Similar problems may also occur within firms, since different organizational sub-units (such as the marketing division vs. the manufacturing division) and corresponding decision making units pursue different performance criteria and information asymmetry must also be taken into account following agency theory [14] [25]. Thus the application of negotiation-based mechanisms even for the coordination of intra-firm supply chains may be appropriate. This conforms to the nexus-of-contract perspective of organization according to the taxonomy of coordination in operations of Whang [41].

[4] Consequently, the crucial question is about “making the pie bigger” instead of “fighting to get the largest piece of a pie of fixed size”. The latter problem is usually termed as zero-sum single-issue bargaining. Going back to fundamental concepts by Nash [29] [30] and Rubinstein [33], bargaining constitutes a main part of the game theory literature [27] [28] [32].

firm which may allow for re-sequencing. The essential question is about the common job sequence at the coupling point (at delivery); see Figure 1. This operations management problem must be solved regularly subject to actual data such as external demand and capacity constraints.

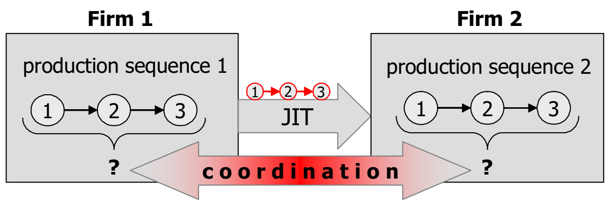

Figure 1. The coordination of production sequences in a two-stage supply chain

Given a set of jobs $J = \{1, \dots, n\}$, a production (job) sequence is defined by the permutation $\Pi = [\pi(1), \dots, \pi(n)>$, with $\pi(i)$ denoting the job that takes the i-th position in the schedule. Each firm evaluates production sequences by some scalar objective function. The firms pursue different goals, which are usually conflicting. Considering two different scenarios A and B, we assume that firm 1 aspires a cost efficient production schedule, while firm 2 is mainly interested in the eventual completion time of the jobs (customer focus).

In scenario A, firm 1 aims for minimizing sequence-dependent setup costs, assuming that each of the jobs has to be processed without interruption on a single machine. Given setup cost data τ_{ij}, $1 \leq i, j \leq$ n, $i \neq j$, the objective function is defined as

$$f(\Pi) = \sum_{i=1}^{n-1} \tau_{\pi(i),\pi(i+1)} .$$

Concerning firm 2, we consider a continuous flow-shop characteristic of the production system with the objective of minimizing the average completion time of the jobs [10]. In a flow-shop setting each job has to be processed in an identical order on a given number of m machines. Each machine can process only one job at a time. The parameters t_{ij}, $1 \leq i \leq n$, $1 \leq j \leq m$, denote the processing time of job i on machine j. For continuous flow-shop scheduling problems the processing of each job must be continuous, which means that there must not be any waiting times between the processing of any consecutive tasks of this job. To allow processing of a job without interruption on all machines, the order in which the jobs are processed on a machine is the same for all machines (assuming non-zero processing times), so that a permutation of the n jobs defines the production schedule. Continuous processing of a job generally determines an inevitable delay d_{ik}, $1 \leq i \leq n$, $1 \leq k \leq n$, $i \neq k$, on the first machine between the start of job i and the start of job k when job k is processed directly after job i. This delay may be computed as

$$d_{ik} = \max_{1 \leq j \leq m} \left\{ \sum_{h=1}^{j} t_{ih} - \sum_{h=2}^{j} t_{k,h-1} \right\} .$$

The objective function is defined as

$$f(\Pi) = \frac{1}{n} \left(\sum_{i=2}^{n} (n+1-i) d_{\pi(i-1),\pi(i)} + \sum_{i=1}^{n} \sum_{j=1}^{m} t_{ij} \right) .$$

One should note that, first, the latter term of this objective function is constant and thus may be neglected, and, second, the scaling factor of $1/n$ may be dropped. Ideally such positive affine transformations, which are somehow arbitrary, should not affect the outcome of the coordination of the production sequences.

In scenario B, we still consider that firm 1 aims for minimizing sequence-dependent setup costs. On the other hand, we suppose that the goal of firm 2 is minimization of weighted job tardiness [5] [6]. Assuming that each job i has to be processed without interruption on a single machine, a problem instance is characterized by processing times p_i, positive weights w_i, and due dates r_i. In dependence on the job sequence Π, the resulting completion times C_i and the resulting tardiness $T_i = \max \{ C_i - r_i, 0 \}$ of each job i can be calculated. The objective function is defined as

$$f(\Pi) = \sum_{i=1}^{n} w_i T_i .$$

When restricting the consideration to one firm which aims for determining an ideal contract for itself, all of the introduced problems are *NP*-hard. Therefore, the problem of collaborative planning between two firms (i.e., determining a mutually beneficial contract with regard to the production sequence), taking into account asymmetric information and opportunistic behavior, is complex. In the next section, we discuss a general approach that provides the necessitated collaborative planning functionality.

3. Negotiating complex contracts

Automated negotiations have recently received considerable attention in research literature [2] [13] [16] [17] [22] [37]. While a literature survey is beyond the scope of this paper, our approach may be broadly classified as follows. We aim for prescriptive insights into the full automation[5] of bilateral (one-to-one) negotiations about complex contracts with multiple interdependent

[5] This is in contrast to partial automation in the sense of "negotiation support tools" or "e-negotiation media", where the course of the negotiation mainly depends on personal decisions that are based on supporting information systems.

issues. These issues should be assessed concurrently, exploiting win-win opportunities in the sense of integrative negotiations.

The proposed negotiation procedure was inspired by Klein et al. [20]. It is generic with regard to the kind of contract under consideration, but we assume that the set of possible contracts can be formally specified. The supply chain sequencing scenarios introduced in Section 2 serve for exemplification and experimental evaluation (see Section 4). The contract space of the set of permutations of n objects leads to multiple interdependent issues in the form of object positions or corresponding predecessor-successor relationships. Obviously, such issues must be considered simultaneously.

The negotiation procedure may be conceptualized and implemented as a system of interacting autonomous software agents [13] [21] [31]. These agents adhere to a formal negotiation protocol on the basis of an initial conditioning by their principals.

3.1 Negotiation protocol

We assume that there are two agents a_1 and a_2, which want to reach agreement about a complex contract that determines the collaboration in question. The set of possible contracts is defined by a formal contract space $C = \{c \mid c \text{ is a possible contract}\}$, over which the agents have already gained agreement. The contract space is usually defined in an implicit way and, due to its size, does not allow for an exhaustive search. The agents assess contracts according to their preferences by utility functions $f_1: C \to \mathrm{R}$ and $f_2: C \to \mathrm{R}$, respectively, which represent private information. Each agent aims for maximizing its utility.

Taking into account asymmetric information and opportunistic behavior of the agents, a mediator m supports the negotiation process. The actions of the mediator are transparent. Besides the contract space, the mediator has neither any specific knowledge about utility functions nor any special trust relationships with the agents. The mediator repeatedly generates new candidate contracts which are accepted or rejected by the agents according to specific strategies. For this the agents do not need an explicit representation of the contract space and the utility functions, but are only required to be capable of assessing potential contracts compared with each other.

The mediator accepts a candidate contract (i.e., this candidate contract c' becomes the contract c that the agents agree on after the particular round) if and only if both agents have signaled acceptance. After some termination criterion is met (e.g., with respect to a maximum number of rounds or a maximum computation time) the contract eventually agreed on constitutes the negotiation outcome. The generic negotiation protocol is shown in the sequence diagram in Figure 2. This negotiation protocol is rather simple, yet its effectiveness depends on two essential variation points, namely the generation of candidate contracts by the mediator and the acceptance criteria of the agents.

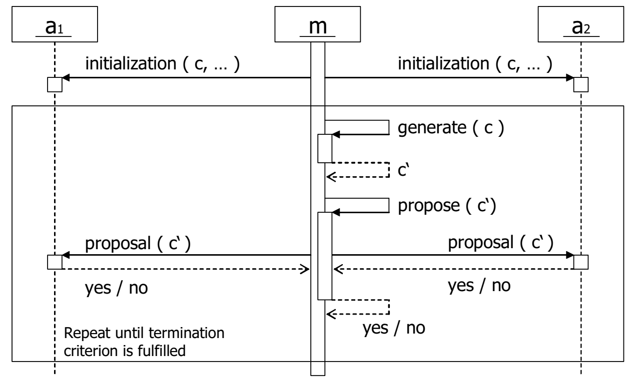

Figure 2. Negotiation protocol

The crucial difference between the described negotiation procedure and multi-criteria optimization is due to the asymmetric information concerning the utility functions. In particular, the course of the negotiation, which may be regarded as a search process, rests on plain yes/no responses by the agents having private information about the specific criteria. This precludes the application of typical multi-criteria combinatorial optimization procedures [7] [12]. In Section 4, a major concern is whether and to what extent this loss of information (with independent decision making units which opportunistically pursue their own goals) deteriorates the efficiency and social welfare of the negotiation outcome.

3.2 Generation of candidate contracts

The generation of candidate contracts is due to the mediator, which has no information about the utility functions. A simple option is some random generation of candidate contracts (e.g., selecting members of the set of the permutations of n objects according to a uniform probability distribution). Presuming that the negotiation process successively produces contracts of higher quality, a better option may be to introduce small changes to the current contract (e.g., a shift of some object to another position in the sequence). Such a progressive modification (perturbation) corresponds to the iterative selection of random neighbor solutions (in the terms of local search) or random mutations (in the terms of evolutionary algorithms).

There are different options for generating candidate contracts in a more "intelligent" way. On the one hand, one may facilitate some kind of learning by the mediator. Namely, the mediator may observe the negotiation process trying to infer certain regularities (e.g., of the kind that both agents have mostly rejected candidate contracts where some object is shifted to the back part of the sequence), which may allow for focusing the generation procedure accordingly. The recombination of fragments of traversed contracts (like for evolutionary algorithms) is also conceivable. On the other hand, one may allow that the agents (perhaps alternately) propose candidate contracts, which requires an adaptation of the negotiation protocol.

3.3 Acceptance criteria

The "natural" behavior of a selfish agent is to myopically accept a new candidate contract *c'* if and only if it is not worse than the current contract *c*. This *greedy acceptance criterion* A_i^g of an agent a_i is defined as follows:

$$A_i^g(c,c') := \begin{cases} \text{Yes} & \text{if } f_i(c') \geq f_i(c) \\ \text{No} & \text{otherwise} \end{cases}$$

However, the mutual application of the greedy acceptance criterion by both agents ("hill climbers" [20]) may result in a negotiation process that gets stuck rather soon. This risk is well known in the context of local search methods, which may get stuck in local optima of inferior quality. One approach to overcome such local optima is to partly accept utility deteriorations. In particular, the classical ("Boltzmann") acceptance criterion of simulated annealing [18] may be applied as the acceptance criterion of an agent ("annealer" [20]), which leads to the following *cooperative acceptance criterion*:

$$P(A_i^k(c,c') = \text{Yes}) := \begin{cases} 1 & \text{if } f_i(c') \geq f_i(c) \\ e^{(f_i(c') - f_i(c))/T} & \text{otherwise} \end{cases}$$

The probability *P* of acceptance of a candidate contract *c'* decreases with a larger (negative) utility deterioration. This probability is computed using a positive control parameter *T* (temperature), which is gradually reduced according to a cooling schedule so that the probability of accepting deteriorating moves decreases in the course of the process.

The crucial problem with regard to the application of the cooperative acceptance criterion is about sensibly determining and adapting the temperature parameter.[6] An adequate mechanism should be generally effective concerning a high quality negotiation outcome. Some kind of fairness must be provided. Ideally, no agent should be able to profit by untruthful utterances or corresponding distortions of its utility function. For example, the negotiation process should be invariant with regard to positive affine transformations of the utility functions. To facilitate a fair mechanism, we introduce acceptance probabilities that must be fulfilled by both agents. In particular, two parameters p_1 and p_2, $p_1 > p_2$, define the ratio of the number of candidate contracts that shall be accepted by an agent for the beginning phase and for the final phase of the negotiation process, respectively. (Such phases may take, e.g., 10% of the whole negotiation process.) Obliging both agents to show the same willingness to accept deteriorating candidate contracts means that we enforce particular degrees of cooperation. This may be regarded as a fair negotiation rule.

For each agent, appropriate temperature values T_1 and T_2 depend on the actual contract space and the utility function. Conditionally determining these values may be accomplished by the following mechanism: The agents simulate the negotiation process separately from each other for a number of rounds ("trial run"); in each round the acceptance decision of the other agent is randomly chosen according to the predefined acceptance probability. For the trial run each agent records the utility deteriorations that occur, which results in a list of negative values [Δ_1, Δ_2, …, Δ_k]. This allows the conditional computation of the expected value of the number of accepted deteriorating candidate contracts, which leads to the following equivalence with regard to determining suitable temperature parameters:

$$\sum_{i=1}^{k} e^{\Delta_i / T_j} = k \times p_j \quad \text{for } j = 1, 2$$

These equivalence relations have unambiguous solutions, which can be calculated by an iterative approximation procedure.

A common technique for adapting the temperature parameter throughout the search process is the application of a geometric cooling schedule. That is, after each round the temperature is multiplied by a parameter α, $0 < \alpha < 1$: $T := \alpha \times T$. This parameter can be calculated in dependence on the temperature parameters and the number *r* of rounds of the negotiation process:

$$\alpha := \sqrt[r-1]{T_2 / T_1}$$

4. Computational experience

The assessment of negotiation processes depends on the pursued criteria [34]. A basic desirable property is Pareto optimality (efficiency) of the final contract. That is, there should be no other contract that improves the utility of one agent without making the other agent worse off. But there are usually a large number of Pareto optimal contracts. Under the questionable assumption that

[6] This point is not elaborated in [20]. Even the majority of the applied research literature on simulated annealing seems to neglect this problem, restricting the analysis to the application of some simple rules of thumb.

an interpersonal comparison of utility values is possible (for instance by means of a monetary equivalence that relates the utility functions to each other), maximization of social welfare may constitute a meaningful criterion for the selection among Pareto optimal contracts. However, maximization of social welfare may be at odds with some perception of fairness. For instance, a socially optimal contract might unfairly allocate the joint gain due to the collaboration completely to one agent. From a one-sided perspective, contracts may be assessed on the basis of the deviation from a hypothetical optimal contract for one agent.[7] Further criteria are incentive-compatibility (i.e., an agent should not be able to profit from untruthful utterances provided that this conforms to the negotiation protocol), computational tractability, and low requirements for information revelation (in particular with respect to the privacy of utility functions).

Each of the following negotiation experiments is made up of 30 collaboration instances for scenario A and 25 collaboration instances for scenario B. In each case, a collaboration instance is constituted by a specific pair of problem instances for firm 1 and firm 2. Each negotiation process for some collaboration instance is repeated ten times with different seed values for the pseudo-random value generator, which affects the initial solutions (random permutations) and the specific responses from the probabilistic acceptance criteria. Aggregate results for some negotiation experiment are given as average values over 300 or 250 specific results, respectively. We are particularly interested in average percentage deviations from reference results for each firm (in the sense of lower bounds), which are determined assuming that a firm can decide about the contract in isolation. Due to the *NP*-hardness of the considered problems the computation of optimal solutions is considered intractable. We approximate these reference results as the best outcome from applying different state-of-the-art metaheuristics (iterated local search, simulated annealing, and reactive tabu search, each for a maximum computation time of 10 seconds) to the single objective problem. We apply the shift of some object to another position in the sequence as the procedure to generate candidate contracts. Computation times are given on the basis of using a Pentium 4 processor with 1.8 GHz.

4.1 Results for scenario A

We use problem instances with 100 jobs from the benchmark scheduling data sets of Taillard [38]. For firm 2, the problem instances ta061–ta090 are treated as continuous flow-shop scheduling problems. For firm 1, the delay values d_{ik} are taken as setup costs. We use 30 collaboration instances by matching different problem instances for firm 1 and firm 2. In particular, we use the pairs ($i+j$, $i+((j+2) \bmod 10)$) for $i \in \{61, 71, 81\}$ and $j=0,\ldots,9$; the tuple values denote the instance numbers from the benchmark data set.

Table 1 provides the results for the case that one firm dominates the negotiation process by forcing its decisions upon the other firm. To allow for this, we adapt the acceptance criterion of the mediator by equating it with the outcome of the acceptance criterion of the dominating firm (i.e., the dominated firm has no influence at all in the negotiation process). The results clearly reflect the power of the firm that can enforce its decisions. However, the enforcing agent profits from applying the cooperative acceptance criterion, which allows overcoming local optima. While resulting outcomes may be Pareto optimal, the social welfare may be far from optimum as the interest of one firm is completely neglected and thus opportunities for win-win situations are mostly thrown away.

Table 1. Average results for the case that one firm can enforce its decisions (scenario A)

Average deviations		100,000 rounds		1,000,000 rounds	
p_1=0.2; p_2=0.004		f_1	f_2	f_1	f_2
Firm / agent 1 holds power	greedy	5.0%	68.0%	4.9%	68.1%
	cooperative	3.4%	67.7%	2.2%	67.8%
Firm / agent 2 holds power	greedy	64.3%	5.0%	64.2%	5.0%
	cooperative	64.0%	3.0%	64.0%	2.0%

Table 2 and Table 3 show the results of combining agents that employ different acceptance criteria. The parameter values of p_1=0.2 und p_2=0.04 have been set on the basis of preliminary experiments for a similar application [9]. Agents that follow a cooperative strategy perform trial runs with 10% of the rounds of the actual negotiation process. Computation times are about one second for a negotiation process of 100,000 rounds.

Table 2. Average negotiation results for scenario A (100,000 rounds)

Average deviations		**Firm / agent 2 (flowtime)**			
100,000 rounds		greedy		cooperative	
p_1=0.2; p_2=0.04		f_1	f_2	f_1	f_2
Firm / agent 1 (setup costs)	greedy	23.0%	23.9%	11.3%	36.4%
	cooperative	38.9%	10.0%	19.7%	15.1%

Table 3. Average negotiation results for scenario A (1,000,000 rounds)

Average deviations		**Firm / agent 2 (flowtime)**			
1,000,000 rounds		greedy		cooperative	
p_1=0.2; p_2=0.04		f_1	f_2	f_1	f_2
Firm / agent 1 (setup costs)	greedy	23.0%	23.9%	7.4%	42.9%
	cooperative	48.4%	6.2%	18.5%	11.6%

[7] However, assessing such deviations is problematic due to similar reasons as in defining social welfare. In particular, adding a constant value to some utility function can arbitrarily distort the interpersonal comparison of percentage deviations.

With respect to the depicted deviations, all combinations but pairing two greedy agents are Pareto optimal. A greedy agent that is combined with a cooperative agent obtains high quality results at the cost of the cooperative agent. Combining two cooperative agents leads to an outcome that dominates the outcome obtained by two greedy agents. When using the average deviation for both agents as a measure of social welfare, the combination of two cooperative agents leads to the best overall results. With regard to comparing these results with a hypothetical integrated optimization model, we applied state-of-the-art metaheuristics (see above) to an aggregate objective function which represents the corresponding average percentage deviation. The outcome of the negotiation process deviates by only approximately 1.5% from the results of an integrated optimization model (which is not discussed further in this paper). Taking the restricted information revelation into account, this outcome is rather surprising.

Comparing the results of Table 2 and Table 3 (100,000 vs. 1,000,000 rounds) confirms the expectation that the negotiation process usually gets stuck rather soon when combining two greedy agents. With a longer negotiation process, a greedy agent that is combined with a cooperative agent obtains even better results at the cost of the cooperative agent, since the cooperative agent operates longer with large acceptance probabilities which is exploited by the greedy agent. Figure 3 depicts the typical course of the negotiation process (concerning the average deviation) for different strategy combinations. Clearly most of the benefit is realized and assigned in the first part of the negotiation process.

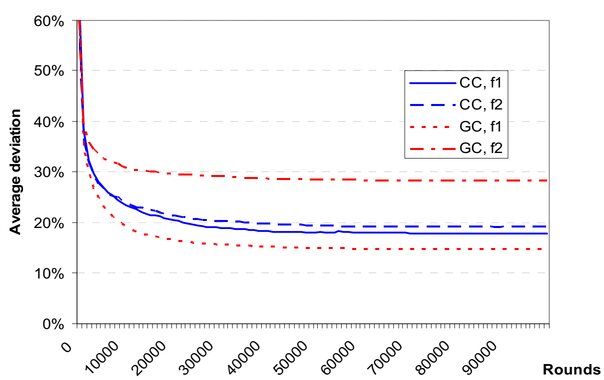

Figure 3. Typical course of the negotiation process for different strategy combinations (CC: both agents apply the cooperative acceptance criterion; GC: agent 1 applies the greedy acceptance criterion, agent 2 applies the cooperative acceptance criterion)

4.2 Results for scenario B

For firm 1, we use the traveling salesman problem generator described by Johnson and McGeoch [15].[8] We interpret Euclidian distance values of geometric instances in two dimensions as setup costs. In particular, we generated five different problem instances with 100 objects each with clustering (ten clusters), which features some regularity (roughly in the sense of ten different product families). For firm 2, we used weighted tardiness problem instances with 100 jobs from [6]. We selected five instances (number 86 up to 90) which are characterized by a fairly balanced distribution of the due dates.

We follow the experimental setup from the previous section. On first glance, the characteristics of the problem instances for scenario B lead to quite different results in comparison to the deviations for scenario A. Table 4 depicts the results for the case that one firm dominates the negotiation process. The outcome clearly reflects the power of the firm that can enforce its decisions, while the dominated firm must bear huge deviations. This effect is mainly due to the “gradiental topology” of the contract space regarding the utility functions. Namely, for firm 1 the geometric characteristics of the instances result in large variations of the costs of predecessor-successor pairs in the sequence. As a result some arbitrary sequence may be a lot worse than a high quality sequence. In effect we have a collaboration scenario where both firms must give in more in comparison to scenario A.

Table 4. Average results for the case that one firm can enforce its decisions (scenario B)

Average deviations p_1=0.2; p_2=0.004		100,000 rounds f_1	f_2	1,000,000 rounds f_1	f_2
Firm / agent 1 holds power	greedy	35.2%	777.5%	34.6%	776.6%
	cooperative	9.9%	763.6%	6.3%	774.6%
Firm / agent 2 holds power	greedy	711.0%	2.4%	710.7%	2.3%
	cooperative	714.2%	1.9%	708.4%	2.3%

The results that are depicted in Table 5 and Table 6 also show larger deviations than for scenario A. However, the basic structure of the outcome is the same. That is, overall best results are obtained when combining two cooperative agents. The average outcome dominates the combination of two greedy agents. When combining two agents with different strategies the greedy agent profits at the cost of the cooperative agent.

[8] This generator is available for download at the website of the “8th DIMACS Implementation Challenge: The Traveling Salesman Problem” (http://www.research.att.com/~dsj/chtsp/).

Table 5. Average negotiation results for scenario B (100,000 rounds)

Average deviations 100,000 rounds p_1=0.2; p_2=0.04		Firm / agent 2 (tardiness) greedy f_1	greedy f_2	cooperative f_1	cooperative f_2
Firm / agent 1 (setup costs)	greedy	188.6%	100.4%	109.4%	174.5%
	cooperative	336.5%	12.2%	157.8%	44.4%

Table 6. Average negotiation results for scenario B (1,000,000 rounds)

Average deviations 1,000,000 rounds p_1=0.2; p_2=0.04		Firm / agent 2 (tardiness) greedy f_1	greedy f_2	cooperative f_1	cooperative f_2
Firm / agent 1 (setup costs)	greedy	188.4%	100.2%	84.4%	216.0%
	cooperative	462.0%	4.0%	164.3%	34.3%

Considering the unbalanced deviations of firm 1 and firm 2, it seems that firm 2 is better off. But in general one cannot directly compare these values because of the problem of interpersonally relating utility functions. Only after defining some aggregate utility function (social welfare) further efforts to counterbalance the deviations might make sense. In general, savings realized by coordinated decision making (as by the applied negotiation scheme) may vary widely depending upon individual objectives and restrictions.

Since both cooperative agents, by design, roughly accept the same number of candidate contracts, the outcome of the negotiation may be regarded as fair (in spite of one firm acquiring production sequences with much smaller deviations from an ideal contract than the other firm). Whether this outcome is adequate or not depends on normative decisions with regard to the understanding of fairness and the interpersonal transfer of utility. From a practical point of view, the introduction of monetary side payments into the contract space may provide a means of handling utility transfer within the considered negotiation framework.

5. Game-theoretic analysis

In Section 4, we experimentally examined the outcome of negotiation processes without contemplating about the strategic behavior of the agents. Considering the mutual choice of the acceptance criteria as a strategic decision situation in the sense of non-cooperative game theory, the results shown in Tables 2, 3, 5, and 6 define corresponding payoff tables (with smaller deviations representing larger utility values). Under the assumption of common knowledge about the ordinal relations in the payoff table (notwithstanding information asymmetry concerning the utility functions), the application of the dominance criterion from game theory results in both agents behaving greedily and thus obtaining the worst outcome. Namely, if an agent considers every choice of the other agent in turn, and selects a best "answer" for each choice in order to maximize its own utility value, it leads to the selection of the greedy acceptance criterion as the dominant individual strategy in every case. Therefore, the combination of two greedy agents represents an unambiguous Nash equilibrium, which unfortunately corresponds to a socially dominated outcome. This situation resembles the classic prisoner's dilemma [1], as has already been observed in [20].

Thus the problem arises how to resolve this unfortunate strategic decision situation. Considering the negotiation process as an iterative game, with each round representing a separate game, may pose as an option since the prisoner's dilemma may be overcome in iterative games [1]. However, it may take many rounds to reliably detect defective behavior, during which most of the benefit might have already been realized and assigned to a greedy agent; see Figure 3. Klein et al. [20] propose to adapt the negotiation protocol by moving the actual acceptance criterion to the mediator (i.e., applying an "annealing mediator"). However, this comes at the cost of complicating the negotiation protocol and requires that the agents reveal more information to enable the mediator to roughly assess the quality of a candidate contract.[9]

Our approach to overcome the prisoner's dilemma is as follows. The mediator basically defines mandatory acceptance rates for the agents in the course of the negotiation process. The procedure as described in Section 3.3 may be applied without change. Thus both agents are enabled to determine suitable temperature values and cooling schedules to attain the defined acceptance probabilities throughout the negotiation process. The mediator can easily verify that both agents conform to the requirements for cooperative behavior, which is quite fair (as both agents must give in equally often) and eventually leads to the high quality negotiation outcome due to two cooperative agents. The conscientious application of the cooperative acceptance criterion (as described in Section 3.3) may actually represent incentive compatible behavior (concerning the honest revelation of private information); since otherwise an agent would have to partly accept worse contracts with a higher probability. However, a formal proof of such a conjecture might be beyond reach, as this would require demonstrating that the Boltzmann criterion is the quasi best probabilistic acceptance criterion that one may think about.

[9] Klein et al. [20] propose that the agents classify responses as strong or weak. Yet it is not quite clear how the agents should achieve a suitable classification and whether this scheme is generally robust and effective.

6. Conclusions

We examined the coordination of production sequences between two independent decision making units within two-stage supply chains. Because of the intractability of such collaboration scenarios and corresponding contract spaces with multiple interdependent issues, simple bargaining concepts or auction-based mechanisms are not suitable. However, automated negotiation procedures may provide an effective building-block mechanism to support supply chain coordination.

The main contributions of this paper are, first, the extension of the negotiation procedure of Klein et al. [20] concerning a specific design of a general and fair cooperative acceptance criterion, and, second, the application and validation of the approach for different collaboration scenarios. The design of the cooperative acceptance criterion allows the definition of negotiation rules that can be verified to force both agents to collaborate in a cooperative manner (which solves the prisoner's dilemma). Moreover, the proposed concept implicitly prescribes the division of the benefits from collaboration among the parties.

The applied negotiation mechanism is quite simple and general: no exploitation of domain-specific knowledge, only plain yes/no responses by the agents which assess pairs of contracts in relation to each other (in particular, neither cardinal utterances nor some argumentation for, or critique on, specific contract features are available, which prevents a more focused generation of contract candidates), minor information revelation, and strategic behavior of selfish agents. Nevertheless, the experimental results are very promising. This is in contrast to the common expectation that simple negotiation processes which are only based on a random generation of potential contracts and minor information revelation might be inefficient.[10]

Evidently, the high-quality outcome may be the result of the rather simple application scenarios. Thus the considered approach must be validated for realistic coordination problems with contract spaces of higher complexity. Respective research might provide more insights regarding the generality of the key findings. In particular, it has to be examined for which kind of scenarios the ordinal relations in the payoff table are the same as observed in this paper. In general, if the combination of two cooperative agents does not always lead to the best outcome, the problem arises how to properly devise strategies in new collaboration scenarios.

There are various other opportunities for further research work. The mentioned options for extending the negotiation process might be followed (e.g., enhancing the role of the mediator with regard to a more intelligent generation of candidate contracts), which probably comes at the cost of stronger requirements for information revelation. The negotiation protocol might be extended for multilateral negotiations. The application within market matching scenarios is also conceivable through computation of affinity measures of all potential partners on a marketplace by provisional negotiation experiments; these assessments may then be used to determine effective assignments.

[10] „If agents can only accept or reject others' proposals, then negotiation (and especially negotiation over objects that are multi-dimensional) can be very time consuming and inefficient since the proposer has no means of ascertaining why the proposal is unacceptable, nor whether the agents are close to an agreement, nor in which dimension/direction of the agreement space it should move next." ([13], p. 203)

References

[1] R. Axelrod, The Evolution of Cooperation, Basic Books, New York, 1984.

[2] M. Bichler, G.E. Kersten, and S. Strecker, "Towards a structured design of electronic negotiations", Group Decision and Negotiation 12, 2003, pp. 311–335.

[3] G.P. Cachon, "Supply chain coordination with contracts", in: A.G. de Kok and S.C. Graves (eds.), Handbooks in Operations Research and Management Science, Vol. 11: Supply Chain Management: Design, Coordination and Operation, North-Holland, Amsterdam, 2004.

[4] E.K. Clemons, S.P. Reddi, and M.C. Row, "The impact of information technology on the organization of economic activity: the ´move to the middle` hypothesis", Journal of Management Information Systems 10 (2), 1993, pp. 9–35.

[5] R.K. Congram, C.N. Potts, and S.L. van de Velde, "An iterated dynasearch algorithm for the single-machine total weighted tardiness scheduling problem", INFORMS Journal on Computing 14, 2002, pp. 52–67.

[6] H.A.J. Crauwels, C.N. Potts, and L.N. van Wassenhove, "Local search heuristics for the single machine total weighted tardiness scheduling problem", INFORMS Journal on Computing 10, 1998, pp. 341–350.

[7] M. Ehrgott and X. Gandibleux, "A survey and annotated bibliography of multiobjective combinatorial optimization", OR Spektrum 22, 2000, pp. 425–460.

[8] S.S. Erengüc, N.C. Simpson, and A.J. Vakharia, "Integrated production/distribution planning in supply chains: An invited review", European Journal of Operational Research 115, 1999, pp. 219–236.

[9] A. Fink, „Automatisierte Koordinationsmechanismen zur Ausgestaltung komplexer Verträge im Rahmen mediatisierter Verhandlungsprozesse" (in German), in: W. Uhr, W. Esswein, E. Schoop (eds.), Wirtschaftsinformatik 2003 / Band I, Medien – Märkte – Mobilität, Springer, Berlin Heidelberg, 2003, pp. 281–301.

[10] A. Fink and S. Voß, "Solving the continuous flow-shop scheduling problem by metaheuristics", European Journal of Operational Research 151, 2003, pp. 400–414.

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

[12] A. Jaszkiewicz, Multiple Objective Metaheuristic Algorithms for Combinatorial Optimization. Habilitation thesis, Politechnika Poznanska, Poznan, 2001.

[13] N.R. Jennings, P. Faratin, A.R. Lomuscio, S. Parsons, M. Woolridge, and C. Sierra, "Automated negotiation: Prospects, methods and challenges", Group Decision and Negotiation 10, 2001, pp. 199–215.

[14] M. Jensen and W. Meckling, "Theory of firm: Managerial behavior, agency costs, and ownership structure", Journal of Financial Economics 3, 1976, pp. 305–360.

[15] D.S. Johnson and L.A. McGeoch, "Experimental analysis of heuristics for the STSP", in: G. Gutin and A. Punnen (eds.), The Traveling Salesman Problem and Its Variations, Kluwer, Dordrecht, 2002, pp. 369–443.

[16] G.E. Kersten, "The science and engineering of e-negotiation: Review of the emerging field", InterNeg Reports INR05/02, Draft paper, October 5, 2002.

[17] G.E. Kersten, "The science and engineering of e-negotiation: An introduction", Proceedings of the 36th Hawaii International Conference on Systems Sciences, IEEE Computer Society Press, Los Alamitos, 2003 (CD ROM, 10 pages).

[18] S. Kirkpatrick, C.D. Gelatt Jr., and M.P. Vecchi, "Optimization by simulated annealing", Science 220, 1983, pp. 671–680.

[19] T.D. Klastorin, K. Moinzadeh, and J. Son, "Coordinating orders in supply chains through price discounts", IIE Transactions 34, 2002, pp. 679–689.

[20] M. Klein, P. Faratin, H. Sayama, and Y. Bar-Yam, "Negotiating complex contracts", Group Decision and Negotiation 12, 2003, pp. 111–125.

[21] S. Kraus, Strategic Negotiation in Multiagent Environments, MIT Press, Cambridge MA, 2001.

[22] A.R. Lomuscio, M. Wooldridge, and N.R. Jennings, "A classification scheme for negotiation in electronic commerce", Group Decision and Negotiation 12, 2003, pp. 31–56.

[23] T.W. Malone, "Modeling coordination in organizations and markets", Management Science 33, 1987, pp. 1317–1332.

[24] T.W. Malone, J. Yates, and R.I. Benjamin, "Electronic markets and electronic hierarchies", Communications of the ACM 30, 1987, pp. 484–497.

[25] P. Milgrom and J. Roberts, Economics, Organization and Management, Prentice Hall, Englewood Cliffs NJ, 1992.

[26] M. Moses and S. Seshadri, "Policy mechanisms for supply chain coordination", IIE Transactions 32, 2000, pp. 245–262.

[27] A. Muthoo, Bargaining Theory with Applications, Cambridge University Press, Cambridge MA, 1999.

[28] S. Napel, Bilateral Bargaining: Theory and Applications, Lecture Notes in Economics and Mathematical Systems 518, Springer, Berlin Heidelberg, 2002.

[29] J.F. Nash, "The bargaining problem", Econometrica 18, 1950, pp. 155–162.

[30] J.F. Nash, "Two-person cooperative bargaining games", Econometrica 21, 1953, pp. 128–140.

[31] J.S. Rosenschein and G. Zlotkin, Rules of Encounter – Designing Conventions for Automated Negotiation among Computers, MIT Press, Cambridge MA, 1994.

[32] A.E. Roth (ed.), Game-theoretic Models of Bargaining, Cambridge University Press, Cambridge MA, 1985.

[33] A. Rubinstein, "Perfect equilibrium in a bargaining model", Econometrica 50, 1982, pp. 97–109.

[34] T.W. Sandholm, "Distributed rational decision making", in: G. Weiss (ed.), Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence, MIT Press, Cambridge MA, 1999, pp. 201–258.

[35] A.M. Sarmiento and R. Nagi, "A review of integrated analysis of production-distribution systems", IIE Transactions 31, 1999, pp. 1061–1074.

[36] J.F. Shapiro, Modeling the Supply Chain, Duxbury, Pacific Grove, 2001.

[37] M. Ströbel and C. Weinhardt, "The Montreal taxonomy for electronic negotiations", Group Decision and Negotiation 12, 2003, pp. 143–164.

[38] E.D. Taillard, "Benchmarks for basic scheduling instances", European Journal of Operational Research 64, 1993, pp. 278–285.

[39] D.J. Thomas and P.M. Griffin, "Coordinated supply chain management", European Journal of Operational Research 94, 1996, pp. 1–15.

[40] A.A. Tsay, S. Nahmias, and N. Agrawal, "Modeling supply chain contracts: A review", in: S. Tayur, R. Ganeshan, and M. Magazine (eds.), Quantitative Models for Supply Chain Management, Kluwer, Boston, 1999, pp. 299–336.

[41] S. Whang, "Coordination in operations: A taxonomy", Journal of Operations Management 12, 1995, pp. 413–422.