

Agents Preferences in Decentralized Task Allocation

Mark Hoogendoorn^{1,2} and Maria L. Gini²

Abstract. The ability to express preferences for specific tasks in multi-agent auctions is an important element for potential users who are considering to use such auctioning systems. This paper presents an approach to make such preferences explicit and to use these preferences in bids for reverse combinatorial auctions. Three different types of preference are considered: (1) preferences for particular durations of tasks, (2) preferences for certain time points, and (3) preferences for specific types of tasks. We study empirically the tradeoffs between the quality of the solutions obtained and the use of preferences in the bidding process, focusing on effects such as increased execution time. We use both synthetic data as well as real data from a logistics company.

1 Introduction

Auctions are used in multi-agent systems, among other things, to perform allocation of tasks (see e.g. [13] and [14]). Such reverse auctions, where the buyer is the auctioneer, can be of a combinatorial type, allowing for bidding on bundles of tasks. Sandholm [12] notes that reverse auctions are not economically efficient because optimal bundling depends on suppliers preferences, which traditionally cannot be expressed. Enabling the agents to express the preferences of their users is an important requirement for actual companies and people to use agents for bidding.

In this paper we propose a concrete preference function to be used by an agent to express preferences over tasks. This function expresses preferences for specific properties of tasks and it is used in a decentralized task allocation setting. We introduce a bidding algorithm, where an agent bids on its most preferred tasks that are feasible given its current commitments. This algorithm uses a pricing mechanism which depends on the actual cost to perform the tasks and on the preference for the task. The influence of preferences on the price can be varied by setting a parameter (look at the role of the parameter p in the algorithm in Section 3.5).

Using this algorithm, we investigate the impact of preferences upon other aspects of task execution, such as execution time. We use both synthetic as well as real data from a logistics company.

This paper is organized as follows. First, the auctioning system used throughout the paper is introduced in Section 2. Section 3 introduces a function to express preferences and a bidding algorithm based upon such preferences. Experiments to evaluate the bidding algorithm and to study the trade-off between preferences and efficiency of task execution are presented in Section 4. Section 5 discusses related work, and finally, Section 6 concludes the paper.

¹ Vrije Universiteit Amsterdam, Department of Artificial Intelligence, De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands, email: mhoogen@cs.vu.nl

² University of Minnesota, Department of Computer Science and Engineering, 200 Union Street SE, Minneapolis, MN 55455, United States of America, email: gini@cs.umn.edu

2 The MAGNET System

The approach we present exploits some unique features of the MAGNET [4] system that allows autonomous agents to negotiate over coordinated tasks with precedence and time constraints. The MAGNET system consists of: (1) a *customer* agent, which puts tasks up for auction. The tasks have time constraints and other restrictions; (2) *suppliers* agents, which bid on the tasks and execute them if awarded; and (3) the MAGNET market server, which keeps track of the activities of the agents and of the auctions. The main interactions between agents in the MAGNET system are as follows:

- A customer agent issues a *Request for Quotes* (RFQ) which specifies the tasks, their precedence relations, and a time line for the bidding process. For each task, a time window specifies the earliest time the task can start and the latest time the task can end.
- Supplier agents submit bids. A bid includes one or more tasks, a price, the portion of the price to be paid as a non-refundable deposit, and the estimated duration and time window for task execution. Bids reflect supplier resource availability and constrain the customer's scheduling process.
- The customer agent decides which bids to accept. Each task needs to be mapped to one bid and the constraints of all awarded bids must be satisfied in the final schedule. In MAGNET the customer can choose from a collection of winner-determination algorithms (A*, IDA* [2], simulated annealing, and integer programming [3]).
- The customer agent awards bids and specifies the work schedule.

3 Preference Algorithm

In the bidding algorithm we propose, price is used as a mechanism to express preferences for tasks. Preferences in our case can be a combination of the following: (1) a preference for tasks of a particular duration (e.g. I hate performing very short tasks), (2) a preference for tasks at particular times during the day (e.g. I love getting up early in the morning, so give me tasks that ought to start early in the morning), and (3) a preference for particular types of tasks (e.g. I really hate to perform a task like that).

We show how to express these preferences and how to combine them. The preference for a task is referred to as ϕ_{task} , which we express using a real number in the interval $[0,1]$. Hereby, $\frac{1}{2}$ indicates a neutral preference, 0 is not preferred, and 1 is fully preferred. Since humans typically do not think in terms of a number when specifying preferences, we provide for each of the preference types covered a more intuitive formulation, as explained next. The specifics of how preferences are computed could be adapted for different domains, while keeping the approach.

3.1 Preferences for Duration

Let the preference to perform tasks of a certain duration be an integer. Such an integer can indicate either a minimum or a maximum duration (i.e. d_{min} , d_{max}). Let d_{min} be the minimum duration you want a task to last, i.e. you want the task to last longer than d_{min} . Durations below d_{min} are not preferred. If the duration is precisely d_{min} your preference is $\frac{1}{2}$, i.e. neutral. Let d_{close} be an integer that indicates how longer than d_{min} you want the task to last to be fully preferred. Tasks with duration in the range $[d_{min}, d_{min}+d_{close}]$ are more preferred than neutral, but not fully preferred. Any duration longer than $d_{min}+d_{close}$ is fully preferred. Then the preference $\phi_{duration}$ of a task with duration d_{task} can be calculated as follows:

- if there is a preference for minimum duration d_{min} :
 $d_{task} \geq d_{min}$: $\phi_{duration,task} = \frac{1}{2} + \min(\frac{1}{2} \times \frac{d_{task}-d_{min}}{d_{close}}, \frac{1}{2})$
 $d_{task} < d_{min}$: $\phi_{duration,task} = \max(\frac{d_{task}}{d_{min}} - \frac{1}{2}, 0)$
- if there is a preference for maximum duration d_{max} :
 $d_{task} \leq d_{max}$: $\phi_{duration,task} = \frac{1}{2} + \min(\frac{1}{2} \times (\frac{d_{max}-d_{task}}{d_{close}}), \frac{1}{2})$
 $d_{task} > d_{max}$: $\phi_{duration,task} = \max(\frac{d_{max}}{d_{task}} - \frac{1}{2}, 0)$

3.2 Preferences for Time Points

Let the preference for particular time points be indicated by a time of the day (e.g. 6.30 a.m.). Such a preference can indicate that the time needs to be before a particular time point t_{before} , or after a time point t_{after} . Let t_{close} indicate a time which is considered close to a particular time point. Again, the preference for a task which is precisely at the specified time point t_{before} or t_{after} is $\frac{1}{2}$, i.e. preference neutral. The preference for a given start time t_{task} can now be calculated as follows (note that for calculations using time points these are represented in seconds of the day):

- if a preference has been set for a task time before t_{before} :
 $t_{task} \leq t_{before}$: $\phi_{time,task} = \frac{1}{2} + \min(\frac{1}{2} \times (\frac{t_{before}-t_{task}}{t_{close}}), \frac{1}{2})$
 $t_{task} > t_{before}$: $\phi_{time,task} = \max(\frac{t_{before}}{t_{task}} - \frac{1}{2}, 0)$
- if a preference has been set for a task time after t_{after} :
 $t_{task} \geq t_{after}$: $\phi_{time,task} = \frac{1}{2} + \min(\frac{1}{2} \times \frac{t_{task}-t_{after}}{t_{close}}, \frac{1}{2})$
 $t_{task} < t_{after}$: $\phi_{time,task} = \max(\frac{t_{task}}{t_{after}} - \frac{1}{2}, 0)$

3.3 Preferences for Tasks

The last way to express preferences is for particular types of tasks. Let $type_{task}$ be the type for a given task. The type of a task is specified by means of a certain range of integers, whereby integers are ordered based upon similarity of the tasks. For example, if the tasks are represented on the interval $[0, 100]$, then the task identified with 1 is completely different from the task identified with 100, but has great similarity with the task identified with 2. Let the preferred tasks include a certain range of tasks $[type_{lower}, type_{upper}]$. Furthermore, let $type_{close}$ be an integer that expresses when a task is close to another task. The preference is calculated as follows:

- if $(type_{lower} \leq type_{task}) \wedge (type_{upper} \geq type_{task})$
 $\phi_{type,task} = 1$;
- if $type_{lower} > type_{task}$
 $\phi_{type,task} = \max(\frac{type_{close}}{type_{lower}-type_{task}}, 0)$
- if $type_{upper} < type_{task}$
 $\phi_{type,task} = \max(\frac{type_{close}}{type_{task}-type_{upper}}, 0)$

3.4 Combining Preferences

The preferences specified above are usually combined. We use a weighted sum of the preferences, setting the weight to 0 if a preference is not expressed.

$$\phi_{task} = w_{duration} \times \phi_{duration,task} + w_{time} \times \phi_{time,task} + w_{type} \times \phi_{type,task}, \text{ where } w_{duration} + w_{time} + w_{type} = 1$$

3.5 Bidding Algorithm with Preference for Tasks

We assume that the supplier agent owns a single resource with a particular capability (with which, of course, a number of different task types can be performed, as explained earlier). Furthermore, the resource has an availability slot (i.e. a begin and end time) as well as a particular $type_{begin}$ when the resource is initially setup and an end $type_{end}$ at which the use of the resource needs to end. The supplier agent maintains a schedule of the tasks planned for its resource.

We now present a bidding algorithm that takes preference values ϕ_{task} into account. The algorithm is a greedy algorithm, supplier agents try to bid upon as many tasks as feasible to maximize the usage of their resource. The algorithm uses a parameter, p , to vary the influence of the preference upon the eventual price bid.

The tasks within an RFQ are first ordered based upon their preference. If some tasks have identical preferences, they are ordered according to the earliest start time specified in the RFQ for the tasks included. We assume that there exists a function $switch_time: TASK_TYPE \times TASK_TYPE \rightarrow DURATION$ that calculates the switching time from one task type to another (when it can be performed on the resource). Furthermore, $performance_time: TASK_TYPE \rightarrow DURATION$ expresses the time needed to perform the task.

Bidding Algorithm

For the bidding algorithm, let $latest_end_time_{previous}$ be the latest end time of the previous task in the current schedule of the resource (or the schedule start time in case no such task exists). Let $type_{previous}$ be the type of the previous task (or the start type in case of no prior task), $latest_start_time_{next}$ be the latest start time of the next task (or the schedule end time in case no such task exists), and $type_{next}$ be the type of the next task (or the schedule end time in case no such task exists).

For each preference ordered task:

Check if task (*current*) can be done using the resource.

If yes, see if it fits in the current schedule (see below).

From the beginning of the schedule and for each empty slot in the schedule do:

If the task fits in the current empty slot in the schedule then insert the task in the bid,

add its time parameters to the schedule, and compute the price of the bid (see below)

else if $latest_end_time_{current} > latest_end_time_{next}$ then continue with the next slot
else continue with the next task.

To see if the task fits in the schedule, check if the following holds:

$$\begin{aligned} & ((latest_end_time_{previous} + switch_time(type_{previous}, type_{current})) \\ & \leq latest_start_time_{current}) \wedge \\ & ((latest_start_time_{next} - switch_time(type_{current}, type_{next}) - \\ & performance_time(type_{current})) \\ & \geq earliest_start_time_{current}) \wedge \end{aligned}$$

$$\begin{aligned} & [(latest_end_time_{previous} + switch_time(type_{previous}, type_{current})) \\ & \leq (latest_start_time_{next} - switch_time(type_{current}, type_{next}) \\ & - performance_time(type_{current}))] \end{aligned}$$

The price of the bid is computed as follows (note the parameter p):

$$\begin{aligned} price_{task} = & (1 + (p \times (1 - \phi_{task}))) \times \\ & [switch_time(type_{previous}, type_{current}) + \\ & switch_time(type_{current}, type_{next},) + \\ & performance_time(type_{current})] \end{aligned}$$

We have shown earlier how to calculate the value of ϕ_{task} for the different type of preferences. This price equation assumes a certain standard price for each minute of time spent. In case these costs vary, the cost per minute can be included as an additional parameter.

4 Experimental Setup

We now describe the effect of adjusting the parameter p in the bidding algorithm defined above. Furthermore, we study the effect of the preferences on the duration of task execution, which is an indicator of how efficiently the tasks are being performed. Together these form the utility function of the suppliers. Of course it is expected that having more preferences awarded will result in a less efficient execution. We are interested in assessing the severeness of these effects. We performed experiments using synthetic data, and experiments using a real dataset obtained from a trucking company.

4.1 Experimental Setup with Synthetic Data

We start by describing the parameters in the setup with synthetic data, and specify the actual settings used. There are many parameters that can influence the results. Many of them influence the difficulty of the task allocation problem in general. These include:

1. The number of tasks to be allocated.
2. The number of resources available.
3. The ratio between the resources required to perform the tasks and the availability of those resources (e.g. one resource might be more scarce than another). This also includes the specification of duration of tasks, switching time, and initial resource settings.
4. The tightness of the time windows specified in the tasks. Wider time windows allow more flexible scheduling of tasks, therefore finding a solution is easier.

The preference value itself is influenced by other parameters, including the following:

1. The parameter setting for the preference functions (e.g. what is considered to be a close by task, the stricter this norm is, the more easily preferences can be met).
2. The variation of tasks that exist (i.e. more variation means that it will be easier to get your preferences met).

Finally, other parameter settings can be varied, such as the number of iterations, and the value of the parameter p , which is used in the bidding algorithm to determine the influence of preferences on price.

4.1.1 Parameter Settings Used

We set the parameters of the preference functions and the variety of tasks to fixed values. This means that the preference function itself

remains constant over time, so that the influence of the parameter p is the only variation regarding the preference function.

We used several variations of the difficulty of the task allocation throughout the experiments. In particular we considered a market where more than sufficient resources are available (*overflow*) versus a market where resources are insufficient (*shortage*). Furthermore, the tightness of the time windows was varied by either setting them very tight or setting them wide. More precisely, the following parameters have been used to affect tightness of tasks:

1. the number of tasks was fixed to 10.
2. the number of resources available varied between 12 (tight market) and 50 (plenty of resources available).
3. the ratio between the required resources to perform tasks and the availability of those resources was fixed. We had three types of resources, each generated with an equal probability. The number of different tasks per resource was set to 9999. The maximum time to change from one task to another was set to 100 minutes. Task types were generated in a random fashion with an equal probability as well.
4. The tightness of the time windows specified in the tasks was varied between just sufficient time to perform the task to twice the time needed plus two full hours.

The parameter setting for the preferences are set so that initially the preference for tasks is around 60%, equally divided over the different preferences. Each of the agents is assigned one type of preference at random. The parameter p varied between 0 and 5.

4.1.2 Results

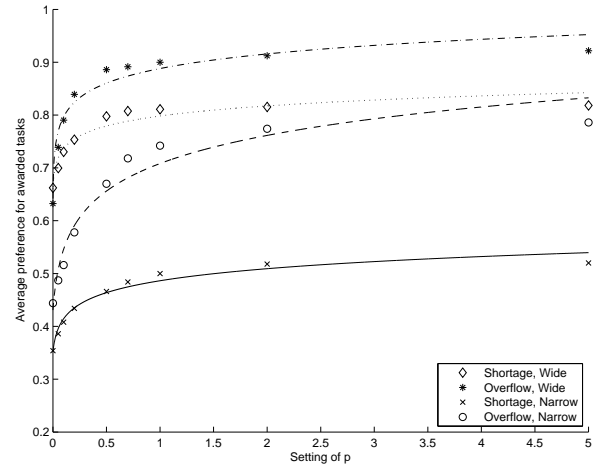


Figure 1. Preferences met for varying values of p

Figure 1 shows the average preferences for tasks with varying p for the different market types and time window settings. As can be seen, the easiest way to get the preferences met is the overflow market with wide time windows. The most difficult is the shortage market with tight time windows. The curves of the shortage market are less steep compared to the overflow market. The influence of the time windows on the average preference value is that the curve is basically lower by a certain constant value. The shape of the curve does not change for varying time window settings (i.e. in both the shortage market and

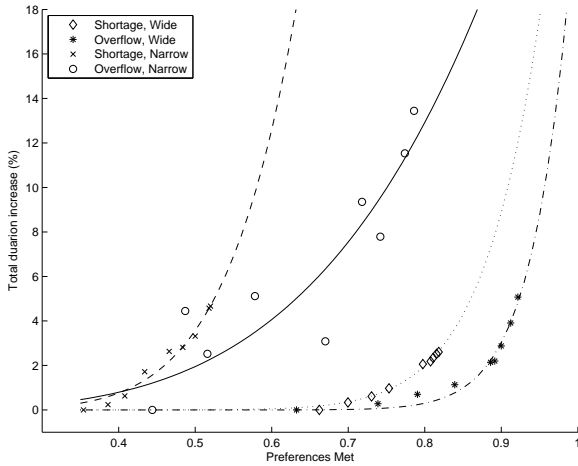


Figure 2. Preferences met versus increase in duration

the overflow market, the shape of the curve is the same for narrow and wide time windows).

Figure 2 shows the average preference for tasks on the x -axis and the increase in the average duration to perform the tasks on the y -axis. This clearly shows the trade-off between preferences awarded and the efficiency of task execution. All curves look similar (an x^2 type shape) except for the point where the huge increase starts, which varies for the different types of markets. The only exception is the curve of the overflow market with narrow time windows. In this case the results are less stable compared to the other results. The curve with the lowest preference value, after which a steep increase is observed, is the one in the shortage market with narrow time windows. This makes sense because there is hardly any room for allocating tasks to other agents. The curve with the highest point is the overflow market with wide time windows, in which there is plenty of space to express preferences and get them awarded.

4.2 Trucking data

Besides synthetic data, we tested our approach using a real company dataset from the trucking domain. The dataset consists of a number of container transports that need to take place. Tasks require a certain transportation from one zip code (the pickup location) to an intermediate location (the delivery location), ending at a third location (the return location). Therefore, a task description does not consist of one integer specifying the task, as before, but of three integers.

Furthermore, each task is associated with a certain early start time and a particular deadline at which the container needs to be returned at the return location. In addition to the containers that require transportation, the dataset also specifies which trucks are available. These can carry one container at a time (so only one type of resource is available), and have a certain availability slot when the truck becomes available, and when the truck needs to be returned. A location is also specified where the truck starts, and where it has to end. This nicely maps to the algorithm specified.

The performance time is now defined as the time to go from the pickup to the delivery location, plus the time to go from the delivery location to the return location. The switching time is no longer an artificial time, but it is the actual driving time from one zip code to another. The only artificial data which we have generated are the preferences of the various trucks. This is done according to the method

mentioned for the synthetic data. Finally, the preference for type of tasks is the average of the three different integers included in the task description (i.e. pickup, delivery, and return location).

4.2.1 Results

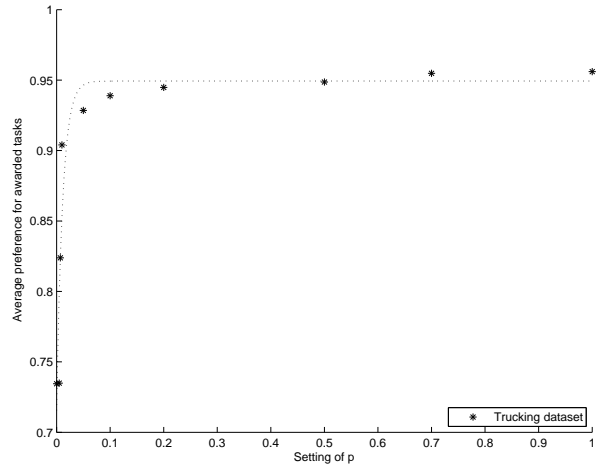


Figure 3. Preferences met for trucking dataset, for varying values of p

Figure 3 shows how the value of p affects the average preference for tasks. It can be seen that the value of p required to increase the average preference significantly is much lower than for the random dataset. Furthermore, the limit seems to be comparable with the overflow market with wide time windows.

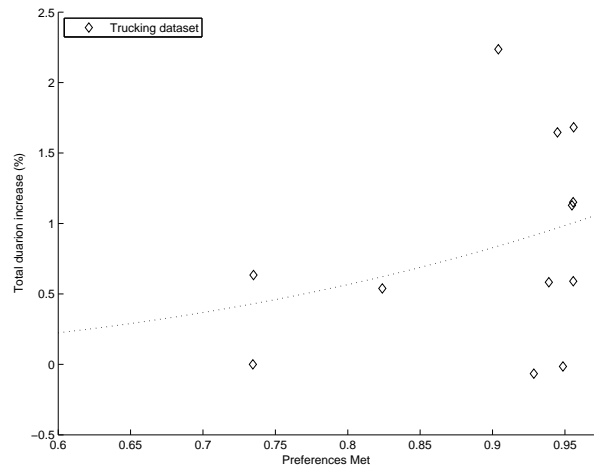


Figure 4. Average task preference versus duration of performing the tasks

Figure 4 shows the preference value versus the average duration increase, i.e. the trade-off between preferences met, and the efficiency of execution. It can be seen that there is hardly any correlation between the average preference value of the trucks and the average increase in duration. This is of course very good news for the trucking company because this means they can award drivers their preferences without increasing the total driving time. This is assuming that preferences are equally divided amongst the truckers, as in the experimental setup.

5 Related Work

In the field of combinatorial auctions, a lot of attention has been devoted to finding out the exact preference for particular bundles of tasks (see e.g. [5] and [11]). In general a certain preference for each of the bundles is assumed, but no detail is given on how the bidder comes to such a preference value. In this paper we introduce a preference function that allows for a more intuitive specification of preferences, thereby taking multiple aspects of the tasks into account. Preferences for different aspects of the tasks are combined using a weighted average to produce a single preference value.

In research on preference elicitation, typically the impact on selling is addressed, but the precise influence of preferences upon the quality of the solution is not. In this paper, we show how the allocation of tasks in a decentralized fashion directly influences the quality of the solution, and we explore the relationship between the average preference of tasks and the solution quality. In [7] an approach for scheduling a meeting between agents is proposed, which takes into account the preferences of the agents. The relationship between such preferences and the quality of the solution is addressed, but the problem is not studied from the perspective of combinatorial auctions.

Task allocation can also be performed from a centralized perspective, using preferences as soft constraints. See, for example, [9] for an approach to consider preferences in decision making. There are decentralized variants of constraint optimization, but the agents in our case are not necessarily cooperative. In the field of planning and scheduling preferences have been considered as well. Languages have been developed that allow for the specification of preferences and soft constraints (see e.g. [8]).

The logistic domain we use for our experiments has been researched for quite some time (see e.g. [10]), mainly focusing on calculating optimal solutions from a centralized perspective. For instance, in [6] the problem addressed is to find optimal routes for transportation orders of a large set of users. Orders have to be picked up and delivered at specific locations, within a given time window, and using a limited number of trucks. The solution proposed is centralized, and it is used to support a human dispatcher.

The current trend in logistics requires an even more distributed setting because of the use of fourth party logistics (4PL) [1]. 4PL companies sign contracts with large companies to arrange their entire transportation demand. These companies, however, do not have sufficient resources on their own to arrange all these transports and therefore distribute many of those tasks to other (partner) companies. Centralized calculation might no longer be feasible due to lack of complete information (availability of resources is too sensitive for a company to communicate) as well as the complexity of calculating an optimal solution within a short period (time is crucial in the business).

6 Conclusions

We have presented an approach to specify preferences for tasks in a combinatorial auction setting. Allowing users to specify such preferences is essential for them to use auctions and to increase the economic efficiency of reverse auctions, as reported, for instance, in [12]. We propose a preference function and use it in a bidding algorithm where bids on non-preferred tasks have a higher price.

We evaluated our approach in two ways, first by rigorously testing it with synthetic data. Several parameters have been varied, namely the tightness of the time windows, within a certain schedule, and the relative availability of resources. It was shown that it was easiest to

get preferences awarded in for markets with wide time windows. The trade-off between meeting preferences and overall execution time has been studied in depth. We have shown that the overall execution time is influenced most in the case of the overflow market, due to the fact that in the shortage market there are hardly any alternatives at hand and therefore, although the agent might not prefer a task, it will still get its bid awarded. The curves observed tend to have the same shape when the time window setting changes but the market type remains the same. For different market types, the curves vary in steepness.

Besides testing with synthetic data, we have also used a real company dataset from the trucking domain. We have shown that the bidding algorithm is effective in awarding suppliers more preferred tasks. The influence of this preference on the overall solution quality was not observed using the real dataset. Hence, in this setting the preferences being met have much less influence on the efficiency of the solution found. For future work, it would be interesting to find out whether other real datasets would show the same results as the dataset used in this paper. Furthermore, exploring how well the companies can express their preferences using these functions would be interesting as well.

Acknowledgments: Partial support is gratefully acknowledged from NSF under grant IIS-0414466.

REFERENCES

- [1] P. Briggs. The hand-off: the future of outsourced logistics may be found in the latest buzzword [fourth party logistics]. *Canadian Transportation Logistics*, 102(5):18, 1999.
- [2] J. Collins, G. Demir, and M. Gini. Bidtree ordering in IDA* combinatorial auction winner-determination with side constraints. In J. Padgett, O. Shehory, D. Parkes, N. Sadeh, and W. Walsh, editors, *Agent Mediated Electronic Commerce IV*, volume LNAI2531, pages 17–33. Springer-Verlag, 2002.
- [3] J. Collins and M. Gini. An integer programming formulation of the bid evaluation problem for coordinated tasks. In B. Dietrich and R. 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.
- [4] J. Collins, W. Ketter, and M. Gini. A multi-agent negotiation testbed for contracting tasks with temporal and precedence constraints. *Int'l Journal of Electronic Commerce*, 7(1):35–57, 2002.
- [5] W. Conen and T. Sandholm. Preference elicitation in combinatorial auctions. In *Proc. First Int'l Conf. on Autonomous Agents and Multi-Agent Systems*, volume 1, pages 168–169, Bologna, Italy, July 2002.
- [6] K. Dorer and M. Calisti. An adaptive solution to dynamic transport optimization. In *Proc. Fourth Int'l Conf. on Autonomous Agents and Multi-Agent Systems*, pages 45–51, 2005.
- [7] M. Franzin, E. Freuder, F. Rossi, and R. Wallace. Multi-agent meeting scheduling with preferences: efficiency, privacy loss, and solution quality. In *Proc. of AAI Workshop on Preference in AI and CP*, 2002.
- [8] A. Gerevini and D. Long. Preferences and soft constraints in PDDL3. In *Proc. ICAPS Workshop on Planning with Preferences and Soft Constraints*, 2006.
- [9] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Wiley, 1976.
- [10] T. Magnanti. Combinatorial optimization and vehicle fleet planning: Perspectives and prospects. *Networks*, 11:179–214, 1981.
- [11] D. C. Parkes. Auction design with costly preference elicitation. *Annals of Mathematics and Artificial Intelligence*, 44:269–302, 2005.
- [12] T. Sandholm. Expressive commerce and its application to sourcing: How we conducted \$35 billion of generalized combinatorial auctions. *AI Magazine*, 28(3):45–58, Fall 2007.
- [13] R. G. Smith. The contract net protocol: High level communication and control in a distributed problem solver. *IEEE Trans. Computers*, 29(12):1104–1113, December 1980.
- [14] W. Walsh and M. Wellman. A market protocol for decentralized task allocation and scheduling with hierarchical dependencies. In *Proc. of 3th Int'l Conf on Multi-Agent Systems*, 1998.