Improving the Performance of the Winner Determination Process in Multi-Agent Contracting*

Alexander Babanov
Department of Computer Science and Engineering
Department of Economics
University of Minnesota
babanov@cs.umn.edu

Wolfgang Ketter
Department of Computer Science and Engineering
University of Minnesota
ketter@cs.umn.edu

February 8, 2004

*This work was conducted in and as a part of the MAGNET framework. Hereafter we refer to "our group" as to the MAGNET research group leaded by Professor Maria Gini in Artificial Intelligence, Robotics and Vision Lab (AIRVL), University of Minnesota.

1 Introduction

We propose a methodology to improve the performance of the winner determination process in the MAGNET (Multi-AGent NEgotiation Testbed) research project. MAGNET agents participate in a first-price sealed-bid combinatorial auction over compound plans with precedence relations and time constraints. After the auction ends, agents solve a winner determination problem for the feasible combinations of bids submitted to the auction and award bids from the best combination. Currently there are two known implementations of the winner determination process: one is based on a Simulated Annealing (SA) and the other employs modified Integer Programming (IP) formulation. However, both SA and IP methods do not scale satisfactory with the number of bids and have solution quality issues that are common for "brute-force" solution concepts not designed atop the market information analysis. As the scalability problem effectively prevents the current system from exploring realistically large and sophisticated markets, we face a need to either find a better algorithm for the winner determination or improve the performance of the existing ones.

We envision our **long-range goal** in perfecting the technology developed inside the MAGNET project to make it applicable to a large variety of automated and mixed-initiative auction-based markets. One of the most important steps on the way to achieve this goal is to improve the performance of the winner determination process, both in terms of its speed and quality, where by quality we assume some function of the number of rejected bids and of the cost to the suppliers of committing to these bids. In the domain of our long-range goal, we define the **objective of this proposal** as to improve the speed and quality of the existing SA and IP algorithms by incorporating the analysis and usage of the market information. Our **central hypothesis** is that the Expected Utility (EU) Theory is a viable choice of the underlying economic concept. This hypothesis is well supported by our preliminary research results and, if justified, will also provide us with the optimization criterion that summarizes risk and profit in a single intuitive measure. In the event we find the central hypothesis not viable, we would be able to use the acquired knowledge to make a more experienced choice among many other economic concepts.

The **rationale** of our overall research effort is that the broad usage of the MAGNET technology will increase the efficiency of the current markets by shifting the burden of the market exploration, auction handling and preliminary decision analysis from a human user to a network of intelligent agents. The variety of the prospective applications includes auction-based and auction-ready markets with repeated interactions and complex time-sensitive tasks, ranging from vacation package sales to construction project subcontracting. The completion of the current proposal goals will allow for the exploration of the realistically populated and sophisticated market, thus bringing us significantly closer to the ultimate goal.

Our team is **best suited** for undertaking the proposed project: first of, because the architecture of our underlying MAGNET system is by far the only presently available multi-agent architecture that is designed to support combinatorial auctions over contracts with precedence relations and time constraints. Another important reason is that our team conducted the extensive preliminary research to take on the proposed objectives, and it possesses the breadth of skills in Artificial Intelligence, Computational Methods and Economic Theory required to successfully achieve a positive outcome.

We plan to test our central hypothesis and complete the goal of this proposal by pursuing the following three **specific aims**:

1. **Design a test suite** for measuring the performance gain due different RFQ generation techniques. This test suite should not only estimate raw speed increase, but also capture the notion of quality. The latter refers, for example, to the above mentioned case with the loss of customer agent's credibility due to its repeated rejection of the excessively large number of the submitted bids. Our group performed preliminary work on testing the speed of SA and IP algorithms, the next step is to use this work together with widely available combinatorial auction test suites to design comprehensive speed testing part of the suite. The quality testing part of the proposed suite will be designed from scratch, since no applicable suite is presently available. It is also important to make sure that the test suite is not inherently biased toward one or another solution concept, hence the test suite should also contain a self-testing part.

2. **Incorporate EU notion** into the existing MAGNET framework. This part requires intense theoretical study and will be completed in the following order: first, we will modify current SA and IP algorithms to make use of EU maximization as well as change current RFQ generation to EU-based one. Next, we will bootstrap the modified MAGNET system with the information from a simple evolutionary market model to perform comprehensive tests of the proper operation. Finally, we will create a system for collecting and analyzing the dynamic market information relevant to EU and perform operation tests with the help of human subjects.

3. **Thoroughly test** the EU-based system from the second part against the old "brute-force" approaches with the use of the test suite from the first part. In this part of the proposed research we will use test information to make all necessary adjustments to make the best of the new system performance. Upon completion of this final part we will have the complete implementation of the EU-based solution concept as well as the proof of this concept and the test suite that can be reused later to tune the existing or design new concepts.

The proposed project will have **significant** academic and economic impact. On the academic side, it will create the **innovative** approach to the winner determination problem in combinatorial auctions with scheduling constraints that is based on the dynamic market information analysis. On the economic side, it will provide a way of utilizing the effective yet currently under-used combinatorial auction concept in the mixed-initiative environment that facilitates human user decision process by relaxing the burden of market exploration, auction handling and initial decision making.

2 Related Work

Auctions are becoming the predominant mechanism for agent-mediated electronic commerce [12]. AuctionBot [26] and eMEDIATOR [24] are among the most well known examples of multi-agent auction systems. The determination of winners of combinatorial auctions [19] is hard. Dynamic programming [22] works for small sets of bids, but does not scale and imposes significant restrictions on the bids. Algorithms such as Bidtree [24] and CASS [11] reduce the search complexity, but their criterion to select bids is just price. Our bids include a time window for each task, and so bid selection cannot be separated from scheduling.

A set of optimal and approximate methods, along with a test set for algorithm evaluation, was published by Fujishjima et al [11]. Hoos and Boutilier [15] describe a stochastic local search approach to solving combinatorial auctions, and characterize its performance with a focus on time-limited situations. A key element of their approach involves ranking bids according to expected revenue; it's hard to see how this could be adapted to the MAGNET domain with temporal and precedence constraints, and without free disposal. Andersson et al [1] describe an Integer Programming approach to the winner determination problem in combinatorial auctions. Nisan [20] extended this with an analysis of bidding languages for combinatorial auctions. More recently, Sandholm [23] has described an improved winner-determination algorithm called CABOB that uses a combination of linear programming and branch-and-bound techniques. It is not clear how this technique could be extended to deal with the temporal constraints in MAGNET, although the bid-graph structure may be of value.

Leyton-Brown et al [16] suggest a way of constructing a universal test suite for winner determination algorithms in combinatorial auctions. Although their work does not speculate on cases with both precedence and time constraints and, thus, is not directly applicable to the MAGNET framework, it nevertheless provides necessary ground-works for testing the speed of the algorithms.

Expected Utility Theory [21] is a mature, yet controversial, field of Economics, that attracted many supportive as well as critical studies, both theoretical [17, 18] and empirical [25, Jullien00?]. We believe that the expected utility and related concepts will attract more attention in the relation to automated auctions soon, in particular, because they suggest a practical way of describing risk estimations and temporal preferences.

3 Background

The MAGNET system is designed to support multiple agents in negotiating contracts for tasks with temporal and precedence constraints [10]. MAGNET proposes to increase the efficiency of the current markets, by shifting the burden of the market exploration, auction handling and preliminary decision analysis from a human user to a network of heterogeneous and self-interested agents. We distinguish between two agent roles, the Consumer and the Supplier, see figure 1. It is possible that one agent performs both roles by composing smaller packages of bids to sell them again. A consumer agent needs resources outside its direct control to complete multi-task plan with precedence and time constraints. Supplier agents may offer to provide resources in response to consumer requests. After that the agents participate in a first-price sealed-bid combinatorial auction.

Customer agents pursue their goals by formulating and presenting Requests for Quotes (RFQs) to Supplier agents through a market infrastructure [9]. The RFQ specifies, among other information, the description of tasks, precedence relations, time windows for the tasks in the plan, and specifying the timeline for the bidding process. Customer agents attempt to satisfy their goals for the least net cost, where cost factors can include not only bid prices, but also goal completion time and risk factors. More precisely, these agents are attempting to maximize the utility function of some user, as discussed in detail in [4]. Supplier agents attempt to maximize the value of the resources under their control by submitting bids in response to those RFQs, specifying what tasks they are able to undertake, when they are available to perform those tasks, and at what price.

In the following we concentrate on the decision process a customer agent needs to go through in order to generate an RFQ. We study in particular the problem of how to specify in the RFQ the

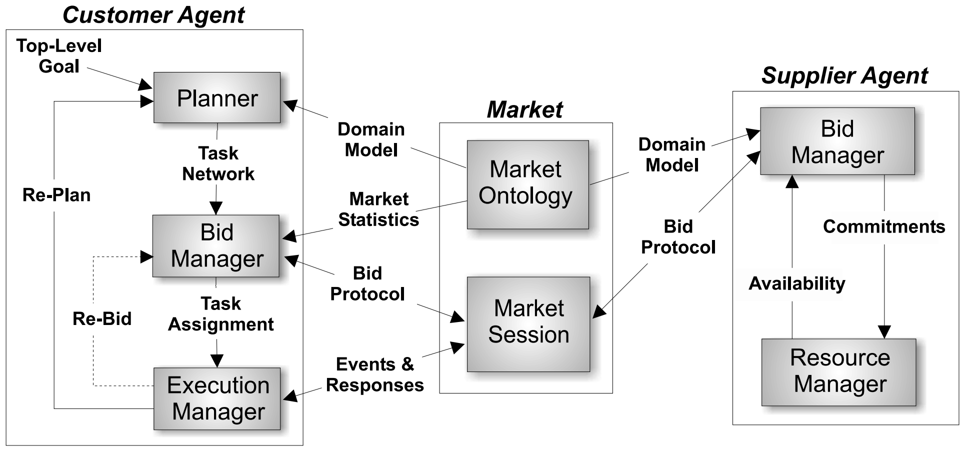

Figure 1: The MAGNET architecture.

time windows for the different tasks. Currently we calculate the duration for the time windows with the help of the Critical Path (CPM) algorithm [14]. This algorithm first walks the directed graph of tasks and precedence constraints forward to compute the earliest start time for each task, and then backward from the goal time to compute the latest possible finish and start times for each task. The entire plan's minimum duration is called the *makespan*, the difference between the deadline and the latest early finish time is called the plan's *total slack*. The *Critical Path* is defined as a sequence of tasks which have the least slack; to calculate the critical path we use CPM as well.

Time windows are set by determining the makespan of the plan and multiplying it by a "slack" factor of 1.2, then "relaxing" the time windows for individual tasks allow some overlap. The final result is that individual tasks are given time windows of at least 125% of their expected times.

After the customer agent did send out the RFQ to the market, the supplier agents can bid on the tasks issued in the RFQ. The time period a supplier agent has to deliberate about bidding for the tasks starts after receiving the RFQ and finishes at the bid deadline. In this time period the supplier agents may submit bids for certain tasks. After receiving bids from the supplier agents the customer agent has to find and to award the best bid combination. The time for the customer agent to deliberate and solve the winner determination problem ends at the bid award deadline. In this optimization process each task in the task-network needs to get covered with exactly one bid from the overall feasible combinations of incoming bids; the optimization criterion is price. In the following figure 2 you can see the described interaction between the customer (contractor) and the supplier agent(s). All interactions between the customer and the supplier agent have to go through the market.

There are several issues in the current way of creating an RFQ. First, with the current approach we only include some market information, and therefore it does not mirroring the real world. At the moment there are prefixed methods which allocate automatically, mentioned above, 20% slack of the entire task duration. This gives us a bigger choice set, which allows for better optimization results, but at the same time the crux of the winning bundle determination problem is, that the consideration of each extra bid adds to the run-time of the optimization procedure. When the number of bids increases, the search time for the optimization algorithms (SA and IP) to run

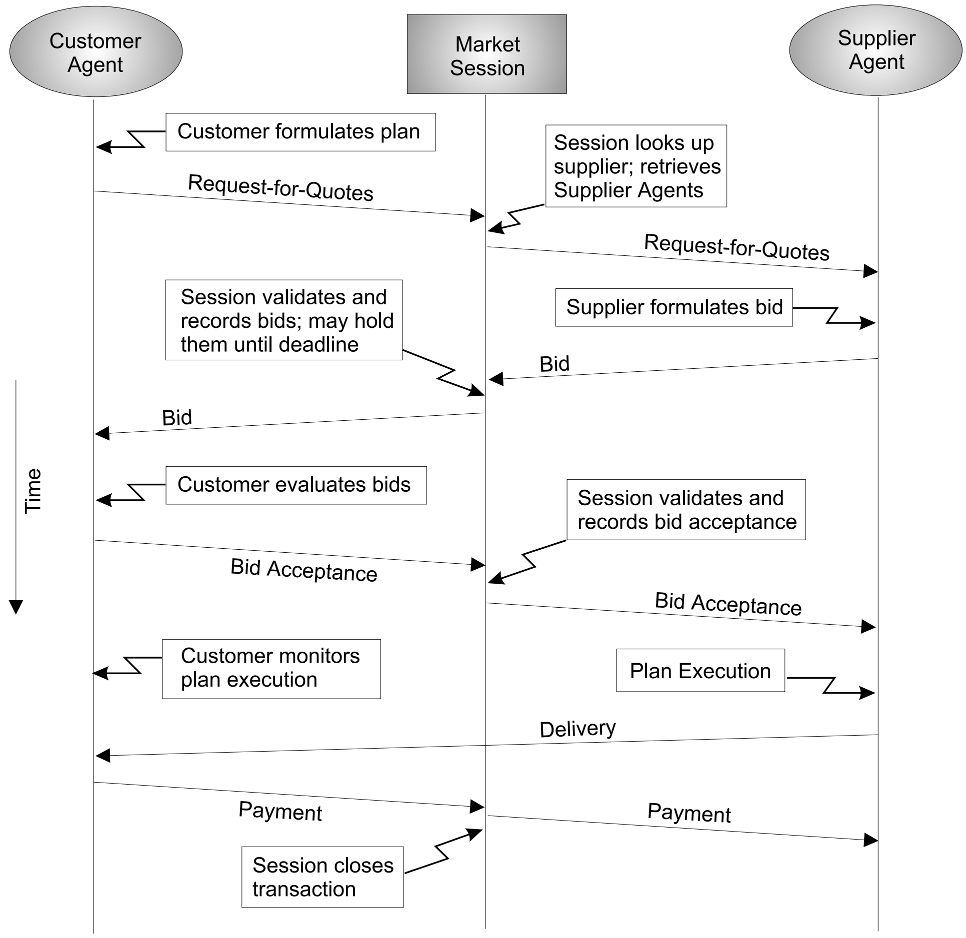

Figure 2: Customer-Supplier Interaction in a Typical Contracting Market Session.

grows exponentially [3]. It is of paramount importance to solve this scalability problem, since this prevents currently MAGNET to be used in large markets. Another less explicit problem with that is, that every extra bid over the absolute minimum that covers all tasks also adds one more rejected bid. Ultimately, a large percentage of rejections will reduce the customer agent's credibility. If the customer agent looses credibility, this will result on the supplier side in either increased cost of the submitted bids, or the submitted bids may have very wide time windows or the suppliers do not bid at all on tasks for this particular customer in the future.

This tension between issuing an RFQ that will solicit the maximum number of bids and reduce costs, and one that will guarantee the feasibility of any plan constructed with the resulting bids is discussed in [8]. We assume that suppliers will bid depending of their current resource commitments, and therefore larger time windows will result in more bids and better utilization of resources, which will result in lower prices. However, an RFQ with overlapping time windows will make the process of winner determination much more complex [6].

An additional factor to be considered is the financial exposure of the customer agent [5]. We assume non-refundable deposits are paid to secure awarded bids, and payments for each task are made as the tasks are completed. The payoff for the customer agent occurs only at the completion of the plan. Once a task starts and, in case it is successfully completed in the time period specified by the contract, the customer is liable for its full cost, regardless of whether in the meantime the plan as a whole has been abandoned due to a failure on some other branch of the plan.

If a task is not completed by the supplier, the customer agent is not liable for its cost, but this failure can have a devastating effect of other parts of the plan. Having slack in the schedule increases the probability that tasks will be successfully completed or that there will be enough time to recover if one of the tasks fails. However, slack extends the completion time and so reduces the payoff. In made-to-order products speed is the essence and taking extra time might prevent a supplier from getting a contract. This complicates the selection of which bids to accept. The lowest cost combination of bids and the tightest schedule achievable is not necessarily the preferable schedule because it is more likely to be brittle.

Risk can also be reduced by consolidating tasks with fewer suppliers. Suppliers can bid on "packages" composed of subsets of tasks from the RFQ. In general, the customer is better off from a risk standpoint if it takes these packages, assuming that the supplier is willing to be paid for the whole package at the time of its completion. In some cases, the customer may be willing to pay a premium over the individual task prices in order to reduce risk. The advantage of doing this is greater toward the end of the plan than near the beginning, since at that point the customer has already paid a significant part of the tasks. Having a greater financial exposure provides an additional incentive to reduce risk.

We are overall interested in improving the performance of the winner determination process. This process is strongly influenced by the problem of setting the optimal time window for the RFQ [2], in this paper we show how to use Expected Utility Theory to determine the time windows for tasks in the task network, so that bids that are close to these time windows form the most preferred risk-payoff combinations for the customer agent. This gives us the opportunity to calculate the slack allocation for each task in the RFQ. With this approach we increase the probability of getting better bids in, with better we mean the number of conflicting bids will be relatively small to the approach what we currently doing. The time to create an appropriate RFQ will be longer, but this will pay off overall later in the winner determination process, since we receive more appropriate (less conflicting) bids to cover the tasks, and so we will decrease the final optimization time.

We also used Expected Utility Theory in [5], but in this work we were mostly concerned with computing the marginal expected utility of completing successfully all the tasks within the duration promised.

4 Proposed Work

4.1 Testing for Performance

This part of the proposed work is based on the preliminary research on SA and IP algorithms performance [3]. We will extend the existing test cases to fit EU framework, also we will work on adopting realistic distribution generation techniques from [16] with the emphasis on the sample task network creation and bid generation.

4.2 Testing for Quality

The most complex and innovative concept in the proposed test suite is the notion of quality of the RFQ generation. Naturally, the quality of a single RFQ should depend on the statistical distribution of the bids that are solicited by this RFQ. The quality of some particular RFQ generation procedure should, in turn, depend on the quality of RFQs created by this procedure for different task networks in various market situations. We believe that it is not possible to summarize customer agent's criteria of what makes one RFQ better than other in a single comprehensive measure, unless we place severe restrictions on the market structure. For example, a type of RFQ that maximizes customer agent's expected utility in a single auction, by offering wide time windows and rejecting most of incoming bids, will most surely lead to the devastating results, if being used in repeated market interactions.

The logic above suggests that we need a set of tests that address different parts of RFQ generation. To facilitate test creation we summarized the set of the concepts in RFQ generation and relations between these concepts in Figure 3. Here nodes are various characteristics and criteria of the process and edges are direct relations between them. The only "sink" in this complicated scheme of 15 nodes and more than 30 direct relationships is the customer's expected profit, all other concepts propagate (often, indirectly) their influence to many other concepts. Note that the number of incoming bids and the customer's credibility are two most highly connected concepts in this graph.

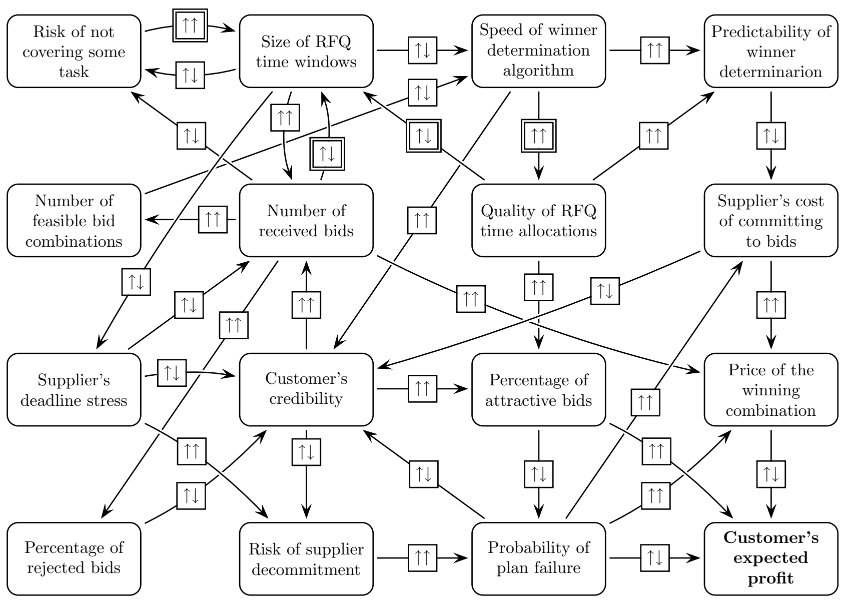

Figure 3: Relationships between performance criteria.

The link label ↑↑ denotes that two concepts are positively correlated while the symbol ↑↓ stands for the negative correlation. For example, increasing number of incoming bids leads to the increasing number of feasible bid combinations and to the decreasing risk of not getting at least one bid for each task in the plan. Next, single-edged box surrounding link label means that this link is "natural", i.e. defined by the market. Complimentary to that, double-edged box shows the algorithmic (or "artificial") relation. For example, increasing size of RFQ window is expected to attract more bids, because it gives a chance to suppliers who otherwise won't be able to participate. In the same time, an RFQ generator may use the increasing number of incoming bids to shrink time windows in attempt to increase the speed of the winner determination algorithm.

The following list represents some basic test criteria that we plan to use in the test suite, the design and analysis of the more sophisticated criteria will be performed as a part of the research:

1. The ratio of the number of feasible bid combination to the total number of combinations. This and many of the following tests will utilize Monte-Carlo methods to prevent test time from growing exponentially with the number of bids.

2. The ratio of the "attractive" bid combinations to the number of all feasible combinations. Here by "attractive" we imply those combination that promise the expected payoff above some threshold in term of the maximum payoff.

3. The structure of the incoming bids distribution in terms of correlation with other bids, likelihood of forming feasible bid combinations, price and risk, etc.

4. The dynamics of the credibility (or reputation building) of a customer agent using one or another RFQ generation mechanism.

4.3 EU-based RFQ Generation

This part of the proposed work is based on the preliminary results in [2]. In particular, we will elaborate on using the EU notion to find optimal placement and size of time windows in an RFQ. The intuition behind this process is represented in Figure 4. In this figure the left graph represents all combinations of bids that were solicited by CPM-based approach, with the gray area denoting "attractive" bid combinations and the circle denoting the best combination. The middle part shows how the knowledge of the market information and integrated EU-maximization help shifting RFQ time windows to capture more "attractive" combinations. The right part of the graph illustrates the idea that with the better placement of the time windows it is possible to shrink these windows in size in order to increase the ratio of attractive bids among all feasible combinations and, thus, decrease the search time.

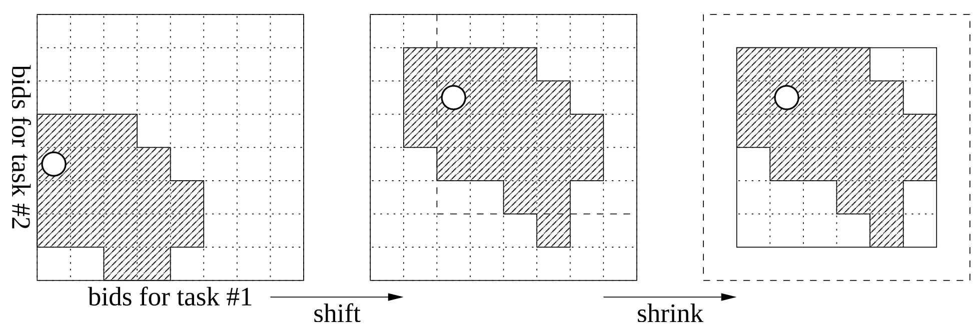

Figure 4: Adjustments to RFQ time windows.

4.4 Bootstrapping EU-based System

We plan to gradually increase the complexity of the market information that is used by the system. The two reasons for that are: first, it is advisable to start with the information that is fully controlled and analytically simple; second, it is cost-efficient to test the system within artificially constructed markets before introducing human involvement. The following list represents the proposed stages of building up the complexity of information inputs and the degree of human involvement:

1. Bootstrapping the model with theoretic distributions, like, Weibull distribution for task completion times and uniform distribution for task start times.
2. Usage of the evolutionary market model to obtain and use more realistic market information, including supplier availability, price distributions and risks.
3. Performing lab test with the help of human subjects (students).
4. Integrating the system in the real-market as a transparent entity that collects and analyses the information by observing human decision maker.
5. Integrating the system as a full-scale mixed-initiative decision maker in the real market.

5 Assessment

As the positive result of the proposed research we expect to observe EU-based RFQ generation techniques to perform significantly better than CPM-based techniques in the major part of individual test. Another condition of the successful research completion is that our test suite is applicable to the analysis of generic RFQ generation techniques and winner determination algorithms.

6 Time Table

We plan to complete the proposed work in a matter of one year. The detailed schedule highly depends on the complexity of the quality notion and the resulting size of the test suite. We are planning to clarify the detailed schedule after finishing work on the quality issues. As for now, we plan to complete test suite creation in the next 4 months, EU integration in the following 6 months and the comprehensive test runs in the last 2 months.

7 Summary of Contributions and Impact of Proposed Research

We will extend automated negotiation beyond the buying and selling of individual goods and services, into the realm of complex supply-chain coordination issues.

7.1 Benefits of Proposed Research to Society

Buyer-supplier relationships are becoming more dependent on factors such as quality, delivery performance, flexibility as opposed to just cost [13], and these must be taken into account in automated negotiation. Our work will add a new dimension to business-to-business interactions, by adding the ability to automate the negotiation and execution of complex contracts among multiple suppliers.

We believe that the study we are proposing is critical for opening electronic marketplaces to a wider audience, meaning more people can participate in the electronic market. Our mixed-initiative system will allow a human decision maker to more efficiently process larger amounts of information with the help of intelligent agents. Our approach to solve the winner determination problem in combinatorial auctions will include dynamic market information analysis. By combining intelligent agents and combinatorial auctions we will increase the efficiency of nowadays markets.

7.2 Educational and Academic Impact

Electronic commerce supported by software agents will have an increasingly role in our future economy and there is a great need to train the future pool of information technology workers in this area. There is a potential for sharing our findings with the Business School of the University of St. Thomas, as well with the Carlson School of Management at the University of Minnesota.

Our approach is innovative, since the winner determination problem in combinatorial auctions with precedence and time constraints is an important problem, that has yet not received much attention in the scientific community.

8 Results from Previous Support

The major results from the previous supported work are:

- Most of the key-features in the MAGNET testbed architecture as proposed previously have been implemented and are ready to be released. This includes prototypes of the customer and supplier agent, as well as the market, interaction and bidding protocols, and a graphical user interface.
- Development of the winner determination algorithms in context of MAGNET with precedence and time constraints. One type of the algorithms is based on a Simulated Annealing [7] and the other employs modified Integer Programming formulation [6]. We developed a test suite [3] to measure the speed of the existing algorithms and found out that the used optimization algorithms do not scale satisfactory with the number of bids.

9 Technical Report

9.1 Bid Compatibility Test

The notation for RFQ time window for the task $n \in N$ is:

- *Early start time* $w_n^{\text{es}} \in \left[t_n^{\text{es}}, t_n^{\text{ls}}\right]$, no bid should start earlier than that.
- *Late start time* $w_n^{\text{ls}} \in \left[w_n^{\text{es}}, t_n^{\text{ls}}\right]$, no bid should start later than that. This parameter is defined to be equal to w_n^{lf} (see below) for the CPM-based algorithm.
- *Late finish time* $w_n^{\text{lf}} \in \left[w_n^{\text{ls}}, t_n^{\text{lf}}\right]$, no bid can finish later than that.

The notation for some set bids for tasks $n \in N$ is:

- *Early start time* $b_n^{\text{es}} \in \left[w_n^{\text{es}}, w_n^{\text{ls}}\right]$, this bid execution cannot start earlier than that.
- *Late start time* $b_n^{\text{ls}} \in \left[b_n^{\text{es}}, w_n^{\text{ls}}\right]$, this bid execution cannot start late than that.
- *Duration* $b_n^{\text{d}} \geq 0$, such that $b_n^{\text{ls}} + b_n^{\text{d}} \leq w_n^{\text{lf}}$.

We perform the bid compatibility test by solving the following system of linear inequalities for all $n \in N$:

$$
\begin{array}{rcl}
b_n^{\text{s}} & \leq & b_n^{\text{es}} \\
-b_n^{\text{s}} & \leq & -b_n^{\text{ls}} \\
b_n^{\text{s}} - b_m^{\text{s}} & \leq & -b_n^{\text{d}}, \quad m \in S_1(n)
\end{array}
$$

For our sample garage building task network it will look as shown in Figure 5. Note that the creation of this system is easily automated and its dimensions grow slowly with the size of the problem.

This test can be used as follows:

1. To find the percentage of feasible combinations of bids among all possible combinations.
2. To plot a histogram of the proportion of bids that visit some ratio of feasible combinations. It makes sense to distinguish between bids for different tasks in this test, i.e. build separate histograms for each of N tasks.

$$
\begin{bmatrix}
-1 & & & & & \\
1 & & & & & \\
1 & -1 & & & & \\
1 & & -1 & & & \\
1 & & & -1 & & \\
 & -1 & & & & \\
 & 1 & & & & \\
 & 1 & & & & -1 \\
 & & -1 & & & \\
 & & 1 & & & \\
 & & 1 & & -1 & \\
 & & & -1 & & \\
 & & & 1 & & \\
 & & & 1 & -1 & \\
 & & & & -1 & \\
 & & & & 1 & \\
 & & & & 1 & -1 \\
 & & & & & -1 \\
 & & & & & 1
\end{bmatrix}_{2N+\sum_{n\in N} S_1(n)\times N}
\times
\begin{bmatrix}
b_1^{\mathrm{s}} \\ b_2^{\mathrm{s}} \\ \vdots \\ b_3^{\mathrm{s}}
\end{bmatrix}_{N\times 1}
\leq
\begin{bmatrix}
-b_1^{\mathrm{es}} \\ b_1^{\mathrm{ls}} \\ -b_1^{\mathrm{d}} \\ -b_1^{\mathrm{d}} \\ -b_1^{\mathrm{d}} \\ -b_2^{\mathrm{es}} \\ b_2^{\mathrm{ls}} \\ -b_2^{\mathrm{d}} \\ -b_3^{\mathrm{es}} \\ b_3^{\mathrm{ls}} \\ -b_3^{\mathrm{d}} \\ -b_4^{\mathrm{es}} \\ b_4^{\mathrm{ls}} \\ -b_4^{\mathrm{d}} \\ -b_5^{\mathrm{es}} \\ b_5^{\mathrm{ls}} \\ -b_5^{\mathrm{d}} \\ -b_6^{\mathrm{es}} \\ b_6^{\mathrm{ls}}
\end{bmatrix}_{2N+\sum_{n\in N} S_1(n)\times 1}
$$

Figure 5: System of inequalities for the bid compatibility test.

9.2 Bid Generation

We use the following algorithm to generate a set of bids for task $n \in N$:

1. Generate bid early start times $\left\{b_{n,1}^{\mathrm{es}}, \ldots, b_{n,m_n}^{\mathrm{es}}\right\} \in \left[w_n^{\mathrm{es}}, w_n^{\mathrm{ls}}\right]$ using a stationary Poisson process with frequency f_n:

$$
\begin{aligned}
u &\leftarrow U(0,1) \\
b_{n,j}^{\mathrm{es}} &\leftarrow -\frac{1}{f_n}\ln u + \begin{cases} w_n^{\mathrm{es}}, & j = 1 \\ b_{n,j-1}^{\mathrm{es}} & \text{otherwise} \end{cases}
\end{aligned}
$$

2. Generate bid durations as:
$$b_{n,j}^{\mathrm{d}} \leftarrow \mathrm{Weibull}\left(\alpha_n, \beta_n\right)$$

3. Drop bids that do not fit in RFQ time window, i.e. if
$$b_{n,j}^{\mathrm{es}} + b_{n,j}^{\mathrm{d}} > w_n^{\mathrm{lf}}$$

4. Generate bid late finish times as:

$$b_{n,j}^{\mathrm{lf}} \leftarrow b_{n,j}^{\mathrm{es}} + b_{n,j}^{\mathrm{d}} + |\mathrm{N}(0, \sigma_n)|$$

5. Restrict late finish times from above by w_n^{lf}.

9.3 List of Attachments

- `calcProbs.m`: Matlab function to calculate payoffs and corresponding probabilities (see algorithm in EU paper).
- A graph of cumulative probabilities of plan completion for near optimal (local maximum) schedule. This schedule is optimal for the given relative task ordering. The globally maximizing schedule for this case can be found in EU paper.
- The corresponding graph of certainty equivalent changes with deviations of one task at a time from the locally maximizing schedule.
- `generateBids.m`: Matlab function to generate bids using the algorithm described in the preceding section.
- A sample distribution of generated bids made under assumption of equal bid frequencies for the EU-based RFQ with no overlaps. We should be able to derive optimal overlaps from the certainty equivalent graph mentioned above, however it requires custom-build non-linear equation solver, due to the presence of the constant value intervals in certainty equivalent functions.

References

[1] Arne Andersson, Mattias Tenhunen, and Fredrik Ygge. Integer programming for combinatorial auction winner determination. In *Proc. of 4th Int'l Conf on Multi-Agent Systems*, pages 39–46, July 2000.

[2] Alexander Babanov, John Collins, and Maria Gini. Risk and expectations in a-priori time allocation in multi-agent contracting. In *Proc. of the First Int'l Conf. on Autonomous Agents and Multi-Agent Systems*, volume 1, pages 53–60, Bologna, Italy, July 2002.

[3] J. Collins and M. Gini. Performance of winner determination search in combinatorial auctions with time constraints. *Submitted to the First International Joint Conference on Autonomous Agents and Multi-Agent Systems*, July 2002.

[4] John Collins, Corey Bilot, Maria Gini, and Bamshad Mobasher. Mixed-initiative decision support in agent-based automated contracting. In *Proc. of the Fourth Int'l Conf. on Autonomous Agents*, pages 247–254, June 2000.

[5] John Collins, Corey Bilot, Maria Gini, and Bamshad Mobasher. Decision processes in agent-based automated contracting. *IEEE Internet Computing*, 5(2):61–72, March 2001.

[6] John Collins and Maria Gini. An integer programming formulation of the bid evaluation problem for coordinated tasks. In Brenda Dietrich and Rakesh V. Vohra, editors, *Mathematics of the Internet: E-Auction and Markets*, volume 127 of *IMA Volumes in Mathematics and its Applications*, pages 59–74. Springer-Verlag, New York, 2001.

[7] John Collins, Rashmi Sundareswara, Maria Gini, and Bamshad Mobasher. Bid selection strategies for multi-agent contracting in the presence of scheduling constraints. In A. Moukas, C. Sierra, and F. Ygge, editors, *Agent Mediated Electronic Commerce II*, volume LNAI1788. Springer-Verlag, 2000.

[8] John Collins, Maksim Tsvetovat, Rashmi Sundareswara, Joshua Van Tonder, Maria Gini, and Bamshad Mobasher. Evaluating risk: Flexibility and feasibility in multi-agent contracting. In *Proc. of the Third Int'l Conf. on Autonomous Agents*, pages 350–351, May 1999.

[9] John Collins, Maxim Tsvetovat, Bamshad Mobasher, and Maria Gini. MAGNET: A multi-agent contracting system for plan execution. In *Proc. of SIGMAN*, pages 63–68. AAAI Press, August 1998.

[10] John Collins, Ben Youngdahl, Scott Jamison, Bamshad Mobasher, and Maria Gini. A market architecture for multi-agent contracting. In *Proc. of the Second Int'l Conf. on Autonomous Agents*, pages 285–292, May 1998.

[11] Yuzo Fujishima, Kevin Leyton-Brown, and Yoav Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Proc. of the 16th Joint Conf. on Artificial Intelligence*, pages 548–553, 1999.

[12] Robert H. Guttman, Alexandros G. Moukas, and Pattie Maes. Agent-mediated electronic commerce: a survey. *Knowledge Engineering Review*, 13(2):143–152, June 1998.

[13] S. Helper. How much has really changed between US manufacturers and their suppliers. *Sloan Management Review*, 32(4):15–28, 1991.

[14] Frederick S. Hillier and Gerald J. Lieberman. *Introduction to Operations Research*. McGraw-Hill, 1990.

[15] Holger H. Hoos and Craig Boutilier. Solving combinatorial auctions using stochastic local search. In *Proc. of the Seventeen Nat'l Conf. on Artificial Intelligence*, pages 22–29, 2000.

[16] Kevin Leyton-Brown, Mark Pearson, and Yoav Shoham. Towards a universal test suite for combinatorial auction algorithms. In *Proc. of ACM Conf on Electronic Commerce (EC'00)*, pages 66–76, Minneapolis, MN, October 2000.

[17] Mark J. Machina. Choice under uncertainty: Problems solved and unsolved. *Journal of Economic Perspectives*, 1(1):121–154, 1987.

[18] Mark J. Machina. Dynamic consistency and non-expected utility models of choice und er uncertainty. *Journal of Economic Literature*, 27(4):1622–1668, December 1989.

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

[20] Noam Nisan. Bidding and allocation in combinatorial auctions. In *Proc. of ACM Conf on Electronic Commerce (EC'00)*, pages 1–12, Minneapolis, Minnesota, October 2000. ACM SIGecom, ACM Press.

[21] John W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32:122–136, 1964.

[22] Michael H. Rothkopf, Alexander Pekeč, and Ronald M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.

[23] Tuomas Sandholm. An algorithm for winner determination in combinatorial auctions. In *Proc. of the 16th Joint Conf. on Artificial Intelligence*, pages 524–547, 1999.

[24] Tuomas Sandholm. Approaches to winner determination in combinatorial auctions. *Decision Support Systems*, 28(1-2):165–176, 2000.

[25] V. Kerry Smith and William H. Desvousges. An empirical analysis of the economic value of risk changes. *Journal of Political Economy*, 95(1):89–114, February 1987.

[26] Peter R. Wurman, Michael P. Wellman, and William E. Walsh. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. In *Second Int'l Conf. on Autonomous Agents*, pages 301–308, May 1998.