Communication Management Using Abstraction in Distributed Bayesian Networks *

Jiaying Shen   Victor Lesser
Department of Computer Science
University of Massachusetts
Amherst, MA 01003-4610
{jyshen, lesser}@cs.umass.edu

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence

General Terms

DESIGN, PERFORMANCE

Keywords

Abstraction, Communication, Decentralized MDP

1. INTRODUCTION

In complex distributed applications, such as distributed interpretation, a problem is often decomposed into a set of subproblems and each subproblem is distributed to an agent who will be responsible for solving it. The existence of interactions between subproblems means that the agents cannot simply solve the subproblems individually and then combine local solutions together. In such systems, the amount of communication among agents may be very significant in order to guarantee global optimality or even global consistency . Thus, "satisficing" approaches have been developed that trade off optimality for reduced communication [2]. An important characterization of such distributed protocols is how much communication is required and the likelihood that the solution will be the same as that generated by an optimal centralized algorithm which uses all available information.

Shen et al. [3] took the satisficing approach to the next step by designing a parameterized algorithm where one can predict, for a desired confidence level in the final solution, the expected amount of communication the agents need. They studied these issues in terms of a two layered Distributed Bayesian Network, as shown in Figure 1. A decentralized Markov Decision Process (DEC-MDP) [1] can be constructed from the Bayesian Network structure to find a joint communication policy that minimizes the expected communication cost. The agents use only the necessary amount of communication to achieve the required level of solution quality.

*This material is based upon work supported by the National Science Foundation under Grant No. IIS-0219606. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

AAMAS'05, July 25-29, 2005, Utrecht, Netherlands.

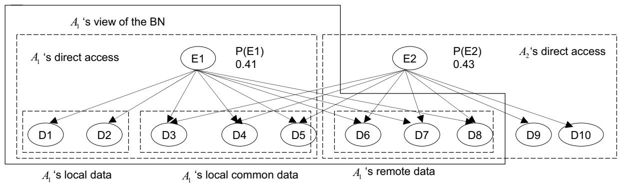

Figure 1: There are two events E_1 and E_2. Data $D_1, D_2, ...D_{10}$ are distributed between two agents. A_1 has access to $D_1, ...D_5$ and is responsible for solving E_1, while A_2 can see only $D_6, ...D_{10}$ and is responsible for solving E_2. The value of E_1 is dependent on A_1's data and vice versa. The objective is for A_1 and A_2 to decide what E_1 and E_2 are with required confidence while minimizing the expected communication cost.

In this paper, we introduce an abstraction layer into the Distributed Bayesian Network as a way of carrying more useful information in transmitted data to further reduce the number of messages that need to be sent. An algorithm is developed to automatically generate appropriate abstraction data, which reduces the expected communication cost necessary to achieve the required confidence level. Techniques are introduced to effectively incorporate this abstraction data set into the DEC-MDP framework. It is shown that the appropriate addition of abstraction data actions simplifies the DEC-MDP while reducing the expected communication cost.

2. GENERATING THE ABSTRACTION LAYER

We need to find an appropriate abstraction layer from the existing BN that, when transmitted from the remote agent, more efficiently conveys the necessary information to facilitate the local agent's problem solving. In other words, this abstraction layer, when acquired, should be able to reduce the expected communication necessary to achieve the required local confidence level. We achieve this goal by developing an algorithm that automatically generates an abstraction layer given a value combination of an agent's local data and the desired confidence level. The basic idea behind the algorithm is to find a set of logic expressions consisting of the remote data such that if at least one of the expressions is true the required confidence is reached. When given a BN and a desired confidence level, the agent generates an abstraction layer for each raw data value combination and adds them to its action options for the states that have the corresponding raw data values. The expanded DEC-MDP can be solved to generate a communication strategy. We call this approach the *all data action selection* approach.

We compared the performance of the *all action selection* approach and the *raw data action selection* approach. The cost of

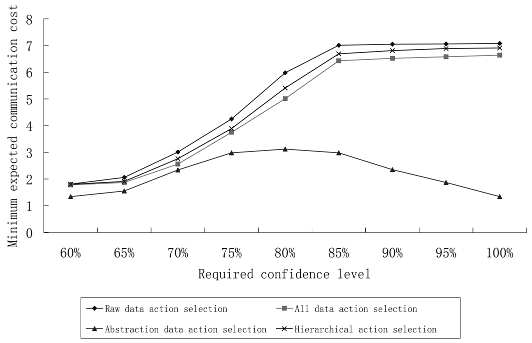

Figure 2: A comparison of the minimum expected communication cost given different action selections

sending a piece of abstraction data was equal to the cost of sending raw data. We used the Iterative Algorithm introduced in [3] to solve the DEC-MDP. We ran experiments on 100 problem structures with 2 high level events and 10 raw data (5 local to each agent) for different confidence levels. All of the networks were fully connected, which means that for both agents to have the complete evidence, 10 pieces of data needed to be transmitted. Figure 2 shows a comparison of the minimum expected communication cost generated by both systems. Column (a) in Table 1 shows the percent improvement in the expected communication cost when transmitting abstraction data in addition to the raw data. As shown, the *all data action selection* has a noticeable improvement over the *raw data action selection* approach. This illustrates that the addition of the abstraction data does help reduce the communication cost required.

3. HIERARCHICAL ACTION SELECTION

Introducing the communication actions that transmit the values of the new abstraction data leads to both a larger state space and a larger action space for the generated DEC-MDP. Column (c) in Table 1 shows the average time the *all data action selection* approach took to solve the DEC-MDP, where the average time needed for the *raw data action selection* approach equals 1.00.

First we examine the case where the agents only transfer the abstraction data between them. We call this approach the *abstraction data action selection* approach. While the size of the DEC-MDP generated is often much smaller than that of the original DEC-MDP, one major drawback of this approach is that we can no longer guarantee that the required confidence level can be reached.

We seek to combine the advantages of the *all data action selection* and the *raw data action selection* approaches to save time on solving the DEC-MDP as well as guarantee the required confidence level. We achieve this by restricting legal actions for different states. Only when the acquisition of all of the abstraction data cannot achieve the desired confidence level will an agent start to acquire the raw data in order to get the necessary information. We call this approach the *hierarchical action selection* approach.

Figure 2 compares the performance of the different approaches we have discussed. On average, the *hierarchical action selection* approach outperforms the *raw data action selection* approach. However, there are cases where the *hierarchical action selection* approach requires more communication than the *raw data action selection* does. In those BNs, it is often the case that there is a low likelihood of any of the abstraction data being true. Column (b) in Table 1 shows the amount of improvement in the minimum communication cost the *hierarchical action selection* approach gains over the *raw data action selection* approach. The pattern is similar

required confidence	(a)	(b)	(c)	(d)
60%	1.66%	0.55%	1.25	0.89
65%	9.22%	7.28%	1.53	0.77
70%	14.95%	8.31%	1.49	0.69
75%	11.76%	8.47%	1.52	0.63
80%	16.22%	9.53%	1.61	0.59
85%	8.27%	4.56%	1.41	0.65
90%	7.52%	3.40%	1.21	0.77
95%	6.80%	2.41%	1.10	0.82
100%	6.21%	2.40%	1.09	0.87

Table 1: Performance given different action selections compared to transmitting only raw data. (a) Expected communication cost improvement of *all action selection*. (b) Expected communication cost improvement of *hierarchical action selection*. (c) Time needed to solve the DEC-MDP for *all action selection* normalized by that for raw data action selection. (d) Time needed to solve the DEC-MDP for *hierarchical action selection* normalized by that for raw data action selection.

to that of column (a). Column (d) in Table 1 shows the average time needed to solve the DEC-MDP for the *hierarchical action selection* approach normalized by the average time needed for the *raw data action selection* approach. It achieves substantial savings. Even though the *hierarchical action selection* approach does not reduce the size of the action space compared to the *all action selection* approach, it does reduce the number of legal actions available to any given state. This also decreases the size of the state space because H, the communication history, has fewer possibilities. These two factors combined together contribute to the time savings, and the larger the network is, the more substantial the savings should be.

4. CONCLUSIONS

In this paper we investigated the techniques of transferring abstraction data in addition to raw data in Distributed Bayesian Networks to reduce the required communication cost. Both the improvement in the minimum expected communication cost and the time savings in solving the DEC-MDP make the *hierarchical action selection* an attractive approach to our problem, especially for the systems which require a mid-ranged confidence level. This work allows us to look at the use of abstraction to reduce communication cost from a formal perspective. We predict that the savings of the *hierarchical action selection* approach shown in this paper will be more significant for larger networks. An important extension to this work is to introduce multiple levels of abstraction, which may help reduce the difficulty in scaling up the system.

5. REFERENCES

[1] D. Bernstein, R. Givan, N. Immerman, and S. Zilberstein. The complexity of decentralized control of markov decision processes. *Mathematics of Operations Research*, 27(4):819–840, November 2002.

[2] N. Carver and V. Lesser. Domain monotonicity and the performance of local solutions strategies for CDPS-based distributed sensor interpretation and distributed diagnosis. *International Journal of Autonomous Agents and Multi-Agent Systems*, 6:35–76., 2003.

[3] J. Shen, V. Lesser, and N. Carver. Minimizing Communication Cost in a Distributed Bayesian Network using a Decentralized MDP. In *Proceedings of Second International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2003)*, volume AAMAS03, pages 678–685, Melbourne, AUS, July 2003. ACM Press.