

Final Performance Report

The Value of Communication in Decentralized Planning and Control

AFOSR Agreement Number F49620-03-1-0090

Shlomo Zilberstein, Principal Investigator
Computer Science Department
140 Governors Drive
University of Massachusetts
Amherst, MA 01003-9264

August 2006

1 Overview

This project produced a mathematical approach to the design and analysis of decentralized control with particular focus on management of communication in an uncertain environment. Target applications of this framework involve multiple decision makers that could exchange a limited amount of information because communication is either costly or risky. The main accomplishments of the project include the formulation of a general mathematical model to study such problems called decentralized POMDP; conducting an extensive complexity analysis of this model; developing the first policy iteration algorithm for solving general problems formalized as decentralized POMDPs; developing several efficient algorithms for solving special classes of decentralized POMDPs that arise in practice; developing an effective myopic approach to communication in collaborative multi-agent systems; and laying the foundations for developing learning techniques for communication. The report describes these research accomplishments and provides references to published papers and PhD dissertations that include detailed descriptions of the results.

2 Summary of Research Accomplishments

2.1 Formalizing the problem of decentralized control and analyzing its complexity

Decentralized control of cooperative systems captures the operation of a group of decision makers that share a single global objective. The difficulty in solving optimally such problems arises when the agents lack full observability of the global state of the system. In earlier work, we proved the general problem to be intractable (NEXP-complete). In this project, we identified classes of decentralized control problems whose complexity ranges between NEXP and P. In particular, we studied problems characterized by independent transitions, independent observations, and goal-oriented objective functions. Two algorithms

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE August 30, 2006	3. REPORT TYPE AND DATES COVERED Final Performance Report 01/02/2003 - 05/31/2006	
4. TITLE AND SUBTITLE The Value of Communication in Decentralized Planning and Control			5. FUNDING NUMBERS AFOSR Agreement Number F49620-03-1-0090	
6. AUTHOR(S) Shlomo Zilberstein Department of Computer Science University of Massachusetts				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Massachusetts 140 governors Drive Amherst, MA 01003			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Grant & Contract Administration 70 Butterfield Terrace University of Massachusetts Amherst MA 01003-9242			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION / AVAILABILITY STATEMENT			12b. DISTRIBUTION CODE	
13. ABSTRACT (<i>Maximum 200 Words</i>) <p>This project produced a mathematical approach to the design and analysis of decentralized control with particular focus on management of communication in an uncertain environment. Target applications of this framework involve multiple decision makers that could exchange a limited amount of information because communication is either costly or risky. The main accomplishments of the project include the formulation of a general mathematical model to study such problems called decentralized POMDP; developing the first policy iteration algorithm for solving general problems formalized as decentralized POMDPs; developing several efficient algorithms for solving special classes of decentralized POMDPs that arise in practice; developing an effective myopic approach to communication in collaborative multi-agent systems; and laying the foundations for developing learning techniques for communication. The report describes these research accomplishments and provides references to published papers and PhD dissertations that include detailed descriptions of the results.</p>				
14. SUBJECT TERMS			15. NUMBER OF PAGES	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	

have been developed to solve optimally useful classes of goal-oriented decentralized processes in polynomial time. We have also studied information sharing among the decision-makers, which can improve their performance. We identified three ways in which agents can exchange information: indirect communication, direct communication and sharing state features that are not controlled by the agents. Our analysis has shown that introducing direct or indirect communication does not change the worst-case complexity of the problem. These results provide a better understanding of the complexity of decentralized control problems that arise in practice and facilitate the development of planning algorithms for these problems. A journal paper describing these results was published in the Journal of Artificial Intelligence Research [7].

2.2 Bounded policy iteration for decentralized POMDPs

We developed a bounded policy iteration algorithm for infinite-horizon decentralized POMDPs. Policies are represented as joint stochastic finite-state controllers, which consist of a local controller for each agent. We also let a joint controller include a correlation device that allows the agents to correlate their behavior without exchanging information during execution. This leads to improved performance, particularly when the number of memory states is small. The algorithm uses a fixed amount of memory, and each iteration is guaranteed to produce a controller with value at least as high as the previous one for all possible initial state distributions. This work was presented at the ICAPS 2005 Workshop on Multi-Agent Planning and Scheduling [9] and at the 2005 International Joint Conference on Artificial Intelligence [14]. A journal paper on this topic is in preparation. Other relevant publications include: [5,8,12,15,17,18,19].

2.3 Solving efficiently special classes of DEC-POMDPs

We identified a class of problems called *transition-independent* MDPs, that captures effectively the control problem of many real-world applications of multi-agent systems. The general class consists of independent collaborating agents that are tied up by a global reward function that depends on both of their execution histories. For example, when two autonomous robots are deployed, each with its own mission, there could be important interactions between the activities they perform. This framework allows us to model activities that are either *complementary* or *redundant*. We developed a novel algorithm for solving this class of problems and examined its properties. This result was the first effective technique for solving optimally a class of decentralized MDPs. The work was presented at AAMAS 2003 and won the Best Paper Award [1]. A journal paper describing these results was published in the Journal of Artificial Intelligence Research [6].

2.4 Solving DEC-MDPs with event-driven interactions

Decentralized MDPs provide a powerful formal framework for planning in multi-agent systems, but the complexity of the general model limits its usefulness. We developed algorithms that restrict the interactions between the agents to a structured, event-driven dependency. These dependencies can model locking a shared resource or temporal enabling constraints, both of which arise frequently in practice. We have shown that the complexity of this class of problems is no harder than exponential in the number of states and doubly exponential in the number of dependencies. Since the number of dependencies is much smaller than the number of states for many problems, this is significantly better than the doubly exponential complexity of DEC-MDPs. We have also demonstrated how the *coverage-set* algorithm we developed can solve problems in this class optimally or approximately. Experimental work indicates that this solution technique is significantly faster than a naive policy search. These results were presented at AAMAS 2004 [3].

2.5 Developing myopic approaches for multi-agent communication

Choosing when to communicate is a fundamental problem in multi-agent systems. In previous work, we have shown that this problem becomes particularly hard when communication is constrained and each agent has different partial information about the overall situation. Although computing the exact value of communication is intractable, it can be estimated using a standard myopic assumption. However, this assumption—that communication is only possible at the present time—introduces error that can lead to poor agent behavior. We examined specific situations in which the myopic approach performs poorly and demonstrated an alternate approach that relaxes the assumption to improve performance. The results open up new research directions that could produce effective method for value-driven communication policies in multi-agent systems. This work received the Best Paper Award at the Intelligent Agent Technology Conference in September 2005 [16].

2.6 Learning to communicate in decentralized systems

Learning to communicate is an emerging challenge in AI research. It is known that agents interacting in decentralized, stochastic environments can benefit from exchanging information. Multiagent planning generally assumes that agents share a common means of communication; however, in building robust distributed systems it is important to address potential miscoordination resulting from misinterpretation of messages exchanged. We have laid the foundations for studying this problem, examining its properties analytically and empirically in a decision-theoretic context [2]. Solving the problem optimally is often intractable, but our approach enables agents using different languages to converge upon coordination over time. This work has been presented at AAMAS 2004 [4], BISFAI 2005 [10], a AAAI 2005 workshop [11], and at the International Joint Conference on Artificial Intelligence [13]. A journal paper describing these results is scheduled to appear in the Journal of Autonomous Agents in Multi-Agent Systems [20].

3 Personnel

In addition to the PI, five graduate students, Martin Allen, Chris Amato, Raphen Becker, Zhengzhu Feng, and Daniel Bernstein, and one postdoctoral research fellow, Dr. Claudia Goldman, were involved in this project. Daniel Bernstein, who completed his PhD in 2005, continued to work on the project as a postdoctoral research fellow until the completion of the project.

4 Publications

Note: The publications are available for download at: <http://rbr.cs.umass.edu/>

4.1 PhD Dissertations

1. Daniel S. Bernstein. “Decentralized Control of Markov Decision Processes: Algorithms and Complexity Analysis.” PhD Dissertation, Computer Science Department, University of Massachusetts Amherst, 2005. (**Nominated for the ACM Best Dissertation Award**)
2. Zhengzhu Feng. “Exploiting Problem Structures in Solving Partially Observable Markov Decision Processes.” PhD Dissertation, Computer Science Department, University of Massachusetts Amherst, 2005.

3. Raphen Becker. “Exploiting Structure in Decentralized Markov Decision Processes.” PhD Dissertation, Computer Science Department, University of Massachusetts Amherst, 2006.

4.2 Journals and Conferences

1. R. Becker, S. Zilberstein, V. Lesser, and C. V. Goldman. “Transition-Independent Decentralized Markov Decision Processes.” *Proceedings of the Second International Conference on Autonomous Agents and Multi-Agent Systems*, Melbourne, Australia, July, 2003. **(Best Paper Award)**
2. C. V. Goldman and S. Zilberstein. “Optimizing Information Exchange in Cooperative Multi-Agent Systems.” *Proceedings of the Second International Conference on Autonomous Agents and Multi-Agent Systems*, Melbourne, Australia, July, 2003.
3. R. Becker, S. Zilberstein, and V. Lesser. “Decentralized Markov Decision Processes with Event-Driven Interactions.” *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi-Agent Systems*, New York City, July 2004.
4. C. V. Goldman, M. Allen, and S. Zilberstein. “Decentralized Language Learning Through Acting.” *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi-Agent Systems*, New York City, July 2004.
5. E. A. Hansen, D. S. Bernstein, and S. Zilberstein. “Dynamic Programming for Partially Observable Stochastic Games.” *Proceedings of the Nineteenth National Conference on Artificial Intelligence*, San Jose, California, July 2004.
6. R. Becker, S. Zilberstein, V. Lesser, and C. V. Goldman. “Solving Transition-Independent Decentralized MDPs.” *Journal of Artificial Intelligence Research*, 22:423-455, 2004.
7. C. V. Goldman and S. Zilberstein. “Decentralized Control of Cooperative Systems: Categorization and Complexity Analysis.” *Journal of Artificial Intelligence Research*, 22:143-174, 2004.
8. Z. Feng and S. Zilberstein. “Region-Based Incremental Pruning for POMDPs.” *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence*, Banff, Canada. 2004.
9. D.S. Bernstein, E.A. Hansen, and S. Zilberstein. “Bounded Policy Iteration for Decentralized POMDPs.” ICAPS 2005 Workshop on Multiagent Planning and Scheduling, Monterey (ICAPS-05), California, 2005.
10. M. Allen, C.V. Goldman, and S. Zilberstein. “Learning to Communicate in Decentralized Systems.” Eighth Binennial Israeli Symposium on the Foundations of AI (BISFAI-05), Haifa, Israel, 2005.
11. M. Allen, C.V. Goldman, and S. Zilberstein. “Learning to Communicate in Decentralized Systems.” AAAI 2005 Workshop on Multi-Agent Learning, Pittsburgh, Pennsylvania, 2005.
12. D. Szer, F. Charpillet, and S. Zilberstein. “MAA*: A Heuristic Search Algorithm for Solving Decentralized POMDPs.” Proceedings of the Twenty First Conference on Uncertainty in Artificial Intelligence (UAI-05), Edinburgh, Scotland, 2005.
13. M. Allen, C.V. Goldman, and S. Zilberstein. “Language Learning in Multi-Agent Systems.” Poster presented at the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI-05), Edinburgh, Scotland, 2005.

14. D.S. Bernstein, E.A. Hansen, and S. Zilberstein. “Bounded Policy Iteration for Decentralized POMDPs.” Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI-05), Edinburgh, Scotland, 2005.
15. Z. Feng and S. Zilberstein. “Efficient Maximization in Solving POMDPs.” *Proceedings of the Twentieth National Conference on Artificial Intelligence*, Pittsburgh, Pennsylvania, July 2005.
16. R. Becker, V. Lesser, and S. Zilberstein. “Analyzing Myopic Approaches for Multi-Agent Communication.” Proceedings of Intelligent Agent Technology (IAT-05), Compiègne, France, 2005. (**Best Paper Award**)
17. C. Amato, D. S. Bernstein, and S. Zilberstein. “Finding Optimal POMDP Controllers Using Quadratically Constrained Linear Programs.” *Ninth International Symposium on Artificial Intelligence and Mathematics*, Ft. Lauderdale, Florida, 2006.
18. C. Amato, D. S. Bernstein, and S. Zilberstein. “Solving POMDPs Using Quadratically Constrained Linear Programs.” Poster in *Proceedings of the Fifth International Joint Conference on Autonomous Systems and Multiagent Systems*, Hakodate, Japan, 2006.
19. C. Amato, D. S. Bernstein, and S. Zilberstein. “Optimal Fixed-Size Controllers for Decentralized POMDPs.” *AAMAS 2006 Workshop on Multi-Agent Sequential Decision Making in Uncertain Domains (MSDM)*, Hakodate, Japan, 2006.
20. C. V. Goldman, M. Allen, and S. Zilberstein. “Learning to Communicate in a Decentralized Environment.” To appear in *Journal of Autonomous Agents and Multi-Agent Systems*, 2006.

5 Interactions and Transitions

The project team was very active in several conferences, symposia, panels, and journals. Three of the students, Daniel Bernstein, Zhengzhu Feng, and Raphen Becker, completed their PhD dissertations in 2005 and 2006. Bernstein was recognized as the Best Graduating PhD in Computer Science in 2005 and his dissertation was nominated for the ACM Best Dissertation Award.

5.1 Conference Organization

1. **2nd International Joint Conference on Autonomous Agents and Multi-Agent Systems**
PI served on the program committee.
2. **5th Workshop on Game-Theoretic and Decision-Theoretic Agents**
PI served on the program committee.
3. **3rd International Joint Conference on Autonomous Agents and Multi-Agent Systems**
PI served on the senior program committee.
4. **6th Workshop on Game-Theoretic and Decision-Theoretic Agents**
PI served on the program committee.
5. **ECAI 2004 Workshop on Multi-Agent Markov Decision Processes: Theories and Models**
PI served on the program committee.

6. **15th International Conference on Automated Planning and Scheduling (ICAPS-05)**
The PI served on the program committee of ICAPS-05, which took place in Monterey, California, in June 2005. He is also a member of the ICAPS Executive Council, which oversee the conference series.
7. **8th Binennial Israeli Symposium on the Foundations of AI (BISFAI-05)**
Claudia Goldman served as the symposium co-chair.
8. **20th National Conference on Artificial Intelligence (AAAI-05)**
PI served on the senior program committee. Daniel Bernstein served as a member of the program committee.
9. **4th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-05)**
PI served on the senior program committee. Daniel Bernstein and Claudia Goldman served as members of the program committee.
10. **Workshop on Game-Theoretic and Decision-Theoretic Agents (GTDT-05)**
PI served on the program committee.
11. **19th International Joint Conference on Artificial Intelligence (IJCAI-05)**
PI served on the program committee.
12. **9th International Symposium on Artificial Intelligence and Mathematics (AIMATH-06)**
The PI served as the chair of the program committee of AI & Math 2006, which took place in Fort Lauderdale, Florida, in January 2006. Daniel Bernstein served as the publicity chair of the symposium. Claudia Goldman served as a program committee member.
13. **AAMAS 2006 Workshop on Sequential Decision making in Uncertain Multiagent Domains**
PI served on the program committee.
14. **International Conference on Automated Planning & Scheduling**
PI served on the program committee.
15. **21st National Conference on Artificial Intelligence**
PI served on the senior program committee.
16. **AAAI 2006 Workshop on Learning for Search**
PI served on the program committee.
17. **22nd National Conference on Artificial Intelligence**
The PI serves as a member of the senior program committee of AAAI-07, which will take place in Vancouver, British Columbia, Canada, in July 2007.

5.2 Conference Participation

1. **18th International Joint Conference on Artificial Intelligence (IJCAI 2003)** [2 participants; presenting one tutorial and two papers]
2. **2nd International Joint Conference Autonomous Agents and Multi Agent Systems (AAMAS 2003)** [3 participants; presenting three papers]

3. **5th Workshop on Game-Theoretic and Decision-Theoretic Agents (GTDT 2003)** [1 participant; presenting one paper]
4. **6th European Workshop on Reinforcement Learning (EWRL 2003)** [1 participant; presenting one paper]
5. **20th International Conference on Automated Planning & Scheduling (ICAPS 2003)** [1 participant; presenting one tutorial and one paper]
6. **14th International Conference on Automated Planning and Scheduling (ICAPS-04)** [3 participants; presenting one tutorial and one workshop paper]
7. **3rd International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-04)** [5 participants; presenting two technical papers, one poster, and one workshop paper]
8. **19th National Conference on Artificial Intelligence (AAAI-04)** [2 participants; presenting one technical paper and one workshop paper]
9. **15th International Conference on Automated Planning and Scheduling (ICAPS-05)** [2 participants; presenting one workshop paper]
10. **8th Binennial Israeli Symposium on the Foundations of AI (NISFAI-05)** [2 participants; presenting one technical paper and a plenary invited talk given by the PI]
11. **20th National Conference on Artificial Intelligence (AAAI-05)** [4 participants; presenting one technical paper and one workshop paper]
12. **19th International Joint Conference on Artificial Intelligence (IJCAI-05)** [2 participants; presenting one technical paper and one poster]
13. **9th International Symposium on Artificial Intelligence and Mathematics (AIMATH-06)** [2 participants; presenting two papers]
14. **5th International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-06)** [3 participants; presenting a technical paper, a poster and one workshop paper]
15. **16th International Conference on Automated Planning and Scheduling (ICAPS-06)** [1 participant; presenting a tutorial]
16. **21th National Conference on Artificial Intelligence (AAAI-06)** [4 participants; presenting one paper]

5.3 Editorial Boards

1. The PI serves on the editorial board of two prominent journals: *Journal of Artificial Intelligence Research* and *Journal of Autonomous Agents and Multi-Agent Systems*. The PI and other members of the project team reviewed numerous papers for these and other leading journals.

5.4 Other Interactions

1. The PI maintains close collaboration ties between his lab and the MAIA group at INRIA, Nancy, France. To advance this collaboration, INRIA has provided funding for exchange of students and short visits. In Fall 2005, a graduate student from France, Daniel Szer, visited our group and contributed to our research efforts and to the objectives of this AFOSR project. In 2006, using additional NSF funding, one graduate student working on the project spent the summer at INRIA.

6 Inventions and Patent Disclosures

None.

7 Honors and Awards

1. Two papers listed in Section 4 [references 1 and 14] received Best Paper Awards from prestigious conferences.
2. Daniel Bernstein named Best Graduating PhD in Computer Science in 2005.
3. Daniel Bernstein, UMass Graduate School Fellowship, 2004-05
4. Daniel Bernstein, Finalist for an IBM fellowship, 2004-05
5. Zhengzhu Feng, First place in a track of the Probabilistic Planning Competition, 2004
6. Christopher Amato, Massachusetts Space Grant Summer Fellowship, 2004
7. Zhengzhu Feng, NASA RIACS Summer Student Research Program Fellowship, 2003
8. Zhengzhu Feng, passed the departmental portfolio exam with distinction, 2002
9. Zhengzhu Feng, UMass Graduate School Fellowship, 2002-03
10. Daniel Bernstein, NASA Graduate Student Researchers Program Fellowship, 2001-04