
Asking the right question: Risk and expectation in multiagent contracting

ALEXANDER BABANOV, JOHN COLLINS, AND MARIA GINI

Department of Computer Science and Engineering, University of Minnesota, Minneapolis, Minnesota 55455, USA

(RECEIVED April 22, 2002; ACCEPTED January 12, 2003)

Abstract

In this paper we are interested in the decision problem faced by an agent when requesting bids for collections of tasks with complex time constraints and interdependencies. In particular, we study the problem of specifying an appropriate schedule for the tasks in the request for bids. We expect bids to require resource commitments, so we expect different settings of time windows to solicit different bids and different costs. The agent is interested in soliciting “desirable” bids, where desirable means bids that can be feasibly combined in a low-cost combination that covers the entire collection of tasks. Since the request for bids has to be issued before the agent can see any bids, in this decision process there is a probability of loss as well as a probability of gain. This requires the decision process to deal with the risk posture of the person or organization on whose behalf the agent is acting. We describe a model based on Expected Utility Theory and show how an agent can attempt to maximize its profits while managing its financial risk exposure. We illustrate the operation and properties of the model and discuss what assumptions are required for its successful integration in multiagent contracting applications.

Keywords: Automated Auctions; Expected Utility; Multiagent Contracting; Optimization; Risk Estimation

1. INTRODUCTION

E-commerce technology has the potential to benefit society by reducing the cost of buying and selling and by opening new market opportunities. We envision an auction-based approach to the management of agile and dynamic supply chains, in which autonomous, self-interested agents negotiate on behalf of organizations and individuals to organize coordinated activities. This is an area in which the potential payoff is very high, given the projected size of the business to business and make to order e-commerce markets.

More production processes are being outsourced to outside contractors, making supply chains longer and more convoluted. This increased complexity is compounded by increasing competitive pressure and accelerated production schedules that demand tight integration of all processes. Finding potential suppliers is only one step in the process of producing goods. Time dependencies among operations make scheduling a major factor. A late delivery of a part

may produce a cascade of devastating effects. Unfortunately, current auction-based systems do not have any notion of time. Handling auctions for tasks with time constraints is beyond the capabilities of current e-commerce systems.

We present the results of a study of how an autonomous agent can maximize its profits while predicting and managing its financial risk exposure when requesting bids for tasks with complex time constraints. We show how this can be done by specifying appropriate time windows for tasks when soliciting bids and by using received bids effectively in building a final work schedule.

2. MAGNET, A MULTIAGENT NEGOTIATION TESTBED FOR CONTRACTING TASKS WITH TEMPORAL AND PRECEDENCE CONSTRAINTS

This study is a part of the MAGNET (multiagent negotiation testbed) research project (Collins et al., 2002). MAGNET agents participate in first-price, sealed-bid combinatorial auctions over collections of tasks with precedence relations and time constraints. MAGNET promises to increase the efficiency of current markets by shifting

Reprint requests to: Alexander Babanov, Department of Computer Science and Engineering, University of Minnesota, 4-192 EE/CSci, 200 Union Street SE, Minneapolis, MN 55455. E-mail: babanov@cs.umn.edu

much of the burden of market exploration, auction handling, and preliminary decision analysis from human decision makers to a network of heterogeneous agents.

We distinguish between two agent *roles*, the *Customer* and the *Supplier* (see Fig. 1). A customer has a set of tasks to be performed, with complex dependencies among the tasks. When a customer lacks the resources to carry out its own tasks, it may solicit resources from suppliers by presenting a *request for quotes* (RFQ) through an agent-mediated market. Supplier agents may offer to provide the requested resources or services, for specified prices, over specified time periods. Once the customer receives bids, it evaluates them to select an optimal set of bids and create a work schedule. This paper deals with decision problems in the Bid Manager component of the Customer Agent.

This is a schematic outline of the main interactions among agents:

- A customer issues an RFQ that specifies tasks, their precedence relations, and a timeline for the bidding process. For each task, a time window is specified giving the earliest time the task can start and the latest time it can end.
- Suppliers submit bids. A bid includes a set of tasks, a price, a portion of the price to be paid as a nonrefundable deposit, and estimated duration and time window data that reflect supplier resource availability and constrain the customer’s scheduling process.
- The customer decides which bids to accept. Each task needs to be mapped to exactly one bid (i.e., no free disposal; Nisan, 1999), and the constraints of all awarded bids must be satisfied in the final work schedule.
- When the customer awards a bid, it pays a deposit and specifies the work schedule.

- When the supplier completes a task, the customer pays the remainder of the price.
- If the supplier fails to complete a task, the price is forfeited and the deposit must be returned to the customer. A penalty may also be levied for nonperformance, or a leveled-commitment protocol (Sandholm, 1996) may be used. The customer decides whether to handle the failure by replanning or rebidding the failed task(s).

2.1. A motivating example

As an example, imagine that we need to construct a garage. Figure 2 shows the tasks needed to complete the construction. The tasks are represented in a *task network*, where links indicate precedence constraints. The first decision we are faced with is how to sequence the tasks in the RFQ and how much time to allocate to each of them. For instance, we could reduce the number of parallel tasks, allocate more time to tasks with higher variability in duration or for which there is a shortage of laborers, or allow more slack time.

A sample RFQ is shown in Figure 2. Note that the time windows in the RFQ do not need to obey the precedence constraints; the only requirement is that the accepted bids obey them. We assume that the supplier is more likely to bid, and submit a lower-cost bid, if it is given a greater flexibility in scheduling its resources. It is up to the customer to find a bid combination that forms a feasible schedule.

2.2. Experiences and observations

We have shown (Collins et al., 2001) that the time constraints specified in the RFQ can affect the customer’s outcome in two major ways:

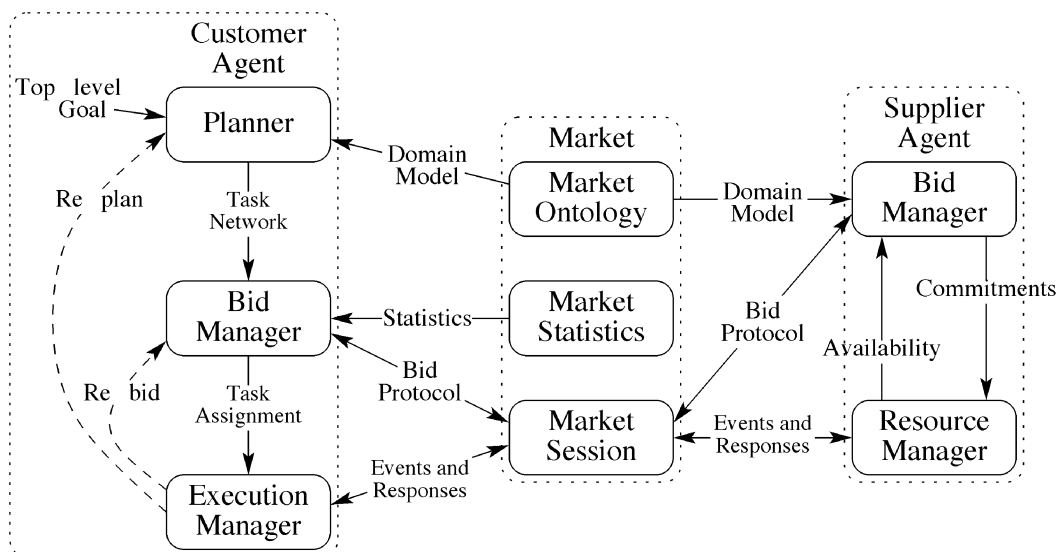


Fig. 1. The MAGNET architecture.




Fig. 2. A task network example and a corresponding RFQ.

1. by affecting the number, price, and time windows of bids. We assume that bids will reflect supplier resource commitments, and therefore larger time windows will result in more bids and better utilization of resources, in turn leading to lower prices (Collins, 2002). However, an RFQ with overlapping time windows makes the process of winner determination more complex (Collins, 2002). Another less obvious problem is that every extra bid over the minimum needed to cover all tasks adds one more rejected bid. Ultimately, a large percentage of rejections will reduce the customer agent’s credibility, which, after repeated interactions in the market, will result in fewer bids and/or higher costs.
2. by affecting the financial exposure of the customer agent (Collins et al., 2001). We assume nonrefundable deposits are paid to secure awarded bids, and payments for each task are made as the tasks are completed. The payoff for the customer occurs only at the completion of all the tasks. Once a task is completed in the time period specified, the customer is liable for its full cost, regardless of whether other tasks have failed in the meantime. If a task is not completed by the supplier, the customer is not liable for its cost, but this failure can ruin other parts of the plan. Slack in the schedule increases the probability that tasks will be completed or that there will be enough time to recover if any fail. However, slack extends the completion time and so reduces the payoff. In many business situations, the speed is the key; the value of the final payoff may drop off very steeply with time.

The agent needs to issue the RFQ before having received any bid, so the process of deciding how to schedule the different tasks and how much time to allocate to each task involves a probability of loss, as well as a probability of gain. This requires the decision process to deal with the risk posture of the person or organization on whose behalf the agent is acting.

The agent can use information available from the market on expected costs, probability of completion of tasks within a time window, and expected numbers of bidders to guide its decision on how to sequence the tasks and how much time to allocate to each.

In the next section we propose a principled method for generating RFQs that takes into account the agent’s risk posture and available market statistics to produce a schedule that optimizes the agent’s expected utility.

3. EXPECTED UTILITY APPROACH

In this section we describe a new approach to the construction of optimal RFQs that employs the Expected Utility Theory to reduce the likelihood of receiving unattractive bids while maximizing the number of bids that are likely to be awarded. This approach was originally suggested in our previous work (Babanov, Collins, et al., 2002). In this work we extend it and pay special attention to the relation between the size of RFQ time windows and the number of expected bids by investigating the balance between the quantity and the quality of expected bids.

3.1. Terminology

A *task network* (see Fig. 2) is a tuple $\langle N, < \rangle$ of a set N of individual tasks and strict partial ordering on them, such that for any $i, j \in N, i < j$ implies that task i immediately precedes task j . We also use N to denote the number of tasks where appropriate.

A task network is characterized by a *start time* (t^s) and a *finish time* (t^f), which delimit the interval of time when tasks can be scheduled. The placement of task n in the schedule is characterized by *task n start time* (t_n^s) and *task n finish time* (t_n^f), subject to the following constraints:

$$t^s \leq t_m^f \leq t_n^s, \quad \forall m \in P_1(n) \quad \text{and} \quad t_n^f \leq t_m^s \leq t^f, \quad \forall m \in S_1(n),$$

where $P_1(n)$ is the a set of *immediate predecessors* of n , $P_1(n) = \{m \in N | m < n\}$, and $S_1(n)$ is defined similarly to be the set of *immediate successors* of task n .

The *probability of task n completion* by time t , conditional on the ultimate successful completion of task n , is distributed according to the cumulative distribution function (CDF) $\Phi_n = \Phi_n(t_n^s; t)$, $\lim_{t \rightarrow \infty} \Phi_n(t_n^s; t) = 1$. Observe that Φ_n is defined to be explicitly dependent on the t_n^s . To see the rationale, consider the probability of successful mail delivery in x days for packages that were mailed on different days of a week.




Fig. 3. The unconditional distribution for the successful completion probability.

There is an associated unconditional *probability of success* $p_n \in [0, 1]$ characterizing the percentage of tasks that are successfully completed given infinite time (see Fig. 3). In the empirical support of work we assumed Weibull probability distribution for Φ_n ; however, the form of the distribution is not tied in the theory. In fact, we expect that the success probabilities will be derived from the available market information.

Task n bears an associated *cost*.¹ We assume the total cost of task n has two parts: a deposit, which is paid when the bid is accepted, and a cost c_n , which is due some time after successful completion of n . Because deposits are assumed to be paid up front, the amount does not change between schedules and we can assume without loss of generality the sum of deposits to be 0.

There is a single *final reward* V scheduled at the plan finish time t^f and paid conditional on all tasks in N being successfully completed by that time.

For each cost and reward, there is an associated *rate of return*² q_n that is used to calculate the *discounted present value* (PV) for payoff (c_n) due at time t as

$$PV(c_n; t) := c_n(1 + q_n)^{-t}.$$

We associate the rate of return q with the final payoff V .

3.2. Expected utility and certainty equivalent

We represent the customer agent’s preferences over payoffs by the von Neumann–Morgenstern utility function u (Mas–Colell et al., 1995). We further assume that the absolute risk-aversion coefficient $r := -u''/u'$ of u is constant for any value of its argument; hence, u can be represented as follows:

$$u(x) = -\exp\{-rx\} \text{ for } r \neq 0 \text{ and } u(x) = x \text{ for } r = 0.$$

It is imperative to note here that we do not compare utility values directly; the counterintuitive (i.e., decreasing in mon-

etary terms) form of the utility for $r < 0$ is a trade-off for simple notation.

We assume that a future state of the world can be described in terms of potential money transfers and the corresponding probabilities. Accordingly, we define *gamble* to be a set of payoff–probability pairs $G = \{(x_i, p_i)\}_i$ s.t. $p_i > 0, \forall i$ and $\sum_i p_i = 1$. The expectation of the utility function over a gamble G is the *expected utility* (EU):

$$EU[G] := \sum_{(x_i, p_i) \in G} p_i u(x_i).$$

The *certainty equivalent* (CE) of a gamble G is defined as the single payoff value whose utility matches the expected utility of the entire gamble G , that is, $u(CE[G]) := EU[G]$. Hence, under our assumptions

$$CE(G) = \frac{-1}{r} \log \sum_{(x_i, p_i) \in G} p_i \exp\{-rx_i\} \text{ for } r \neq 0$$

and

$$CE(G) = \sum_{(x_i, p_i) \in G} p_i x_i \text{ for } r = 0.$$

Our evaluation criterion is based upon comparing CE values, because they represent money transfers in certain and current money. Because of this interpretation the CE values, unlike utilities, can be compared across various risk aversities and alternative schedules, even between different plans. Naturally, the agent will not be willing to accept gambles with a negative certainty equivalent and higher values of the certainty equivalent will correspond to more attractive gambles.

To illustrate the concept, Figure 4 shows how the CE depends on the risk aversity r of an agent. In this figure we consider a gamble that brings the agent either 100 or nothing with equal probabilities. Agents with positive r values are risk averse; those with negative r values are risk loving. Agents with risk aversity close to zero, that is, almost risk neutral, have a CE equal to the gamble’s weighted mean (50).

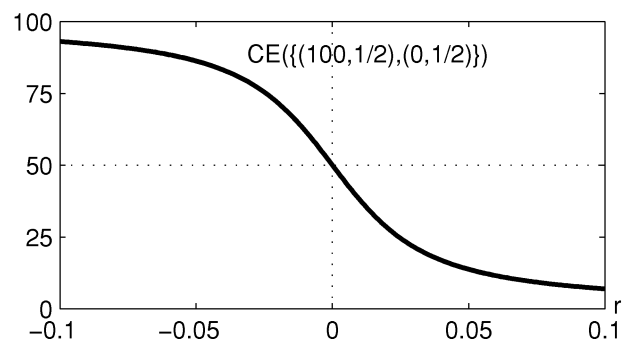


Fig. 4. The CE of a simple gamble as a function of the risk aversity.

¹Hereafter we use the words “cost” and “reward” to denote some monetary value, while referring the same value as “payment” or “payoff” whenever it is scheduled at some time t .

²The reason for having multiple q_n values is that individual tasks can be financed from different sources, thus affecting task scheduling.




Fig. 5. The CE maximizing time allocations for the plan in Figure 2 for $r = -0.01$ (left) and $r = 0.02$ (right).

3.3. Cumulative probabilities

To compute the certainty equivalent of a gamble we need to determine a schedule for the tasks and compute the payoff probability pairs.

We assume that the c_n for task n is scheduled at t_n^f , so its present value³ \tilde{c}_n is

$$\tilde{c}_n := c_n(1 + q_n)^{-t_n^f}.$$

We define the conditional probability of task n success as

$$\tilde{p}_n := p_n \Phi_n(t_n^s; t_n^f).$$

We also define the *precursors* of task n as a set of tasks that finish before task n starts in a schedule, that is,

$$\tilde{P}(n) := \{m \in N | t_m^f \leq t_n^s\}.$$

The unconditional probability that task n will be completed successfully is

$$\tilde{p}_n^c = \tilde{p}_n \times \prod_{m \in \tilde{P}(n)} \tilde{p}_m.$$

That is, the probability of successful completion of every precursor and of task n itself are considered as independent events. The reason this is calculated in such form is because, if any task in $\tilde{P}(n)$ fails to be completed, there is no need to execute task n .

The probability of receiving the final reward V is therefore

$$\tilde{p} = \prod_{n \in N} \tilde{p}_n^c.$$

³Hereafter we use the tilde to distinguish variables that depend on the current task schedule, while omitting corresponding indices for the sake of notational simplicity.

3.4. Example and discussion

To illustrate the definitions above, let us return to the task network in Figure 2 and consider the sample task schedules shown in Figure 5. In this figure the x axis is time and the y axis shows both the task numbers and the cumulative distribution of the unconditional probability of completion (compare to Fig. 3). Circle markers show t_n^s . Crosses indicate both t_n^f and success probabilities \tilde{p}_n (numbers next to each point). Square markers denote that the corresponding task cannot span past this point due to precedence constraints. The thick part of each CDF shows the time allocated to each task.

The customer agent needs a way of collecting the market information necessary to build and use the probability model. The probability of success is relatively easy to observe in the market. This is the reason for introducing the cumulative probability of success Φ_n and probability of success p_n , instead of the average project life span or probability of failure. Indeed, it is rational for the supplier to report a successful completion immediately in order to maximize the present value of a payment. In addition, it is rational not to report a failure until the last possible moment, because of a possibility of earning the payment by rescheduling, outsourcing, or fixing the problem in some way.

3.5. Gamble calculation algorithm and maximization

Given a schedule like the one shown in Figure 5, we need to compute the payoff probability and then maximize the CE for the gamble. Writing an explicit description of the expected utility as a function of gambles is overly complicated and relies on the order of task completions. Instead, we propose a simple recursive algorithm that creates these gambles. We then maximize the CE over the space of all feasible schedules and the corresponding gambles.

```

Algorithm:  $G \leftarrow \text{calcGamble}(T, D)$ 
  Requires:  $T$  "tasks to process,"  $D$  "processed tasks"
  Returns:  $G$  "subtree gamble"
 $M \leftarrow \{m \in T \mid \bar{P}(m) \subset D\}$ 
if  $M \neq \emptyset$  "it's a branch"
   $n \leftarrow \text{first}\{M\}$  "according to some ordering"
   $T \leftarrow T \setminus \{n\}$ 
   $G \leftarrow \emptyset$ 
   $E \leftarrow \text{calcGamble}(T, D)$  "follow  $\dots \rightarrow \bar{n}$  path"
  forall  $(x, p) \in E$ 
     $G \leftarrow G \cup \{(x, p \times (1 - \bar{p}_n))\}$ 
  endfor
   $I \leftarrow \text{calcGamble}(T, D \cup \{n\})$  "follow  $\dots \rightarrow n$  path"
  forall  $(x, p) \in I$ 
     $G \leftarrow G \cup \{(x + \tilde{c}_n, p \times \bar{p}_n)\}$ 
  endfor
  return  $G$  "subtree is processed"
else "it's a leaf"
  if  $N = D$  "all tasks are done"
    return  $\{(V, 1)\}$ 
  else "some task failed"
    return  $\{(0, 1)\}$ 
  endif
endif
endif
    
```

In the first call, the algorithm receives a "todo" task list $T = N$ and a "done" task list $D = \emptyset$. All the subsequent calls are recursive. To illustrate the idea behind this algorithm, we refer to the payoff-probability trees in Figure 6. These two

trees were built for the time allocations in Figure 5 and reflect the precursor relations for each case.

Considering the more sequential schedule on the right, we note that with probability $1 - \bar{p}_1$ ask 1 fails, and the customer agent does not pay or receive anything and stops the execution (path $\bar{1}$ in the right tree). With probability $\bar{p}_1^c = \bar{p}_1$ the agent proceeds with task 3 (path 1 in the tree). In turn, task 3 either fails with probability $\bar{p}_1 \times (1 - \bar{p}_3)$, in which case the agent ends up stopping the plan and paying a total of c_1 (path $1 \rightarrow \bar{3}$), or it is completed with the corresponding probability $\bar{p}_3^c = \bar{p}_1 \times \bar{p}_3$. In the case where both 1 and 3 are completed, the agent starts both 2 and 4 in parallel and becomes liable for paying c_2 and c_4 , respectively, even if the other task fails (paths $1 \rightarrow 3 \rightarrow 2 \rightarrow \bar{4}$ and $1 \rightarrow 3 \rightarrow \bar{2} \rightarrow 4$). If both 2 and 4 fail, the resulting path in the tree is $1 \rightarrow 3 \rightarrow \bar{2} \rightarrow \bar{4}$ and the corresponding payoff-probability pair is framed in the figure.

The algorithm's complexity is $O(2^{K-1} \times N)$, where K is the maximum number of tasks that are scheduled to be executed in parallel. Reducing the complexity of *calcGamble* is critical, because it will be executed in the inner loop of any CE maximization procedure, unless we somehow fix precursor relations and, consequently, a tree structure. In commercial projects the ratio K/N is likely to be low, because few of these exhibit a high degree of parallelism. Our preliminary experiments allow us to conclude that the K/N ratio is lower for risk-averse agents (presumably, business-people) than for risk lovers (gamblers). These two considerations may reduce the need for a faster algorithm, although additional work to improve the algorithm is planned.




Fig. 6. Payoff-probability trees for the time allocations in Figure 5.

3.6. Experimental results

We have conducted a set of experiments on CE maximization in a variety of task networks. Some of the results for our reference six-task network are summarized in Figure 7. In this figure, the y axis shows 11 different risk-aversity r settings, and the bottom x axis is time t in the plan. The rounded horizontal bars in each of 11 sections denote time allocations for each of six tasks, with task 1 on top. Sections $r = -0.01$ and $r = 0.02$ correspond to schedules in Figure 5a and b, respectively. Finally, the vertical bars show the maximum CE values on the top x axis for each r setting.

Let us examine the relative placement of time allocations as a function of r . For this example we chose two tasks, 3 and 4, which have similar positions in the task network. Task 3 (black horizontal bars) has a lower probability of success in the infinite horizon than task 4 (white bars), as well as a higher variance of the probability of success distribution. In addition, the cost of task 3 is twice the cost of task 4.

Given this setup, we observed four distinct cases in the experimental data:

1. Risk-loving agents tend to schedule tasks in parallel and late in time in order to maximize the present value

of the expected difference between reward and payoffs to suppliers. This confirms the intuitive observation from Figure 4 that risk lovers lean toward receiving high, risky payoffs rather than low, certain payoffs.

2. Risk neutral and minimally risk-averse agents place risky task 3 first to make sure that any failure does not happen too far into the project. Note that they still keep task 2 in parallel so in case 2 fails, they are liable for paying the supplier of task 4 on success. One can consider those agents as somewhat optimistic.
3. Moderately risk-averse agents try to dodge the situation above by scheduling task 3 after task 2 is finished. These agents are willing to accept the plan, but their expectations are quite pessimistic.
4. Highly risk-averse agents shrink the task 1 interval to zero, thus “cheating” to avoid covering any costs. One may interpret this as a way of signaling a refusal to accept the plan. Indeed, the assignment of zero duration to a precursor-less task ensures zero probability of completion and hence zero CE, even in the cases where any nondegenerate schedule has a negative CE value.




Fig. 7. The CE maximizing schedules and CE values for the plan in Figure 2 and $r \in [-0.03, 0.07]$.




Fig. 8. The $CE(t_n^s)$ graphs for the corresponding reasonable schedules in Figure 5.

4. GENERATION OF RATIONAL RFQ

In the previous sections we have shown a way of generating a CE maximizing schedule of task execution, which we hereafter refer to as the *reasonable schedule*. For a chosen risk-aversity value and known market statistics, the reasonable schedule ensures the highest expected quality of the bids that satisfy it. By *quality*, we assume some function of the expectations over the cost, the probability of successful completion, and the profitability of the incoming bids in their feasible combinations with other bids.

The reasonable schedule, however, cannot serve as a rational RFQ, because it is unlikely that bids will be available to cover precisely the same intervals as mandated by the CE maximizing schedule. In order to construct a viable RFQ using the reasonable schedule as a basis, the customer agent might choose to lower its expectations of the bid quality to some level by widening the RFQ time windows around the time windows in the reasonable schedule, thus increasing⁴ the expected number of incoming bids. In this section we discuss criteria that allow for rationalizing the selection among all such RFQs.

We approach the individually rational (i.e., agent-dependent) RFQ generation as follows:

1. Measure the sensitivity of the expected bid quality to deviations from the CE maximizing schedule.
2. Derive the relationship between the quality of incoming bids and the size of RFQ time windows.
3. Choose a rational quality–quantity combination.

In addition, we search for a solution concept that generates viable RFQs and is comprehensible to a human user of the system.

⁴At least to some extent, because there is a fair chance that the number of the incoming bids will cease to increase whenever RFQ time windows become too large to inspire confidence on the part of suppliers.

4.1. CE Sensitivity to schedule changes

We propose measuring the sensitivity of CE by investigating how CE values change with variations of a single task n start time (t_n^s) in the reasonable schedule. For the sake of brevity the resulting dependency of CE values is denoted by $CE(t_n^s)$. Figure 8 shows $CE(t_n^s)$, $n = 1 \dots 6$ for our six-task sample problem for $r = -0.01$ and 0.02 , respectively. In the figure, the y axis of each horizontal stripe n represents the percentage of the maximum CE value as t_n^s varies, the x axis represents t_n^s , and the horizontal lines with circle and cross ends show the corresponding reasonable schedules.

Tasks 1, 3, and 5 in the right graph are relatively restrictive to the start times of the bids that can be bundled with the reasonable bids without considerably impairing the resulting bundle’s value. However, the fact that task 2 in the right graph is more flexible does not guarantee that it will attract a higher number of bids, because the latter depend both on the size of the corresponding time window and on the market properties of the task: resource availability, number of prospective bidders, seasonal changes, and so forth.

We assert that for the purpose of creating a rational RFQ, it is admissible to choose time windows based on the sensitivity of CE to deviations of a single time constraint from the reasonable schedule. The rationale is that the relations between tasks are already encapsulated in the calculations of CE, so the change of one constraint will approximate the rescheduling of several related tasks in the neighborhood of the reasonable schedule.

4.2. Quality versus quantity

Observe that the time window for the task n , $\{t_n^s | CE(t_n^s) \geq x\}$, grows as the lowest expected CE value (x) decreases. The relation between these two variables for the tasks 3 and 4 of the test problem is shown in Figure 9. The corresponding relation between the lowest expected CE value and the expected number of bids as a function of the window size is




Fig. 9. The relationship between the RFQ window size (units of time on x axis) and the lowest admissible percentage of the maximum CE value.

shown in Figure 10. In the right graph we assumed, for the sake of example, that the supply of task 3 in the market is higher than task 4, hence the difference in relative positions of task 3 and task 4 graphs in the two figures.

The graph in Figure 10 is an example of the relation between the quality and the quantity of bids for which we are searching. Indeed, the only independent variable in this graph is t_n^s . The quantity of bids depends on the size and positions of RFQ time windows that, in turn, depend on the decision about the lowest acceptable CE value. The quality of bids is a function of the RFQ choice and the properties of the plan. Finally, it is expected that the customer agent will prefer a point on the graph to any point below and to the left of it; thus, the best choice should lie on the graph.

4.3. Rational quality–quantity choice and RFQ

We illustrate the decision process of the customer agent in Figure 11, where the customer agent’s preferences over

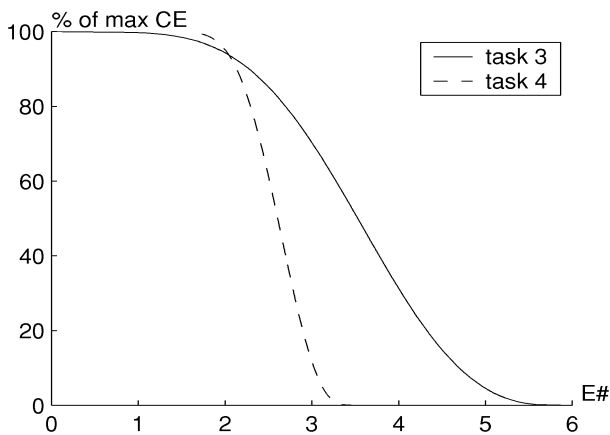


Fig. 10. The relationship between the expected number of bids (x axis) and the lowest admissible percentage of the maximum CE value.

quality–quantity combinations are represented by a family of indifference curves and the graph of underlying quality–quantity relationship is as derived in the previous section. Each indifference curve shows quality–quantity pairs that are equivalent from the agent’s point of view. In particular, points A, B, and C are considered to be equally attractive. The intuition is that, although in point A the agent receives a much smaller number of bids to compare with C and is exposed to a higher risk of not covering some tasks, this is offset by a positive effect of a lower percentage of bid rejections on the agent’s reputation. In addition, the winner determination problem is exponentially easier to solve (Collins, 2002) for the lower expected number of bids in point A.

For all points below the maximum expected quality line (Fig. 11, dashed line) agent’s preferences increase in the direction of point M. Thus, a curve through point D is preferred over one through point C and even more so over one through point E. The rational choice belongs to the intersection of the quality–quantity graph and the highest indifference curve (point B).

After the rational choice of the quality–quantity combinations for all tasks in the plan is revealed, we proceed with constructing the RFQ time windows. The choice of early start time (t_n^{es}) and late start time (t_n^{ls}) are determined by the value of the reciprocal of the CE (t_n^s) at the minimum admissible CE choice for task n . The late finish time (t_n^{lf}) is chosen to be at the reasonable time window length distance from t_n^{ls} . Figure 12 shows two sample RFQs for the garage building example. The start time intervals ($[t_n^{es}, t_n^{ls}]$), late finish time t_n^{lf} , and the corresponding reasonable schedules are presented in the figure.

Our choice of the RFQ may not be optimal in the quantitative sense. However, it is individually rational for the customer agent, fast to compute, and, arguably, easily grasped by a human user of the system. It should be emphasized here that the choice of the RFQ is based on the uncertain market information; hence, any quantitatively “optimal” solution is itself a compromise.

5. OPEN ISSUES AND FURTHER RESEARCH

In this section we outline two major issues that arise when we employ the expected utility approach to generate rational RFQs. The first issue concerns the CE maximization in the domain with temporal and precedence constraints. The second issue is the assessment of the EU approach and, ultimately, the MAGNET system itself in the absence of the real-world data for the domain of interest.

5.1. Multiple local maxima

One of the most important issues related to CE maximization is the presence of multiple local maxima of CE, even in cases where task networks are fairly simple. We argue that this property is partially due to the relative positioning of the tasks off the critical path. Any two tasks that are not




Fig. 11. A quality–quantity graph with three (—) indifference curves; (---) the maximum expected quality line.

ordered by the precedence constraints can be scheduled in three ways: parallel and two sequential. Scheduling tasks in parallel increases the probability of successful completion, whereas sequential scheduling minimizes overall payments in case one of the tasks fails. In cases where extra slack allows for sequential scheduling, it turns out that parallel and sequential positionings of two independent individual tasks lead to similar resulting CE values.

To illustrate the issue, we constructed a sample task network with two parallel tasks. Task 1 has a higher variance of completion time probability and lower probability of success than task 2; everything else is the same. The resulting graph of CE is shown in Figure 13. There are three local maxima with positive CE values in this figure: one that corresponds to task 2 being scheduled first in sequential order, another on the right side corresponding to task 1 being first, and yet another representing both tasks being scheduled at time 0 and executed in parallel. The number of local maxima grows considerably with the number of the tasks that are not restricted by the precedence relationship.

5.1.1. Domain and algorithm properties

The following list shows the properties of the domain that influence the search algorithm design:

- local maxima are due to different scheduling order of tasks off the critical path;
- groups of local maxima have similar CE values; and
- an RFQ based on the global maximum can be overly restrictive.

The properties of the domain frame the properties of the search algorithm that we design to fit this domain. Namely, the search algorithm must be able to test different orderings of tasks, should know how to explore groups of similar local maxima, and, whenever possible, should provide alternative schedules with CE values close to the global maximum.

We propose a search algorithm based on the ideas of simulated annealing (Reeves, 1993) and genetic algorithms (Forrest, 1993). The algorithm will combine the stochastic, temperature-driven nature of simulated annealing with the simultaneous search space exploration of genetic algorithms. In this section we describe the proposed algorithm in more detail and explain the rationale of its design.

5.1.2. Search algorithm

The proposed search algorithm explores several alternative schedules in parallel. The initial set of alternatives can be generated in many ways: random generation, hill climb-




Fig. 12. The rational RFQs for the corresponding reasonable schedules in Figure 5 and the maximum CE percentages of 80, 95, 50, 70, 50, and 90%. (■) start time intervals, (□) late finish times, and (○—×) corresponding reasonable schedules.

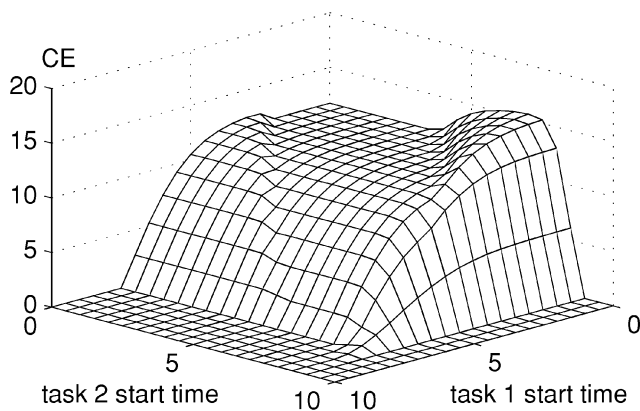


Fig. 13. The local maxima for two parallel tasks.

ing from a random schedule, critical path method, and so forth. The execution of the algorithm proceeds in steps by randomly applying one of the six transformation rules to each alternative schedule. The probabilities of choosing some rule can be adjusted to adjust the algorithm's behavior over a wide range.

Figure 14 illustrates the algorithm for the case of three pairwise independent tasks. In this figure the columns represent three consecutive states of the algorithm, and each column lists several alternative schedules. Arrows and letters next to the columns denote various transformations from the following list:

Distortion alters start and finish times of one or several tasks and adjusts the time windows of all related tasks to maintain precedence constraints (1 → 5, Fig. 14). Distortion mimics the basic step of the simulated annealing algorithm in continuous $2N$ -dimensional space of task start and finish times.

Gradient following is one or more steps of a generic numerical maximization method (5 → 9). This step is very computationally intensive, because it requires many calls to the calcGamble algorithm in the process of calculating numerical derivatives. By varying the relative probabilities of the distortion and gradient following, we may choose the stochastic properties of the proposed approach.

Shuffling changes relative scheduling of two or more tasks wherever it is permitted by the precedence constraints. It can switch ordering of the tasks (6 → 10), change the sequential ordering to parallel, or reschedule parallel tasks to be executed sequentially (4 → 8). The major role of shuffling is to explore local maxima that have similar CE values due to different scheduling of tasks off the critical path.

Explosion adds a copy of the subject schedule to the list of alternatives (2 → 6, 7). Explosion compliments shuffling by allowing for simultaneous exploration of the

groups of similar schedules. We may choose to decrease the rate of explosions with the annealing temperature to focus on improving the current set of solutions after the search space has been explored to some extent.

Implosion merges two similar⁵ schedules in one. It helps to reduce computational expenses from crowding several alternative schedules around one maximum (7, 8 → 11). The rate of implosions will change in the direction opposite to the rate of explosions.

Removal eliminates alternatives that do not score well relative to others (3 → \emptyset). This transformation takes care of the schedules that are stuck in local maxima with low CE values. The rate of removals grows as the annealing temperature decreases.

Each of the first five transformations is tested against the simulated annealing temperature rule whenever it leads to a decrease in the CE value. In case it is discarded, other transformations are chosen at random and applied until one of them increases the CE or passes the temperature rule.

The probabilities of transformations and the details of the proposed search algorithm's properties are subject to further research. However, it is reasonable to believe that the comprehensive study of the RFQ generation mechanism is only possible in the dynamic market environment. In the next section we outline the approach to the large-scale testing of the MAGNET system that we are presently researching and that will provide us with the necessary data.

5.2. Evolutionary framework for large-scale testing

We plan to devote efforts to thoroughly test the suggested CE-based approach of rational RFQ generation. In particular, we are interested in testing how well individual agents interact in a populated market. The major goals of this part of the study would be the following:

- to provide the statistical data necessary for the evaluation of the theoretical assumptions and derivations;
- to facilitate the understanding of the nuances of CE-based RFQ generation and suggest improvements to the theory and implementation; and
- to study the relative performance of agents in a simulated market, developing an understanding of the properties of automated and mixed-initiative combinatorial auction-based trading societies.

The most compelling approach would be to gather a rich set of statistical data from a commerce domain. That has not proven to be feasible for two reasons. First, few industrial organizations are sufficiently open to expose the type of data you would need to do that, and we would need data

⁵Similarity is a function of the distance between two schedules, as between two points in the $2N$ -dimensional time space.




Fig. 14. The two steps of the search algorithm execution.

from multiple organizations in a single market. Second, data is gathered to serve a purpose, and our experience tells us that when you attempt to apply existing data to a new purpose, it frequently turns out to be full of inconsistencies and methodological problems.

In lieu of using real industry data, we are designing a large-scale test suite atop an abstract domain with controllable statistics, based on the evolutionary approach to economic simulation. Evolutionary frameworks have been used extensively in economics (Nelson, 1995; Rode, 1997; Testfatsion, 2001). The framework will allow us to tune the market by tweaking the frequency of issuing RFQs and will allow for the dynamic introduction of new supplier strategies without imposing any assumptions on the nature of strategies. We will later extend the framework to support trade games played by human subjects. This will be an especially useful tool for teaching and a tool to explore strategic behaviors and study the emergence of cooperation (Axelrod, 1984; Axelrod, 1997).

We suggest detailed reasoning behind our choice of the evolutionary approach and describe experimental results from a pilot model of a trading agent society in our related research (Babanov, Ketter, et al., 2002).

6. RELATED WORK

Expected Utility Theory (Pratt, 1964) is a mature field of economics that has attracted many supportive and critical studies, both theoretical (Machina, 1987, 1989) and empirical (Smith & Desvousges, 1987; Jullien & Salanié, 2000). We believe that expected utility will play an increasing role in automated auctions, because it provides a practical way of describing risk estimations and temporal preferences.

Despite the abundance of work in auctions (McAfee & McMillan, 1987), limited attention has been devoted to auctions over tasks with complex time constraints and interdependencies. Parkes and Ungar (2001) propose a method

to auction a shared track line for train scheduling. The problem is formulated with mixed integer programming with many domain-specific optimizations. Bids are expressed by specifying a price to enter a line and a time window. Time slots are used in Wellman et al. (2001), where a protocol for decentralized scheduling is proposed. The study is limited to scheduling a single resource. MAGNET agents deal with multiple resources.

Most work in supply-chain management is limited to hierarchical modeling of the decision making process, which is inadequate for distributed supply chains, where each organization is self-interested rather than cooperative. Walsh et al. (2000) propose a protocol for combinatorial auctions for supply chain formation, using a game-theoretical perspective. They allow complex task networks but do not include time constraints. MAGNET agents must also ensure the scheduling feasibility of the bids they accept and evaluate risk. Agents in MASCOT (Sadeh et al., 1999) coordinate scheduling with the user, but there is no explicit notion of payments or contracts and no explicit statement of the criteria for accepting/rejecting a bid. Their major objective is to show policies that optimize schedules locally (Kjenstad, 1998), whereas our objective is to optimize the customer's utility.

Different heuristics for scheduling are proposed in Dutta et al. (2001). The strategies are intended for supplier agents that are trying to adjust their schedules to win new awards. In the work presented here, we are concerned with the scheduling that customers need to do before requesting bids.

In MAGNET the agents interact with each other through a market. The market infrastructure provides a common vocabulary, collects statistical information that helps agents estimate costs, schedules, and risks, and acts as a trusted intermediary during the negotiation process. The market also acts as a matchmaker (Sycara et al., 1997), allowing us to ignore the issue of how agents will find each other. For a survey on the use of intelligent agents in manufacturing, see Shen and Norrie (1999).

The determination of winners of combinatorial auctions (Rothkopf et al., 1998) is difficult. Dynamic programming (Rothkopf et al., 1998) works well for small sets of bids, but it does not scale and imposes significant restrictions on the bids. Algorithms such as CABOB (Sandholm et al., 2001), Bidtree (Sandholm, 2000), and CASS (Fujishima et al., 1999) reduce the search complexity. Reeves et al. (2001) use auction mechanisms to “fill in the blanks” in prototype declarative contracts that are specified in a language based on Courteous Logic Programming (Grosz et al., 1999). These auctions support bidding on many attributes other than price, but the problem of combining combinatorial bids with side constraints is not addressed.

Combinatorial auctions are becoming an important mechanism not only for agent-mediated e-commerce (Guttman et al., 1998; Wurman et al., 1998; Sandholm, 1999) but also for the allocation of tasks to cooperative agents (e.g., Dias & Stentz, 2000; Hunsberger & Grosz, 2000).

Hunsberger and Grosz (2000) combinatorial auctions for the initial commitment decision problem, which is the problem an agent must solve when deciding whether to join a proposed collaboration. Their agents have precedence and difficult temporal constraints. However, to reduce search effort, they use domain-specific *roles*, a shorthand notation for collections of tasks. In their formulation, each task type can be associated with only a single role. MAGNET agents are self-interested, and there are no limits to the types of tasks they can decide to do. In Glass and Grosz (2000) scheduling decisions are made, not by the agents, but instead by a central authority. The central authority has insight to the states and schedules of participating agents, and agents rely on the authority for supporting their decisions.

Leyton-Brown et al. (2000) suggest a way of constructing a universal test suite for winner determination algorithms in combinatorial auctions. Their work does not include cases with precedence and time constraints and, thus, is not directly applicable to the MAGNET framework. It nevertheless provides well-understood test cases for comparing the performance of algorithms.

7. CONCLUSIONS

Auction mechanisms are an effective approach to negotiation among groups of self-interested economic agents. We are particularly interested in situations where agents need to negotiate over multiple factors, including not only price, but task combinations and temporal factors as well.

We have shown how an agent can use information about the risk posture of its principal, along with market statistics, to formulate RFQs that optimize the trade-off between risk and value and increase the quality of the bids received. This requires deciding how to sequence tasks and how much time to allocate to each of them. Bids closest to the specified time windows are the most preferred risk-payoff combinations.

The work described here is a part of a larger effort at the University of Minnesota that aims to study how autonomous or semiautonomous agents can be used in complex commerce-oriented domains.

ACKNOWLEDGMENTS

We gratefully acknowledge the partial support of this research by the National Science Foundation (NSF/IIS-0084202).

REFERENCES

- Axelrod, R. (1997). *The Complexity of Cooperation*. Princeton, NJ: Princeton University Press.
- Axelrod, R.M. (1984). *The Evolution of Cooperation*. New York: Basic Books.
- Babanov, A., Collins, J., & Gini, M. (2002). Risk and expectations in a-priori time allocation in multi-agent contracting. *Proc. First Int. Conf. Autonomous Agents and Multi-Agent Systems*, vol. 1, pp. 53–60, Bologna, Italy.
- Babanov, A., Ketter, W., & Gini, M. (2002). An evolutionary framework for large-scale experimentation in multi-agent systems. In *Toward an Application Science: MAS Problem Spaces and Their Implications to Achieving Globally Coherent Behavior*. Bologna, Italy.
- Collins, J. (2002). *Solving combinatorial auctions with temporal constraints in economic agents*. PhD Thesis. University of Minnesota.
- Collins, J., Bilot, C., Gini, M., & Mobasher, B. (2001). Decision processes in agent-based automated contracting. *IEEE Internet Computing* 5(2), 61–72.
- Collins, J., Ketter, W., & Gini, M. (2002). A multi-agent negotiation tested for contracting tasks with temporal and precedence constraints. *International Journal of Electronic Commerce* 7(1), 35–57.
- Dias, M.B., & Stentz, A. (2000). A free market architecture for distributed control of a multirobot system. *Sixth Int. Conf. Intelligent Autonomous Systems*, pp. 115–122, Venice, Italy.
- Dutta, P.S., Sen, S., & Mukherjee, R. (2001). Scheduling to be competitive in supply chains. *IJCAI workshop on E-Business and the Intelligent Web*.
- Forrest, S. (1993). Genetic algorithms: Principles of natural selection applied to computation. *Science* 261, 872–878.
- Fujishima, Y., Leyton-Brown, K., & Shoham, Y. (1999). Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. *Proc. 16th Int. Joint Conf. Artificial Intelligence*, pp. 548–553.
- Glass, A., & Grosz, B.J. (2000). Socially conscious decision-making. *Proc. Fourth Int. Conf. Autonomous Agents*, pp. 217–224.
- Grosz, B.N., Labrou, Y., & Chan, H.Y. (1999). A declarative approach to business rules in contracts: Courteous logic programs in XML. *Proc. ACM Conf. Electronic Commerce (EC'99)*, pp. 68–77. New York: ACM.
- Guttman, R.H., Moukas, A.G., & Maes, P. (1998). Agent-mediated electronic commerce: A survey. *Knowledge Engineering Review* 13(2), 143–152.
- Hunsberger, L., & Grosz, B.J. (2000). A combinatorial auction for collaborative planning. *Proc. 4th Int. Conf. Multi-Agent Systems*, pp. 151–158. Boston: IEEE Computer Society Press.
- Jullien, B., & Salanié, B. (2000). Estimating preferences under risk: The case of racetrack bettors. *The Journal of Political Economy* 108(3), 503–530.
- Kjenstad, D. (1998). *Coordinated supply chain scheduling*. PhD Thesis. Norwegian University of Science and Technology.
- Leyton-Brown, K., Pearson, M., & Shoham, Y. (2000). Towards a universal test suite for combinatorial auction algorithms. *Proc. ACM Conf. Electronic Commerce (EC'00)*, pp. 66–76, Minneapolis, MN.
- Machina, M.J. (1987). Choice under uncertainty: Problems solved and unsolved. *The Journal of Economic Perspectives* 1(1), 121–154.
- Machina, M.J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *The Journal of Economic Literature* 27(4), 1622–1668.
- Mas-Colell, A., Whinston, M.D., & Green, J.R. (1995). *Microeconomic Theory*. New York: Oxford University Press.
- McAfee, R., & McMillan, P.J. (1987). Auctions and bidding. *Journal of Economic Literature* 25, 699–738.

- Nelson, R.R. (1995). Recent evolutionary theorizing about economic change. *Journal of Economic Literature* 33(1), 48–90.
- Nisan, N. (1999). Bidding and allocation in combinatorial auctions. *Proc. ACM Conf. Electronic Commerce (EC'00)*, pp. 1–12, Minneapolis, MN. New York: ACM.
- Parke, D.C., & Ungar, L.H. (2001). An auction-based method for decentralized train scheduling. *Proc. Fifth Int. Conf. Autonomous Agents*, pp. 43–50, Montreal. New York: ACM.
- Pratt, J.W. (1964). Risk aversion in the small and in the large. *Econometrica* 32, 122–136.
- Reeves, C.R. (1993). *Modern Heuristic Techniques for Combinatorial Problems*. New York: Wiley.
- Reeves, D.M., Wellman, M.P., & Grosz, B.N. (2001). Automated negotiation from declarative contract descriptions. *Proc. Fifth Int. Conf. Autonomous Agents*, pp. 51–58. Montreal. New York: ACM.
- Rode, D. (1997). *Market efficiency, decision processes, and evolutionary games*. Carnegie Mellon University.
- Rothkopf, M.H., Pekeč, A., & Harstad, R.M. (1998). Computationally manageable combinatorial auctions. *Management Science* 44(8), 1131–1147.
- Sadeh, N.M., Hildum, D.W., Kjenstad, D., & Tseng, A. (1999). MASCOT: An agent-based architecture for coordinated mixed-initiative supply chain planning and scheduling. *Workshop on Agent-Based Decision Support in Managing the Internet-Enabled Supply-Chain, Agents '99*, pp. 133–138.
- Sandholm, T. (1999). An algorithm for winner determination in combinatorial auctions. *Proc. 16th Int. Joint Conf. Artificial Intelligence*, pp. 524–547.
- Sandholm, T. (2000). Approaches to winner determination in combinatorial auctions. *Decision Support Systems* 28(1–2), 165–176.
- Sandholm, T., Suri, S., Gilpin, A., & Levine, D. (2001). CABOB: A fast optimal algorithm for combinatorial auctions. *Proc. 17th Int. Joint Conf. Artificial Intelligence*, Seattle, WA, pp. 1102–1108.
- Sandholm, T.W. (1996). *Negotiation Among Self-Interested Computationally Limited Agents*. PhD Thesis. University of Massachusetts, Amherst.
- Shen, W., & Norrie, D.H. (1999). Agent-based systems for intelligent manufacturing: A state-of-the-art survey. *Knowledge and Information Systems* 1(2), 129–156.
- Smith, V.K., & Desvousges, W.H. (1987). An empirical analysis of the economic value of risk changes. *The Journal of Political Economy* 95(1), 89–114.
- Sycara, K., Decker, K., & Williamson, M. (1997). Middle-agents for the Internet. *Proc. 15th Joint Conf. Artificial Intelligence*, pp. 578–583.
- Tesfatsion, L. (2001). *Agent-Based Computational Economics: Growing Economies from the Bottom Up*. ISU Economics Working Paper No. 1. Ames, IA: Iowa State University.
- Walsh, W.E., Wellman, M., & Ygge, F. (2000). Combinatorial auctions for supply chain formation. *Proc. ACM Conf. Electronic Commerce (EC'00)*, pp. 260–269.
- Wellman, M.P., Walsh, W.E., Wurman, P.R., & MacKie-Mason, J.K. (2001). Auction protocols for decentralized scheduling. *Games and Economic Behavior* 35, 271–303.
- Wurman, P.R., Wellman, M.P., & Walsh, W.E. (1998). The Michigan Internet AuctionBot: A configurable auction server for human and software agents. *Second Int. Conf. Autonomous Agents*, pp. 301–308.

Alexander Babanov is a graduate student in economics and a PhD student in computer science at the University of Minnesota. Prior to joining the University of Minnesota, he obtained his MS degree in applied mathematics at the Moscow Institute of Physics and Technology. His research interests include applications of economic concepts, such as uncertainty and evolutionary game theory, to multiagent systems.

John Collins is an Assistant Professor of computer science at the University of Minnesota. Prior to beginning his academic career, he spent 30 years in industry in product development and industrial research. Dr. Collins' research interests are primarily in artificial intelligence, specifically in the decision problems of autonomous agents.

Maria Gini is a Professor in computer science and engineering at the University of Minnesota. She has been a Research Associate at the Politecnico of Milan, Italy, and a Visiting Research Associate at Stanford University. Her contributions span robot motion planning to robot coordination and multiagents negotiations. She is on the Editorial Board of *Autonomous Robots* and *Integrated Computer-Aided Engineering*. She was the General Chair for the 7th International Conference on Intelligent Autonomous Systems and for the 1st International Conference on Autonomous Agents and Multi-Agent Systems, both in 2002, and the Editorial Program Co-Chair for the 4th International Conference on Autonomous Agents in 2000.