NASA GSRP TRAINING GRANT

FINAL SUMMARY OF RESEARCH

Grant Title:

A Decision-Theoretic Approach to Autonomous Planetary Rover Control

#NGT 2-52276

Principal Investigator:

Shlomo Zilberstein

University of Massachusetts - Amherst
Department of Computer Science
140 Governors Drive, Rm 100
Amherst MA 01003

Period Covered:

July 1, 2001 – July 30, 2004

NASA GSRP Final Report

Daniel S. Bernstein

September 20, 2004

This is a final report for the NASA GSRP Fellowship that supported my graduate studies from June 2001 to June 2004. I plan to graduate in 2005 from the University of Massachusetts with my Ph.D. in computer science. The tentative title for my thesis is "Decentralized Control of Markov Decision Processes," and my advisor is Shlomo Zilberstein. I am currently seeking employment as a professor at a research-oriented university. My achievements for the period during which I held the fellowship are listed below. I accomplished the objectives stated in my original and renewal proposals, in addition to completing other projects.

Research Projects

- *Decentralized Control of Markov Decision Processes*, University of Massachusetts Amherst, 2000-present.
  - We studied the complexity of decentralized control of Markov decision processes, and developed algorithms for finding optimal control policies.
- *Scheduling Contract Algorithms*, University of Massachusetts Amherst, 2001-present.
  - We developed an optimal method for scheduling runs of a contract anytime algorithm (one that takes the deadline as input) in situations where the deadline is unknown, multiple problem instances must be solved, and a multi-processor machine is available.
- *Planetary Rover Control as a Markov Decision Process*, University of Massachusetts Amherst and NASA Ames Research Center, 2000-present.
  - We used the Markov decision process framework to formalize and solve problems in planetary rover control.
- *Adaptive Peer Selection*, University of Massachusetts Amherst, 2001-present.
  - We used reinforcement learning to maximize the expected download speed for a client in a peer-to-peer file sharing system.

Publications

Refereed Journal Publications

- Daniel S. Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. The Complexity of Decentralized Control of Markov Decision Processes. *Mathematics of Operations Research*, 27(4):819-840, 2002.

Refereed Conferences

- Eric A. Hansen, Daniel S. Bernstein, and Shlomo Zilberstein. Dynamic Programming for Partially Observable Stochastic Games. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence*, San Jose, California, 2004.

- Daniel S. Bernstein, Lev Finkelstein, and Shlomo Zilberstein. Contract Algorithms and Robots on Rays: Unifying Two Scheduling Problems. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, Acapulco, Mexico, 2003.

- Daniel S. Bernstein, Theodore J. Perkins, Shlomo Zilberstein, and Lev Finkelstein. Scheduling Contract Algorithms on Multiple Processors. In *Proceedings of the 18th National Conference on Artificial Intelligence*, 702-706, Edmonton, Canada, 2002.

- Daniel S. Bernstein and Shlomo Zilberstein. Reinforcement Learning for Weakly-Coupled MDPs and an Application to Planetary Rover Control. In *Proceedings of the 6th European Conference on Planning*, Toledo, Spain, 2001.

Refereed Workshops and Symposia

- Daniel S. Bernstein, Zhengzhu Feng, Brian Levine, and Shlomo Zilberstein. Adaptive Peer Selection. In *Proceedings of the International Workshop on Peer-to-Peer Systems*, Berkeley, California, 2003.

- Shlomo Zilberstein, Richard Washington, Daniel S. Bernstein, and Abdel-Illah Mouaddib. Decision-Theoretic Control of Planetary Rovers. In *Proceedings of the Dagstuhl Workshop on Plan-based Control of Robotic Agents*, Wadern, Germany, 2001.

- Daniel S. Bernstein, Shlomo Zilberstein, Richard Washington, and John L. Bresina. Planetary Rover Control as a Markov Decision Process. In *Proceedings of the 6th International Symposium on Artificial*

Intelligence, Robotics, and Automation in Space, Montreal, Canada, 2001.

- Daniel S. Bernstein, Shlomo Zilberstein, Richard Washington, and John L. Bresina. Planetary Rover Control as a Markov Decision Process. In *Proceedings of the AAAI Spring Symposium on Game Theoretic and Decision Theoretic Agents*, Stanford, California, 2001.

Presentations

Invited Talks

- Dynamic Programming for Partially Observable Stochastic Games. *Massachusetts Institute of Technology*, Cambridge, Massachusetts, June 2004.
- Dynamic Programming for Partially Observable Stochastic Games. *NASA Ames Research Center*, Moffett Field, California, March 2004.
- Coordination and Control of Collaborative Agents Using Decentralized POMDPs. *INFORMS Annual Meeting*, Atlanta, Georgia, October 2003.

Talks at Conferences and Workshops

- Generalized Dynamic Programming for Decentralized POMDPs. *Sixth European Workshop on Reinforcement Learning*, Nancy, France, September 2004.
- Contract Algorithms and Robots on Rays: Unifying Two Scheduling Problems. *18th International Joint Conference on Artificial Intelligence*, Acapulco, Mexico, August 2003.
- Adaptive Peer Selection. *2nd International Workshop on Peer-to-Peer Systems*, Berkeley, California, February 2003.
- Reinforcement Learning for Fast Peer-to-Peer File Transfers. *State-Action-Reward Day in New England*, Amherst, Massachusetts, May 2002.
- Planetary Rover Control as a Markov Decision Process. *6th International Symposium on Artificial Intelligence, Robotics, and Automation in Space*, Montreal, Canada, June 2001.

POSTER PRESENTATIONS AT CONFERENCES AND WORKSHOPS

- Scheduling Contract Algorithms on Multiple Processors. *18th National Conference on Artificial Intelligence*, Edmonton, Canada, July 2002.
- Reinforcement Learning for Weakly-Coupled MDPs and an Application to Planetary Rover Control. *ICML Workshop on Hierarchy and Memory in Reinforcement Learning*, Williamstown, Massachusetts, June 2001.

Professional Activities

- Program committee member for *Autonomous Agents and Multi-Agent Systems*, 2003-2004.
- Organizing committee member for *State-Action-Reward Day in New England*, 2002.
- Referee for *Artificial Intelligence Journal* (2001).
- Referee for various conferences: *European Conference on Artificial Intelligence* (2002), *European Symposium on Algorithms* (2001), *International Joint Conference on Artificial Intelligence* (2003), *National Conference on Artificial Intelligence* (2002), *Neural Information Processing Systems* (2003).