

Bidders with Hard Valuation Problems

Kate Larson
Computer Science Department
Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213
klarson@cs.cmu.edu

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213
sandholm@cs.cmu.edu

ABSTRACT

Auctions provide efficient and distributed ways of allocating goods and tasks among agents. In this paper we study optimal strategies for computationally limited agents, where agents must use their limited computing resources to determine valuations for (bundles of) the items being auctioned. Agents are free to compute on any valuation problems including their opponents'. The deliberation actions are incorporated into the agents' strategies and different auction settings (both single-item and combinatorial) are analyzed in order to determine equilibrium strategies. We show that in some auction mechanisms, but not others, in equilibrium the bidders compute on others' problems as well. We show that the model of bounded rationality (limited *or* costly computing) impacts the agents' equilibrium strategies and so must be considered when designing mechanisms for computationally limited agents.

Categories and Subject Descriptors

I.2.11 Multiagent systems

General Terms: Theory

1. INTRODUCTION

Auctions are useful mechanisms for allocating items in multi-agent systems. There is a large body of research that has focused on designing different auction mechanisms and determining optimal bidding strategies for agents participating in auctions [4]. The bulk of auction theory assumes that bidders' valuations for the items being auctioned are given *a priori*. In many applications, however, agents must expend significant effort to determine their valuations.

Limitations on agents' computational capabilities can mean that agents are not able to compute their valuations optimally. This can impact the bidding strategies that are chosen to be played. It turns out that the equilibrium for rational agents does not generally remain an equilibrium for computationally limited agents.

This paper investigates the effect of computational limitations on agents strategies, where agents are participating in different classical auction mechanisms. The agents' limitations take the form of deadlines. Agents can compute freely up until a certain point in

time, at which point all computing must cease. Before their deadlines, agents are free to compute on any valuation problem they wish, including their opponents' valuation problems. Information that the agents gain about valuations by computing can be used in forming bidding strategies. We present the concept of *strategic computing*, and analyze four classical single-item auction mechanism and one multi-item auction to determine under what circumstances strategic computing will occur in equilibrium.

2. ROLE AND CONTROL OF COMPUTING

Agents must decide how to put their allowed computing time to best use. They can decide to compute on their own problems in order to determine their own valuations or may compute on competitors' problems in order to learn what sort of bids they may submit. These decisions are based on the results of the agent's computing and on what the agent believes other agents are doing.

Computing can play different roles. In particular, we are interested in settings where agents either compute to *improve* their valuations or else compute to *refine* their valuations.

When agents compute to improve their valuations, we assume that they have available to them *anytime algorithms* for solving optimization problems on how to use the item up for auction. As the agent devotes more computing resources on a valuation problem, the solution quality improves. As better solutions are found for the valuation problem, the agent is able to better tailor its bid so as to optimize its expected utility, where the utility of an agent is defined to be the difference between its computed value for the item and the amount that it pays for the item.

In other settings agents refine their valuations. They do this by acquiring information which improves their knowledge about what the true valuation of an item is. For example, over time an agent may learn what the true quality of the item up for auction is, or may gain information about possible uses for the item. An agent can use this information in the bid formation process, to improve the agent's expected utility.

Agents are provided with *performance profile trees* that can be used to determine how to optimally use their limited computational resources. Performance profile trees represent the effects of computing time on solution quality and is used by agents to decide how to compute at each step in the process, based on results of computing so far.

Figure 1 is an example of a performance profile tree. There are two different types of nodes in a tree; value nodes and random nodes. Value nodes store the value of the solution computed so far while random nodes occur when ever a random number is needed to chart the path of the algorithm run. The edges in the trees are all labeled. An edge emanating from a value node is labeled with the probability of reaching the child given the parent was reached.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'02, July 15-19, 2002, Bologna, Italy.

Copyright 2002 ACM 1-58113-480-0/02/0007 ...\$5.00.

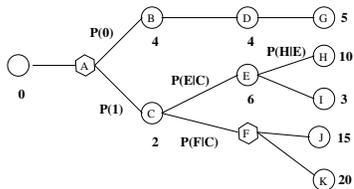


Figure 1: A performance profile tree.

An edge emanating from a random node is labelled with a number and the probability that the number was drawn. Once an agent has reached a node in the tree, it need only consider algorithm paths that continue into the subtree rooted at the node that was reached.

There has been considerable work on performance profile based deliberation control for single agents [6, 1], but often simplifying assumptions have been used. This is fine for single agent settings as long as agents behave reasonably well. However, in multiagent settings any deviation from full normativity can be catastrophic as each agent's strategy may depend on other agents' strategies. On the other hand, the performance profile tree is a fully normative model for control of computing [2]. It allows conditioning on results obtained by computing as well as conditioning on different aspects of the problem instances.

Performance profile trees can capture uncertainty in computing from different sources including uncertainty in the problem instance, and uncertainty from using randomized algorithms. Any performance profile tree that models uncertainty is called *stochastic*, while a performance profile tree where there is no uncertainty as to what result will be obtained after any amount of time spent computing is called *deterministic*.

Agents use their computational resources in different ways. They can compute on their own problems in order to obtain better valuations. They can also compute on their opponents' problems in an attempt to gather information about the bids that the opponents may be submitting. Both uses of the computational resources allow agents to submit better bids, however, we make a distinction between these two types of computation. The first we simply call *computation*. The second we call *strategic computation*. Strategic computing can be further refined to *weak strategic computation* and *strong strategic computation*.

DEFINITION 1. *If an agent uses part of its computational resources to compute on another agent's valuation problem, then the agent is performing strong strategic computation. If an agent does not use its limited computational resources to compute on another agent's valuation problem, but does use information obtained from the opponent's performance profile to devise a strategy, then the agent is performing weak strategic computation.*

3. RESULTS

The auction mechanism, the type of performance profiles, and the role of computing can all influence the strategies of computationally limited agents. Four single-item auctions and a multiple item auction were analyzed to see whether any form of strategic deliberation would occur in equilibrium. The results are summarized in Table 1. For first-price and Dutch auctions, strategic computing can occur in equilibrium, and the type of strategic computing depends on the performance profiles. For the English and Vickrey auctions, no form of strategic computing occurs in equilibrium. This is independent of the performance profiles. In the generalized Vickrey auction (GVA), the performance profiles determine the type of strategic computing that may occur in equilibrium.

Other models of computationally limited agents have been studied. In particular, Larson and Sandholm studied a model where

Auction Mechanism	Performance Profile Type	Strategic Computing?		
		None	Weak	Strong
First Price	deterministic		✓	
	stochastic		✓	✓
Dutch	deterministic		✓	
	stochastic		✓	✓
English	deterministic	✓		
	stochastic	✓		
Vickrey	deterministic	✓		
	stochastic	✓		
GVA	deterministic		✓	
	stochastic		✓	✓

Table 1: The occurrence and type of strategic computing depends on the auction mechanism and the type of performance profiles. Whether the agents compute to improve their valuations or to refine their valuations does not change the equilibrium strategies for each auction and performance profile type.

bidding agents have unlimited but costly computational resources [3]. In particular, they showed that if agents incur a cost from computing then in English and Vickrey auctions strong strategic computing can occur in Nash equilibrium. Table 2 shows whether strong strategic computation occurs in equilibrium for both models of bounded rationality in the situation where the performance profiles are stochastic.

	Auction mechanism	Rational agents counterspeculate?	Strategic Computing?	
			Limited	Costly
Single item	English	no	no	yes
	Vickrey	no	no	yes
	First Price	yes	yes	yes
	Dutch	yes	yes	yes
Multiple items	GVA	no	yes	yes

Table 2: When does strategic deliberation occur? The answer depends on the model of limited computation being used.

From this result, one can conclude that *how* agents are computationally limited deeply impacts their equilibrium strategies. While previous work has looked at which auction mechanisms might be better for computationally limited agents [5], our results show that the model of how agents are bounded in their rationality should be included in any game theoretic analysis.

4. REFERENCES

- [1] M. Boddy and T. Dean. Deliberation scheduling for problem solving in time-constrained environments. *Artificial Intelligence*, 67:245–285, 1994.
- [2] K. Larson and T. Sandholm. Bargaining with limited computation: Deliberation equilibrium. *Artificial Intelligence*, 132(2):183–217, 2001.
- [3] K. Larson and T. Sandholm. Costly valuation computation in auctions: Deliberation equilibrium. In *TARK*, pages 169–182, Siena, Italy, July 2001.
- [4] P. Milgrom. Auctions and bidding: A primer. *Journal of Economic Perspectives*, 3(3):3–22, 1989.
- [5] D. C. Parkes. Optimal auction design for agents with hard valuation problems. In *Agent-Mediated Electronic Commerce Workshop at IJCAI99*, Stockholm, Sweden, 1999.
- [6] S. Russell and E. Wefald. Principles of metareasoning. *Artificial Intelligence*, 49:361–395, 1991.