Available online at www.sciencedirect.com
SCIENCE DIRECT
Electronic Notes in Theoretical Computer Science
ELSEVIER Electronic Notes in Theoretical Computer Science 119 (2005) 51–65
www.elsevier.com/locate/entcs

New Algorithms for Solving Simple Stochastic Games

Rafał Somla[1],[2]

Computing Science Department, Uppsala University, Box 337, SE-75105, Uppsala, Sweden

Abstract

We present new algorithms for determining optimal strategies for two-player games with probabilistic moves and reachability winning conditions. Such games, known as simple stochastic games, were extensively studied by A.Condon [2,3]. Many interesting problems, including parity games and hence also mu-calculus model checking, can be reduced to simple stochastic games. It is an open problem, whether simple stochastic games can be solved in polynomial time.
Our algorithms determine the optimal expected payoffs in the game. We use geometric interpretation of the search space as a subset of the hyper-cube $[0,1]^N$. The main idea is to divide this set into convex subregions in which linear optimization methods can be used. We show how one can proceed from one subregion to the other so that, eventually, a region containing the optinal payoffs will be found. The total number of subregions is exponential in the size of the game but, in practice, the algorithms need to visit only few of them to find a solution.
We believe that our new algorithms could provide new insights into the difficult problem of determining algorithmic complexity of simple stochastic games and other, equivalent problems.

Keywords: infinite graph games, parity games, simple stochastic games, finding optimal strategies, successive approximation.

1 Introduction

Many problems studied in computer science have an elegant presentation in a form of two-player graph games with various winning conditions. This includes verification of open components, controller synthesis and also theory of alternating automata. Hence a question of finding efficient algorithms for solving graph games, i.e., for deciding which player possess a winning strategy and

[1] Email: rsomla@it.uu.se
[2] This work is supported by Games for Processes project (VR grant 621–2002–455).

doi:10.1016/j.entcs.2004.07.008

possibly finding this strategy, becomes important. The problem was extensively studied for games with a wide range of winning conditions: from simple safety/reachability objectives to ω-regular ones expressed by either Büchi/co-Büchi or most general parity conditions [9,15,14,4,8,13,5]. A long standing open question in this area is whether it is possible to solve games with parity winning conditions in polynomial time. Through known reductions [12], a positive answer to this question would also mean that the mu-calculus model checking can be done in polynomial time.

In this paper we focus on *simple stochastic games* [2]. These are two-player, turn-based games with random moves. The objective in the game is to reach a final position (a sink) with the best possible associated payoff. Thus, rather than looking for a winning strategy, we want to find an *optimal strategy*, that is a strategy which guarantees the best expected payoff for a player.

We are interested in this kind of games because, on the one hand, other important graph games, like parity and mean-payoff games, can be easily reduced to simple stochastic games. On the other hand, simple stochastic games are instances of general stochastic games which have a rich and well developed theory. We believe that this link between an old area of operational research and the current studies can provide new insights into the problem of the complexity of graph games. Simple stochastic games are also interesting as a model for open, probabilistic components. Efficient algorithms for solving simple stochastic games can be used for verification of such components or even for synthesis of components meeting given specification.

Over the years, many algorithms has been proposed, which solve (simple) stochastic games. Many of them were later shown to be incorrect [3]. The correct ones, usually don't have any satisfactory complexity analysis. The two main methods used in these algorithms are the strategy improvement method and solving the local optimality equations.

Strategy improvement was developed by Hoffman and Karp for general stochastic games [7]. In this method an initial strategy for one of the players is improved in each iteration by switching it at positions at which choices are not locally optimal.

The other method is based on solving a system of constraints for the optimal expected payoffs in the game, which we call local optimality equations. Having optimal payoffs, one can easily reconstruct optimal strategies. For one-player games, the local optimality equations are linear, hence such games can be solved in polynomial time using linear programming techniques [6]. For two-player games the constraints are no longer linear and thus other methods are used, usually some form of iterative approximation.

We propose two algorithms which are based on the second method. For a