

Finite-Memory Strategy Synthesis for Robust Multidimensional Mean-Payoff Objectives

Yaron Velner

The Blavatnik School of Computer Science, Tel Aviv University

Abstract

Two-player games on graphs provide the mathematical foundation for the study of reactive systems. In the quantitative framework, an objective assigns a value to every play, and the goal of player 1 is to minimize the value of the objective. In this framework, there are two relevant synthesis problems to consider: the quantitative analysis problem is to compute the minimal (or infimum) value that player 1 can assure, and the boolean analysis problem asks whether player 1 can assure that the value of the objective is at most ν (for a given threshold ν). Mean-payoff expression games are played on a multidimensional weighted graph. An atomic mean-payoff expression objective is the mean-payoff value (the long-run average weight) of a certain dimension, and the class of mean-payoff expressions is the closure of atomic mean-payoff expressions under the algebraic operations of max, min, numerical complement and sum. In this work, we study for the first time the strategy synthesis problems for games with robust quantitative objectives, namely, games with mean-payoff expression objectives. While in general, optimal strategies for these games require infinite-memory, in synthesis we are typically interested in the construction of a finite-state system. Hence, we consider games in which player 1 is restricted to finite-memory strategies, and our main contribution is as follows. We prove that for mean-payoff expressions, the quantitative analysis problem is computable, and the boolean analysis problem is inter-reducible with Hilbert's tenth problem over rationals — a fundamental long-standing open problem in computer science and mathematics.

Categories and Subject Descriptors F.4.3 [Formal Languages]: Decision problems

General Terms Theory.

Keywords Quantitative synthesis, multidimensional quantitative games on graphs, mean-payoff games.

1. Introduction

In the classical framework of boolean formal verification, a program may only violate or satisfy a given specification, and in the

framework of synthesis, the task is to automatically construct a program that satisfies the specification. The boolean framework does not discriminate between programs that satisfy a given specification, and consequently, it may produce (or verify) unreasonable implementations.

In the recent years, there is an emerging line of research that aims to measure the quality of a program with quantitative metrics, e.g., [1, 3, 5, 7, 8, 11]. The quantitative verification problem asks how well a program satisfies a given specification, and the synthesis task is to construct the optimal program with respect to a specification.

Quantitative verification and synthesis problems are modelled by infinite-duration games over weighted graphs. In these games, the set of vertices is partitioned into player-1 and player-2 vertices; initially, a pebble is placed on an initial vertex, and in every round, the player who owns the vertex that the pebble resides in, advances the pebble to an adjacent vertex. This process is repeated forever and gives rise to a *play* that induces an infinite sequence of weights (or weight vectors), and a quantitative objective assigns a value to every play (or equivalently to every infinite sequence of weights).

The classical work on these games only considered games with single objectives, such as minimizing the long run average weight, or minimize the sum of weights. In order to have robust quantitative specifications, it is necessary to investigate games on graphs with multiple (and possibly conflicting) objectives. Typically, multiple objectives are modeled by multidimensional weight functions (e.g., [1, 4, 5, 10]), and the outcome of a *play* is a vector of values. In the boolean setting, the goal of player 1 is to satisfy a boolean condition on the values (with respect to a threshold vector). For example, player 1 needs to assure that the average response time (*rt*) of an arbiter is at most 2.4 and that the average energy consumption (*ec*) is below 7. In the quantitative setting, the outcome of a play is a unique (real) value, and the goal of player 1 is to minimize the value of the play. A multiple objective specification is modelled by algebraic operations on single objectives. In the example above, we define the quantitative objective $\max(rt - 2.4, ec - 7)$, and a non-positive value to the quantitative objective implies that the boolean objective is satisfied. In the general case, an objective is determined either by the projection of the weight function to one dimension, or it is formed by algebraic operation on two (or more) objectives. In the literature, the common and natural algebraic operations are min, max, *numerical complement* (multiplication by -1) and sum. We note that when the goal is to minimize the value of the objective, then the first three operations generalize the boolean disjunction, conjunction and negation. A class of quantitative objectives is *robust* if it is closed under the four algebraic operations. So far, the only known class of robust quantitative objectives that has an effective algorithm for the model checking problem (that is, for solving one-player games) is the class of *mean-payoff expressions* [6], which is the closure of one-dimensional

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CSL-LICS 2014, July 14–18, 2014, Vienna, Austria.
Copyright © 2014 ACM 978-1-4503-2886-9...\$15.00.
<http://dx.doi.org/10.1145/2603088.2603096>

mean-payoff (long-run average of the weights) objectives to the four algebraic operation. For example, for an infinite sequence of vectors $a = a_1, a_2, \dots \in (\mathbb{R}^3)^\omega$ the objective

$$E(a) = \text{LimAvg}_1(a) + \min(\text{LimAvg}_2(a), -\text{LimAvg}_3(a))$$

is a mean-payoff expression (where LimAvg_i is the long-run average of dimension i) and $E((1, 2, 6)^\omega) = 1 + \min(2, -6) = -5$.

In the quantitative setting, there are two relevant synthesis problems: (i) the *quantitative analysis problem* is to compute the optimal (infimum) value that a player-1 strategy can assure; and (ii) the *boolean analysis problem* is to determine whether player 1 can assure a value of at most ν to the objective (for a given ν). From the perspective of synthesis, these problems are most important when player 1 is restricted to finite-memory strategies (in Example 4 we show that infinite-memory strategies may yield a better value for player 1, hence the restriction to finite-memory strategies may affect the analysis of the synthesis problem).

For mean-payoff expressions, optimal finite-memory strategies may not always exist. Hence, the quantitative analysis problem is to compute the greatest lower bound on the minimal value that player 1 can assure. We note that since all model checking problems (i.e., the quantitative generalization of the emptiness, universality and language inclusion) are decidable for mean-payoff expression, then the computability of the quantitative analysis will give us an effective algorithm to synthesize ϵ -optimal finite-memory strategies, and if the boolean analysis problem were decidable, then we would have an algorithm that construct the corresponding player-1 strategy.

Our contribution. In this paper, we consider for the first time the synthesis problem for a robust class of quantitative objectives. We prove computability for the quantitative synthesis problem, and we show that the boolean analysis problem is inter-reducible with Hilbert’s tenth problem over rationals ($H10(\mathbb{Q})$), which is a fundamental long-standing open question in computer science and mathematics. We show that the problem is inter-reducible with $H10(\mathbb{Q})$ even when both players are restricted to finite-memory strategies, and we show that there is a fragment of mean-payoff expressions that is $H10(\mathbb{Q})$ -hard when one or both players are restricted to finite-memory strategies, but decidable when both players may use infinite-memory strategies.

Our main technical contribution is the introduction of a general scheme that lifts a one-player game solution (equivalently, a model checking algorithm) to a solution for a two-player game (when player 1 is restricted to finite-memory strategies). The scheme works for a large class of quantitative objectives that have certain properties (which we define in Subsection 2.2).

Related work. The class of mean-payoff expressions was introduced in [6], and the decidability of the model checking problems (which correspond to one-player games) was established. A simpler and more efficient algorithm for mean-payoff expression games was given in [15]. Mean-payoff games on multidimensional graphs were first studied in [9]. In these games the objective of player 1 was to satisfy a conjunctive condition (in the terms of this paper, the objective was a maximum of multiple one-dimensional objectives). In [16], decidability for an objective that is formed by the min and max operators was established. But the proof can not be extended to include the numerical complement operator, and it does not scale for the case that player 1 is restricted to finite-memory strategies.

Structure of the paper. In the next section we give the basic definitions for quantitative games and we define a class of quantitative objectives that have special properties. In Sections 3 and 4 we give a generic solution for the synthesis problem of quantitative objectives that satisfies the special properties (and an overview of the solution is given in Subsection 2.3). In Section 5 we show that mean-payoff expressions satisfy the special properties and the main results of the

paper follow. Due to lack of space, some of the proofs were omitted, and the full proofs are available in the technical report [17].

2. Games with Quantitative Objectives

In this section we give the formal definitions for quantitative objectives and games on graphs with quantitative objectives (Subsection 2.1). We define four special properties of quantitative objectives (Subsection 2.2), and we give an informal overview for the two-player game solution of games with quantitative objectives that satisfy the special properties (Subsection 2.3).

2.1 Quantitative games on graphs

Quantitative objectives. In this paper we consider directed finite graphs with a k -dimensional weight function that assigns a vector of rationals to each edge. A *quantitative objective* is a function that assigns a value to every infinite sequence of weight vectors. Formally an objective is a function $obj : (\mathbb{R}^k)^\omega \rightarrow \mathbb{R}$. A simple example for quantitative objective is to consider a one-dimensional weight function and an objective that assigns to each infinite path the maximal weight that occurs infinitely often in the path. An objective $obj : (\mathbb{R}^k)^\omega \rightarrow \mathbb{R}$ is called *prefix-independent* if for every $a_1 \in (\mathbb{R}^k)^*$ and $a_2 \in (\mathbb{R}^k)^\omega$ it holds that $obj(a_1 a_2) = obj(a_2)$.

Algebraic operations over quantitative objectives. The quantitative counterpart of the boolean operations of disjunction, conjunction and negation are the max, min and numerical complement operators (numerical complement is multiplication by -1). The sum operator, which does not have a boolean counterpart, is also very natural operator in the framework of quantitative objectives. For two quantitative objectives obj_1 and obj_2 , the quantitative objective $\text{op}(obj_1, obj_2)$ (for $\text{op} \in \{\min, \max, \text{sum}\}$) assigns to every infinite sequence of weights $\ell \in (\mathbb{R}^k)^\omega$ the value $\text{op}(obj_1(\ell), obj_2(\ell))$, and the numerical complement of obj_1 assigns the value of $-obj_1(\ell)$.

Robust quantitative objectives A class of quantitative objectives $\mathcal{O}$ is *robust* if it is closed under the algebraic operations of min, max, sum and numerical complement. Formally, a class of objectives $\mathcal{O}$ is robust, if for every two objectives $obj_1, obj_2 \in \mathcal{O}$ the four quantitative objectives $obj_{\min}, obj_{\max}, obj_{\text{sum}}$ and obj^- are in $\mathcal{O}$ (such that for every $\ell \in (\mathbb{R}^k)^\omega$ and $\text{op} \in \{\min, \max, \text{sum}\}$: $obj_{\text{op}}(\ell) = \text{op}(obj_1(\ell), obj_2(\ell))$ and $obj^-(\ell) = -obj_1(\ell)$). We note that in [6, 8], Chatterjee et al. gave a broader definition for robustness of quantitative objectives, but since the concrete objectives that we consider in this paper are robust according to both definitions, we prefer to use the narrower (and simpler) notion of robustness.

Games on graph. A game graph is a directed graph $G = (V = V_1 \cup V_2, v_0, E, w : E \rightarrow \mathbb{Q}^k)$, where V is the set of vertices; V_i is the set of *player i* vertices; v_0 is the *initial vertex*; $E \subseteq V \times V$ is the set of edges; and $w : E \rightarrow \mathbb{Q}^k$ is a multidimensional weight function (e.g., see Figure 1). A *play* is an infinite sequence of rounds. In the first round a pebble is placed on the initial vertex and in every round the player who owns the vertex of the pebble advances the pebble to an adjacent vertex. Hence, a play corresponds to an infinite path in the graph that begins in v_0 and the labeling of the play is the corresponding infinite sequence of weight vectors. A game graph is a *one-player game* if only one of the players has a vertex with out-degree more than one.

Strategies. A *strategy* is a recipe for determining the next move based on the *history* of the play. A *player- i strategy* is a function $\sigma : V^* V_i \rightarrow V$, such that for every finite path π that ends in vertex v we have $(v, \sigma(\pi)) \in E$. A strategy has *finite memory* if it can be implemented by a Moore machine $(M, m_0, \alpha_n, \alpha_u)$, where M is




Figure 1. Game graph G . Player 1 owns the round vertices.




Figure 2. Game graph G according to strategy σ .

a finite set of memory states, m_0 is the initial memory state, $\alpha_u : M \times V \rightarrow M$ is the update function, and $\alpha_n : M \times V_i \rightarrow V$ is the next vertex function. If a play prefix is in state v_i and memory state M , then the strategy choice for the next vertex is $v = \alpha_n(M, v_i)$ and the memory is updated to $\alpha_u(M, v_i)$. A strategy is *memoryless* if it depends only in the current location of the pebble. Formally a player- i memoryless strategy is a function $\sigma : V_i \rightarrow V$. (We note that a memoryless strategy is also a finite-memory strategy.)

We denote the set of all player- i strategies by $\mathcal{S}_i$ and we denote the set of all player- i finite memory strategies by $\mathcal{FM}_i$.

Game graph according to a finite-memory strategy. For a game graph $G = (V = V_1 \cup V_2, E, w)$ and a player-1 finite-memory strategy $\sigma = (M, m_0, \alpha_u, \alpha_n)$, we denote the *game graph according to strategy σ* by G^σ , and we define it as follows:

- The vertices of G^σ are the Cartesian product $V \times M$; player- i vertices are $V_i \times M$; and the initial vertex of G^σ is (v_0, m_0) .
- For a player-1 vertex (v, m) , the only successor vertex is $(\alpha_n(v, m), \alpha_u(v, m))$. For a player-2 vertex (v, m) the set of successor vertices is $\{(u, n) \mid (v, u) \in E \text{ and } \alpha_u(v, m) = n\}$.

We note that the out-degree of all player-1 vertices is one, and thus G^σ is a *one-player game graph*. The main property of graphs according to a finite-memory strategy is that every infinite path in G^σ corresponds to a play that is consistent with σ in G . A game graph according to a memoryless strategy is a special case of games according to finite-memory strategies. In this case, the game graph is obtained from G by removing all the player's out-edges that are not chosen by the memoryless strategy.

Example 1. Consider the game graph from Figure 1 and consider a player-1 strategy σ that in vertex v_2 moves the pebble to v_3 if v_2 was visited an odd number of times and otherwise it moves the pebble to v_1 . For example, in the first time that v_2 is visited, player 1 moves the pebble to v_3 , in the second time he will move the pebble to v_1 , in the third time to v_3 and so on. The strategy σ requires one bit of memory (i.e., $M = \{0, 1\}$), and G^σ is illustrated in Figure 2 (the labeling of the nodes represents the memory state). In G^σ all the choices are done by player 2.

Values of strategies and games. A tuple (σ, τ) of player-1 and player-2 strategies (respectively) uniquely defines a play $\pi_{\sigma, \tau}$ in a




Figure 3. $\text{CONVEX}(A, B, C, D, E, F, G)$ is the polygon $ABDEG$.

given graph. For a game graph G , a quantitative objective obj and a tuple of strategies (σ, τ) we denote $Val_{\sigma, \tau} = obj(\pi_{\sigma, \tau})$. In this paper, we assume that player 1 wishes to minimize¹ the value of the quantitative objective, and we define the *value* of a player-1 strategy σ to be $Val_\sigma = \sup_{\tau \in \mathcal{S}_2} Val_{\sigma, \tau}$. (Intuitively, this is the maximal value that player 2 can achieve against strategy σ .) The *minimal value of a game* is defined as $\inf_{\sigma \in \mathcal{FM}_1} Val_\sigma$. Intuitively, the minimal value of a game is the minimal value that player 1 can ensure by a finite-memory strategy.

Quantitative and boolean analysis For a given game graph, objective, and a rational threshold $r \in \mathbb{Q}$: The *quantitative analysis* task is to compute the minimal value of the game that can be enforced by a finite-memory strategy. The *boolean analysis* task is to decide whether there is a player-1 finite-memory strategy σ for which $Val_\sigma \leq r$. That is, whether player 1 can assure a value of at most r for the objective.

Boolean games and winning strategies. A *boolean game* is a game on graph equipped with a winning condition $W \subseteq E^\omega$ (that is, a winning condition is a set of infinite paths). A play π is *winning* for player 1 if $\pi \in W$, and a strategy σ is a player-1 *winning strategy* if for every player-2 strategy τ we have $\pi_{\sigma, \tau} \in W$. For a quantitative objective obj and a threshold $\nu \in \mathbb{R}$ we denote by (obj, ν) the boolean winning condition $\{\pi \in E^\omega \mid obj(\pi) \leq \nu\}$.

2.2 One-player game solution

In this paper, we consider objectives that have special properties for their one-player game solution and we present a general scheme that lifts a one-player game solution into a two-player game solution. To formally define the special properties of the solutions, we give the next definitions.

Definitions and notions for weighted graphs. Let $G = (V, E, w : E \rightarrow \mathbb{Q}^k)$ be a k -dimensional weighted graph. The *weight vector* of a finite path $\pi = e_1 \dots e_n$ is $w(\pi) = \sum_{i=1}^n w(e_i)$ and the *average weight* of a path is $Avg(\pi) = \frac{w(\pi)}{|\pi|}$. For a set of finite paths $\Pi = \{\pi_1, \dots, \pi_n\}$ we denote $Avg(\Pi) = \{Avg(\pi_1), \dots, Avg(\pi_n)\}$. We denote the set of simple cycles in G by $C(G)$, and we abbreviate $Avg(G) = Avg(C(G))$. For a finite set of vectors $V = \{v_1, v_2, \dots, v_n\} \in \mathbb{R}^k$, we denote $\text{CONVEX}(V) = \{\sum_{i=1}^n \alpha_i v_i \mid \sum_{i=1}^n \alpha_i = 1 \text{ and } \alpha_1, \dots, \alpha_n \geq 0\}$ (see Figure 3). We abbreviate $\text{CONVEX}(G) = \text{CONVEX}(Avg(G))$. An m -dimensional *simplex* is the set $\mathcal{S}(m) = \{(x_1, \dots, x_m) \in \mathbb{R}^m \mid x_i \geq 0 \wedge \sum_{i=1}^m x_i = 1\}$. The *simplex interior* is $SI(m) = \{(x_1, \dots, x_m) \in \mathbb{R}^m \mid x_i > 0 \wedge \sum_{i=1}^m x_i = 1\}$, and the *rational interior of a simplex* is $QSI(m) = SI(m) \cap \mathbb{Q}^m$. When m is clear from the context we abbreviate $\mathcal{S}(m)$, $SI(m)$ and $QSI(m)$ with $\mathcal{S}$, SI and QSI (respectively).

Solution for one-player game with special properties. A solution for a one-player quantitative game is a function f that assigns to every one-player game graph G the maximal value that the player can achieve in graph G . We note that for prefix-independent objectives, a function f' that assigns to each *strongly connected*

¹ Since we consider robust objectives, then the same results hold when player-1 goal is to maximize the value of the objective.

graph its maximal value uniquely defines the solution function f (since the value of f is the maximal value of f' over all the strongly connected component of the graph). In this paper, we will consider only prefix-independent objective, hence, we define the special properties of a solution for strongly connected graphs. The special properties that we consider are:

1. *First-order definable.* For every $n \in \mathbb{N}$ there is a first-order formula $\zeta_n(\bar{x}_1, \dots, \bar{x}_n, y)$ over $\langle \mathbb{R}, =, <, +, \times \rangle$ such that for every graph G with $\text{Avg}(G) = \{\bar{x}_1, \dots, \bar{x}_n\}$ we have $f(G) = y$ if and only if $\zeta_n(\bar{x}_1, \dots, \bar{x}_n, y)$ holds. In addition we require ζ_n to be computable from n . In the sequel, we write $y = \zeta_n(\bar{x}_1, \dots, \bar{x}_n)$ instead of $\zeta_n(\bar{x}_1, \dots, \bar{x}_n, y)$.
2. *Monotone in CONVEX(G).* If for two (strongly connected) graphs H and G we have $\text{CONVEX}(G) \subseteq \text{CONVEX}(H)$, then $f(G) \leq f(H)$. As a consequence, we get that for a k -dimensional objective, f is a function from $(\mathbb{R}^k)^*$ to $\mathbb{R}$, and that $f(G) \equiv g(\text{CONVEX}(G))$ for some function $g : (\mathbb{R}^k)^* \rightarrow \mathbb{R}$. Hence, by abusing the notation, we sometime write $f(\text{CONVEX}(G))$ instead of $f(G)$.
3. *Continuous function.* f is a continuous function. Formally, if f is the solution for a k -dimensional objective, then for every $n \in \mathbb{N}$ the function $\zeta_n : (\mathbb{R}^k)^n \rightarrow \mathbb{R}$ is a continuous function, i.e., for every $\epsilon > 0$ there exists $\delta > 0$ such that for every two vectors $A, B \in (\mathbb{R}^k)^n$ with $|A - B| < \delta$ it holds that $|\zeta_n(A) - \zeta_n(B)| < \epsilon$.

We will show computability for the quantitative analysis problem for objectives that have a solution that satisfies the above three properties. We also consider a fourth special property, and we will show decidability for the boolean analysis problem for objectives that have a solution that satisfies all four properties.

4. *Fourth property.* A solution $f = \{\zeta_1, \dots, \zeta_n, \zeta_{n+1}, \dots\}$ satisfies the fourth property if the next problem is decidable (for the set $\{\zeta_1, \dots, \zeta_n, \zeta_{n+1}, \dots\}$):
 - Input: a threshold $\nu \in \mathbb{Q}$ and a set of n matrices $A_1, \dots, A_n$, where A_i is a $k \times m_i$ matrix for some $m_i \in \mathbb{N}$.
 - Task: determine if the inequality $\zeta_n(A_1 \cdot x_1, \dots, A_n \cdot x_m) \leq \nu$ subject to $x_i \in \text{QSI}(m_i)$ is feasible (note that the result of the multiplication $A_i \cdot x_i$ is a vector of size k).

In the next example we demonstrate the above properties.

Example 2. Consider the two-dimensional one-player solution function $f(G) = \max_{(x,y) \in \text{CONVEX}(G)} [\max(x + y + 10, -x + y + 10, \min(-x + y - 10, x + y - 10))]$. We demonstrate that f is first-order definable by giving the explicit formula for ζ_2 , that is, the formula for a (strongly connected) graph with only two simple cycles with average weights (x_1, y_1) and (x_2, y_2) .

$$\begin{aligned} \zeta_2(x_1, y_1, x_2, y_2, r) \equiv & \\ \forall \alpha_1, \alpha_2, x, y & (\alpha_1 \geq 0 \wedge \alpha_2 \geq 0 \wedge (\alpha_1 + \alpha_2 = 1) \wedge (x = \\ & \alpha_1 x_1 + \alpha_2 x_2) \wedge (y = \alpha_1 y_1 + \alpha_2 y_2)) \rightarrow \\ & r \geq \max(x + y + 10, -x + y + 10, \min(-x + y - 10, x + \\ & y - 10)) \\ \wedge \\ \exists \alpha_1, \alpha_2, x, y & (\alpha_1 \geq 0 \wedge \alpha_2 \geq 0 \wedge (\alpha_1 + \alpha_2 = 1) \wedge (x = \\ & \alpha_1 x_1 + \alpha_2 x_2) \wedge (y = \alpha_1 y_1 + \alpha_2 y_2)) \wedge \\ & r = \max(x + y + 10, -x + y + 10, \min(-x + y - 10, x + \\ & y - 10)) \end{aligned}$$

(Technically max and min are not in $\langle \mathbb{R}, <, +, \times \rangle$, but they are trivially definable in this vocabulary.) Clearly if for two graphs we have $\text{CONVEX}(G_1) \subseteq \text{CONVEX}(G_2)$, then $f(G_1) \leq f(G_2)$ (hence, f is monotone), and ζ_2 is obviously a continuous function (and in general ζ_n is also continuous). Hence, f satisfies Properties 1-3. In Figure 4 we illustrate the geometrical interpreta-




Figure 4.

tion of Property 2, namely, the fact that the value of f depends only in $\text{CONVEX}(G)$. The equality $\max(x + y + 10, -x + y + 10, \min(-x + y - 10, x + y - 10)) = 0$ is represented by the thick line. The points that are connected by the dotted line represent the weights of the simple cycles of a strongly connected graph G_1 and the points that are connected by the dashed line represent the weights of the simple cycles of a strongly connected graph G_2 . The reader can see that $\text{CONVEX}(G_1)$ is below the thick line and $\text{CONVEX}(G_2)$ intersects with it. Hence, $f(G_1) < 0$ and $f(G_2) > 0$.

2.3 Informal overview of the solution for two-player games

The key notion for our solution is games according to strategies. When a one-player solution f is given, the boolean analysis problem amounts to determining whether there is a finite-memory strategy σ such that for every strongly-connected component (SCC) $S \in G^\sigma$ it holds that $f(\text{CONVEX}(S)) \leq \nu$. In Lemma 4 we show that w.l.o.g we may assume that for any σ the graph G^σ is strongly connected. Hence, in Section 3 we investigate the set $\{\text{CONVEX}(G^\sigma) \mid \sigma \in \mathcal{FM}_1\}$ and obtain a computable characterization for it. In Section 4 we exploit the properties of the one-player solution and the results of Section 3 and we obtain a first-order formula over **rationals** that computes the values that player 1 can enforce. We use the fact that f is continuous to show that the formula has the same infimum over rationals and reals, and hence, due to Tarski's Theorem the infimum value is computable. We also show that if Property 4 holds, then one can effectively determine whether the formula has an assignment that gives a value of at most ν . In Section 5 we apply these results for mean-payoff expressions. We show that their one-player solution satisfies Properties 1-3, and that it satisfies property 4 if and only if $H10(\mathbb{Q})$ is decidable.

3. CONVEX Cycles Problem

In this section, we consider the next problem:

- Problem 1 (CONVEX cycles problem).**
- Input: a k -dimensional game graph G and a set of k -dimensional vectors $\mathcal{V}$.
 - Task: determine whether there is a player-1 finite-memory strategy σ such that $\text{CONVEX}(G^\sigma) \subseteq \text{CONVEX}(\mathcal{V})$. (We call such strategy a realizing strategy.)

We first present the solution for the above problem, and then we show how to find all the sets of vectors for which there is a realizing player-1 finite-memory strategy.

The solution for Problem 1 relies on the next lemma.

Lemma 1. For a game graph G and a set of vectors $\mathcal{V}$, there exists a player-1 finite-memory strategy σ for which $\text{CONVEX}(G^\sigma) \subseteq \text{CONVEX}(\mathcal{V})$ iff for every player-2 memoryless strategy τ_i there exists a player-1 finite-memory strategy σ_i such that $\text{CONVEX}((G^{\tau_i})^{\sigma_i}) \subseteq \text{CONVEX}(\mathcal{V})$.

Proof. The proof for the direction from left to right is trivial (since any cycle in $(G^{\tau_i})^\sigma$ is also a cycle in G^σ).

Our proof for the converse direction is inspired by [12], and the key intuition of the proof is the following. Let v be a player-2 vertex, with two out-edges e_1 and e_2 , and let $G_1 = G - \{e_1\}$ and $G_2 = G - \{e_2\}$. Suppose that player 1 has two finite-memory strategies σ_1 and σ_2 such that $\text{CONVEX}(G_i^{\sigma_i}) \subseteq \text{CONVEX}(\mathcal{V})$ (for $i = 1, 2$). Then player 1 can combine the two strategies over $G - \{e_1\}$ and $G - \{e_2\}$, and he can obtain a finite-memory strategy σ , such that each simple cycle in G^σ is a convex combination of cycles from $G_1^{\sigma_1}$ and $G_2^{\sigma_2}$, and hence, $\text{CONVEX}(G^\sigma) \subseteq \text{CONVEX}(\mathcal{V})$. Hence, either player 1 has a realizable strategy for G or he does not have realizable strategy for $G_1 = G - \{e_1\}$ or for $G_2 = G - \{e_2\}$. Since this holds for every player-2 state, the proof follows.

In order to formally prove the key intuition we claim that if player 1 has a realizable strategy σ_i against any player-2 memoryless strategy, then he has a realizable strategy σ that satisfies $\text{CONVEX}(G^\sigma) \subseteq \text{CONVEX}(\mathcal{V})$, and we prove the claim by induction on the number of player-2 vertices with out-degree greater than one. The base case, where all of player-2 vertices have out-degree one, is trivial. For the inductive step, let us assume that there is a player-2 vertex v with out-edges e_1 and e_2 (if there is no such vertex, then we are in the base case). For $i = 1, 2$, let G_i be $G - \{e_i\}$. If player-2 has a violating memoryless strategy in either G_1 or G_2 , then surely this is also a violating memoryless strategy for G , and the claim follows. Otherwise, we construct a realizable player-1 strategy in G in the following way. For $i = 1, 2$, let σ_i be a finite-memory player-1 realizable strategy in G_i . If in σ_1 (resp., σ_2), the vertex v is unreachable then it is surely a winning strategy also for G . Otherwise, there exists a memory state m such that (m, v) is a vertex in $G_1^{\sigma_1}$, and we denote by σ'_1 the strategy that is obtained by changing σ_1 initial memory state to m . We construct σ in the following way. The memory structure of σ is a tuple $(M_1, M_2, \{1, 2\})$ where M_1 is the memory structure of σ'_1 , M_2 is the memory structure of σ_2 , and the third value in the tuple indicates if we are playing according to σ'_1 or σ_2 . At the beginning of a play, σ decides according to σ_2 (and updates M_2 accordingly). If σ decides according to σ_2 and edge e_1 is visited, then σ decides according to σ'_1 (and updates M_1 accordingly), until edge e_2 is visited, and then σ again decides according to σ_2 , and so on. We note that σ is a finite-memory strategy, and that any simple cycle in G^σ is a composition of simple cycles from $G_1^{\sigma'_1}$ and $G_2^{\sigma_2}$. Hence, the average weight of any simple cycle in G^σ is $\lambda x_1 + (1 - \lambda)x_2$ for some $\lambda \in [0, 1]$ and $x_i \in \text{Avg}(G_i) \subseteq \text{CONVEX}(\mathcal{V})$. And thus, a convex combination of x_1 and x_2 is also in $\text{CONVEX}(\mathcal{V})$, and we get that $\text{CONVEX}(G^\sigma) \subseteq \text{CONVEX}(\mathcal{V})$. Therefore, σ is a realizing strategy and the proof is completed. $\square$

We now wish to characterize all the sets of vectors that have a realizing strategy. For this purpose we give the next definition. For a player-2 memoryless strategy τ , let Π_e^τ be the (finite) set of Eulerian cyclic paths in G^τ , that is Π_e^τ contains only cyclic paths that visit every edge at most once. For every path $\pi \in \Pi_e^\tau$, let $c_1, \dots, c_t$ be the simple cycles that occur in π and we associate a $t \times k$ matrix A_π to every path π such that the i -th column of the matrix is $\text{Avg}(c_i)$. We observe that

$$\{\text{Avg}(\pi) \mid \pi \text{ is a cyclic path in } G^\tau\} = \bigcup_{\pi \in \Pi_e^\tau} \{A_\pi \cdot x \mid x \in \text{QST}\}$$

The next lemma shows how to compute the realizable sets of vectors.

Lemma 2. Let G be a k -dimensional graph, and let $\tau_1, \dots, \tau_\ell$ be the (finitely many) player-2 memoryless strategies in G . A set of




Figure 5. Game graph G .




Figure 6.

vectors $\mathcal{V} \subseteq \mathbb{R}^k$ is realizable if and only if there exist $x_1, \dots, x_\ell \in \mathbb{R}^k$ such that $x_i \in \bigcup_{\pi \in \Pi_e^{\tau_i}} \{A_\pi \cdot x \mid x \in \text{QST}\}$ (for every $i \in \{1, \dots, \ell\}$) and $\text{CONVEX}(x_1, \dots, x_\ell) \subseteq \text{CONVEX}(\mathcal{V})$.

Proof. First we characterize the realizable vector sets when a player-2 memoryless strategy τ is given, that is, we characterize the realizable vectors in a one-player game. A finite-memory strategy σ in a one-player graph G^τ is an ultimately periodic infinite path, and $(G^\tau)^\sigma$ is a lasso shaped graph with exactly one cycle. The cycle of $(G^\tau)^\sigma$ is obviously a cyclic path in G^τ , and thus $\mathcal{V}$ is realizable in G^τ iff there is a cyclic path π in G^τ with $\text{Avg}(\pi) \in \text{CONVEX}(\mathcal{V})$.

Hence, by Lemma 1, we get that $\mathcal{V}$ is realizable iff for every player-2 memoryless strategy τ_i there is a cyclic path π_i in G^{τ_i} with $\text{Avg}(\pi_i) \in \text{CONVEX}(\mathcal{V})$. Since such witness π_i exists iff there exists $x_i \in \bigcup_{\pi \in \Pi_e^{\tau_i}} \{A_\pi \cdot x \mid x \in \text{QST}\}$ with $\text{Avg}(\pi_i) = x_i$, then the proof is completed. $\square$

In the next example we illustrate the geometrical interpretation of Lemma 2.

Example 3. Consider the game graph G in Figure 5, where the box vertices are controlled by player 2. Player 2 has two possible memoryless strategies, namely, τ_1 that follows the edge $v_0 \rightarrow v_1$ and τ_2 that follows $v_0 \rightarrow v_4$. In G^{τ_1} the set of Eulerian cyclic paths $\Pi_e^{\tau_1}$ contains all cyclic sub-paths of the Eulerian cyclic path $v_1 \rightarrow v_2 \rightarrow v_3 \rightarrow v_1$. Hence, the average weight of any infinite lasso path in G^{τ_1} is a convex combination of $\text{Avg}(v_1 \rightarrow v_2 \rightarrow v_3 \rightarrow v_1)$, $\text{Avg}(v_2 \rightarrow v_3 \rightarrow v_2)$ and $\text{Avg}(v_3 \rightarrow v_3)$ (points D, F and E in Figure 6). In G^{τ_2} , an Eulerian cyclic path is either a sub-path of $v_4 \rightarrow v_5 \rightarrow v_4$ or the path $v_6 \rightarrow v_5 \rightarrow v_4$. Hence, the average weight of any infinite lasso path is either a convex combination of $\text{Avg}(v_4 \rightarrow v_5 \rightarrow v_4)$ and $\text{Avg}(v_5 \rightarrow v_5)$ (points A and B in Figure 6), or it is $\text{Avg}(v_6 \rightarrow v_5 \rightarrow v_4)$ (point C in Figure 6). By Lemma 2, we get that a set of vectors $\mathcal{V}$ is realizable if and only if $\text{CONVEX}(\mathcal{V})$ intersects with the polygon DEF and with either the line AB or with the point C (or with both).

4. Generic Solution for Games with Quantitative Objectives

In this section we solve the quantitative analysis problem for games with quantitative objectives that satisfy Properties 1-3 and we solve the boolean analysis problem for objectives that satisfy Properties 1-4. We first give a conceptual (i.e., not always computable) solution for the boolean analysis problem, and then extend the solution for the quantitative analysis problem.

An equivalent formulation for the boolean analysis problem is to ask whether for a game graph G and a threshold ν there is a player-1 (finite-memory) strategy σ such that the one-player solution over G^σ is at most ν . By the third property (convex monotonicity), it is enough to determine whether there is σ such that for every SCC S of G^σ it holds that $f(\text{CONVEX}(S)) \leq \nu$ (where f is the solution for the one-player game). However, we first show how to determine whether there is σ such that $f(\text{CONVEX}(G^\sigma)) \leq \nu$ and only then solve the original problem.

Lemma 3. *Let f be a one-player solution that satisfies Properties 1-3. Then $\inf_{\sigma \in \mathcal{F}\mathcal{M}_1} f(\text{CONVEX}(G^\sigma))$ is computable (when the input is a game graph G). If f also satisfies Property 4, then the problem of determining whether there is a player-1 finite-memory strategy σ such that $f(\text{CONVEX}(G^\sigma)) \leq \nu$ is decidable (when the input is G and ν).*

Proof. Let $\tau_1, \dots, \tau_m$ be all player-2 memoryless strategies in G (note that m is at most exponential in $|G|$). By Lemma 2, and by the monotonicity of f , there is a player-1 strategy σ that satisfies $f(\text{CONVEX}(G^\sigma)) \leq \nu$ if and only if there are matrices $A_{\pi_1}, \dots, A_{\pi_m}$ and vectors $x_1, \dots, x_m$ such that $\pi_i \in \Pi_e^{\tau_i}$ and $x_i \in \mathcal{QSI}$ and $\zeta_m(A_{\pi_1} \cdot x_1, \dots, A_{\pi_m} \cdot x_m) \leq \nu$. For every τ_i , the set $\Pi_e^{\tau_i}$ is finite (and at most of exponential size). Hence, we can enumerate all m -tuples of $\Pi_e^{\tau_1} \times \dots \times \Pi_e^{\tau_m}$ and check if for at least one tuple there is a solution to the inequality $\zeta_m(A_{\pi_1} \cdot x_1, \dots, A_{\pi_m} \cdot x_m) \leq \nu$. If f satisfies Property 4, then for a given $\pi_1, \dots, \pi_m$ we can effectively check if the inequality is satisfiable. Hence, we can effectively determine whether there is σ such that $f(\text{CONVEX}(G^\sigma)) \leq \nu$. Moreover, for a given $\pi_1, \dots, \pi_m$ the expression $\inf_{x_1, \dots, x_m \in \mathcal{QSI}} \zeta_m(A_{\pi_1} \cdot x_1, \dots, A_{\pi_m} \cdot x_m)$ is first-order definable (recall that ζ_m is first-order definable) over $\langle \mathbb{R}, <, +, \times \rangle$ (note that x_i ranges over $\mathcal{QSI}$ and not over $\mathcal{QSI}$) and therefore, by Taski's Theorem [13] its value is computable. Since ζ_m is continuous we get that $\inf_{x_1, \dots, x_m \in \mathcal{QSI}} \zeta_m(A_{\pi_1} \cdot x_1, \dots, A_{\pi_m} \cdot x_m) = \inf_{x_1, \dots, x_m \in \mathcal{QSI}} \zeta_m(A_{\pi_1} \cdot x_1, \dots, A_{\pi_m} \cdot x_m)$. Finally, by Lemma 2, we have that $\inf_{\sigma \in \mathcal{F}\mathcal{M}_1} f(\text{CONVEX}(G^\sigma)) = \inf_{\pi_1 \in \Pi_e^{\tau_1}, \dots, \pi_m \in \Pi_e^{\tau_m}} \inf_{x_1, \dots, x_m \in \mathcal{QSI}} \zeta_m(A_{\pi_1} \cdot x_1, \dots, A_{\pi_m} \cdot x_m)$, and since $\Pi_e^{\tau_i}$ is finite we get that $\inf_{\sigma \in \mathcal{F}\mathcal{M}_1} f(\text{CONVEX}(G^\sigma))$ is computable. $\square$

Before presenting the algorithm for the boolean analysis problem we recall the (standard) definitions of *winning regions* and *attractors*. Let G be a game graph with an initial vertex v_0 , and let v be an arbitrary vertex in G . We denote by (G, v) the game graph that is formed from G by changing the initial vertex to v . We say that a vertex v is in *player-1 winning region* (denoted by Win_1) if player 1 wins in (G, v) (that is, player 1 has a finite-memory strategy that assures a value at most ν to the objective). The *player-1 attractor set* of a vertex v (denoted by $\text{Attr}_1(v)$) contains all the vertices from which player 1 can force reachability to v (after finite number of rounds). It is well known that the attractor set of a vertex is computable (even in linear time) and that player 1 can force reachability by a finite-memory strategy (in fact, even by a memo-

ryless strategy). The next remark shows another important property of attractors and winning regions.

Remark 1. *Let G be a game graph over a boolean objective that is formed by a quantitative objective with a solution function f and a threshold ν , and let v be a vertex in G . Then for every vertex $u \notin \text{Attr}_1(v)$, if σ is a finite-memory player-1 strategy for (G, u) , then σ is a winning strategy in $(G - \text{Attr}_1(v), u)$.*

Proof. We denote $H = G - \text{Attr}_1(v)$ and we observe that $(H, u)^\sigma$ is a subgraph of $(G, u)^\sigma$. Hence, for every SCC $S \in H$ there is a corresponding SCC $S' \in G$ such that $f(\text{CONVEX}(S')) \leq \nu$. Since $\text{CONVEX}(S) \subseteq \text{CONVEX}(S')$ and by the monotonicity of f we get that $f(\text{CONVEX}(S)) \leq \nu$ and therefore σ is a winning strategy in (H, u) . $\square$

Algorithm 1 computes player-1 winning region, and we prove its correctness in Lemma 4

Algorithm 1 Player-1 winning region computation for quantitative objectives

```

WINNINGREGION( $G, f, \nu$ )
  for  $v \in G$  do
    if  $\exists \sigma$  s.t.  $f(\text{CONVEX}((G, v)^\sigma)) \leq \nu$  then
       $W \leftarrow \text{Attr}_1(v)$ 
       $W \leftarrow W \cup \text{WINNINGREGION}(G - \text{Attr}_1(v), f, \nu)$ 
    return  $W$ 
  end if
end for
return  $\emptyset$ 

```

Lemma 4. *Algorithm 1 computes player-1 winning region.*

Proof. We first prove that in every step of the algorithm, if a vertex $u \in W$, then $u \in \text{Win}_1$. We prove the assertion by considering the next three cases: (i) There is a strategy σ for which $f(\text{CONVEX}((G, u)^\sigma)) \leq \nu$. In this case, for every SCC $S \in (G, u)^\sigma$ we have that $\text{CONVEX}(S) \subseteq \text{CONVEX}((G, u)^\sigma)$ and by the monotonicity of f we get that $f(\text{CONVEX}(S)) \leq \nu$. Hence, $u \in \text{Win}_1$. (ii) There is a vertex v and a strategy σ s.t. $f(\text{CONVEX}((G, v)^\sigma)) \leq \nu$ and $u \in \text{Attr}_1(v)$. In this case, v is in player-1 winning region and therefore the attractor of v is also in Win_1 . (iii) For some vertex v we have $u \in \text{WINNINGREGION}(G - \text{Attr}_1(v), f, \nu)$ and $f(\text{CONVEX}((G, v)^\sigma)) \leq \nu$ for some strategy σ . By a simple induction on the size of the graph we get that u is in player-1 winning region for the game graph $G - \text{Attr}_1(v)$. The following strategy is a winning strategy for (G, u) : (a) play according to the winning strategy over $G - \text{Attr}_1(v)$; (b) if the pebble is in vertex v , then play according to σ . Hence, if $u \in W$, then $u \in \text{Win}_1$ and we get that $\text{Win}_1 \supseteq W$.

In order prove the converse direction, we first prove that if $\text{Win}_1 \neq \emptyset$, then $W \neq \emptyset$. Indeed, if $v \in \text{Win}_1$, then for some strategy σ we have that for every SCC $S \in (G, v)^\sigma$ it holds that $f(\text{CONVEX}(S)) \leq \nu$. Let S' be a terminal SCC in $(G, v)^\sigma$ and let (u, m) be a vertex in S' (where u is a vertex in G and m is a memory state of σ). Let σ' be the strategy that is formed by changing σ initial memory state to m . Then $(G, u)^{\sigma'} = S'$, and therefore $f(\text{CONVEX}((G, u)^{\sigma'})) \leq \nu$. Hence, the **if** condition in the **for** loop is satisfied at least once, and $W \neq \emptyset$. We are now ready to prove that $\text{Win}_1 \subseteq W$. Towards a contradiction we assume the existence of $u \in (\text{Win}_1 - W)$. By the definition of Algorithm 1 it follows that there is a subgraph $H \subseteq G$ such that $u \in H$ and the algorithm returns $\emptyset$ when it runs over H . Hence, player-1 winning region in H is empty (namely, $u \notin \text{Win}_1$ over game graph H).

and by Remark 1 we get that $u \notin \text{Win}_1$ in game graph G and the contradiction follows. Thus $\text{Win}_1 \subseteq W$. $\square$

We present a similar algorithm for the computation of quantitative analysis of quantitative objectives. For this purpose we extend the notion of winning regions to quantitative objectives by defining *value regions*. For a threshold ν we say that a vertex v is in ν *value region* (denoted by $VR(\nu)$) if $\inf_{\sigma \in \mathcal{FM}_1} \sup_{\tau \in \mathcal{S}_2} \text{Val}_{\sigma, \tau} = \nu$ (when the initial vertex of the game is v). Algorithm 2 computes value regions by a call to $\text{VALUEREGION}(G, f, -\infty)$, and its correctness follows by the same arguments as in the proof of Lemma 4. We note that if f satisfies Properties 1-3, then by Lemma 3, there is an effective procedure to compute $\inf_{\sigma \in \mathcal{FM}_1} f(\text{CONVEX}((G, v)^\sigma))$ (hence, Algorithm 2 can be effectively executed) and if f satisfies Properties 1-4, then by the same lemma we get that there is a procedure to determine whether $f(\text{CONVEX}((G, v)^\sigma)) \leq \nu$ (hence, Algorithm 1 can be effectively executed). Hence, we get the main result of this section.

Theorem 1. *Let f be the one-player solution of a quantitative objective.*

- *If f satisfies Properties 1-3, then the corresponding quantitative analysis problem is computable.*
- *If f satisfies Properties 1-4, then the corresponding boolean analysis problem is decidable.*

We note that Theorem 1 provides a recipe for the construction of ϵ -optimal strategies. If the infimum value that player 1 can achieve is ν , then the process that enumerates all $\sigma \in \mathcal{FM}_1$ and halts if the one-player solution of G^σ is at most $\nu + \epsilon$ will always terminate. Similarly, if the boolean analysis problem is decidable, then it is possible to effectively construct a finite-memory strategy that assures the corresponding threshold (we first check if such a strategy exists, and if it does exist, then we enumerate all finite-memory strategies until we find a strategy σ such that the solution for G^σ is at most ν).

5. Games with Mean-Payoff Expression Objectives

In this section, we give the formal definition of mean-payoff expressions and we use the results of Section 4 to analyze games with mean-payoff expressions. In Subsection 5.1 we define mean-payoff expressions and show that optimal strategies may require infinite memory. In Subsection 5.2 we analyze mean-payoff expressions.

5.1 Mean-payoff expression objectives

The class of mean-payoff expressions is the closure of single dimension mean-payoff objectives under the algebraic operations of $\min, \max, \text{sum}$ and numerical complement. Formally, for an infinite sequence of reals $\rho = a_1, a_2, \dots \in \mathbb{R}^\omega$, we denote $\text{LimInfAvg}(\rho) = \liminf_{n \rightarrow \infty} \frac{a_1 + \dots + a_n}{n}$ and $\text{LimSupAvg}(a_1, a_2, \dots) = \limsup_{n \rightarrow \infty} \frac{a_1 + \dots + a_n}{n}$. For an infinite sequence of vectors $\rho = v_1, v_2, \dots \in (\mathbb{R}^k)^\omega$ we denote the projection of ρ to the i -th dimension by ρ_i , and we denote $\text{LimInfAvg}_i(\rho) = \text{LimInfAvg}(\rho_i)$ and $\text{LimSupAvg}_i(\rho) = \text{LimSupAvg}(\rho_i)$. An *atomic expression* over $\mathbb{R}^k$ is either LimInfAvg_i or LimSupAvg_i . If E_1 and E_2 are expressions, then $-E_1$, $\max(E_1, E_2)$, $\min(E_1, E_2)$ and $\text{sum}(E_1, E_2)$ are also expressions. For a sequence $\rho \in (\mathbb{R}^k)^\omega$ and an expression E , the value of $E(\rho)$ is $\text{LimInfAvg}_i(\rho)$ if $E = \text{LimInfAvg}_i$, $\text{LimSupAvg}_i(\rho)$ if $E = \text{LimSupAvg}_i$, $\text{op}(E_1(\rho), E_2(\rho))$ if $E = \text{op}(E_1, E_2)$ (for $\text{op} \in \{\min, \max, \text{sum}\}$) and $-E_1(\rho)$ if $E = -E_1$. Over $\mathbb{R}^2$, a possible expression is $E = \min(\text{LimInfAvg}_1, \text{LimSupAvg}_1 + \text{LimInfAvg}_2) +$

$\max(\text{LimInfAvg}_1, \text{LimSupAvg}_2)$, and the value of E for the sequence $(-1, 1)^\omega$ is $\min(-1, -1 + 1) + \max(-1, 1) = 0$.

We say that an expression E is of *normal form* if (i) the numerical complement does not occur in E ; and (ii) for every dimension i , there is at most one occurrence of an atomic expression $A_i \in \{\text{LimInfAvg}_i, \text{LimSupAvg}_i\}$; and (iii) $E = \max(E_1, \dots, E_\ell)$, where E_i is a *max-free* expression (that is, the $\max$ operator does not occur in E_i). The next simple lemma shows that w.l.o.g we may consider only games over normal form expressions.

Lemma 5. *For every k -dimensional weighted graph G with a weight function w and an expression E , we can effectively construct an m -dimensional weight function w' and a normal form expression F such that every infinite path in G gets the same value according to (E, w) and according to (F, w') .*

Proof. We can easily overcome the restriction on the number of atomic expressions per dimension by creating several *copies* of the same dimension (that is, additional dimensions with weights that are identical to the original dimension). We can create an equivalent numerical complement free expression by the following recursive process. If $E = -\text{LimInfAvg}_i$ (respectively, $E = -\text{LimSupAvg}_i$), then we multiply all the weights in dimension i by -1 and define $F = \text{LimSupAvg}_i$ (resp. $F = \text{LimInfAvg}_i$). F is equivalent to E since $\text{LimInfAvg}(a_1, a_2, \dots) = -\text{LimSupAvg}(-a_1, -a_2, \dots)$. If $E = -\text{op}(E_1, E_2)$, then we recursively change the weights and construct normal form expressions F_1 and F_2 that are equivalent to $-E_1$ and $-E_2$, and return the normal form expression $F = \text{op}(F_1, F_2)$. And we similarly handle the expression $E = \text{op}(E_1, E_2)$. Finally, if we have a numerical complement free expression E , then we construct an equivalent expression $F = \max(F_1, \dots, F_\ell)$, where F_i is a max free expression, by the following recursive procedure: If E is an atomic expression, then we return $F = \max(E, E)$. If $E = \text{op}(E_1, E_2)$, then we recursively construct two expressions F_1 and F_2 , such that F_i is equivalent to E_i and $F_1 = \max(G_1, \dots, G_r)$, $F_2 = \max(H_1, \dots, H_q)$ (where H_i and G_i are max-free expressions), and we return $F = \max_{i \in \{1, \dots, r\}, j \in \{1, \dots, q\}} \{\text{op}(G_i, H_j)\}$. $\square$

Hence, in the rest of the paper we will assume w.l.o.g that all the expressions are of normal form. The next example shows that optimal strategies for mean-payoff expressions may require infinite memory.

Example 4. *Consider the game graph in Figure 7 and the expression $E = \max(\text{LimInfAvg}_1, \text{LimInfAvg}_2)$. In this game graph there is only one vertex that is controlled by player 1 and two self-loop edges, namely e_1 with $w(e_1) = (9, 1)$ and e_2 with $w(e_2) = (1, 9)$. We first observe that any finite-memory strategy gives a value of at least 5 to E . Indeed, a finite-memory strategy induces an ultimately periodic path π with $\text{LimInfAvg}(\pi) = \alpha w(e_1) + (1 - \alpha)w(e_2)$ for some $\alpha \in [0, 1] \cap \mathbb{Q}$. Hence, $E(\pi) = \max(9\alpha + 1 - \alpha, \alpha + 9 - 9\alpha) = \max(8\alpha + 1, 9 - 8\alpha)$, and the minimum value for E is obtained when $\alpha = \frac{1}{2}$ and we get that the minimal value for E is 5. We now describe a player-1 infinite-memory strategy that gives a value of at most 2 to E . The strategy is simple. It follows e_2 as long as the average weight in the first dimension is more than 2, then it follows e_1 as long as the average weight in the second dimension is more than 2, and this process is repeated forever (i.e., e_2 is followed for a while, then e_1 and so on). Clearly, in the formed path π the average weight of the first dimension is at most 2 for infinitely many prefixes of π . Hence $\text{LimInfAvg}_1(\pi) \leq 2$, and by the same arguments $\text{LimInfAvg}_2(\pi) \leq 2$. Thus, $E(\pi) \leq 2$, and we establish the fact that optimal strategies may require infinite-memory strategies (in this example, the presented infinite-memory strategy is not optimal,*

Algorithm 2 Value region computation for quantitative objectives. The algorithm invokes $\text{VALUEREION}(G, f, -\infty)$.

$\text{VALUEREION}(G, f, \text{ValLowerBound})$

```

if  $G \neq \emptyset$  then
  for  $v \in G$  do
     $I[v] \leftarrow \max(\inf_{\sigma \in \mathcal{FM}_1} f(\text{CONVEX}((G, v)^\sigma)), \text{ValLowerBound})$ 
  end for
   $u \leftarrow \text{argmin}_{v \in G} I[v]$  {Choose  $u$  s.t  $I[u] = \min_{v \in G} I[v]$ }
   $\text{VR}(I[u]) \leftarrow \text{VR}(I[u]) \cup \text{Attr}_1(u)$  {Add  $\text{Attr}_1(u)$  to the value region of  $I[u]$ }
  return  $\text{VALUEREION}(G - \text{Attr}_1(u), f, I[u])$  {Continue the computation recursively. The new lower bound is  $I[u]$ .}
end if

```




Figure 7.

but we demonstrated that the best finite-memory strategy does not give an optimal value).

5.2 Synthesis of a finite-memory controller for mean-payoff expression objectives

In this subsection we apply Theorem 1 to mean-payoff expression objectives. We first prove that the solution for mean-payoff expressions satisfies Properties 1-3, and thus the quantitative analysis problem is computable for mean-payoff expression games. We then show that the boolean analysis problem is inter-reducible with Hilbert's tenth problem over rationals ($H10(\mathbb{Q})$) by showing that an effective algorithm for $H10(\mathbb{Q})$ implies that mean-payoff expressions satisfy Property 4, and by a reduction from $H10(\mathbb{Q})$ to mean-payoff expression games.

One-player games were solved in [6] and in [15]. We present our solution from [15] to establish properties of the one-player solution. For an expression E and a one-player game (G, v_0) , that is, a game over graph G with initial vertex v_0 , we say that a threshold ν is *feasible* if the player has a strategy that achieves a value at least ν (we recall that in the one-player setting, the player aim to maximize the value of the objective). The *max-free* constraints were presented in [14] (Section A.4), and they describe the feasible thresholds of a max-free expression (a threshold ν is feasible if the one-player can achieve a value of at least ν).

Definition 1 (Max-free constraints). *Let G be a strongly-connected k -dimensional game graph, and we recall that $C(G)$ is the set of simple cycles of G . Let E be a max-free expression such that the first j dimensions of G occur in E as *lim-inf* (and the others as *lim-sup*). We define a variable X_c^i for every simple cycle c and index $i \in \{j+1, \dots, k\}$, and we define a vector of variables $\bar{r} = (r_1, \dots, r_{2k})$. Then the max-free constraints for threshold $\nu \in \mathbb{Q}$ are*

1. $\sum_{c \in C(G)} X_c^i \text{Avg}_m(c) \geq r_m$ for every $i \in \{j+1, \dots, k\}$ and $m \in \{1, \dots, j, i\}$
2. $\sum_{c \in C(G)} X_c^i = 1$ for every $i \in \{j+1, \dots, k\}$
3. $X_c^i \geq 0$ for every $i \in \{j+1, \dots, k\}$ and $c \in C(G)$
4. $M_E \times \bar{r} \geq (0, \dots, 0, \nu)^T$

where M_E is a matrix that is independent of the graph, and computable from E . (We note that in [14], the first type of constraints was $\sum_{c \in C(G)} X_c^i w_m(c) \geq r_m$, where w_m is the projection of w to the m -th dimension, and the second type of constraints was $\sum_{c \in C(G)} |c| X_c^i = 1$. It is straight forward to verify that the constraints are equivalent — in terms of feasibility. In addition, the fourth constraint was presented as $M_E \times \bar{r} \geq \bar{b}_\nu$; but

the proof of Lemma 7 in [14] implies that $\bar{b}_\nu = (0, \dots, 0, \nu)^T$.) We proved in [15] that a threshold ν is feasible if and only if the corresponding max-free constraints are feasible. For a max-free expression E , a strongly-connected graph G and a threshold ν , we denote the max-free constraints by $\text{MFC}(E, G, \nu)$ and we observe that for a (normal-form) mean-payoff expression $E = \max(E_1, \dots, E_\ell)$ and a strongly-connected graph G , the solution function for the one-player game is $f(G) = \max\{\nu \in \mathbb{R} \mid \exists i \in \{1, \dots, \ell\} \text{ s.t } \text{MFC}(E_i, G, \nu) \text{ is feasible}\}$. By the definition of the max-free constraints, it easily follows that the solution is a function that is first-order definable and continuous (i.e., it satisfies Properties 1 and 3). In the next Lemma we prove that the solution also satisfies the second property.

Lemma 6. *Let E be a mean-payoff expression over k dimensions, and let f be its one-player solution function. Then for every two strongly-connected graphs G and H : if $\text{CONVEX}(H) \subseteq \text{CONVEX}(G)$, then $f(H) \leq f(G)$.*

Proof. Since we assume that $E = \max(E_1, \dots, E_n)$, where E_i is a max-free expression, it is enough to prove that if a threshold ν is feasible in H for the max-free expression E_i , then it is also feasible in G . Let $c_1^G, \dots, c_n^G$ and $c_1^H, \dots, c_m^H$ be the simple cycles of G and H respectively. We note that since $\text{CONVEX}(H) \subseteq \text{CONVEX}(G)$, then for every convex combination $x_1, \dots, x_m$, there is a convex combination $y_1, \dots, y_n$ such that $\sum_{i=1}^m x_i \text{Avg}(c_i^H) = \sum_{i=1}^n y_i \text{Avg}(c_i^G)$. Hence, a solution for the max-free constraints over graph H induces a solution for the max-free constraints over G (by replacing, in the inequalities of constraints 1 over graph H , every convex combination of cycles of H by the corresponding convex combination of cycles of G).

Thus, every threshold that is feasible for H is also feasible for G , and the proof follows. $\square$

Hence, the one-player solution function of mean-payoff expressions satisfies Properties 1-3 and the next theorem follows.

Theorem 2. *The quantitative analysis problem for mean-payoff expression games (where player 1 is restricted to finite-memory strategies) is computable.*

We now show that the solution for one-player mean-payoff expression games satisfies Property 4 if and only if $H10(\mathbb{Q})$ is decidable. We first prove the direction from right to left.

Lemma 7. *If $H10(\mathbb{Q})$ is decidable, then mean-payoff expressions satisfy the fourth property.*

Proof. Let G be an arbitrary strongly connected graph with n simple cycles, let $C(G) = \{c_1, \dots, c_n\}$ be its set of simple cycles, let $E = \max(E_1, \dots, E_m)$ be a mean-payoff expression (where E_i is a max-free expression), and let ν be a rational threshold. We recall that the sentence $(\zeta_n(\text{Avg}(c_1), \dots, \text{Avg}(c_n)) \leq \nu)$ is equivalent to the statement: "For every $y > \nu$ and $i \in \{1, \dots, m\}$, the constraints $\text{MFC}(E_i, G, y)$ are infeasible." By the definition

of the max-free constraints, when the set $Avg(G)$ is fixed the above statement is easily reduced to the infeasibility of m linear systems. Motzkins Transposition Theorem (e.g., Theorem 1 in [2]) gives a witness to the infeasibility of a set of linear inequalities. We use Lagrange four-square Theorem to construct a Diophantine equation that has a rational root if and only if the witness exists. We show that the construction works also when $Avg(G)$ is not fixed, i.e., when $Avg(G) = \{A_1\bar{x}_1, \dots, A_1\bar{x}_n\}$ for some n matrices $A_1, \dots, A_n$ and n vectors of rational variables $\bar{x}_1, \dots, \bar{x}_n \in \mathbb{QSL}$. (The details of the construction are given in the technical report [17].) Hence, if $H10(\mathbb{Q})$ is decidable, then mean-payoff expressions satisfy the fourth property. $\square$

We now prove the reduction from $H10(\mathbb{Q})$ to the boolean analysis of mean-payoff expression games, and we show that there is a reduction even for a simpler subclass of mean-payoff expressions. An expression E is sum-free and $LimInfAvg$ -only if only the min and max operators occur in E and all the atomic expressions in E are of the form $LimInfAvg_i$. (In addition, the numerical complement operator also does not occur). The next lemma shows that the boolean analysis problem for sum-free $LimInfAvg$ -only expressions is $H10(\mathbb{Q})$ -hard.

Lemma 8. *If the boolean analysis problem is decidable for sum-free $LimInfAvg$ -only expressions, then $H10(\mathbb{Q})$ is decidable.*

Proof. We only present a rough and informal sketch of the proof. The full proof is given in the technical report [17]. We first show a reduction from $H10(\mathbb{Q})$ to the problem of finding a rational solution for two set of variables $Q = \{q_1, \dots, q_n\}$ and $P = \{p_1, \dots, p_n\}$ and a set of constraints, each of them is of the form: (i) $\sum_{i \in I} \alpha_i q_i \leq 0$, for some $I \subseteq \{1, \dots, n\}$; or (ii) $\sum_{i \in I} \alpha_i p_i \leq 0$, for some $I \subseteq \{1, \dots, n\}$; or (iii) $q_i p_j = q_k p_\ell$ for some $i, j, k, \ell \in \{1, \dots, n\}$; subject to $q_i, p_i > 0$. We then show a reduction from the boolean analysis problem to the above problem. We illustrate the reduction by showing the construction for the set of constraints $\{q_1 - 2q_2 \leq 0, 2p_1 - 3p_2 \leq 0, p_1 q_1 = p_2 q_2\}$. For these constraints we build a game graph G that is illustrated in Figure 8. In the figure we explicitly show only part of the weight vectors and only part of the dimensions. The initial vertex of G is s_0 and this vertex is a player-2 vertex (and the rest are player-1 vertices). The objective of the game is the expression $E = \max(1, \min(2, 3), \min(4, 5))$, where i stands for $LimInfAvg_i$ (for $i = 1, 2, 3, 4, 5$), and the threshold is 0. In G , player 2 has only two memoryless strategies, namely $\tau_1 = s_0 \rightarrow b_1$ and $\tau_1 = s_0 \rightarrow a_1$. We rely on Lemma 4 and show that player 1 has a winning strategy if and only if there is a finite path π_1 that visits only the self-loops of b_2 and a path π_2 that visits only the self-loops of a_2 such that every vector in $v \in CONVEX(Avg(\pi_1), Avg(\pi_2))$ satisfies the winning condition, i.e., if $v = (v_1, v_2, v_3, v_4, v_5)$, then $\max(v_1, \min(v_2, v_3), \min(v_4, v_5)) \leq 0$. We observe that for any such path π_1 it holds that $Avg(\pi_1) = q_1(3, 1, 0, -1, 0) + q_2(-2, 0, -1, 0, 1)$ and similarly $Avg(\pi_2) = p_1(2, -1, 0, 1, 0) + p_2(-1, 0, 1, 0, -1)$ for some positive rational q_1, q_2, p_1, p_2 . We further observe that if $q_1 - 2q_2 > 0$, then $Avg(\pi_1)$ is positive in the first dimension, and thus there is a vector $v \in CONVEX(Avg(\pi_1), Avg(\pi_2))$ that gives a positive value to the expression. Hence, it must hold that $q_1 - 2q_2 \leq 0$ and similarly $2p_1 - 3p_2 \leq 0$. Moreover, we prove that if $p_1 q_1 \neq p_2 q_2$, then there exists $v \in CONVEX(Avg(\pi_1), Avg(\pi_2))$ that is positive either in dimensions 2 and 3 or in dimensions 4 and 5. Hence, it must hold that $p_1 q_1 = p_2 q_2$ and the proof follows. $\square$

The next theorem summarizes the results of Lemmas 8 and 7




Figure 8.

Theorem 3. *The boolean analysis problem for mean-payoff expression games (when player 1 is restricted to finite-memory strategies) is inter-reducible with $H10(\mathbb{Q})$, and it is $H10(\mathbb{Q})$ -hard even for sum-free $LimInfAvg$ -only expressions.*

We also consider the case where both players are restricted to finite-memory strategies. In this setting, the quantitative analysis problem is to compute $\inf_{\sigma \in \mathcal{FM}_1} \sup_{\tau \in \mathcal{FM}_2} Val_{\sigma, \tau}$. The boolean analysis problem is to determine whether player 1 has a finite-memory strategy that assures a value of at most ν against any player-2 finite-memory strategy.

Theorem 4. *When both players are restricted to finite-memory strategies: (i) the quantitative analysis problem for mean-payoff expression games is computable; (ii) the boolean analysis problem for mean-payoff expression games is inter-reducible with $H10(\mathbb{Q})$, and it is $H10(\mathbb{Q})$ -hard even for sum-free $LimInfAvg$ -only expressions.*

Proof. We present a sketch of the proof. The full proof is given in the technical report [17]. Informally, when both players are restricted to finite-memory strategies, the outcome of a play is an ultimately periodic path, and thus we may assume that all the atomic expressions are of the form $LimInfAvg$ (because for periodic paths we have $LimInfAvg(\pi) = LimSupAvg(\pi)$). We also show that if all the atomic expressions are $LimInfAvg$ and player 1 is restricted to finite-memory strategies, then player 2 can achieve a value greater than ν if and only if he can do it with a finite-memory strategy, and the proof follows. $\square$

As a final remark, we note that while the boolean analysis for sum-free $LimInfAvg$ -only expressions is $H10(\mathbb{Q})$ -hard when player 1 is restricted to a finite-memory strategy (and also when both players are restricted to finite-memory strategies), the next lemma shows that the problem is decidable when both players may use arbitrary strategies.

Lemma 9 (Theorem 5 in [16]). *When both players may use arbitrary strategies, the boolean analysis of sum-free $LimInfAvg$ -only expression games is decidable.*

Proof. The proof follows from Theorem 5 in [16] due to the fact that there is an immediate translation from sum-free $LimInfAvg$ -only expressions to the $\bigvee \bigwedge \text{MeanPayoffInf}^{\leq}(\nu)$ objectives that were defined in [16]. $\square$

6. Discussion and Future Work

In this work we studied the synthesis of finite-memory strategies for games with robust multidimensional mean-payoff objectives, and we obtained two main results. The first is a positive result, namely, the computability of the quantitative analysis problem. The second has a negative flavour, and it shows that the boolean analysis is inter-reducible with Hilbert's Tenth problem over rationals.

From a practical point of view, the positive result is the most interesting, since for the first time (to the best of our knowledge) a recipe is given for computing ϵ -optimal finite-memory strategies for a robust class of quantitative objectives. A future work is to investigate whether the construction of these ϵ -optimal strategies is feasible, both in terms of memory size and computational complexity. From the theoretical point of view, the negative result is a bit surprising since it suggests that the boolean analysis is harder than the optimization problem, and in computer science typically there is a naive reduction from optimization problems to the corresponding decision problems. However, in our case, the optimization computes only the greatest upper bound, and since optimal finite-memory strategies need not exist, then the reduction fails. In fact, the hardness result suggests that it is even $H10(\mathbb{Q})$ -hard to determine whether an optimal strategy exists. A future work is to investigate games in which player 1 may use arbitrary infinite-memory strategies. Additional direction for future work is to consider more general algebraic structures over multidimensional mean-payoff objectives.

Acknowledgments

I would like to thank Prof. Alexander Rabinovich for many discussions on this work, and in particular for (i) discussions on Hilbert tenth problem and Tarski's Theorem; and (ii) suggesting to present a modular and abstract solution.

References

- [1] R. Alur, A. Degorre, O. Maler, and G. Weiss. On omega-languages defined by mean-payoff conditions. *FOSSACS*, 2009.
- [2] A. Ben-Israel. Motzkins transposition theorem, and the related theorems of farkas, gordan and stiemke. *Encyclopaedia of Mathematics*, 2001.
- [3] U. Boker, K. Chatterjee, T. A. Henzinger, and O. Kupferman. Temporal specifications with accumulative values. In *LICS*, 2011.
- [4] T. Brázdil, V. Brozek, K. Chatterjee, V. Forejt, and A. Kucera. Two views on multiple mean-payoff objectives in markov decision processes. In *LICS*, 2011.
- [5] T. Brázdil, K. Chatterjee, A. Kucera, and P. Novotný. Efficient controller synthesis for consumption games with multiple resource types. In *CAV*, 2012.
- [6] K. Chatterjee, L. Doyen, H. Edelsbrunner, T. A. Henzinger, and P. Rannou. Mean-payoff automaton expressions. In *CONCUR*, 2010.
- [7] K. Chatterjee, L. Doyen, and T. A. Henzinger. Expressiveness and closure properties for quantitative languages. *LICS*, 2009.
- [8] K. Chatterjee, L. Doyen, and T. A. Henzinger. Quantitative languages. *ACM Trans. Comput. Log.*, 2010.
- [9] K. Chatterjee, L. Doyen, T. A. Henzinger, and J.-F. Raskin. Generalized mean-payoff and energy games. In *FSTTCS*, 2010.
- [10] K. Chatterjee, M. Randour, and J.-F. Raskin. Strategy synthesis for multi-dimensional quantitative objectives. In *CONCUR*, 2012.
- [11] M. Droste and G. Rahonis. Weighted automata and weighted logics on infinite words. In *Developments in Language Theory*, 2006.
- [12] E. Kopczyński. Half-positional determinacy of infinite games. *ICALP*, 2006.
- [13] A. Tarski. A decision method for elementary algebra and geometry. 1951.
- [14] Y. Velner. The complexity of mean-payoff automaton expression. *CoRR*, 2011.
- [15] Y. Velner. The complexity of mean-payoff automaton expression. *ICALP*, 2012.
- [16] Y. Velner and A. Rabinovich. Church synthesis problem for noisy input. In *FOSSACS*, 2011.
- [17] Y. Velner. Finite-memory strategy synthesis for robust multidimensional mean-payoff objectives. *CoRR*, arXiv:1209.1007v2, 2014.