Dynamic Multiagent Resource Allocation: Integrating Auctions and MDPs for Real-Time Decisions

Hadi Hosseini
David R. Cheriton School of Computer Science
University of Waterloo
200 University Avenue West
Waterloo, ON N2L 3G1
h5hossei@uwaterloo.ca

Abstract

Multiagent resource allocation under uncertainty raises various computational challenges in terms of efficiency such as intractability, communication cost, and preference representation. To date most approaches do not provide efficient solutions for dynamic environments where temporal constraints pose particular challenges. We propose two techniques to cope with such settings: auctions to allocate fairly according to preferences, and MDPs to address stochasticity. This research seeks to determine the ideal combination between the two methods to handle wide range of allocation problems with reduced computation and communication cost between agents.

Problem Statement

The main goal in Multiagent Resource Allocation (MARA) is to distribute a set of resources among a set of intelligent agents in order to respect the preferences of the agents and to maximize some measure of global utility. Resource allocation can benefit from various multiagent system approaches such as market mechanisms including auctions (Wellman et al. 2001) and negotiation protocols (Sandholm 1998) and stochastic planning techniques (Dolgov and Durfee 2006). Dynamic auctions (Seuken, Cavallo, and Parkes 2008) are shown to be effective in dealing with uncertainty while preserving the properties of a VCG mechanism; however, they require an optimal solution, which is intractable in large scale problems. We seek to develop novel techniques to enable efficient solutions in stochastic environments where temporal constraints between resources may arise.

Progress to Date

Resource planning under uncertainty in multiagent environments scales exponentially with the number of agents and resources. Solving large MDPs in multiagent systems becomes NEXP-complete and computationally expensive, and in fact, multiagent Markov decision process (MMDP) and decentralized MDP (Dec-MDP) approaches become infeasible (Bernstein et al. 2002) when dealing with a large number of agents and resources.

We approach the dynamic stochastic resource scheduling problem for domains in which the success of each task or agent is dependent stochastically on its ability to obtain a sequence of resources over time. We formulate this problem as a decentralized Markov decision process (Dec-MDP), in which each consumer models its decision-making process by a single MDP that has a utility function over its success, considering the uncertain effects of resources on success. Agent MDPs also model the ability of the agent to obtain resources by exposing their expected value related to that resource to a centralized decision maker using an iterative auction procedure. The actions of the consumer agents are then to bid for resources in this sealed-bid iterative auction, or to wait until resources become free. Our approximate solution for solving Dec-MDPs benefits from a simple sealed bid auction mechanism to coordinate the decentralized MDPs in order to maximize the social welfare. Auctions introduce indirect communication and help to reduce communication costs.

In this research, we are motivated by the challenge of effectively allocating resources to patients in hospital settings. This domain is potentially rich for evaluation purposes: the non-deterministic steps in the medical procedures require stochastic analysis for decision making to take randomness of patient behaviors into account.

Patients should go through a series of medical tasks (acquiring equipment or doctors that are possibly in demand by other patients) in a certain (partial) *order* to complete their medical pathways. An effective allocation of such resources to the agents requires coping with stochastic, temporal, and dynamic elements in order to increase the overall health of patients (and minimize possible health state deterioration).

A cooperative market-based allocation mechanism to allocate different resources based on a linear health-based cost function was proposed (Paulussen et al. 2003) that uses rounds of negotiations to find a final setting that minimizes costs for all the patients, without considering independencies between tasks and stochastic nature of the environment. A related approach (Dolgov and Durfee 2006) focuses on solving the resource allocation problem and stochastic planning problem simultaneously by having agents reporting their local preferences through an MDP to the auctioneer and solving the winner determination problem. This algorithm is a one-shot allocation solution with static number of agents and no re-allocation of resources is allowed afterwards. Thus, it is not applicable to environments with dynamic agents or resources because the formulation should be modified by each change to compute a new allocation. In a preliminary paper, we have shown (Hosseini, Hoey, and Cohen 2011) that our Dec-MDP approach copes with dynamic allocation of resources under uncertainty (as it uses multiagent stochastic planning) and converges fast with low communication cost (as it uses an auction mechanism), while addressing temporal preferences.

Proposed Plan for Research

The first step will be to conduct more extensive validation, in order to improve the results. The primary aim is to then offer a more generalized approach for domains beyond health.

Combinatorial auctions (CAs) allow bidders to express their true preferences and valuations over the *combination of items (or bundles)* (in MARA, a full preference model over all possible sets of resources) that results in more effective allocations in terms of total economic efficiency (Cramton, Shoham, and Steinberg 2006). I plan to study combinatorial bidding mechanisms and bidding languages (Boutilier and Hoos 2001; Nisan 2000) to deal with situations where agents need to expose a richer preference model in order to optimize the allocation process. MDPs can potentially be used in combinatorial bidding processes in which agents reveal their valuations over a bundle of items by reasoning on the future expected outcomes. This can reduce the complexity of the bidding model in combinatorial bidding process by eliminating the paths that could lead into smaller or negative expected outcomes. I plan to theoretically study the use of MDPs as stochastic mechanisms for preference representation in CAs.

In auction-based mechanisms, the assumption is that auctioneers are reliable and trustworthy, although they may lack sufficient information about the requesting agents (as the communication is just through passing bids). However, from another perspective resource agents can negotiate over consumer agents (patients in the healthcare scenario) in some negotiation rounds and enforce understandings while competing with each other in order to make an agreement over the final Pareto-optimal assignment. One future direction would be for me to study the integration of such techniques to the context of MARA problems.

Another future direction would be defining fairness criterion along with addressing optimality and social welfare. Defining fairness in terms of envy-freeness does not necessarily result in optimal solutions (Sandholm 1999). One interesting direction would be to study envy-reducing strategies in negotiation protocols and GVA mechanisms that can potentially lead into efficient allocation of resources.

In the side of stochastic planning, I plan to apply the auction-based coordination technique to Dec-POMDPs where the states of obtaining resources and progression is partially observable through an agent's observation on environment dynamics, which may include a probability distribution over all the resource requests. This opens up an interesting direction in solving large-scale resource planning problems using Dec-POMDPs.

I anticipate addressing a number of challenging research questions as part of this topic. Included is how to effectively reason about expected future outcomes in a planning horizon model. Another is examining the relative value of regret-based or reward-based bidding for the MDP-based auction component. Finally, the most critical challenge will be to support large-scale resource planning problems where real-time decision making is vital.

In the next 2 years, I will integrate these mentioned techniques to provide a more effective solution for MARA. This will result in dual contributions to the fields of stochastic planning using MDPs (integrating market-oriented coordination solutions) and the design of multiagent auctions (with rich preference representation for bidding and proposals for addressing challenges in solving combinatorial auctions). Researchers in both subfields will thus benefit as we demonstrate how best to combine these techniques into a comprehensive solution.

References

Bernstein, D.; Givan, R.; Immerman, N.; and Zilberstein, S. 2002. The complexity of decentralized control of Markov decision processes. *Mathematics of operations research* 27(4):819–840.

Boutilier, C., and Hoos, H. 2001. Bidding languages for combinatorial auctions. In *International Joint Conference on Artificial Intelligence*, volume 17, 1211–1217. Lawrence Erlbaum Associates LTD.

Cramton, P.; Shoham, Y.; and Steinberg, R. 2006. Introduction to combinatorial auctions.

Dolgov, D., and Durfee, E. 2006. Resource allocation among agents with mdp-induced preferences. *Journal of Artificial Intelligence Research* 27(1):505–549.

Hosseini, H.; Hoey, J.; and Cohen, R. 2011. Multi-agent patient scheduling through auctioned decentralized mdps. In *Proceedings of the 6th INFORMS Workshop on Data Mining and Health Informatics (DM-HI 2011)*, 73–78.

Nisan, N. 2000. Bidding and allocation in combinatorial auctions. In *Proceedings of the 2nd ACM conference on Electronic commerce*, 1–12. ACM.

Paulussen, T.; Jennings, N.; Decker, K.; and Heinzl, A. 2003. Distributed patient scheduling in hospitals. In *International Joint Conference on Artificial Intelligence*, volume 18, 1224–1232. Citeseer.

Sandholm, T. 1998. Contract types for satisficing task allocation. In *Proceedings of the AAAI spring symposium: Satisficing models*, 23–25.

Sandholm, T. 1999. Distributed rational decision making. *Multiagent systems: a modern approach to distributed artificial intelligence* 201–258.

Seuken, S.; Cavallo, R.; and Parkes, D. 2008. Partially-synchronized dec-mdps in dynamic mechanism design. In *AAAI*, volume 8, 162–169.

Wellman, M.; Walsh, W.; Wurman, P.; and MacKie-Mason, J. 2001. Auction protocols for decentralized scheduling* 1,* 2. *Games and Economic Behavior* 35(1-2):271–303.