Local Improvement/Policy Iteration Algorithms

Uri Zwick

Abstract

Policy Iteration algorithms form a simple family of algorithms that can be applied in many different settings, ranging from the relatively simple problem of finding a minimum mean weight cycle in a graph, the more challenging solution of Markov Decision Processes (MDPs), to the solution of 2-player full information stochastic games, also known as Simple Stochastic Games (SSGs).

It was recently shown by Fridmann that the worst case running time of a natural deterministic version of the policy iteration algorithm, when applied to Parity Games (PGs), is exponential. It is still open, however, whether deterministic policy iteration algorithm can solve Markov Decision Processes in polynomial time, and whether randomized policy iteration algorithms can solve Simple Stochastic Games in polynomial time.

The talk will survey what is known regarding policy iteration algorithms and mention many intriguing open problems.