Evaluating Risk: Flexibility and Feasibility in Multi-Agent Contracting

John Collins, Maksim Tsvetovat, Rashmi Sundareswara, Joshua van Tonder, Maria Gini
Department of Computer Science and Engineering, University of Minnesota

Bamshad Mobasher
School of Computer Science, Telecomm., and Information Systems, DePaul University

Abstract

In an automated contracting environment, where a *contractor* agent must negotiate with other self-interested *supplier* agents in order to execute its plans, there is a tradeoff between giving the suppliers sufficient flexibility to incorporate the requirements of the contractor's call-for-bids into their own resource schedules, and ensuring the contractor that any bids received can be composed into a feasible plan. We introduce a bid evaluation process that incorporates cost, task coverage, temporal feasibility, and risk estimation. Using this evaluation process, we describe an empirical study of the tradeoffs between flexibility, plan feasibility, and cost in the context of our MAGNET multi-agent contracting market infrastructure.

1 Introduction

The ability to automate the negotiation of complex contracts among multiple suppliers is an important but not yet achieved feature of multi-agent systems [1, 4, 6]. The University of Minnesota's MAGNET (Multi-Agent Negotiation Testbed) system [3] is an innovative agent-based approach to complex contracting and supply-chain management problems. The MAGNET system comprises a set of agents who negotiate with each other through a market infrastructure designed to support the execution of complex plans among a population of independent, autonomous, heterogeneous, self-interested agents. We call this activity *Plan Execution By Contracting*. A *contractor* agent, in order to fulfill its goals, must contract with other *supplier* agents for all or part of the necessary tasks. The MAGNET system incorporates a simple three step, leveled commitment protocol, with a contractor agent issuing a call-for-bids, suppliers replying with bids, and the contractor accepting the bids it chooses. A detailed description of the MAGNET architecture and its negotiation protocol were presented in [7].

Here we introduce a bid evaluation process for automated contracting that incorporates cost, task coverage, temporal feasibility, and risk estimation. Using this evaluation process, we provide an empirical study of the tradeoffs between flexibility, plan feasibility, and cost in the context of the MAGNET market infrastructure.

2 Bid Generation and Evaluation: Experimental Results

The main goal of this paper is to present a bid evaluation algorithm and describe experiments designed to observe the relationships among numbers of suppliers, the temporal flexibility offered to suppliers, and the price and risk factors experienced by a contractor agent. The experimental setup includes a MAGNET Server, a contractor agent that generates plans, requests bids, and evaluates bids, and a supplier agent that generates and submits bids.

Generation of call-for-bids. The contractor agent starts with a plan, a partially ordered sequence of tasks, and issues a call-for-bids for the tasks that will be contracted out. In the call-for-bids, the contractor specifies for each task (i) an early start time, (ii) a late finish time, and (iii) the set of tasks that must precede it.

Generation of bids. The supplier agent is a test agent that masquerades as an entire community of suppliers. In the full implementation of MAGNET, this will be replaced by heterogeneous and more sophisticated agents. The supplier agent generates bids, selecting an initial task at random, following predecessor and successor links with some probability to add more tasks to the one selected initially. The task duration is a normally distributed random variable. If the duration is longer than the time window in the call-for-bids, the task is dropped from the bid. The negotiation protocol requires that the time parameters in a bid are within the time windows specified in the call-for-bids. For each task, the expected early start time is selected as a uniformly distributed variable between the early start time in the call-for-bids, and a late start time derived by subtracting the task duration from the late finish time specified in the call-for-bids. The late finish time is similarly determined. The price for a task is computed as the product of an initial price, a flexibility discount, and a resource commitment premium. The flexibility discount and resource commitment premium are computed from the ratio of the time allocated to the task over its expected duration. The decommitment penalty that the contractor must pay for backing out of a bid is chosen as a random value uniformly distributed between 0 and the price of the bid. If the supplier is bidding on multiple tasks, the overall price for the bid is discounted by a random value between 1.0 and 0.8. All these values are arbitrarily chosen to create a reasonable spread of prices.

Bid Evaluation. When the bidding deadline is past, the contractor agent evaluates the bids in an attempt to find a combination that minimizes a combination of cost and risk, provides coverage of all tasks, and allows for a feasible schedule. For each bid, the contractor has the option of selecting the entire bid and paying the overall discounted price or selecting a subset of the task bids. Since we expect the combination of plan size, number of bids, and number of tasks per bid to be too large for systematic-search, we use a heuristically-guided simulated-annealing search [5], that includes the following steps:

- *choose a node*: A node is a partial or complete mapping of tasks to bids. Each node has a set of numeric evaluations that represent the costs of the mapped bids, a risk value that estimates recovery cost, and weighted penalties for unmapped tasks and infeasibilities. The nodes are kept in a queue of fixed maximum length, sorted by the node's evaluation, and the choice of which node to expand is made by using an exponentially-distributed value, whose mean is the current annealing temperature (see [5] for details on simulated annealing). As the search proceeds, the annealing temperature is periodically reduced to increase the focus of the search.
- *choose an expansion*: Expansions are made by adding a bid or an individual task to a node. If one or more tasks mapped by the selected bid were already covered by other bids, those bids are removed from the node's mapping. In order to minimize thrashing, bids that have already been used to expand a given node are disallowed, as are bids that are in the *tabu list*, which contains the last 10 expansions in the parentage of the node. In addition, if there are "singleton" bids or tasks for which there is a single bid, bids for those tasks may not be displaced by expansions.
- *evaluate the resulting node*: If a valid node was produced by the expansion, it is tested for coverage, feasibility, cost, and risk. Cost and risk are measured in the same units, while coverage and feasibility involve arbitrary penalty factors.
- *stopping criterion*: The search ends when a number of expansions have been attempted without finding an improvement. The best feasible and covered solution, if any, is returned.

3 Experimental Results

We explored two dimensions of the bid-evaluation problem, the number of bidders and the flexibility in the time-windows for the tasks specified in the call-for-bids. We tested two sets of bidders, one with 10 and one with 30. We used varying values of the time windows specified in the call-for-bids, allowing the tasks to overlap by 10%, 40%, and 70%. For each of the two sets of bidders, we performed 3 runs with varying time windows, and we run each experiment 40 times. A detailed description of the results is in [2].

For the purpose of the experiments, the plans consist of seven tasks, each with an expected duration arbitrarily set to 1 hour, and with a combination of parallel and serial precedence relationships between the tasks.

Our experimental results demonstrate that the advantage of increasing supplier flexibility is dependent on the number of available suppliers. In other words, if the number of suppliers is small, the risk of plan infeasibility outweighs the advantage of added flexibility. On the other hand, if the number of suppliers is large, the more flexible plan specifications result in lower-risk plans.

The probability of finding a feasible, covered plan is reduced dramatically with small numbers of bidders. When time windows are short, the infeasibility test in the bid preparation process results in smaller numbers of bids, with the bids received being composed of fewer tasks. However when the time windows are increased, the effect of larger number of (larger) bids is more than offset by fewer combinations that form temporally feasible plans.

Risk evaluation generally decreases as the time windows are increased. There is a clear relationship between the size of the time windows specified in the call-for-bids and the expected decommitment cost. This is because the bidders are generally composing bids with larger time windows, and even though many of them cannot be composed while preserving precedence relations, the evaluator is able to find combinations that maximize slack early in the plan, where it has the greatest impact on expected decommitment cost.

The increasing cost with larger time windows is due to the same effect that reduces risk: suppliers are specifying larger time windows in the bids, and are applying the resource-commitment premium. In a more realistic situation, bidders could submit both a higher-cost, more flexible and lower-risk bid, along with a lower-cost, less flexible and higher-risk bid, allowing the contractor to make the decision.

4 Conclusions

We have presented a bid evaluation process for automated contracting that incorporates cost, task coverage, temporal feasibility, and risk estimation, and we have provided an empirical study of the tradeoffs between flexibility, feasibility, and cost in the context of the MAGNET system. Our results show that a contractor agent can make tradeoffs between price, feasibility, and risk by adjusting the temporal specifications of the tasks. The nature of these tradeoffs varies with the number of potential suppliers in the market.

References

[1] C. Beam and A. Segev. Automated negotiations: A survey of the state of the art. Technical Report CITM Working Paper 96-WP-1022, Walter A. Haas School of Business, 1997.

[2] J. Collins, M. Tsvetovat, R. Sundareswara, J. V. Tonder, M. Gini, and B. Mobasher. Evaluating risk: Flexibility and feasibility in multi-agent contracting. Technical Report 99-001, University of Minnesota, Dept of CSE, Minneapolis, Minnesota, February 1999.

[3] J. Collins, B. Youngdahl, S. Jamison, B. Mobasher, and M. Gini. A market architecture for multi-agent contracting. In *Proc. of the Second Int'l Conf. on Autonomous Agents*, pages 285–292, May 1998.

[4] R. H. Guttman, A. G. Moukas, and P. Maes. Agent-mediated electronic commerce: a survey. *Knowledge Engineering Review*, June 1998.

[5] C. R. Reeves. *Modern Heuristic Techniques for Combinatorial Problems.* John Wiley & Sons, New York, NY, 1993.

[6] T. W. Sandholm and V. Lesser. On automated contracting in multi-enterprise manufacturing. In *Distributed Enterprise: Advanced Systems and Tools*, pages 33–42, Edinburgh, Scotland, 1995.

[7] E. Steinmetz, J. Collins, M. Gini, and B. Mobasher. An efficient algorithm for multiple-component bid selection in automated contracting. In *Agent Mediated Electronic Trading*, volume LNAI1571, pages 105–125. Springer-Verlag, 1998.