

REPORT DOCUMENTATION PAGE

*Form Approved
OMB No. 0704-0188*

The 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 the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to the Department of Defense, Executive Service Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE (DD-MM-YYYY) 14-01-2009		2. REPORT TYPE FINAL		3. DATES COVERED (From - To) 15-04-2005 - 14/10/2008	
4. TITLE AND SUBTITLE New Algorithms for Collaborative and Adversarial Decision Making in Partially Observable Stochastic Games				5a. CONTRACT NUMBER FA9550-05-1-0254	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Zilberstein, Shlomo				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Massachusetts 140 Governors Drive, CMPSCI Dept. Amherst MA 01003-9264				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Grant and Contract Administration 70 Butterfield Terrace University of Massachusetts Amherst MA 01003				10. SPONSOR/MONITOR'S ACRONYM(S) UMASS-Amherst	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A. APPROVED FOR PUBLIC RELEASE; DISTRIBUTION IS UNLIMITED.					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author, and not necessarily shared by AFOSR.					
14. ABSTRACT The project has produced new computational models and algorithms for coordination, prediction and planning in situations involving multiple decision makers that operate over an extended period of time in either collaborative or adversarial domains. This includes the development of the decentralized partially-observable Markov decision process (DEC-POMDP) model, memory-bounded algorithm for solving finite-horizon DEC-POMDPs, sparse representations of agent strategies using finite-state controllers, bounded policy iteration algorithms for infinite-horizon DEC-POMDPs, and algorithms for solving DEC-POMDPs using non-linear optimization methods. The project produced the best existing exact algorithms for these problems as well as scalable approximation techniques and benchmark problems that are now widely used within the multi-agent systems community. The report describes these research accomplishments and provides references to published papers and PhD dissertations that include detailed descriptions of the results.					
15. SUBJECT TERMS multi-agent systems; planning under uncertainty; Markov decision processes; sequential games; coordination.					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Shlomo Zilberstein
U	U	U	SAR	000	19b. TELEPHONE NUMBER (Include area code) 413-545-4189

Final Performance Report

New Algorithms for Collaborative and Adversarial Decision Making in Partially Observable Stochastic Games

AFOSR Agreement Number FA9550-05-1-0254

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

January 14, 2009

1 Overview

This project has focused on the development of new formal models and algorithms for decision-theoretic planning in multi-agent settings. We have studied both collaborative and adversarial domains in which decision makers interact over time and must base their decisions on incomplete and noisy information about the overall situation. This problem arises in many application domains such as multi-robot coordination, distributed management of servers or a power grid, weapon allocation problems, distributed information gathering as well as the operation of complex human organizations. Our results include both complexity analysis of the formal models and the development of the first set of exact and approximate algorithms for solving these complex decision problems. These new algorithms address some fundamental drawbacks of existing approaches and eliminate the need to make some common simplifying assumptions such as: limiting the approach to just two players or zero-sum games; considering just a few steps in a memoryless environment; assuming that decision makers have perfect information or that they can share information all the time; or assuming that opponents are perfectly rational.

To address this challenge, we have developed a formal framework that integrates game-theoretic solution techniques with partially observable Markov decision processes, a model that is widely used for decision-theoretic planning in artificial intelligence and operations research. We have employed two formal models. A Decentralized Partially Observable Markov Decision Process (DEC-POMDP) is designed for collaborative systems in which all the decision makers share the same objective or utility function. A Partially Observable Stochastic Game (POSG) is a proper extension of a DEC-POMDP that is designed for competitive systems in which decision makers have separate, possibly conflicting, objectives. A comprehensive complexity study of these models has shown that they are intractable (NEXP-complete) even when two agents are involved. Consequently, we have identified useful classes of these general models that have lower complexity. We developed the first exact dynamic programming algorithms for these problems, but not surprisingly, these algorithms can only solve small “toy” problems. Thus, in the last two years of the project, the focus has been on the design of memory-bounded approximation techniques that can produce good results and provide an error bound.

20090323116

The project produced a wide range of approximation methods for solving these hard computational problems. This includes the development of memory-bounded dynamic programming algorithms for solving finite-horizon DEC-POMDPs, sparse representation of agent strategies using finite-state controllers, bounded policy iteration algorithms for infinite-horizon DEC-POMDPs, and the development of algorithms for solving DEC-POMDPs using non-linear optimization methods. The project produced the best existing scalable approximation methods and benchmark problems that are now widely used within the multi-agent systems research community. 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 Developing the first policy iteration algorithm for decentralized POMDPs

We developed the first bounded policy iteration algorithm for infinite-horizon decentralized POMDPs. The algorithm uses stochastic finite-state controllers to represent policies. The solution can include a correlation device, which allows agents to correlate their actions without communicating. This approach alternates between expanding the controller and performing value-preserving transformations, which modify the controller without sacrificing value. We developed two efficient value-preserving transformations: one can reduce the size of the controller and the other can improve its value while keeping the size fixed. Empirical results demonstrate the usefulness of value-preserving transformations in increasing value while keeping controller size to a minimum. Initial papers describing this approach were presented at ICAPS-05 [1] and IJCAI-05 [2]. To broaden the applicability of the approach, we also developed a heuristic version of the policy iteration algorithm, which does not guarantee convergence to optimality. This algorithm further reduces the size of the controllers at each step by assuming that probability distributions over the other agents' actions are known. While this assumption may not hold in general, it helps in practice to produce higher quality solutions in a range of test problems. A comprehensive journal paper on this approach was accepted for publication in the *Journal of AI Research* [27].

2.2 Solving POMDPs using quadratically constrained linear programs

As part of the efforts to develop scalable algorithms for solving partially observable Markov decision processes, a new approach was developed that formulates the problem as a quadratically constrained linear program (QCLP). This representation allows a wide range of powerful nonlinear programming algorithms to be used to find solutions for decentralized POMDPs. Although these solvers do not guarantee global optimality, we got very good results using off-the-shelf optimization software for solving QCLPs. Our approach produced consistent solution quality improvement over the state-of-the-art techniques. We can achieve these better results using smaller policies and less memory, and thus use less computation time than alternative methods. The initial work on this approach was presented at ISAIM-06 [6], AAMAS-06 [7], IJCAI-07 [10] and UAI-07 [15]. A comprehensive journal paper on this approach was accepted for publication in the *Journal of AI Research* [29].

2.3 Solving decentralized decision problems using heuristic search

In collaboration with colleagues from INRIA (France), we have developed a multi-agent variant of A* called MAA*. The algorithm is the first complete and optimal heuristic search algorithm for solving decentralized partially-observable Markov decision problems (DEC-POMDPs) with finite horizon. The algorithm is suitable for computing optimal plans for a cooperative group of agents that operate in a stochastic environment such as multi-robot coordination, network traffic control, or distributed resource allocation. The solution is based on a synthesis of classical heuristic search and decentralized control

theory. Experimental results show that MAA* has significant advantages. This work was presented at the Conference on Uncertainty in Artificial Intelligence (UAI-05) [2].

2.4 Managing costly communication in decentralized systems

Choosing when to communicate is a fundamental problem in multi-agent systems. This problem becomes particularly hard when communication is constrained or costly and each agent has different partial information about the overall situation. We developed a decision-theoretic approach to decide *when* to communicate based on the *value of communication* (VoC). 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 developed an alternate approach that relaxes the assumption to improve performance. The results provide an effective method for value-driven communication policies in a useful class of DEC-POMDPs. A paper describing this approach received the best paper award at the *IEEE/WIC/ACM International Conference on Intelligent Agent Technology* in 2005 [5]. A more comprehensive study of this method is scheduled to appear in *Computational Intelligence* [28]. In related work on communication we have examined ways to manage communication when the agents must learn the communication language while acting [4, 9]. A comprehensive study of the notion of *communication-based decomposition mechanisms*—an approach to simplify a decentralized MDP by breaking it into multiple single-agent problems—was published in the *Journal of AI Research* in 2008 [21].

2.5 Developing memory-bounded dynamic-programming algorithm for DEC-POMDPs

One of the important outcomes of the project is the first memory-bounded dynamic-programming algorithm (MBDP) for solving finite-horizon DEC-POMDPs. The algorithm uses a set of heuristics to identify relevant points of the infinitely large belief space. Using these belief points, it successively selects the best joint policies for each decision horizon. The initial algorithm was presented at IJCAI-07 [11]; an improved version was presented at UAI-07 [14]. We subsequently improved the implementation of the algorithm and its scalability with respect to the number of observations each agent can make. The resulting algorithm is extremely efficient, having linear time and space complexity with respect to the horizon length and the number of observations. Experimental results show that it can handle horizons that are multiple orders of magnitude larger than what was previously possible, while achieving the same or better solution quality in a small fraction of the runtime. To evaluate the effectiveness of these improvements, we introduced a new, larger benchmark problem. Experimental results show that despite the high complexity of decentralized POMDPs, scalable solution techniques such as MBDP perform surprisingly well. A comprehensive journal paper that compares the various solution techniques for DEC-POMDPs appeared in the journal *Autonomous Agents and Multi-Agent Systems* in 2008 [20].

2.6 Solving average-reward decentralized Markov decision processes

We have identified several application domains in which the standard approach to describing the objectives of the decision makers does not work well. The standard approach is based on optimizing discounted cumulative reward, but optimizing average reward is sometimes a more suitable criterion. In these problems, the system operates over an extended period of time and the main objective is to perform consistently well over the long run. The more common discounted reward criterion usually leads to poor long-term performance in such domains. We formalized a class of such problems and analyzed its characteristics, showing that it is NP complete and that optimal policies are deterministic. This analysis provided the foundation for designing two optimal algorithms. Both methods are based on formulating the problem as a mathematical program. Experimental results with a standard problem from the literature illustrate the

efficiency of these new solution techniques. A paper describing this work was presented at IJCAI-07 [12].

2.7 Anytime coordination using separable bilinear programs

For DEC-POMDP problems that exhibit a great degree of independence between the decision makers, we have previously developed an approach called the Coverage Set Algorithm (CSA). Essentially, the agents can be modeled in this case as separate MDPs with an overall reward function that depends on the global state. CSA works by first enumerating the policies of one agent that are best responses to at least one policy of the other agent, that is, policies that are not dominated. Then the algorithm searches over these policies to get the best joint policy for all agents. Empirically, CSA was shown to be quite efficient, solving relatively large problems. It also exhibits good anytime performance: When solving a multi-rover coordination problem, a solution value within 1% of optimal is found within 1% of the total execution time on average. Unfortunately, this is only known in hindsight once the optimal solution is found. Additionally, the algorithm has several drawbacks. It is numerically unstable and its complexity increases exponentially with the number of best-response policies. Runtime varies widely over different problem instances. Finally, the algorithm is limited to a relatively small subclass of distributed coordination problems. As part of this project, we improved this technique in several important ways. First, we presented a reformulation of CSA – using separable bilinear programs – that is more general, more efficient, and easier to implement. We also derived an error bound using the convexity of the best-response function, without relying on the optimal solution. The new algorithm exhibits excellent anytime performance, making it suitable for time-constrained situations. Finally, we derived offline bounds on the approximation error and developed a general method for automatic dimensionality reduction. This work was presented at AAAI-07 [17]. A comprehensive journal paper on this method has been accepted for publication in the *journal of Artificial Intelligence Research* [26].

3 Personnel

In addition to the Principal Investigator, the project personnel includes Prof. Eric Hansen at Mississippi State University and seven graduate students: Martin Allen, Christopher Amato, Daniel Bernstein, Alan Carlin, Akshat Kumar, Marek Petrik, and Sven Seuken. Daniel Bernstein, who completed his PhD in 2005, continued to work on the project as a postdoctoral research fellow for one year.

4 Publications

Note: The publications are available for download at:

<http://anytime.cs.umass.edu/shlomo/Publications.html>

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 in 2005. Received an Honorable Mention for the ICAPS Best Dissertation Award in 2007)
2. Raphen Becker. “Exploiting Structure in Decentralized Markov Decision Processes.” PhD Dissertation, Computer Science Department, University of Massachusetts Amherst, 2006.
3. Martin Allen. “Agent Interactions in Decentralized Environments.” PhD Dissertation, Computer Science Department, University of Massachusetts Amherst, 2008.

4.2 Journals and Conferences

1. D.S. Bernstein, E.A. Hansen, and S. Zilberstein. "Bounded Policy Iteration for Decentralized POMDPs." *ICAPS 2005 Workshop on Multiagent Planning and Scheduling (ICAPS-05)*, Monterey, California, 2005.
2. 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.
3. 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.
4. 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.
5. R. Becker, V. Lesser, and S. Zilberstein. "Analyzing Myopic Approaches for Multi-Agent Communication." *Proceedings of Intelligent Agent Technology (IAT-05)*, Compiègne, France, 2005.
(Received the Best Paper Award)
6. C. Amato, D.S. Bernstein, and S. Zilberstein. "Finding Optimal POMDP Controllers Using Quadratically Constrained Linear Programs." *Proceedings of the Ninth International Symposium on Artificial Intelligence and Mathematics (ISAIA-06)*, Ft. Lauderdale, Florida, January, 2006.
7. C. Amato, D.S. Bernstein, and S. Zilberstein. "Solving POMDPs Using Quadratically Constrained Linear Programs." *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-06)*, Hakodate, Japan, May, 2006.
8. 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 (AAMAS-06)*, Hakodate, Japan, May, 2006.
9. C. V. Goldman, M. Allen, and S. Zilberstein. "Learning to Communicate in a Decentralized Environment." *Autonomous Agents and Multi-Agent Systems*, 15(1):47–90, 2007.
10. C. Amato, D. S. Bernstein, and S. Zilberstein. "Solving POMDPs Using Quadratically Constrained Linear Programs." *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, Hyderabad, India, January, 2007.
11. S. Seuken and S. Zilberstein. "Memory-Bounded Dynamic Programming for DEC-POMDPs." *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, Hyderabad, India, January, 2007.
12. M. Petrik and S. Zilberstein. "Average-Reward Decentralized Markov Decision Processes." *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, Hyderabad, India, January, 2007.
13. C. Amato, A. Carlin, and S. Zilberstein. "Bounded Dynamic Programming for Decentralized POMDPs." *AAMAS 2007 Workshop on Multi-Agent Sequential Decision Making in Uncertain Domains (AAMAS-07)*, Honolulu, Hawaii, May, 2007.

14. S. Seuken and S. Zilberstein. "Improved Memory-Bounded Dynamic Programming for Decentralized POMDPs." *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence (UAI-07)*, Vancouver, British Columbia, July 2007.
15. C. Amato, D.S. Bernstein, and S. Zilberstein. "Optimizing Memory-Bounded Controllers for Decentralized POMDPs." *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence (UAI-07)*, Vancouver, British Columbia, July, 2007.
16. M. Allen and S. Zilberstein. "Agent Influence as a Predictor of Difficulty for Decentralized Problem-Solving." *Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI-07)*, Vancouver, British Columbia, July, 2007.
17. M. Petrik and S. Zilberstein. "Anytime Coordination Using Separable Bilinear Programs." *Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI-07)*, Vancouver, British Columbia, July, 2007.
18. E. Hansen. "Indefinite-Horizon POMDPs with Action-Based Termination." *Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI-07)*, Vancouver, British Columbia, July, 2007.
19. M. Petrik and S. Zilberstein. "A Successive Approximation Algorithm for Coordination Problems." *Proceedings of the Tenth International Symposium on Artificial Intelligence and Mathematics (ISAIM-08)*, Ft. Lauderdale, Florida, 2008.
20. S. Seuken and S. Zilberstein. "Formal Models and Algorithms for Decentralized Decision Making under Uncertainty." *Autonomous Agents and Multi-Agent Systems*, 17(2):190-250, 2008.
21. C.V. Goldman and S. Zilberstein. "Communication-Based Decomposition Mechanisms for Decentralized MDPs." *Journal of Artificial Intelligence Research*, 32:169-202, 2008.
22. E. Hansen. "Sparse Stochastic Finite-State Controllers for POMDPs." *Proceedings of Twenty-Fourth Conference on Uncertainty in Artificial Intelligence (UAI-08)*, Helsinki, Finland, 2008.
23. C. Amato, D.S. Bernstein, and S. Zilberstein. "Optimizing Fixed-Size Stochastic Controllers for POMDPs." *AAAI Workshop on Advancements in POMDP Solvers (AAAI-08)*, Chicago, Illinois, 2008.
24. A. Carlin and S. Zilberstein. "POMDP and DEC-POMDP Point-Based Observation Aggregation." *AAAI Workshop on Advancements in POMDP Solvers (AAAI-08)*, Chicago, Illinois, 2008.
25. M. Allen, M. Petrik, and S. Zilberstein. "Interaction Structure and Dimensionality in Decentralized Problem Solving." Technical Report 08-11, Computer Science Department, University of Massachusetts, 2008.
26. M. Petrik and S. Zilberstein. "A Bilinear Programming Approach for Multiagent Planning." To appear in *Journal of Artificial Intelligence Research*, 2009.
27. D. Bernstein, C. Amato, E. A. Hansen, and S. Zilberstein. "Policy Iteration for Decentralized Control of Markov Decision Processes." To appear in *Journal of Artificial Intelligence Research*, 2009.
28. R. Becker, A. Carlin, V. Lesser, and S. Zilberstein. "Analyzing Myopic Approaches for Multi-Agent Communication." To appear in *Computational Intelligence*, 2009.
29. C. Amato, D. S. Bernstein, and S. Zilberstein. "Optimizing Fixed-Size Stochastic Controllers for POMDPs and Decentralized POMDPs." To appear in *Autonomous Agents and Multi-Agent Systems*, 2009.

5 Interactions and Transitions

The project team was very active in several conferences, symposia, panels, and journals. Three of the students, Daniel Bernstein, Raphen Becker and Martin Allen, have completed their PhD dissertations. Team members were engaged in several international collaborations and received several awards. These interactions, which help disseminate the results of the project, are summarized below.

5.1 Editorial Positions

1. The PI is currently the Associate Editor-in-Chief of the *Journal of Artificial Intelligence Research*, one of the top journals in the field of AI. He has been serving on the editorial board of the journal since 2002.
2. The PI serves on the editorial board of two other journals: *Autonomous Agents and Multi-Agent Systems* and *Annals of Mathematics and Artificial Intelligence*.

5.2 Participation in Conference and Workshop Organization

1. **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 was also a member of the ICAPS Executive Council, which oversees this conference series.
2. **20th National Conference on Artificial Intelligence (AAAI-05)**
The PI and Co-PI served as members of the senior program committee of AAAI-05, which took place in Pittsburgh in July 2005. Daniel Bernstein served as a member of the program committee.
3. **4th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-05)**
The PI served as a member of the senior program committee of AAMAS-05, which took place in Utrecht, The Netherlands, in July 2005. Daniel Bernstein served as a member of the program committee.
4. **Workshop on Game-Theoretic and Decision-Theoretic Agents (GTDT-05)**
The PI served on the program committee of GTDT-05, which took place in Edinburgh, Scotland, in July 2005.
5. **19th International Joint Conference on Artificial Intelligence (IJCAI-05)**
The PI served as a member of the program committee of IJCAI-05, which took place in Edinburgh, Scotland, in August 2005.
6. **9th International Symposium on Artificial Intelligence and Mathematics (AIMATH-06)**
The PI is the chair of the program committee of AI & Math 2006, which will take place in Fort Lauderdale, Florida, in January 2006. Daniel Bernstein serves as the publicity chair of the symposium.
7. **16th International Conference on Automated Planning and Scheduling (ICAPS-06)**
The PI served on the program committee of ICAPS-06, which took place in the Lake District, UK, in June 2006. He is also an officer of the ICAPS Executive Council, which oversees this conference series.
8. **21st National Conference on Artificial Intelligence (AAAI-06)**
The PI served as a member of the senior program committee of AAAI-06, which took place in Boston in July 2006.

9. **AAMAS 2006 Workshop on Sequential Decision Making in Uncertain Domains**
The PI served as a member of the program committee of this workshop, which took place in Hakodate, Japan, in May 2006.
10. **AAAI 2006 Workshop on Learning for Search**
The PI served as a member of the program committee of this workshop, which took place in Boston in July 2006.
11. **22nd National Conference on Artificial Intelligence**
The PI served on the senior program committee of AAAI-07, which took place July 22-26, 2007, in Vancouver, British Columbia, Canada. The Co-PI served on the program committee.
12. **6th International Conference on Autonomous Agents and Multiagent Systems**
The PI and Co-PI served on the senior program committee of AAMAS-07, which took place May 14-18, 2007, in Honolulu, Hawaii.
13. **17th International Conference on Automated Planning and Scheduling**
The PI and Co-PI served on the program committee of ICAPS-07, which took place September 22-26, 2007, in Providence, Rhode Island.
14. **AAMAS 2007 Workshop on Multi-Agent Sequential Decision Making in Uncertain Domains**
The PI served on the program committee of this workshop, which took place May 14-18, 2007, in Honolulu, Hawaii.
15. **AAMAS 2007 Workshop on Metareasoning in Agent-Based Systems** The PI served on the program committee of this workshop, which took place May 14-18, 2007, in Honolulu, Hawaii.
16. **AAMAS 2007 Workshop on Coordinating Agents' Plans and Schedules**
The PI served on the program committee of this workshop, which took place May 14-18, 2007, in Honolulu, Hawaii.
17. **AAAI 2007 Spring Symposium on Game Theoretic and Decision Theoretic Agents**
The PI served on the program committee of GTDT-07, which took place March 26-28, 2007, Stanford University, California.
18. **10th International Symposium on Artificial Intelligence and Mathematics**
The PI and Co-PI served on the program committee of ISAIM-08, which took place in January 2008, Fort Lauderdale, Florida.
19. **AAMAS 2008 Workshop on Multi-Agent Sequential Decision Making in Uncertain Domains**
The PI served on the program committee of this workshop, which took place in May, 2008, Estoril, Portugal.
20. **7th International Conference on Autonomous Agents and Multi-Agent**
The PI and Co-PI served on the program committee of AAMAS-08, which took place in May, 2008, Estoril, Portugal.
21. **AAAI 2008 Workshop on Metareasoning: Thinking about Thinking**
The PI served on the organizing committee of this workshop, which took place in July, 2008, Chicago, Illinois.
22. **1st International Symposium on Search in Artificial Intelligence and Robotics**
The PI served on the organizing committee of this symposium, which took place in July, 2008, Chicago, Illinois.
23. **23rd National Conference on Artificial Intelligence**
The PI served on the senior program committee of AAAI-08, which took place in July, 2008, Chicago, Illinois. The Co-PI served on the program committee.

24. **18th International Conference on Automated Planning & Scheduling**

Co-PI Eric Hansen served as co-chair of the program committee of ICAPS-08, which took place in September, 2008, Sydney, Australia. The PI served on the program committee.

25. **ICAPS 2008 Workshop on Multiagent Planning**

The PI served as co-organizer of this Workshop, which took place in September, 2008, Sydney, Australia.

5.3 Other Interactions

1. The PI has maintained 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. The PI has participated in a multi-institutional NSF grant that provided additional funding for this collaboration. These activities contributed directly to this project and enabled us to host several visitors from France and to send 3 of the graduate students who worked on this project for internships at INRIA.
2. The PI has served as member and conference liaison of the ICAPS Executive Council. The council oversees the annual ICAPS conference, which is the premier venue for researchers and practitioners in the area of automated planning and scheduling. The PI is currently the President Elect of the organization.

6 Inventions and Patent Disclosures

None.

7 Honors and Awards

1. One of our publications on “Analyzing Myopic Approaches to Multi-Agent Communication” [4] received Best Paper Awards from the *IEEE/WIC/ACM International Conference on Intelligent Agent Technology* in 2005. There were 305 submissions and 55 accepted papers at the conference. One paper received the award.
2. Graduate student Daniel Bernstein has been recognized as the Best Graduating PhD in Computer Science in 2005. One student was selected from each department within the College of Natural Science and Mathematics. Dan was also nominated by the Computer Science Department for the ACM Best Doctoral Dissertation Award.
3. Mark Gruman, an undergraduate student who completed his honors project under the supervision of the PI, was recognized as the Best Graduating Student in AI in 2005. A total of six students were recognized in different areas of computer science.
4. Daniel Bernstein’s Ph.D. Dissertation, that formed the foundation of this project, received an Honorable Mention for the 2007 ICAPS Best Dissertation Award. ICAPS runs the premier conference on Automated Planning and Scheduling. The 2007 award is for dissertations completed in the previous two years. The awards committee noted Daniel Bernstein for “his highly innovative research on planning under uncertainty for multiple agents introducing and characterizing a new framework of decentralized MDPs.”

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The 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 the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to the Department of Defense, Executive Service Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE (DD-MM-YYYY) 14-01-2009		2. REPORT TYPE FINAL		3. DATES COVERED (From - To) 15-04-2005 - 14/10/2008	
4. TITLE AND SUBTITLE New Algorithms for Collaborative and Adversarial Decision Making in Partially Observable Stochastic Games				5a. CONTRACT NUMBER FA9550-05-1-0254	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
6. AUTHOR(S) Zilberstein, Shlomo				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Massachusetts 140 Governors Drive, CMPSCI Dept. Amherst MA 01003-9264				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Grant and Contract Administration 70 Butterfield Terrace University of Massachusetts Amherst MA 01003				10. SPONSOR/MONITOR'S ACRONYM(S) UMASS-Amherst	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A. APPROVED FOR PUBLIC RELEASE; DISTRIBUTION IS UNLIMITED.					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author, and not necessarily shared by AFOSR.					
14. ABSTRACT The project has produced new computational models and algorithms for coordination, prediction and planning in situations involving multiple decision makers that operate over an extended period of time in either collaborative or adversarial domains. This includes the development of the decentralized partially-observable Markov decision process (DEC-POMDP) model, memory-bounded algorithm for solving finite-horizon DEC-POMDPs, sparse representations of agent strategies using finite-state controllers, bounded policy iteration algorithms for infinite-horizon DEC-POMDPs, and algorithms for solving DEC-POMDPs using non-linear optimization methods. The project produced the best existing exact algorithms for these problems as well as scalable approximation techniques and benchmark problems that are now widely used within the multi-agent systems community. The report describes these research accomplishments and provides references to published papers and PhD dissertations that include detailed descriptions of the results.					
15. SUBJECT TERMS multi-agent systems; planning under uncertainty; Markov decision processes; sequential games; coordination.					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)
U	U	U	SAR	000	Shlomo Zilberstein 413-545-4189