Adaptive Decision Support for Structured Organizations: A Case for OrgPOMDPs

(Extended Abstract)

Pradeep Varakantham**, Nathan Schurr*, Alan Carlin*, Christopher Amato*
** - School of Information Systems,Singapore Management University,{pradeepv}@smu.edu.sg,
* - Aptima, Inc., Washington, DC, {nschurr,acarlin,camato}@aptima.com

ABSTRACT

In today's world, organizations are faced with increasingly large and complex problems that require decision-making under uncertainty. Current methods for optimizing such decisions fall short of handling the problem scale and time constraints. We argue that this is due to existing methods not exploiting the inherent structure of the organizations which solve these problems. We propose a new model called the *OrgPOMDP* (Organizational POMDP), which is based on the partially observable Markov decision process (POMDP). This new model combines two powerful representations for modeling large scale problems: hierarchical modeling and factored representations. In this paper we make three key contributions: (a) Introduce the OrgPOMDP model; (b) Present an algorithm to solve OrgPOMDP problems efficiently; and (c) Apply OrgPOMDPs to scenarios in an existing large organization, the Air and Space Operation Center (AOC). We conduct experiments and show that our OrgPOMDP approach results in greater scalability and greatly reduced runtime. In fact, as the size of the problem increases, we soon reach a point at which the OrgPOMDP approach continues to provide solutions while traditional POMDP methods cannot. We also provide an empirical evaluation to highlight the benefits of an organization implementing an OrgPOMDP policy.

Categories and Subject Descriptors

I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, and Search

General Terms

Algorithms, Management

Keywords

POMDPs, Organizations, Decision Support, Uncertainty

1. INTRODUCTION

Cite as: Adaptive Decision Support for Structured Organizations: A Case for OrgPOMDPs (Extended Abstract), Pradeep Varakantham, Nathan Scurr, Alan Carlin and Christopher Amato, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. 1149-1150.

Solving decision problems in uncertain domains with imperfect information is a difficult challenge. These problems include situations with uncertain action effects and only partial information about the current state of the environment. Partially observable Markov decision processes (POMDPs) provide a robust model for representing these problems. While many promising algorithms have been developed [1, 2, 9, 11], scalability to large real-world domains remains an open question.

Recently, work on hierarchical [5, 6, 10] and factored models [3, 4, 7, 8] has shown increased scalability by making use of inherent structure in a problem. These approaches allow the problem to be broken up into more manageable pieces which can be solved more easily by using either a hierarchy of more finely grained problems or factored problem variables which contain sets that are independent of one another. In this paper, we combine the benefits of both approaches by breaking up a large problem into a set of hierarchically related problems, each of which is made up of a factored model. This model is motivated by the need to find the best use of an organization's resources while taking into account the partially observable nature of a domain, leading us to call our model an Organizational POMDP, or OrgPOMDP. The OrgPOMDP's advantage is that it leverages the hierarchical nature of the organization and the structure in dependencies between different levels to compute policies for decision makers at various levels efficiently.

From the perspective of organizations such as Air and Space Operation Center (AOC), the OrgPOMDP is an ideal model to represent (a) organizations' control hierarchy; (b) decision problems (primarily under uncertainty) faced at each level of the control hierarchy; and most importantly (c) the interactions between decision makers at different levels of the hierarchy. Due to such rich representation of the decision problem, an OrgPOMDP policy ensures that an organization reacts to unexpected events in a coherent manner. In fact, we provide empirical evidence illustrating this very aspect in the context of AOC. It is worth noting that the OrgPOMDP model is very general, allowing a large number of hierarchical problems to be represented and solved.

Apart from presenting the OrgPOMDP model, we also introduce a novel algorithm to solve OrgPOMDPs. This algorithm provides methods to exploit the factored and hierarchical structure present in the OrgPOMDP, drastically reducing the solution complexity. Finally, we also show the performance of this solver on scenarios from the AOC domain. These results show that as the complexity of the problem increases, the benefits of the OrgPOMDP approach increase as well.

2. MODEL

To represent the domains of interest in this paper, we introduce an extension to the well known partially observable Markov decision process (POMDP) model which we call the OrgPOMDP, $\mathcal{OP}$. The model is defined as the tuple

$$\mathcal{OP} = (\mathcal{P}, \{c\mathcal{OP}_1, c\mathcal{OP}_2, \cdots\}, \mathcal{SD}, \mathcal{MD}, H)$$

with the following attributes: $\mathcal{P}$ is the standard POMDP tuple, $c\mathcal{OP}_1, c\mathcal{OP}_2, \cdots$ are child OrgPOMDPs, $\mathcal{SD}$ are state dependencies and $\mathcal{MD}$ are model dependencies. As can be noted from the definition above, $\mathcal{OP}$ is recursively defined, thus an OrgPOMDP can be represented as a tree with each "node" in the tree representing an OrgPOMDP. Without loss of generality, we assume Q nodes in this tree and each node is referred to as $\mathcal{OP}_q$.

State space dependencies, $\mathcal{SD}$: These are dependencies from child nodes to their parent nodes that arise due to the dependence between state space features. For instance in AOC type organizations, these dependencies arise because the performance of the organization as a whole (i.e. root $\mathcal{OP}$) depends on overall progress (feature in the state space of root node) achieved on various tasks, which in-turn is computed from the progress achieved by child nodes on subtasks (feature in the state space of child node).

Model dependencies: These are dependency links from a parent node to one of its child nodes. In this paper, we assume that the state and actions of a parent $\mathcal{OP}_q$ could affect all aspects of the child decision problem, except the observations.

Due to these dependencies between different levels of the hierarchy, the OrgPOMDP model is only partially specified.

3. ALGORITHM

We provide an algorithm for fully specifying and solving an OrgPOMDP problem. The key challenge in solving the OrgPOMDP is reasoning with circular dependencies that exist between the parent and child nodes in the hierarchy: (a) The model for the child nodes is constructed based on the actions selected at the parent node; and (b) Because certain features of the state space at the parent nodes are dependent on states at child nodes, the transition probabilities for parent nodes can only be computed by knowing child policies. In this paper, the key idea is to resolve the circular dependency by converting each node in the hierarchy into a fully specified POMDP and solving it. We achieve this in three steps:

(a) We start from the root of the hierarchy and move towards the leaf nodes, while initializing the POMDPs at all nodes with states, actions and observations.

(b) At the leaf nodes of the hierarchy, OrgPOMDP nodes are already full specified POMDPs. The parent nodes for the leaf nodes are not POMDPs and the models at the leaf levels can change based on the state and actions of the parent node (as explained in state and action dependencies). To account for this, we generate and solve all POMDPs corresponding to the set of states and actions of the parent. The policies thus generated are stored and used for computing state transitions for the parent POMDPs. Our first contribution in this paper is in exploiting structure in the domain to reduce the number of possible POMDPs that are generated and solved.

(c) We construct the parent model by using the policies computed at the child (corresponding to all possible state, action pairs). This stage involves simulating the execution of policy for the child and subsequently computing the transition and observation probability functions at the parent.

4. RESULTS

We conducted experiments that demonstrate that our OrgPOMDP approach is both scalable and useful. We have applied OrgPOMDPs to two realistic scenarios used by the Air and Space Operations Center (AOC): Rescue Mission and Organizational Planning. Our results in two domains show that OrgPOMDPs dramatically reduces computation time. In fact, our results show that as we shifted to the more complex domain of planning for the entire organization, we quickly reached a point where the OrgPOMDP could provide optimal solutions, whereas a traditional POMDP could not. The OrgPOMDP's advantage is that it leverages the hierarchical nature of the organization and the structure in dependencies to compute policies for decision makers at various levels efficiently.

5. REFERENCES

[1] C. Amato, B. Bonet, and S. Zilberstein. Finite-state controllers based on mealy machines for centralized and decentralized POMDPs. In *Proceedings of the 24th National Conference on Artificial intelligence*, 2010.

[2] B. Bonet and H. Geffner. Solving POMDPs: RTDP-Bel vs. point-based algorithms. In *International Joint Conference on Artificial Intelligence*, 2009.

[3] C. Boutilier and D. Poole. Computing optimal policies for partially observable decision processes using compact representations. In *American Association of Artificial Intelligence*, 1996.

[4] E. A. Hansen and Z. Feng. Dynamic programming for POMDPs using a factored state representation. In *AIPS*, 2000.

[5] E. A. Hansen and R. Zhou. Synthesis of hierarchical finite-state controllers for POMDPs. In *International Conference on Automated Planning and Scheduling*, 2003.

[6] J. Pineau, N. Roy, and S. Thrun. A hierarchical approach to POMDP planning and execution. In *ICML Workshop on Hierarchy and Memory in Reinforcement Learning*, 2001.

[7] P. Poupart. *Exploiting Structure to Efficiently Solve Large Scale Partially Observable Markov Decision Processes*. PhD thesis, Department of Computer Science, University of Toronto, 2005.

[8] H. S. Sim, K.-E. Kim, J. H. Kim, D.-S. Chang, and M.-W. Koo. Symbolic heuristic search value iteration for factored POMDPs. In *Proceedings of the 23rd National Conference on Artificial intelligence*, 2008.

[9] T. Smith and R. G. Simmons. Point-based POMDP algorithms: Improved analysis and implementation. In *Uncertainty in Artificial Intelligence*, 2005.

[10] G. Theocharous, K. Murphy, and L. P. Kaelbling. Representing hierarchical POMDPs as DBNs for multi-scale robot localization. 2004.

[11] P. Varakantham, R. Maheswaran, T. Gupta, and M.Tambe. Towards efficient computation of quality bounded solutions in POMDPs. In *International Joint Conference on Artificial Intelligence*, 2007.