Simulation-Based Methods for Control and Optimization

Michael Fu and Steve Marcus
ISR collaborators: Hyeong-Soo Chang, Vahid Ramezani, Jiaqiao Hu, Enlu Zhou

30 THIRTY YEARS OF SYSTEMS RESEARCH EXCELLENCE

INSTITUTE FOR SYSTEMS RESEARCH
A. JAMES CLARK SCHOOL OF ENGINEERING

History: ISR brought us together, allowing us to unite our complementary expertise in stochastic simulation and Markov decision processes. We have co-advised numerous ISR Ph.D. students and postdoctoral fellows, who have gone on to academic positions in Seoul, Bangalore, Stony Brook, Illinois, Georgia Tech.

Future work: (led by new postdoctoral fellow joining research group: Dr. Prashanth L.A.)

- New algorithms incorporating risk, based on cumulative prospect theory
- Reinforcement learning techniques
- Linking local search with global optimization

Adaptive Sampling Algorithm for Solving Markov Decision Processes

- Applicable to dynamic stochastic optimization problems that are difficult to model, or that have no analytical model but where simulations are available
- Goal: design & analysis of algorithms to estimate value function efficiently
- Use *adaptive sampling:* decide which action to sample next (from given state at given time) based on bandit model —tradeoff between exploration & exploitation
- Applications: manufacturing (preventive maintenance, capacity expansion in semiconductor fab), financial engineering (pricing, hedging, risk management)

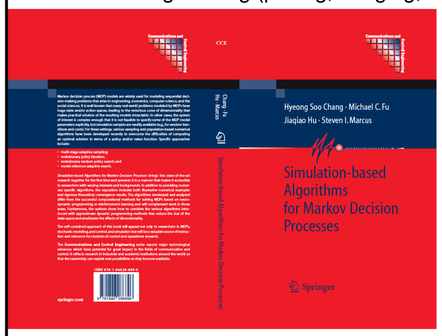

H.S. Chang, M.C. Fu, J. Hu, and S.I. Marcus, "An Adaptive Sampling Algorithm for Solving Markov Decision Processes," *Operations Research*, Vol.53, No.1, 126–139, 2005.
H.S. Chang, M.C. Fu, J. Hu, and S.I. Marcus, "An Asymptotically Efficient Simulation-Based Algorithm for Finite Horizon Stochastic Dynamic Programming," *IEEE Transactions on Automatic Control*, Vol.52, No.1, 89–94, 2007.

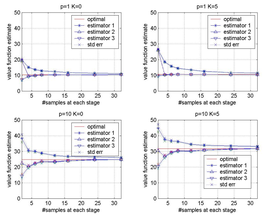

Population Based Evolutionary Approaches for Solving Markov Decision Processes

- Large action space setting; alternative to policy improvement in policy iteration
- Goal: find optimal (or good) policies
- Approach: update *population* of policies as opposed to single policy
- Key ideas:
  - Avoid optimization over entire action space (parallelizable)
  - Monotonicity among elite policies
  - Convergence w.p. 1 to optimal value function

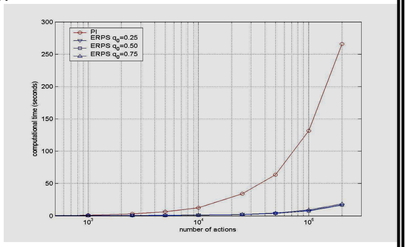

- Two methods of generating elite policy: (i) Policy Switching; (ii) Policy Improvement with Cost Swapping (PICS)

H.S. Chang, H.-G. Lee, M.C. Fu, and S.I. Marcus, "Evolutionary Policy Iteration for Solving Markov Decision Processes," *IEEE Transactions on Automatic Control*, Vol.50, No.11, 1804–1808, 2005.
J. Hu, M.C. Fu, V. Ramezani, and S.I. Marcus, "An Evolutionary Random Search Algorithm for Solving Markov Decision Processes," INFORMS Journal on Computing, Vol.19, No.2, 161–174, 2007.

Simulation-Based Methods for Partially Observed Markov Decision Processes

- Applicable to general continuous-state POMDPs
- Past work almost all for discrete-state POMDPs
- Provable error bounds and computationally efficient
- New simulation-based Projection Particle Filter for state estimation
- Projects belief state onto exponential family
- Successful application to inventory control problem

E. Zhou, M.C. Fu, and S.I. Marcus, "Solving Continuous-State POMDPs via Density Projection," *IEEE Transactions on Automatic Control*, Vol.55, No.5, 1101–1116, 2010.

E. Zhou, M.C. Fu, and S.I. Marcus, "A Particle Filtering Framework for a Class of Randomized Optimization Algorithms," *IEEE Transactions on Automatic Control*, Vol.59, No.4, 1025–1030, 2014.

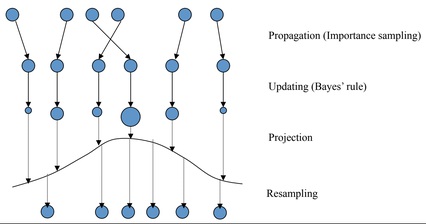

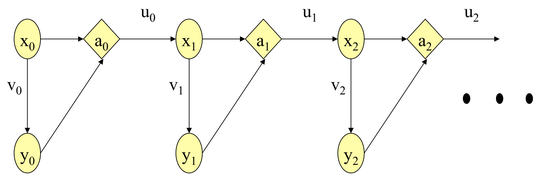

Model Reference Adaptive Search (MRAS) for Global Optimization

- Applicable to general global optimization problems
- Framework for design and analysis of algorithms
- Provably convergent and computationally efficient
- Probability distribution converges to distribution concentrated at the global optimum
- Applications: discrete optimization, continuous optimization, inventory control, buffer allocation, data mining

J. Hu, M.C. Fu, and S.I. Marcus, "A Model Reference Adaptive Search Method for Global Optimization," *Operations Research*, Vol.55, No.3, 549–568, 2007.

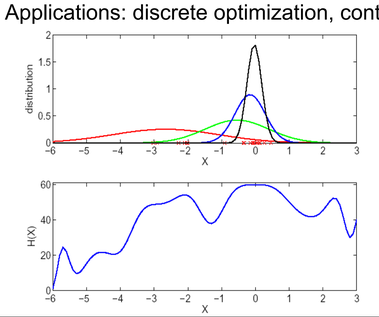

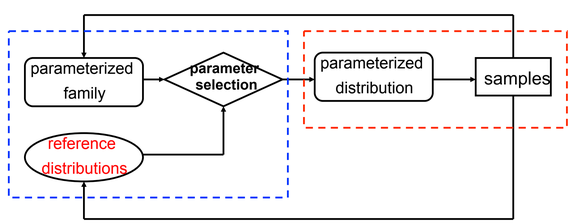

Main Sources of Support: NSF-DMI/CMII, NSF-EECS, AFOSR