Information and Computation 242 (2015) 2–24

Contents lists available at ScienceDirect

Information and Computation

www.elsevier.com/locate/yinco

Qualitative analysis of concurrent mean-payoff games ☆

Krishnendu Chatterjee *, Rasmus Ibsen-Jensen

IST Austria, Austria

ARTICLE INFO

Article history:
Received 11 October 2013
Available online 23 March 2015

ABSTRACT

We consider concurrent games played by two players on a finite-state graph, where in every round the players simultaneously choose a move, and the current state along with the joint moves determine the successor state. We study the most fundamental objective for concurrent games, namely, mean-payoff or limit-average objective, where a reward is associated to each transition, and the goal of player 1 is to maximize the long-run average of the rewards, and the objective of player 2 is strictly the opposite (i.e., the games are zero-sum). The path constraint for player 1 could be qualitative, i.e., the mean-payoff is the maximal reward, or arbitrarily close to it; or quantitative, i.e., a given threshold between the minimal and maximal reward. We consider the computation of the almost-sure (resp. positive) winning sets, where player 1 can ensure that the path constraint is satisfied with probability 1 (resp. positive probability). Almost-sure winning with qualitative constraint exactly corresponds to the question of whether there exists a strategy to ensure that the payoff is the maximal reward of the game. Our main results for qualitative path constraints are as follows: (1) we establish qualitative determinacy results that show that for every state either player 1 has a strategy to ensure almost-sure (resp. positive) winning against all player-2 strategies, or player 2 has a spoiling strategy to falsify almost-sure (resp. positive) winning against all player-1 strategies; (2) we present optimal strategy complexity results that precisely characterize the classes of strategies required for almost-sure and positive winning for both players; and (3) we present quadratic time algorithms to compute the almost-sure and the positive winning sets, matching the best known bound of the algorithms for much simpler problems (such as reachability objectives). For quantitative constraints we show that a polynomial time solution for the almost-sure or the positive winning set would imply a solution to a long-standing open problem (of solving the value problem of turn-based deterministic mean-payoff games) that is not known to be solvable in polynomial time.

1. Introduction

Concurrent games. Concurrent games are played by two players (player 1 and player 2) on finite-state graphs for an infinite number of rounds. In every round, both players independently choose moves (or actions), and the current state along with

☆ The first author was supported by FWF Grant No. P 23499-N23, FWF NFN Grant No. S11407-N23 (RiSE), ERC Start grant (279307: Graph Games), and Microsoft faculty fellows award. Work of the second author supported by the Sino-Danish Center for the Theory of Interactive Computation, funded by the Danish National Research Foundation and the National Science Foundation of China (under the grant 61061130540). The second author acknowledge support from the Center for research in the Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council.

* Corresponding author.
E-mail addresses: krish.chat@ist.ac.at (K. Chatterjee), ribsen@ist.ac.at (R. Ibsen-Jensen).

http://dx.doi.org/10.1016/j.ic.2015.03.009
0890-5401/

the two chosen moves determine the successor state. In *deterministic* concurrent games, the successor state is unique; in *stochastic* concurrent games, the successor state is given by a probability distribution. The outcome of the game (or a *play*) is an infinite sequence of states and action pairs. These games were introduced in a seminal work by Shapley [53], and have been one of the most fundamental and well-studied game models in stochastic graph games. An important sub-class of concurrent games are *turn-based* games, where in each state at most one player can choose between multiple moves (if the transition is stochastic we have turn-based stochastic games, and if the transition is deterministic we have turn-based deterministic games).

Mean-payoff (limit-average) objectives. The most well-studied objective for concurrent games is the *limit-average* (or mean-payoff) objective, where a reward is associated to every transition and the payoff of a play is the limit-inferior (or limit-superior) average of the rewards of the play. The original work of Shapley [53] considered *discounted* sum objectives (or games that stop with probability 1); and concurrent stochastic games with limit-average objectives (or games that have zero stop probabilities) was introduced by Gillette in [39]. The player-1 *value* $\text{val}(s)$ of the game at a state s is the supremum value of the expectation that player 1 can guarantee for the limit-average objective against all strategies of player 2. The games are zero-sum where the objective of player 2 is the opposite. Concurrent limit-average games and many important sub-classes have received huge attention over the last five decades. The prominent sub-classes are turn-based games as restrictions of the game graphs, and *reachability* objectives as restrictions of the objectives. A reachability objective consists of a set U of *terminal* states (absorbing or sink states that are states with only self-loops), and the set U is exactly the set of states where out-going transitions are assigned reward 1 and all other transitions are assigned reward 0. Many celebrated results have been established for concurrent limit-average games and its sub-classes: (1) the existence of values (or determinacy or equivalence of switching of strategy quantifiers for the players as in von-Neumann's min–max theorem) for concurrent discounted games was established in [53]; (2) the existence of values (or determinacy) for concurrent reachability games was established in [36]; (3) the existence of values (or determinacy) for turn-based stochastic limit-average games was established in [48]; (4) the result of Blackwell–Fergusson established existence of values for the celebrated game of Big-Match [5]; and (5) developing on the results of [5] and Bewley–Kohlberg on Puisuex series [4] the existence of values for concurrent limit-average games was established in [49]. The decision problem of whether the value $\text{val}(s)$ is at least a rational constant λ can be decided in PSPACE [26,42]; and is *square-root sum* hard even for concurrent reachability games [35].[1] The algorithmic question of the value computation has also been studied in depth for special classes such as ergodic concurrent games [44] (where all states can be reached with probability 1 from all other states); turn-based stochastic reachability games [28]; and turn-based deterministic limit-average games [33,56,40,10]. The decision problem of whether the value $\text{val}(s)$ is at least a rational constant λ lie in NP $\cap$ coNP both for turn-based stochastic reachability games and turn-based deterministic limit-average games. They are among the rare and intriguing combinatorial problems that lie in NP $\cap$ coNP, but are not known to be in PTIME. The existence of polynomial time algorithms for the above decision questions are long-standing open problems.

Qualitative winning modes. In another seminal work, the notion of *qualitative winning* modes was introduced in [29] for concurrent reachability games. In qualitative winning modes, instead of the exact value computation the question is whether the objective can be satisfied with probability 1 (*almost-sure* winning) or with positive probability (*positive* winning). The qualitative analysis is of independent interest and importance in many applications (such as in system analysis) where we need to know whether the correct behavior arises with probability 1. For instance, when analyzing a randomized embedded scheduler, we are interested in whether every thread progresses with probability 1 [15]. Even in settings where it suffices to satisfy certain specifications with probability $p < 1$, the correct choice of p is a challenging problem, due to the simplifications introduced during modeling. For example, in the analysis of randomized distributed algorithms it is quite common to require correctness with probability 1 (see, e.g., [51,47,54]). More importantly it was shown in [29] that the qualitative analysis for concurrent reachability games can be solved in polynomial time (quadratic time for almost-sure winning, and linear time for positive winning). Moreover the algorithms were discrete graph theoretic algorithms, and the combinatorial algorithms were independent of the precise transition probabilities. Since qualitative analysis is robust to numerical perturbations and modeling errors in the transition probabilities, and admits efficient combinatorial algorithms for the special case of concurrent reachability games, they have been studied in many different contexts such as Markov decision processes and turn-based stochastic games with ω-regular objectives [25,21,22]; pushdown stochastic games with reachability objectives [34,35,9]; and partial-observation games with ω-regular objectives [19,3,2,17,27,14,50,20], to name a few. However, the qualitative analysis for the very important problem of concurrent limit-average games has not been studied before. In this work, we consider qualitative analysis of concurrent limit-average games. We show that the qualitative analysis of concurrent limit-average games is significantly different from and more involved than qualitative analysis of concurrent reachability games.

Relevance of concurrent limit-average games. Besides the mathematical elegance of concurrent limit-average games, they also provide useful modeling framework for system analysis. Concurrent games are relevant in modeling systems with synchronous interaction of components [30,31,1]. Mean-payoff objectives are widely used for performance measure of systems,

[1] The square-root sum problem is an important problem from computational geometry, where given a set of natural numbers $n_1, n_2, \ldots, n_k$, the question is whether the sum of the square roots exceed an integer b. The square root sum problem is not known to be in NP.