Information and Computation 241 (2015) 177–196

Contents lists available at ScienceDirect

Information and Computation

www.elsevier.com/locate/yinco

The complexity of multi-mean-payoff and multi-energy games ☆

Yaron Velner [a], Krishnendu Chatterjee [b,*], Laurent Doyen [c], Thomas A. Henzinger [b], Alexander Rabinovich [a], Jean-François Raskin [d]

[a] *The Blavatnik School of Computer Science, Tel Aviv University, Israel*
[b] *IST Austria (Institute of Science and Technology Austria), Austria*
[c] *LSV, ENS Cachan & CNRS, France*
[d] *Département d'Informatique, Université Libre de Bruxelles (U.L.B.), Belgium*

ARTICLE INFO

Article history:
Received 12 April 2013
Received in revised form 2 September 2014
Available online 5 March 2015

ABSTRACT

In mean-payoff games, the objective of the protagonist is to ensure that the limit average of an infinite sequence of numeric weights is nonnegative. In energy games, the objective is to ensure that the running sum of weights is always nonnegative. Multi-mean-payoff and multi-energy games replace individual weights by tuples, and the limit average (resp., running sum) of each coordinate must be (resp., remain) nonnegative. We prove finite-memory determinacy of multi-energy games and show inter-reducibility of multi-mean-payoff and multi-energy games for finite-memory strategies. We improve the computational complexity for solving both classes with finite-memory strategies: we prove coNP-completeness improving the previous known EXPSPACE bound. For memoryless strategies, we show that deciding the existence of a winning strategy for the protagonist is NP-complete. We present the first solution of multi-mean-payoff games with infinite-memory strategies: we show that mean-payoff-sup objectives can be decided in NP∩coNP, whereas mean-payoff-inf objectives are coNP-complete.

1. Introduction

Graph games and multi-objectives Two-player games on graphs are central in many applications of computer science. For example, in the synthesis problem, implementations of reactive systems are obtained from winning strategies in games with a qualitative objective formalized by an ω-regular specification [30,29,1]. In these applications, the games have a qualitative (Boolean) objective that determines which player wins. On the other hand, games with quantitative objectives which are natural models in economics (where players have to optimize a real-valued payoff) have also been studied in the context of automated design [31,17,32]. In the recent past, there has been considerable interest in the design of reactive systems

☆ Preliminary versions appeared in the *Proceedings of the IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science* (FSTTCS), Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, LIPIcs, 2010, pp. 505–516, and in the *Proceedings of the 14th International Conference on Foundations of Software Science and Computational Structures* (FoSSaCS), Lecture Notes in Computer Science 6604, Springer, 2011, pp. 275–289.

* Corresponding author at: Krishnendu Chatterjee, Am Campus 1, IST Austria, Klosterneuburg 3400, Austria.
E-mail addresses: yaronvel@post.tau.ac.il (Y. Velner), Krishnendu.Chatterjee@ist.ac.at (K. Chatterjee), doyen@lsv.ens-cachan.fr (L. Doyen), tah@ist.ac.at (T.A. Henzinger), rabinoa@post.tau.ac.il (A. Rabinovich), jraskin@ulb.ac.be (J.-F. Raskin).

http://dx.doi.org/10.1016/j.ic.2015.03.001
0890-5401/

that work in resource-constrained environments (such as embedded systems). The specifications for such reactive systems are quantitative, and give rise to quantitative games. In most system design problems, there is no unique objective to be optimized, but multiple, potentially conflicting objectives. For example, in designing a computer system, one is interested not only in minimizing the average response time but also the average power consumption. In this work we study such multi-objective generalizations of the two most widely used quantitative objectives in games, namely, *mean-payoff* and *energy* objectives [19,32,8,3].

Multi-mean-payoff games A *multi-mean-payoff game* is played on a finite weighted game graph by two players. The vertices of the game graph are partitioned into positions that belong to player 1 and positions that belong to player 2. Edges of the graphs are labeled with k-dimensional vectors w of integer values, i.e., $w \in \mathbb{Z}^k$. The game is played as follows. A pebble is placed on a designated initial vertex of the game graph. The game is played in rounds in which the player owning the position where the pebble lies moves the pebble to an adjacent position of the graph using an outgoing edge. The game is played for an infinite number of rounds, resulting in an infinite path through the graph, called a play. The value associated to a play is the mean value[1] in each dimension of the vectors of weights labeling the edges of the play. Accordingly, the winning condition for player 1 is defined by a vector of rational values $v \in \mathbb{Q}^k$ that specifies a threshold for each dimension. A play is winning for player 1 if its vector of mean values is at least v. All other plays are winning for player 2, thus the game is zero-sum. We are interested in the problem of deciding the existence of a winning strategy for player 1 in multi-mean-payoff games. In general infinite memory may be required to win multi-mean-payoff games, but in many practical applications such as the synthesis of reactive systems with multiple resource constraints, the multi-mean-payoff games with finite memory is the relevant problem. Also they provide the framework for the synthesis of specifications defined by mean-payoff conditions [2,10,11], and the synthesis question for such specifications under *regular (ultimately periodic)* words correspond to multi-mean-payoff games with finite-memory strategies. Hence we study multi-mean-payoff games both for general strategies as well as finite-memory strategies.

Multi-energy games In multi-energy games, the winning condition for player 1 requires that, given an initial credit $v_0 \in \mathbb{N}^k$, the sum of v_0 and all the vectors labeling edges up to position i in the play is nonnegative, for all $i \in \mathbb{N}$. The decision problem for multi-energy games asks whether there exists an initial credit v_0 and a strategy for player 1 to maintain the energy nonnegative in all dimensions against all strategies of player 2.

Contributions In this paper, we study the strategy complexity and computational complexity of solving multi-mean-payoff and multi-energy games. The contributions are as follows.

First, we show that multi-energy and multi-mean-payoff games are determined when played with finite-memory strategies. When considering finite-memory strategies, those games correspond to the synthesis question with ultimately periodic words, and they enjoy pleasant mathematical properties like existence of the limit of the mean value of the weights. We also establish that multi-energy and multi-mean-payoff games are not determined for memoryless strategies. Additionally, we show for multi-energy games determinacy under finite-memory coincides with determinacy under arbitrary strategies, and each player has a winning strategy if and only if he has a finite-memory winning strategy. In contrast, we show for multi-mean-payoff games that determinacy under finite-memory and determinacy under arbitrary strategies do not coincide. Moreover, for multi-mean-payoff games when the strategies for player 1 is restricted to finite-memory strategies, the winning set for player 1 remains unchanged irrespective of whether we consider finite-memory or infinite-memory counter strategies for player 2.

Second, we show that under the restriction that either both players play finite-memory or both play memoryless strategies, the decision problems for multi-mean-payoff games and multi-energy games are equivalent.

Third, we study the computational complexity of the decision problems for multi-mean-payoff games and multi-energy games, both for finite-memory strategies and the special case of memoryless strategies. Our complexity results can be summarized as follows. (A) For finite-memory strategies, we provide a nondeterministic polynomial-time algorithm for deciding negative instances of the problems.[2] Thus we show that the decision problems are in coNP. This significantly improves the complexity as compared to the EXPSPACE algorithm that can be obtained by reduction to VASS (vector addition systems with states) [5]. Furthermore, we establish a coNP lower bound for these problems by reduction from the complement of the 3SAT problem, hence showing that the problem is coNP-complete. (B) For the case of memoryless strategies, as the games are not determined, we consider the problem of determining if player 1 has a memoryless winning strategy. First, we show that the problem of determining if player 1 has a memoryless winning strategy is in NP, and then show that the problem is NP-hard even when the weights are restricted to $\{-1, 0, 1\}$ and two dimensions.

Finally, we study the computational complexity of multi-mean-payoff games for infinite-memory strategies. Our complexity results are summarized as follows. (A) We show that multi-mean-payoff games with mean-payoff-sup objectives

[1] The mean value can be either the mean-payoff-sup or the mean-payoff-inf value defined as the limsup (resp., liminf) of the means of the weights of the finite prefixes.

[2] Negative instances are those where player 1 is losing, and by determinacy under finite-memory, where player 2 is winning.