



I|S|T AUSTRIA

Institute of Science and Technology

Quantitative Fair Simulation Gamest

Krishnendu Chatterjee and Thomas A Henzinger and Jan Otop and Yaron Velner

Technical Report No. IST-2014-315-v1+1
Deposited at 05 Dec 2014 11:36
<http://repository.ist.ac.at/315/1/report.pdf>

IST Austria (Institute of Science and Technology Austria)
Am Campus 1
A-3400 Klosterneuburg, Austria

Copyright © 2012, by the author(s).

All rights reserved.

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Quantitative Fair Simulation Games

Krishnendu Chatterjee¹, Thomas A. Henzinger¹, Jan Otop¹, and Yaron Velner²

¹IST Austria

²Tel Aviv University

December 4, 2014

Abstract

Simulation is an attractive alternative for language inclusion for automata as it is an under-approximation of language inclusion, but usually has much lower complexity. For non-deterministic automata, while language inclusion is PSPACE-complete, simulation can be computed in polynomial time. Simulation has also been extended in two orthogonal directions, namely, (1) fair simulation, for simulation over specified set of infinite runs; and (2) quantitative simulation, for simulation between weighted automata. Again, while fair trace inclusion is PSPACE-complete, fair simulation can be computed in polynomial time. For weighted automata, the (quantitative) language inclusion problem is undecidable for mean-payoff automata and the decidability is open for discounted-sum automata, whereas the (quantitative) simulation reduce to mean-payoff games and discounted-sum games, which admit pseudo-polynomial time algorithms.

In this work, we study (quantitative) simulation for weighted automata with Büchi acceptance conditions, i.e., we generalize fair simulation from non-weighted automata to weighted automata. We show that imposing Büchi acceptance conditions on weighted automata changes many fundamental properties of the simulation games. For example, whereas for mean-payoff and discounted-sum games, the players do not need memory to play optimally; we show in contrast that for simulation games with Büchi acceptance conditions, (i) for mean-payoff objectives, optimal strategies for both players require infinite memory in general, and (ii) for discounted-sum objectives, optimal strategies need not exist for both players. While the simulation games with Büchi acceptance conditions are more complicated (e.g., due to infinite-memory requirements for mean-payoff objectives) as compared to their counterpart without Büchi acceptance conditions, we still present pseudo-polynomial time algorithms to solve simulation games with Büchi acceptance conditions for both weighted mean-payoff and weighted discounted-sum automata.

1 Introduction

Language inclusion and simulation. Language inclusion is a central decision question for automata as it subsumes language emptiness, universality, and can express language equivalence. Unfortunately, for non-deterministic automata, the language inclusion problem is PSPACE-complete. To mitigate the high complexity issue, language inclusion is often under-approximated by a stronger notion of *simulation*, which has polynomial-time complexity [1]. The simulation between automata is captured by a two-player game [1, 2, 3], called the *simulation game*. The simulation game proceeds in turns played by Challenger and Simulator; in each turn Challenger produces a transition of the first automaton, which then Simulator tries to match with a transition of the second automaton. A winning strategy for Simulator proves simulation of the first automaton by the second automaton and implies inclusion of the language of the first automaton in the language of second automaton. Intuitively, in the language inclusion game Challenger must play oblivious of the choices of Simulator (i.e., it is a partial-information game where Challenger has partial information), whereas the simulation game strengthens Challenger by allowing him to observe the choices of Simulator (and thereby obtaining a perfect-information game). The game-theoretical characterization allows for natural generalizations of the simulation notion.

Fair and quantitative simulation. The notion of simulation has been extended in two orthogonal directions. The first is to consider Büchi acceptance conditions on the executions of the automata [4], which leads to a *fair-simulation game*. In the fair simulation game, along with the automata, each player is given a Büchi acceptance condition, and Challenger wins if his own acceptance condition is satisfied, and either the acceptance condition of Simulator is violated or Challenger wins the simulation game. While the standard simulation game is a safety game, whose violation can be witnessed in a finite number of steps, to consider infinite computations, one needs to specify which are the desired infinite computations for the simulation relation. The Büchi acceptance conditions

specify the desired set of infinite computations, and thus classical simulation is extended to fair simulation. Note that since we consider non-deterministic automata with Büchi acceptance conditions they can express all ω -regular properties (such as liveness, fairness, and all commonly used specifications in practice). Another extension of the simulation notion is *quantitative simulation* defined on weighted automata [5]. A weighted automaton assigns a real value to every run. For example, a weighted automaton consists of an automaton with an integer-valued weight assigned to every transition, and two classical quantitative objectives are the *mean-payoff* objective that assigns to a run the long-run average of the weights along the run, whereas the *discounted-sum* objective assigns the discounted sum of the weights. A game that characterizes quantitative simulation proceeds as a classical simulation game, but the winning condition for Simulator is stronger. Simulator has to match all transitions picked by Challenger and, moreover, the run defined by transitions picked by Simulator needs to have the value at most the value of the run of Challenger. Recently the analysis of quantitative properties has received a lot of attention to specify quantitative aspects of systems [5, 6, 7], such as resource-consumption, response time etc, and has also been used in synthesis of systems [8, 9]. Weighted automata are central in quantitative verification and synthesis as they provide a natural specification formalism to express quantitative properties.

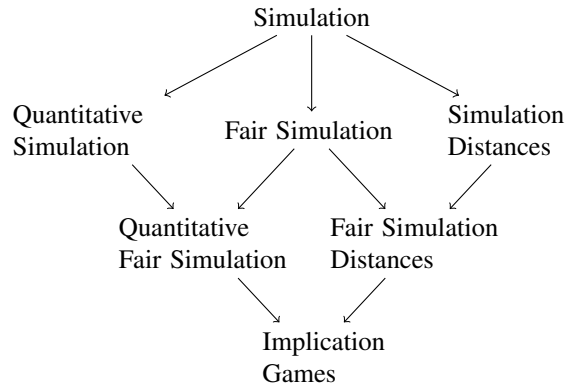
Previous results. The classical simulation game problem is a perfect-information safety game, and can be solved in linear time [10, 11]. The fair-simulation game problem is a perfect-information game where the objective is if a Büchi condition is satisfied, then another Büchi condition should be satisfied, along with a safety condition. Such games can be solved in quadratic time [12, 13]. Note that while both language inclusion and fair trace inclusion problems are PSPACE-complete, the corresponding simulation question can be solved in polynomial time. The quantitative simulation problem for weighted automata with the quantitative objective defined as mean-payoff (resp., discounted-sum) reduces to perfect-information games with mean-payoff (resp., discounted sum) objectives [5]. Using the results for perfect-information mean-payoff and discounted-sum games [14, 15], it follows that there exist pseudo-polynomial time algorithms for the quantitative simulation problem, and if the weights are encoded in unary, then the algorithm is polynomial. This is in sharp contrast to the language inclusion problem for weighted automata, where the problem is undecidable for mean-payoff objectives [16], and the decidability is open for discounted-sum objectives [5]. One very crucial property of the games both for fair simulation as well as quantitative simulation is that in the corresponding perfect-information games each player has memoryless optimal strategies (where a memoryless strategy does not depend on the history of interactions and depends only on the current position of the game). For the game problem for quantitative simulation, we obtain a perfect-information game with the same objective as the objective of the weighted automata.

Our results. In this paper we generalize both fair simulation and quantitative simulation. We consider simulation of weighted automata with Büchi acceptance conditions. We call the resulting perfect-information game *Quantitative Fair Simulation Games (QFSGs)*, which generalize both fair simulation games and games for quantitative simulations. To capture QFSGs we present a framework of new quantitative games called *implication games* defined as follows. In quantitative games an objective is a function from plays to reals, i.e., each play has an outcome, which is a real number. Then, the goal of Player 1 is to minimize the outcome against any strategy of Player 2, who aims to maximize the outcome of the play. Implication games result from imposing additional Boolean conditions on both players, i.e., the value of a play is (i) ∞ if Player 1 fails to satisfy his Boolean condition; (ii) $-\infty$ if Player 1 satisfies his Boolean condition, but Player 2 fails to satisfy his Boolean condition; and (iii) if both players satisfy their Boolean conditions, then the payoff is determined according to the objective function. QFSGs are exactly captured by implication games. However, the implication games are fundamentally different both from fair simulation games as well as games for quantitative simulation. We show that (i) for mean-payoff objectives, in implication games both players require infinite memory strategies for optimality, and optimal strategies exist, whereas (ii) for discounted-sum objectives, for every $\epsilon > 0$, there is a finite-memory ϵ -optimal strategy, but in general, optimal strategies need not exist. This is in sharp contrast to games for quantitative simulation with mean-payoff and discounted-sum objectives which admit memoryless optimal strategies for both players. While QFSGs are more complicated (e.g., due to infinite-memory requirements for mean-payoff objectives) as compared to games for quantitative simulation, we still present pseudo-polynomial time algorithms to solve QFSGs for both weighted mean-payoff and weighted discounted-sum automata. If the weights are encoded in unary, then our algorithm works in polynomial time. In summary, the contributions of this paper are as follows:

1. We introduce implication games.
2. We solve implication games with Büchi objectives for both players and the sum of mean-payoff objectives.
3. We define and solve discounted-sum parity games.
4. We solve implication games with Büchi objectives for both players and the discounted-sum objective.

Another important application: fair simulation distance. Our results for implication games apply not only for QFSGs, but also generalize other important applications. The simulation problem for non-deterministic non-weighted automata has also been generalized to compute simulation distance between automata. In such a setting, if a transition of one automata is not matched, instead of loosing the game immediately, the Simulator must pay a penalty, and the goal of the Simulator is to minimize the average penalty. Quantitative simulation games have been employed to measure similarity properties between automata, such as *correctness*, *robustness* or *coverage* [17]. Again the problem of simulation distances reduce to mean-payoff games. The simulation distance problem was considered for automata without fairness. Since we consider Büchi acceptance conditions, our results for implication games provide the solution to compute *fair simulation distances* (i.e., simulation distances with respect to infinite computations described as Büchi conditions). In summary, our results for implication games on one hand provide a solution for QFSGs, which generalize both fair simulation and quantitative simulation; and on the other hand they provide the solution to compute fair simulation distances, which generalize simulation distances.

Pictorial illustration. The following picture shows the various generalizations. Simulation has been generalized to quantitative simulation (between weighted automata), fair simulation (simulation with respect to Büchi acceptance conditions for non-weighted automata), and simulation distances (to measure distances between non-weighted automata). We generalize quantitative and fair simulation to define QFSGs (for fair simulation for weighted automata). Our solution framework of implications games for QFSGs also apply to compute fair simulation distances (i.e., simulation distances for non-weighted automata with Büchi acceptance conditions). In this paper we will only consider QFSGs and not present details about simulation distances, which has been studied in [17] and reduction was presented to mean-payoff games and adding Büchi acceptance conditions directly leads to implication games.



Organization and technical contributions of the paper. The organization and the technical contributions are as follows:

1. In Section 2 we present notation and notions used throughout the paper.
2. In Section 3 we define implication games, which have not been considered before to the best of our knowledge.
3. In Section 4 we define QFSGs, which correspond to simulation between weighted automata with Büchi acceptance conditions. We identify classes of implication games that correspond to QFSGs.
4. In Section 5, we study implication games motivated by QFSGs for mean-payoff automata with Büchi acceptance condition. In these games both players have to satisfy a Büchi condition, and if the Büchi conditions are satisfied, then the result of the game is determined by the sum of two mean-payoff objectives. We present a reduction from the sum of two mean-payoff objectives to a two dimensional mean-payoff objective, where the protagonist has to ensure a value of at least q in the first dimension and a value of at least p in the second dimension. We then show how to use results of [18] to obtain a pseudo-polynomial time algorithm for our implication games.
5. In Section 6 we study two classes of implication games. First, we consider *discounted-sum parity games* in which Player 1 has to ensure that a parity objective is satisfied and the value of the game is negative. While mean-payoff parity games have been studied before, discounted-sum parity games were not studied before and our result may be of independent interest. Second, we consider implication games that correspond to QFSGs for discounted-sum automata with Büchi acceptance conditions. For discounted-sum parity games

we present a polynomial-time Turing reduction to discounted-sum games and parity games, and for implication games with discounted-sum objectives and Büchi acceptance conditions, we present a polynomial-time Turing reduction to discounted-sum games.

6. We conclude the paper and discuss the future work in Section 7.

2 Preliminaries

2.1 Automata

Words. Given a finite alphabet Σ of letters, a word w is an infinite sequence of letters. For a word w , we define $w[i]$ as the i -th letter of w .

Non-deterministic automata. A (*non-deterministic*) automaton \mathcal{A} is a tuple $(\Sigma, Q, q_0, \delta, F)$, where Σ is the alphabet, Q is a finite set of states, $q_0 \in Q$ is the initial state, $\delta \subseteq Q \times \Sigma \times Q$ is a transition relation, and $F \subseteq Q$ is a set of *accepting* states.

Runs. Given an automaton \mathcal{A} and an infinite word w , a *run* $\pi = \pi[0]\pi[1]\dots$ is an infinite sequence of states such that $\pi[0] = q_0$ and for every $i \in \mathbb{N} \setminus \{0\}$ we have $(\pi[i-1], w[i], \pi[i]) \in \delta$. Given a word w , we denote by $\text{Run}(w)$ the set of all possible runs on w .

Büchi acceptance. The (Büchi) acceptance of words is defined using the accepting states. An (infinite) run π is (*Büchi*) *accepting*, if there exists infinitely many j such that $\pi[j] \in F$. Let $\text{Acc}(w) \subseteq \text{Run}(w)$ denote the set of accepting runs, and a word w is accepted iff $\text{Acc}(w)$ is non-empty. We denote by $\mathcal{L}_{\mathcal{A}}$ the set of words accepted by \mathcal{A} .

Weighted automata. A *weighted automaton* is an automaton whose transitions are labeled by integers \mathbb{Z} . Formally, a weighted automaton \mathcal{A} is a tuple $(\Sigma, Q, q_0, \delta, F, \text{wt})$ such that $(\Sigma, Q, q_0, \delta, F)$ is an automaton and $\text{wt} : \delta \mapsto \mathbb{Z}$. The labels of the transitions are referred to as *weights*.

Semantics of weighted automata. To define the semantics of weighted automata we need to define the value of a run (that combines the sequence of weights of a run to a single value) and the value across runs (that combines values of different runs to a single value). To define values of runs, we will consider *value functions* f that assign real numbers to sequences of integers. Given a word w , every run π of \mathcal{A} on w defines a sequence of weights of successive transitions of \mathcal{A} , i.e., $\text{wt}(\pi) = (\text{wt}(\pi[i-1], w[i], \pi[i]))_{i \in \mathbb{N} \setminus \{0\}}$; and the value $f^{\text{wt}}(\pi)$ of the run π is defined as $f(\text{wt}(\pi))$. The value of a word w assigned by the automaton \mathcal{A} , denoted by $\mathcal{L}_{\mathcal{A}}(w)$, is the infimum of the set of values of all *accepting* runs; i.e., $\inf_{\pi \in \text{Acc}(w)} f^{\text{wt}}(\pi)$, and we have the usual semantics that infimum of an empty set is infinite, i.e., the value of a word that has no accepting runs is infinite. To indicate a particular value function f that defines the semantics, we will call a weighted automaton \mathcal{A} an f -automaton. We consider value functions from VALFUNC defined below.

Value functions. We consider the following functions that aggregate a sequence of integers into a single real number. Let $\bar{a} = (a_i)_{i \geq 1}$ be a sequence of integers, and let $\text{Avg}_k(\bar{a}) = \frac{1}{k} \cdot \sum_{i=1}^k a_i$. We define the following functions:

1. *Limit-average infimum:* $\text{LIMAVGINF}(\bar{a}) = \liminf_{k \rightarrow \infty} \text{Avg}_k(\bar{a})$.
2. *Limit-average supremum:* $\text{LIMAVGSUP}(\bar{a}) = \limsup_{k \rightarrow \infty} \text{Avg}_k(\bar{a})$.
3. *Discounted-sum:* $\text{DISC}_{\lambda}(\bar{a}) = \sum_{i=1}^{\infty} \lambda^i \cdot a_i$, for $0 < \lambda < 1$.

We define $\text{VALFUNC} = \{\text{LIMAVGINF}, \text{LIMAVGSUP}, \text{DISC}_{\lambda}\}$.

Quantitative inclusion. We consider the quantitative variant of the inclusion question. Given an f -automaton \mathcal{A}_1 and a g -automaton \mathcal{A}_2 the *inclusion* question asks whether for every word w we have $\mathcal{L}_{\mathcal{A}_1}(w) \leq \mathcal{L}_{\mathcal{A}_2}(w)$. If it is the case we say that \mathcal{A}_1 *includes* \mathcal{A}_2 . The quantitative inclusion generalizes the (Boolean) inclusion question. Indeed, Boolean automata can be considered as weighted automata that assign the value 1 to rejected words and 0 to accepted words.

2.2 Games

Game arena. A *game arena* \mathcal{G} is a tuple (V, V_1, V_2, E) where (V, E) is a finite graph, (V_1, V_2) is a partition of V into positions of Player 1 and Player 2, respectively. We consider (for technical convenience) that for every position $v \in V$ there is at least one outgoing edge. Given an arena \mathcal{G} and a set $U \subseteq V$ of positions we denote by $\mathcal{G} \upharpoonright U$

the arena restricted to the graph induced by U , i.e., the set of positions is U , and the set of edges $E \cap (U \times U)$. Moreover, if every position in U has an outgoing edge in U , then $\mathcal{G} \upharpoonright U$ is also a game arena.

Game plays. A game on an arena \mathcal{G} is played as follows: a token is placed at a starting position, and whenever the token is at a Player-1 position, then Player 1 chooses an outgoing edge to move the token, and when the token is at a Player-2 position, then Player 2 does likewise. As a consequence we obtain an infinite sequence of positions, which are called plays, and strategies are recipes to extend finite prefix of plays (i.e., the recipes to describe how to move tokens). We formally define them below.

Strategies and plays. Given a game arena \mathcal{G} , a function $\sigma_1 : V^* \cdot V_1 \mapsto V$ (resp., $\sigma_2 : V^* \cdot V_2 \mapsto V$) is a *strategy* for Player 1 (resp., Player 2) on \mathcal{G} iff $\sigma_j(v_0 v_1 \dots v_k) = v$ implies $(v_k, v) \in E$. In other words, given a finite sequence of positions that ends at a Player-1 position (representing the history of interactions), a strategy for Player 1 chooses the next position respecting the edge relationship (to move the token). We denote the set of all strategies for Player 1 (resp., Player 2) on \mathcal{G} by $\mathcal{S}_1[\mathcal{G}]$ (resp., $\mathcal{S}_2[\mathcal{G}]$). A strategy σ_i is *memoryless* iff for all $w, w' \in V^*, v \in V_i$ we have $wv_i, w'v_i \in \text{dom}(\sigma_i)$ implies $\sigma_i(wv_i) = \sigma_i(w'v_i)$. Informally, a memoryless strategy does not depend on the history, but only on the current position. We denote the set of all memoryless strategies for Player 1 (resp., Player 2) on \mathcal{G} by $\mathcal{S}_1[\mathcal{G}, M]$ (resp., $\mathcal{S}_2[\mathcal{G}, M]$). A pair of strategies σ_1, σ_2 on \mathcal{G} , along with a starting position v , defines a *play* $\pi(\sigma_1, \sigma_2, v)$, which is a word over V . The play $\pi(\sigma_1, \sigma_2, v) = v_1 v_2 \dots$ is defined inductively as follows: (a) $v_1 = v$; (b) $v_{i+1} = \sigma_1(v_1 \dots v_i)$ if $v_i \in V_1$; and (c) $v_{i+1} = \sigma_2(v_1 \dots v_i)$ if $v_i \in V_2$. We define $\Pi(\mathcal{G})$ as the set of all plays on \mathcal{G} . Since every position has at least one outgoing edge, every play is indeed infinite.

We consider three types of objectives: Boolean, quantitative and implication. Implication objectives combine Boolean and quantitative objectives and are presented in Section 3.

Boolean objectives. A Boolean objective is a function $\Phi : \Pi(\mathcal{G}) \mapsto \{0, 1\}$. We consider three types of Boolean objectives: tautology, Büchi and parity. Tautology are those objectives Φ_T that for every play return 1. Büchi objectives Φ_B are defined by a subset F of the positions of the arena. Then, $\Phi_B(\pi) = 1$ iff some position from F occurs infinitely often in π . Finally, parity objectives are defined by labellings p of the positions in the arena with natural numbers. The parity objective p is satisfied by π , i.e., $\Phi_P(\pi) = 1$, iff $\liminf \{p(v_0), p(v_1), \dots\}$ is even, i.e., among numbers that appear infinitely often in $p(v_0), p(v_1), \dots$ the minimal one is even.

Winning strategies and winning sets. A strategy σ_1 (resp., σ_2) is *winning* for Player 1 (resp., Player 2) from a position v iff for all strategies σ_2 for Player 2 (resp., all strategies σ_1 for Player 1), the play π defined by σ_1, σ_2 given v satisfies $\Phi(\pi) = 1$ (resp., $\Phi(\pi) = 0$). For a Boolean objective Φ , and $i \in \{1, 2\}$, the winning set for Player i for the objective, denoted by $\text{Win}_i(\Phi)$, is the set of positions v such that there exists a winning strategy for Player i from v . For all Boolean objectives defined above, the winning sets form a partition, i.e., if Φ is the winning objective and $\bar{\Phi}$ its complement, then $\text{Win}_1(\Phi) \cap \text{Win}_2(\bar{\Phi}) = \emptyset$; and $\text{Win}_1(\Phi) \cup \text{Win}_2(\bar{\Phi}) = V$; and there exist memoryless winning strategies for the players from their respective winning set [19, 20].

Quantitative objectives. A *quantitative objective* in general is a Borel measurable function $f : \Pi(\mathcal{G}) \mapsto \mathbb{R} \cup \{-\infty, \infty\}$. Unlike in games with Boolean objectives, quantitative games do not have the winner. Instead, Player 1 (called also Minimizer) plays in a way to construct plays π of a possibly small value $f(\pi)$, whereas Player 2 (called also Maximizer) attempts to maximize $f(\pi)$. The minimal value of the game which Player 1 can ensure (called the lower value) is defined as $\text{val}(f, v) = \inf_{\sigma_1 \in \mathcal{S}_1[\mathcal{G}]} \sup_{\sigma_2 \in \mathcal{S}_2[\mathcal{G}]} f(\pi(\sigma_1, \sigma_2, v))$. Player 2 on the other hand can ensure that the value of the game is at least the upper value, denoted as $\bar{\text{val}}(f, v) = \sup_{\sigma_2 \in \mathcal{S}_2[\mathcal{G}]} \inf_{\sigma_1 \in \mathcal{S}_1[\mathcal{G}]} f(\pi(\sigma_1, \sigma_2, v))$. By Borel determinacy [19], the upper and lower values coincide with respect to f , and is called the value of the game, denoted by $\text{val}(f, v)$. Given a quantitative objective f , we can consider a Boolean objective by imposing a threshold, i.e., given a threshold ν we consider the set of winning plays to be $\{\pi \in \Pi(\mathcal{G}) : f(\pi) \leq \nu\}$, all plays π whose value does not exceed ν .

Optimal and ϵ -optimal strategies. Consider a game arena \mathcal{G} with a quantitative objective f . For a real $\epsilon \geq 0$, a strategy σ^o for Player 1 (resp., Player 2) is ϵ -*optimal* for a position v iff $\sup_{\sigma_2 \in \mathcal{S}_2[\mathcal{G}]} f(\pi(\sigma^o, \sigma_2, v)) - \text{val}(f, v) \leq \epsilon$ (resp., $\text{val}(f, v) - \inf_{\sigma_1 \in \mathcal{S}_1[\mathcal{G}]} f(\pi(\sigma_1, \sigma^o, v)) \geq \epsilon$). A strategy σ is *optimal* iff it is 0-optimal.

Basic quantitative objectives. We consider the following quantitative objectives defined by a labeling $\text{wt} : E \mapsto \mathbb{Z}$ of moves E on \mathcal{G} with integers. Given an arena \mathcal{G} , a labeling wt and a play $\pi = v_0 v_1 \dots$ on \mathcal{G} we define $\text{wt}(\pi)$ as a sequence of integers $\text{wt}(v_0, v_1), \text{wt}(v_1, v_2) \dots$. Given a value function $f \in \text{VALFUNC}$, we consider the quantitative objectives f^{wt} defined as $f^{\text{wt}}(\pi) = f(\text{wt}(\pi))$. For example, limit-average infimum objective is defined on plays by $\text{LIMAVGINF}^{\text{wt}}(\pi) = \text{LIMAVGINF}(\text{wt}(\pi))$.

Determinacy and memoryless optimal strategies. For every $f \in \text{VALFUNC}$ and every labeling wt , the function f^{wt} is Borel measurable, and hence due to Borel determinacy theorem [19], the upper and lower values coincide, and we have a value of the game for every initial position. Moreover, all $f \in \text{VALFUNC}$ admit optimal strategies that are memoryless [21, 14, 15].

Composed quantitative objectives. Consider an arena \mathcal{G} and labellings of moves wt_1 and wt_2 . A *composed quantitative objective* $f + g$ is a quantitative objective defined as follows: $(f + g)(\pi) = f(\text{wt}_1(\pi)) + g(\text{wt}_2(\pi))$.

3 Implication Games

In this section we introduce implication games, and in the following section we argue how simulation of weighted automata with Büchi acceptance conditions can be reduced to implication games, and finally we will present solution for implication games.

Implication games. We define *implication games* that combine Boolean and quantitative objectives. An *implication objective* is a triple (Φ_1, Φ_2, f) , where Φ_1, Φ_2 are Boolean objectives and f is a quantitative objective. Intuitively, Player 1 has to satisfy the Boolean objective Φ_1 and either ensure that Φ_2 is violated or play to minimize the value of $f(\pi)$. Player 2 has to either violate Φ_1 or play to both satisfy Φ_2 and maximize the value of $f(\pi)$. Formally, we define f' as follows: for every $\pi \in \Pi(\mathcal{G})$,

$$f'(\pi) = \begin{cases} f(\pi) & \text{if } \Phi_1(\pi) = \Phi_2(\pi) = 1, \\ \infty & \text{if } \Phi_1(\pi) = 0, \\ -\infty & \text{if } \Phi_1(\pi) = 1 \text{ and } \Phi_2(\pi) = 0. \end{cases}$$

The value of an implication game with the objective (Φ_1, Φ_2, f) is the value of a quantitative game with the objective f' , whenever the latter exists.

Classes of implication objectives. Let $\mathcal{C}_1, \mathcal{C}_2$ be classes of Boolean objectives, and let \mathcal{F} be a class of quantitative objectives. We define the class of implication games with $(\mathcal{C}_1, \mathcal{C}_2, \mathcal{F})$ as the class of all implication games with objectives (Φ_1, Φ_2, f) , where $\Phi_1 \in \mathcal{C}_1, \Phi_2 \in \mathcal{C}_2$ and $f \in \mathcal{F}$.

Values of implication games. The value of the implication game with the objective (Φ_1, Φ_2, f) is the value of the quantitative game with the objective f' . One can construct functions f' such that the quantitative game with the objective does not have the value (if f is not measurable). In the following we show that every reasonable choice of Φ_1, Φ_2 and f leads to a game with the properly defined value. A function $f : \Pi(\mathcal{G}) \mapsto \mathbb{R} \cup \{-\infty, \infty\}$ is Borel iff for every $a \in \mathbb{R}$, the counter-image $f^{-1}[(a, \infty]]$ is Borel in $\Pi(\mathcal{G})$.


Proposition 1. *Let \mathcal{G} be an arena, let Φ_1, Φ_2 be Boolean objectives and let f be a quantitative objective on \mathcal{G} . If Φ_1, Φ_2, f are Borel, then the implication game with the objective (Φ_1, Φ_2, f) has a value, i.e., for every v we have*

$$\inf_{\sigma_1 \in \mathcal{S}_1[\mathcal{G}]} \sup_{\sigma_2 \in \mathcal{S}_2[\mathcal{G}]} f'(\pi(\sigma_1, \sigma_2, v)) = \sup_{\sigma_2 \in \mathcal{S}_2[\mathcal{G}]} \inf_{\sigma_1 \in \mathcal{S}_1[\mathcal{G}]} f'(\pi(\sigma_1, \sigma_2, v)).$$

Proof. If Φ_1, Φ_2, f are Borel, the objective f' is Borel as well. Therefore, for every threshold t , the set of plays such that the value of the play does not exceed t is Borel. By Borel determinacy theorem [19], we have for every threshold t , either Player 1 has a strategy to keep the value of all plays below t against all strategies of Player 2; or Player 2 can enforce that the value of the play is at least t against all strategies of Player 1. Observe that there exists the least t_0 such that Player 2 can enforce that the value of the play is at least t_0 against all strategies of Player 1. On the other hand, for every $t' < t_0$, Player 1 has a strategy to keep the value of all plays below t' against all strategies of Player 2. Thus, t_0 is the value of the game. \square

In classical mean-payoff games [21] (as well as in mean-payoff games with conjunction with a parity objective [22]), the value functions for LIMAVGINF and LIMAVGSUP objectives coincide for every position. In contrast, in the following example we show that for implication games, the value functions for LIMAVGINF and LIMAVGSUP objectives do not coincide in general.

Example 2. *Consider an arena \mathcal{G} depicted below:*




Player 1 owns the position A (depicted as circles) and Player 2 owns B (depicted as squares). Consider an implication game played on \mathcal{G} with objective $(\Phi_1, \Phi_2, \text{LIMAVGINF})$, where Φ_1 states that B is visited infinitely

often, Φ_2 states that A is visited infinitely often, and LIMAVGINF is taken w.r.t. the weight presented in the figure. Intuitively, Φ_1, Φ_2 force Players 1 and 2 to leave their positions, which are beneficial for them in the quantitative aspect. Observe that Player 1 can satisfy Φ_1 and force the game value to be -1 . Indeed, Player 1 maintains a counter i which is increased every time Player 2 moves from B to A . Every time Player 2 moves to the position A , Player 1 stays in A until the partial average of weights decreases below $-1 + \frac{1}{i}$. Then, he moves to the position B . Observe that such a strategy satisfies Φ_1 as Player 1 moves to B after finite time spent in A . Moreover, if Player 2 leaves B infinitely often, then for every $\epsilon > 0$, infinitely often the partial average of the weights are below $-1 + \epsilon$. Therefore, the value of the game is -1 .

However, a similar implication game with an objective $(\Phi_1, \Phi_2, \text{LIMAVGSUP})$ has value 1, since Player 2 can satisfy Φ_2 and as long as Player 1 leaves A infinitely often, Player 2 can force the limit supremum to reach 1 in a similar way to the previous strategy of Player 1 for $(\Phi_1, \Phi_2, \text{LIMAVGINF})$. Player 2 maintains a counter i which is increased every time Player 1 moves from A to B . Every time Player 1 moves to the position B , Player 2 stays in B until the partial average of weights increases above $1 - \frac{1}{i}$. As in the previous case, the value of the game is 1.

Observe that implication games are different than games with Boolean objectives. More precisely, the following example shows that Player 1 can ensure that every play has a negative value, however, the value of the game is 0.

Example 3. Consider an implication game $(\Phi_1, \Phi_2, \text{DISC}_\lambda)$ depicted in the figure below where all positions are owned by Player 2. We define the condition Φ_1 to be satisfied for all plays and Φ_2 to state that B is visited infinitely often.



Player 2 has to eventually leave A and move to B to satisfy Φ_2 . Clearly, every play that eventually moves to B has a negative value. However, Player 2 can stay in A arbitrarily long and achieve values arbitrarily close to 0. Therefore, the supremum (which is the value of the game) is 0.

4 Quantitative Liveness Simulation Games

Language inclusion is a central problem in automata-based verification. Typically, the behavior of the model is described by one automaton, while the other automaton describes traces allowed by the specification. In such a case, the model satisfies the specification iff behaviors of the model are included in the behaviors allowed by the specification. Unfortunately, the quantitative inclusion problem for LIMAVGINF -automata (resp., LIMAVGSUP -automata) is undecidable. The decidability of the quantitative inclusion problem for DISC_λ -automata is open. Still, we can under-approximate (quantitative) inclusion by a more restrictive notion of (quantitative) simulation, which implies inclusion. We define quantitative simulation in terms of games, i.e., we say that an automaton \mathcal{A}_1 is quantitatively simulated by \mathcal{A}_2 iff the player called Simulator wins the following *Quantitative Fair Simulation Game (QFSG)* on $\mathcal{A}_1, \mathcal{A}_2$ against Challenger.

Quantitative Fair Simulation Games (QFSGs). We define a *Quantitative Fair Simulation Game (QFSG)* as follows. Let $\mathcal{A}_1, \mathcal{A}_2$ be weighted automata over the alphabet Σ . For $i \in \{1, 2\}$, let $Q_i, \delta_i, q_{0,i}$ be respectively the set of states, the transition relation and the initial state of \mathcal{A}_i . We define the arena of a QFSG on $\mathcal{A}_1, \mathcal{A}_2$ as $\mathcal{G} = (V, V_1, V_2, E)$, where:

1. $V = V_1 \cup V_2$, where $V_1 = Q_1 \times Q_2, V_2 = Q_1 \times \Sigma \times Q_2$,
2. $E = E_1 \cup E_2$, where $E_1 = \{ \langle (q_1, q_2), (q'_1, a, q_2) \rangle : (q_1, a, q'_1) \in \delta_1, q_2 \in Q_2 \}$ and $E_2 = \{ \langle (q'_1, a, q_2), (q'_1, q'_2) \rangle : q'_1 \in Q_1, (q_2, a, q'_2) \in \delta_2 \}$

Intuitively, $Q_1 \times Q_2$ are Player-1 positions, where Player 1 chooses a letter from the alphabet and a transition of \mathcal{A}_1 ; and then at positions $Q_1 \times \Sigma \times Q_2$ which belong to Player 2 a response of a transition in \mathcal{A}_2 given the chosen letter from Player 1. We will consider the starting position as $v_0 = \langle q_{0,1}, q_{0,2} \rangle$.

We call Player 1 *Challenger* and Player 2 *Simulator*. Challenger plays on \mathcal{A}_1 and picks letters and transitions of \mathcal{A}_1 , while *Simulator* plays on \mathcal{A}_2 and attempts to match transitions of \mathcal{A}_1 with \mathcal{A}_2 transitions. Specifically, at a position (q_1, q_2) Challenger picks a letter a and a transition (q_1, a, q'_1) of \mathcal{A}_1 , which determine the move $\langle (q_1, q_2), (q'_1, a, q_2) \rangle$. Next, at the position (q'_1, a, q_2) Simulator responds with a transition (q_2, a, q'_2)

of \mathcal{A}_2 labeled with the same letter a and moves to a position of Challenger (q'_1, q'_2) . The constructed play $(q_1^0, q_2^0), (q_1^1, a, q_2^0), (q_1^1, q_2^1), \dots$ yields two runs: π_1 of \mathcal{A}_1 and π_2 of \mathcal{A}_2 . Challenger wins the QFSG on automata $\mathcal{A}_1, \mathcal{A}_2$ iff he has a strategy σ_1 and there exists $\epsilon > 0$ such that for every strategy σ_2 of Simulator, the constructed runs satisfy the following: (1) π_1 is an accepting run of \mathcal{A}_1 , and (2) either (a) π_2 is not accepted by \mathcal{A}_2 or (b) the value of the run π_2 in \mathcal{A}_2 is greater than the value of the run π_1 in \mathcal{A}_1 plus ϵ . Simulator wins the QFSG on automata $\mathcal{A}_1, \mathcal{A}_2$ when Challenger does not win the game.

QFSGs vs. automata inclusion. Consider a modification of QFSGs in which Challenger does not distinguish moves of Simulator. The modified game is called a *partial-information* game. In such a game, strategies of Challenger are simply runs of \mathcal{A}_1 , while strategies of Simulator are as in QFSGs. Observe that Challenger has a winning strategy in the modified game precisely when the answer to the quantitative inclusion problem of \mathcal{A}_1 in \mathcal{A}_2 is “no”, i.e., the modified game is equivalent to the quantitative inclusion problem. QFSGs result from giving more power to Challenger, who can observe moves of Simulator, therefore if Simulator wins QFSG then he wins the modified game on the same automata. It follows that QFSG under-approximates the inclusion problem, i.e., if Simulator wins the QFSG, then the answer to the inclusion problem is yes, however, the converse need not hold. We prove this observation formally in the following result:

Theorem 4. *Let $f \in \text{VALFUNC}$ and let $\mathcal{A}_1, \mathcal{A}_2$ be f -automata with Büchi acceptance conditions, and labeling functions wt_1 and wt_2 . Let $q_{0,1}$ (resp., $q_{0,2}$) be the initial state of \mathcal{A}_1 (resp., \mathcal{A}_2). If Challenger does not win the QFSG on automata $\mathcal{A}_1, \mathcal{A}_2$ starting from position $v_0 = \langle q_{0,1}, q_{0,2} \rangle$, then the answer to the quantitative inclusion problem of \mathcal{A}_1 in \mathcal{A}_2 is “yes”.*

Proof. Assume that Challenger does not have a winning strategy from the initial position $\langle q_{0,1}, q_{0,2} \rangle$. Let π_1 be an accepting run of \mathcal{A}_1 . Consider a strategy for Challenger that produces the run π_1 regardless of the actions of Simulator. This strategy is not winning, therefore for every $\epsilon > 0$ there is a strategy of Simulator that produces a run π_2 , which is accepting for \mathcal{A}_2 and $f^{\text{wt}_2}(\pi_2) \leq f^{\text{wt}_1}(\pi_1) + \epsilon$. It follows that for every accepting run π_1 of \mathcal{A}_1 and every $\epsilon > 0$ there is an accepting run of \mathcal{A}_2 such that $f^{\text{wt}_2}(\pi_2) \leq f^{\text{wt}_1}(\pi_1) + \epsilon$. Then, for every word w we have $\inf_{\pi_2 \in \text{Acc}_2(w)} f^{\text{wt}_2}(\pi_2) \leq \inf_{\pi_1 \in \text{Acc}_1(w)} f^{\text{wt}_1}(\pi_1) + \epsilon$, i.e., the answer to the quantitative inclusion problem of \mathcal{A}_1 in \mathcal{A}_2 is “yes”. \square

QFSGs to implication games. The QFSG problem on \mathcal{A}_1 and \mathcal{A}_2 can be reduced to implication games where the first (resp., the second) Boolean objective for the implication games represent the Boolean acceptance condition for \mathcal{A}_1 (resp., \mathcal{A}_2), and the quantitative objective represents the composed quantitative objective. Hence in the following section we present solution of implication games where the Boolean acceptance conditions are Büchi conditions, and the quantitative objectives are sum of mean-payoff objectives or discounted-sum objectives.

5 Mean-Payoff Implication Games

In this section we consider implication games, where the quantitative objective is the sum of two mean-payoff objectives, and the Boolean objectives are Büchi objectives. Our solution will use the notion of attractors that we define below.

Attractors. For a set $U \subseteq V$ we denote by $\text{Attr}_i(U)$ the attractor set of Player i , namely, the set of positions from which Player i can force reachability to a position in U . It is well known that an attractor is computable in linear time [10, 11], and the corresponding winning strategy for the reachability objective is memoryless (which is referred as an *attractor strategy*). An attractor strategy also ensures that the target set U is reached within $|V|$ steps. It is also well known that if \mathcal{G} is a game arena, then for a set of positions $U \subseteq V$ and for Player $i = 1, 2$, the arena $\mathcal{G} \upharpoonright (V \setminus \text{Attr}_i(U))$ is a game arena, i.e., it has at least one move from every position, and all positions of Player i in $V \setminus \text{Attr}_i(U)$ do not have a move to $\text{Attr}_i(U)$.

5.1 Mean-payoff parity games

Mean-payoff parity games, which are a subclass of implication games have been studied in [22]. Formally, mean-payoff parity games are implication games $(\Phi_1, \Phi_2, \text{LIMAVGINF})$ where Φ_1 is a parity objective and Φ_2 is a tautology objective (i.e., it is satisfied for all plays). Intuitively, in a mean-payoff parity game Player 1 has to ensure that the parity condition is satisfied and the mean-payoff objective is above a certain threshold. The objective of Player 2 is the converse. The following results have been showed in [22, 8].


Theorem 5 ([22, 8]). *The following assertions hold: (i) Mean-payoff parity games can be solved in $O(n^d \cdot (m + \text{MPALGO} + \text{PARITYALGO}))$ time, where MPALGO and PARITYALGO denote run-times of algorithms to solve mean-payoff games and parity games, respectively, and n, m , and d represent the number of positions, the number*

of edges, and the number of priorities of the parity objective, respectively. (ii) In every mean-payoff parity game, Player 1 has an infinite-memory optimal strategy, and in general, optimal strategies for Player 1 require infinite memory. (iii) In every mean-payoff parity game, Player 2 has a memoryless optimal strategy. (iv) For every $\epsilon > 0$, there exists a finite-memory ϵ -optimal strategy for Player 1, and ϵ -optimal strategies, for $\epsilon > 0$, for Player 1 in general require memory. However, if there is a finite-memory optimal strategy for Player 1, then there is a memoryless optimal strategy.

5.2 Complexity of winning strategies

In mean-payoff parity games while Player 1 requires infinite memory for optimal strategies, there always exist memoryless optimal strategies for Player 2. In contrast, in implication games with (Büchi, Büchi, LIMAVGINF₁ + LIMAVGINF₂) objectives, we show (in the following two examples) that the optimal strategies, both for Player 1 and Player 2 in general require infinite memory. Similar result also holds if LIMAVGINF is replaced by LIMAVGSUP which we omit for succinctness.

Example 6 (Optimal strategies for Player 2 require infinite memory). Consider an arena \mathcal{G} depicted below:




Player 2 owns both positions. Consider an implication game played on \mathcal{G} with an objective $(\Phi_1, \Phi_2, \text{LIMAVGINF}_1 + \text{LIMAVGINF}_2)$, where Φ_1 is a tautology, Φ_2 states that A is visited infinitely often, and LIMAVGINF_i (for $i = 1, 2$) is taken w.r.t. the weights presented in the figure. Player 2 can achieve value 2 with the following strategy:

- For $m = 1, 2, \dots$:
 - Follow the transition from A to B.
 - Traverse the self-loop of B for m times.
 - Follow the transition from B to A.

However, any finite-memory strategy of Player 2 is an ultimately periodic path $\pi = \pi_0(\pi_1)^\omega$, where π_0 and π_1 are finite paths, and hence $\text{LIMAVGINF}_1(\pi) + \text{LIMAVGINF}_2(\pi) = \text{Avg}_1(\pi_1) + \text{Avg}_2(\pi_1)$. If π_1 contains the position A, then $\text{Avg}_1(\pi_1), \text{Avg}_2(\pi_1) < 1$, and its value is strictly smaller than 2. Otherwise, Φ_2 is not satisfied, and the value of the play is $-\infty$. Hence, with infinite-memory strategy Player 2 can achieve a value of 2, while any finite-memory strategy gives a value strictly smaller than 2.

Example 7 (Optimal strategies for Player 1 require infinite memory). Consider an arena \mathcal{G} depicted below:



Player 1 owns both positions. Consider an implication game played on \mathcal{G} with an objective $(\Phi_1, \Phi_2, \text{LIMAVGINF}_1 + \text{LIMAVGINF}_2)$, where both Φ_1 and Φ_2 are tautologies, and LIMAVGINF_i (for $i = 1, 2$) is taken w.r.t. the weights presented in the figure. Player 1 can achieve value -2 with the following strategy:

- For $m = 1, 2, \dots$:
 - Traverse the self-loop of A until the average weight of the first dimension is smaller than $-1 + \frac{1}{m}$.
 - Follow the transition from A to B.
 - Traverse the self-loop of B until the average weight of the second dimension is smaller than $-1 + \frac{1}{m}$.
 - Follow the transition from B to A.

With this strategy Player 1 achieves $\text{LIMAVGINF}_1 = \text{LIMAVGINF}_2 = -1$. Hence, the value of the play is -2 . However, any finite-memory strategy of Player 1 is an ultimately periodic path $\pi = \pi_0(\pi_1)^\omega$, with $\text{LIMAVGINF}_1(\pi) + \text{LIMAVGINF}_2(\pi) = \text{Avg}_1(\pi_1) + \text{Avg}_2(\pi_1) = 0$. Hence, with infinite-memory strategy Player 1 can achieve a value of -2 , while any finite-memory strategy gives a value 0. Hence optimal strategies (even ϵ -optimal strategies, for $\epsilon > 0$) for Player 1 require infinite memory in general.

In this section we have established that optimal strategies for both players require infinite-memory in general. In games with Boolean combination of multiple mean-payoff objectives, infinite-memory strategies are required for both player for optimality, and the problem is undecidable [23]. In contrast, we show in the following subsections that in implication games with mean-payoff objectives and Büchi acceptance conditions, though both players require infinite memory for optimality, the decision problem is decidable.

5.3 Implication games with objectives (Büchi, Büchi, LIMAVGINF + LIMAVGINF)

We consider implication games defined by (Büchi, Büchi, $\text{LIMAVGINF}_1 + \text{LIMAVGINF}_2$) objective. More formally, we have weight functions wt_1 and wt_2 , and $\text{LIMAVGINF}_1 = \text{LIMAVGINF}^{\text{wt}_1}$ and $\text{LIMAVGINF}_2 = \text{LIMAVGINF}^{\text{wt}_2}$. We start with a few notations.

Notations. Let us consider a threshold ν , and we will use the following notations:

$$\Phi_{\text{sum}}^{\text{inf}}(\nu) = \{\pi : \text{LIMAVGINF}_1(\pi) + \text{LIMAVGINF}_2(\pi) \geq \nu\};$$

i.e., the objective requires the sum of the mean-payoff objectives is at least ν ; and for $i \in \{1, 2\}$,

$$\Phi_i^{\text{inf}}(\nu) = \{\pi : \text{LIMAVGINF}_i(\pi) \geq \nu\};$$


i.e., the objective requires that the individual mean-payoff in dimension i is at least ν . For thresholds x and y , games with objectives that are conjunction $\Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y)$ of two individual mean-payoff objectives are called two-dimensional mean-payoff games (and have been studied in [24]). We also consider the existential version of two-dimensional mean-payoff games that has not been studied before. For a threshold ν , the existential version of two-dimensional mean-payoff objectives is defined as follows:

$$\Phi_{\text{exi}}^{\text{inf}}(\nu) = \{\pi : \exists x, y \in \mathbb{R}. x + y \geq \nu \wedge \pi \in \Phi_1^{\text{inf}}(x) \wedge \pi \in \Phi_2^{\text{inf}}(y)\};$$

i.e., it quantifies existentially over the individual thresholds such that their sum exceeds ν and the individual objectives can be satisfied.

Single-dimensional vs. two-dimensional objectives. Before we proceed with the solution, we discuss the difference between single-dimensional vs. two-dimensional objectives. More precisely, let us consider two weight functions wt_1, wt_2 . We define $\text{wt}_1 + \text{wt}_2$ as the weight function whose weight, for every move, is the sum of weights of wt_1 and wt_2 on that move. In the following example we show that the objectives $B_1 \wedge (\text{co}B_2 \vee \Phi^*(\nu))$, where $\Phi^*(\nu) = \{\pi : \text{LIMAVGINF}^{\text{wt}_1 + \text{wt}_2}(\pi) \geq \nu\}$, and $B_1 \wedge (\text{co}B_2 \vee \Phi_{\text{sum}}^{\text{inf}}(\nu))$ are different. In fact, we show that even the objectives $\Phi^*(\nu)$ and $\Phi_{\text{sum}}^{\text{inf}}(\nu)$ are different.

Example 8. Consider arenas $\mathcal{G}_1, \mathcal{G}_2$ depicted below:



Arenas \mathcal{G}_1 and \mathcal{G}_2 have the same underlying games structures and differ only in labels of moves. All positions are owned by Player 2 and moves in \mathcal{G}_1 are labeled by two weight functions wt_1, wt_2 , while moves in \mathcal{G}_2 are labeled by a single weight $\text{wt}_1 + \text{wt}_2$. The value of any play in \mathcal{G}_2 is 0 as this is the only weight in the arena. However, for every $\nu > -2$, Player 2 can violate the objective $\Phi_{\text{sum}}^{\text{inf}}(\nu)$ in \mathcal{G}_1 . Indeed, the strategy σ_2 achieving $\text{LIMAVGINF}_1(\pi) = \text{LIMAVGINF}_2(\pi) = -1$ is as follows: Let i be a counter initially set to 1.

Step 1. Go to position B and stay there until the value of the partial average in the first dimension is less than $-1 + \frac{1}{i}$.

Step 2. Alternate between A and B until the value of the partial average in the second dimension is less than $-1 + \frac{1}{i}$.

Step 3. Increment i and go to Step 1.

Consider the play π generated by the strategy σ_2 . For every $\epsilon > 0$, the partial average of π in the first (resp., second) dimension is less than $-1 + \epsilon$ infinitely often. On the other hand $\text{LIMAVGINF}_1(\pi), \text{LIMAVGINF}_2(\pi) \geq -1$. Therefore, $\text{LIMAVGINF}_1(\pi) = \text{LIMAVGINF}_2(\pi) = -1$ and for every $\nu > -2$ the play π violates $\Phi_{sum}^{\text{inf}}(\nu)$.

The above example shows that the $\Phi_{sum}^{\text{inf}}(\nu)$ objective cannot be treated simply as a one-dimensional mean-payoff objective. Before we proceed to our solution we present a general fact about prefix-independent objectives, and properties of winning strategies for them.

Tail (prefix-independent) objectives. A Boolean objective φ is a *tail (prefix-independent) objective* if for any finite play π_0 and any infinite play π it holds that $\varphi(\pi) = 1$ if and only if $\varphi(\pi_0 \cdot \pi) = 1$. Observe that Büchi and co-Büchi objectives as well as $\Phi_{sum}^{\text{inf}}(\nu)$, $\Phi_i^{\text{inf}}(\nu)$, and $\Phi_{exi}^{\text{inf}}(\nu)$ are all tail objectives.

Finite-history independent property. Consider a tail objective φ and a winning strategy σ_1 for the objective φ . Consider any finite history h that ends in $\text{Win}_1(\varphi)$, and the strategy that ignores that history h and plays as σ_1 , i.e., for histories $h \cdot h'$ it plays as $\sigma_1(h')$. Since φ is a tail objective, the resulting strategy ensures that if h is executed, then against all strategies of the opponent the resulting play is still winning. In the sequel we will refer to this property as the finite-history independent winning property.

The basic computation problem. For a given threshold ν , the implication game (Büchi, Büchi, $\text{LIMAVGINF}_1 + \text{LIMAVGINF}_2$) is equivalent to a game with the Boolean winning condition $\varphi(\nu) = B_1 \wedge (\text{co}B_2 \vee \Phi_{sum}^{\text{inf}}(\nu))$, where B_1 stands for a Büchi condition over the set B_1 and $\text{co}B_2$ is a co-Büchi condition over the set B_2 . Hence, in this subsection we study a game where Player 1 wishes to satisfy $\varphi(\nu)$, and show how to compute the maximal threshold ν such that Player 1 is the winner with respect to the condition $\varphi(\nu)$ (specifically, we show that such maximal threshold exists and how to compute it).

Solution overview. Intuitively, we show that to find the winner of a game with objective $B_1 \wedge (\text{co}B_2 \vee \Phi_{sum}^{\text{inf}}(\nu))$, it is enough to consider two games independently, one with condition $B_1 \wedge \text{co}B_2$, and the other with condition $B_1 \wedge \Phi_{sum}^{\text{inf}}(\nu)$. If Player 1 wins for one of the objectives, then clearly he wins for the original objective $\varphi(\nu)$. Conversely, we prove that if Player 1 loses in both the games from everywhere, then he also loses for the objective $\varphi(\nu)$ from everywhere. Games with conditions $B_1 \wedge \text{co}B_2$ are well studied [12]. Hence, most of this section is devoted to analyze the winning condition $B_1 \wedge \Phi_{sum}^{\text{inf}}(\nu)$. The analysis is achieved as follows:

1. First, we consider the existential version of two-dimensional mean-payoff games $\Phi_{exi}^{\text{inf}}(\nu)$, and show how to solve them (in Subsection 5.3.1). The solution is obtained by extending the techniques to solve two-dimensional mean-payoff games that were studied in [18].
2. Then we show how to use the solution of the existential version of two-dimensional mean-payoff games to solve $B_1 \wedge \Phi_{sum}^{\text{inf}}(\nu)$ objectives (in Subsection 5.3.2).
3. Finally (in Subsection 5.3.3), we prove that it is enough to consider the two above winning objectives (namely, $B_1 \wedge \text{co}B_2$ and $B_1 \wedge \Phi_{sum}^{\text{inf}}(\nu)$) and present an algorithm that solves the original condition (and computes the maximal value that can be obtained for a given initial position).

5.3.1 Solving games with existential two-dimensional objectives

In this section we study games with the objective $\Phi_{exi}^{\text{inf}}(\nu)$ (i.e., existential version of two-dimensional mean-payoff games). We fix a starting position v_0 for the rest of the section. Given an arena \mathcal{G} , we define $\text{PAIRS}[\mathcal{G}]$ as $\text{PAIRS}[\mathcal{G}] = \{(x, y) : \text{Player 1 has a winning strategy from } v_0 \text{ for } \Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y) \text{ on } \mathcal{G}\}$. In the following lemma we establish that if there exist elements in $\text{PAIRS}[\mathcal{G}]$ with sum at least ν , then there exist such witnesses in the set that have $O(\log(nW))$ description.

Lemma 9. *Let \mathcal{G} be a game arena with n positions and weights from $[-W, W]$. For every threshold ν , if $\text{PAIRS}[\mathcal{G}]$ contains $(x, y) \in \mathbb{R}^2$ with $x + y \geq \nu$, then there exist rationals $(q, p) \in \text{PAIRS}[\mathcal{G}]$ with $p + q \geq \nu$ such that the numerators and denominators of p, q belong to $\{-(nW)^4, \dots, (nW)^4\}$.*

Proof. Since we consider objective $\Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y)$ for Player 1, which is a two-dimensional mean-payoff objective, it follows from [24] that the analysis can be restricted to memoryless counter strategies for Player 2. Given $\sigma_2 \in \mathcal{S}_2[\mathcal{G}, M]$, a memoryless strategy for Player 2 on \mathcal{G} , we denote by \mathcal{G}^{σ_2} the graph resulting from fixing Player-2 choices according to σ_2 . Without loss of generality we consider that for any $\sigma_2 \in \mathcal{S}_2[\mathcal{G}, M]$, (i) the graph \mathcal{G}^{σ_2} contains a cycle with average weight $(-W, -W)$ (note that we can trivially add in a Player-1 position such

a cycle which Player 1 will never choose); and (ii) we further consider that \mathcal{G}^{σ_2} has only one strongly connected component (SCC) and later present the solution for the general case. We denote by \mathbb{C}_{σ_2} the set of simple cycles in the (only) SCC of \mathcal{G}^{σ_2} . For a finite path $\pi = e_1 e_2 \dots e_n$ and $i = 1, 2$ we denote $\text{wt}_i(\pi) = \sum_{j=1}^n \text{wt}_i(e_j)$. Then, we define $\text{Avg}(\pi) = (\frac{\text{wt}_1(\pi)}{n}, \frac{\text{wt}_2(\pi)}{n})$.

By [24], Player 1 can win $\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$ if and only if

$$(q, p) \in \bigcap_{\sigma_2 \in \mathcal{S}_2[\mathcal{G}, M]} \text{CONV}\{\text{Avg}(c) : c \in \mathbb{C}_{\sigma_2}\} \quad (\star)$$

(where $\text{CONV}(X)$ contains all the convex combinations of elements from X). The problem of maximizing the sum $p + q$ subject to (\star) is a linear programming problem, and thus if the constraints of (\star) are feasible and the solution is bounded, then a (possible) solution is a vertex of the convex polygon $\bigcap_{\sigma_2 \in \mathcal{S}_2[\mathcal{G}, M]} \text{CONV}\{\text{Avg}(c) : c \in \mathbb{C}_{\sigma_2}\}$. The constraints are feasible, since Player 1 can surely win for $p, q = -W$, and the solution is bounded by $2 \cdot W$, since Player 1 surely loose when p or q are greater than W . Hence, a solution is obtained in one of the vertices of the polygon. We observe that a vertex in the polygon is either the average weight of some cycle or the intersection of two edges (in \mathbb{R}^2): $(\text{Avg}(c_1), \text{Avg}(c_2))$ and $(\text{Avg}(c_3), \text{Avg}(c_4))$ for some four simple cycles $c_1, c_2, c_3, c_4 \in \bigcup_{\sigma_2 \in \mathcal{S}_2[\mathcal{G}, M]} \mathbb{C}_{\sigma_2}$. We further observe that for any $\sigma_2 \in \mathcal{S}_2[\mathcal{G}, M]$ and any simple cycle c in the graph \mathcal{G}^{σ_2} we have $\text{Avg}(c) \in \{\frac{x}{y} : -nW \leq x \leq nW \wedge 1 \leq y \leq n\}$. Thus, by simple algebra, the intersection of the two edges is a point (a, b) that satisfies $a, b \in \{\frac{r}{\ell} : -(nW)^4 \leq r, \ell \leq (nW)^4\}$. Hence, we may assume that the maximal sum is obtained when p and q are of the above form. Hence, Player 1 can ensure a threshold couple with sum at least ν if and only if he can ensure a threshold couple of the above form.

In the general case each \mathcal{G}^{σ_2} may have more than one SCC. But the same arguments still hold. The maximal value is still the solution of a linear programming problem in which a single SCC is chosen for every \mathcal{G}^{σ_2} , and the optimal combination of SCCs gives the maximal value. \square

Proposition 10. *For an arena \mathcal{G} , an initial position v_0 , and a threshold $\nu \in \mathbb{R}$, if Player 1 does not have a winning strategy for $\Phi_{\text{exi}}^{\text{inf}}(\nu)$ from v_0 , then there exists $\epsilon > 0$ such that Player 1 does not have a winning strategy for $\Phi_{\text{exi}}^{\text{inf}}(\nu - \epsilon)$ from v_0 .*

Proof. Given an arena \mathcal{G} , recall that $\text{PAIRS}[\mathcal{G}] = \{(x, y) : \text{Player 1 can ensure } \Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y) \text{ on } \mathcal{G} \text{ from } v_0\}$. Consider $s = \sup\{x + y : (x, y) \in \text{PAIRS}[\mathcal{G}]\}$. By Lemma 9, s is also the supremum of the set $S = \{p + q : (q, p) \in \text{PAIRS}[\mathcal{G}] \text{ and } p, q \text{ are rationals whose numerators and denominators belong to } \{-(nW)^4, \dots, (nW)^4\}\}$. As S is a finite set, it follows that there are two rationals $(q, p) \in S$ such that $p + q = s$ and Player 1 wins in $\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$. Hence, if for every x, y such that $x + y \geq \nu$ Player 1 loses in $\Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y)$, then it must be the case that $s < \nu$, and the assertion holds for $\epsilon = \frac{\nu - s}{2}$. \square

5.3.2 Solving games with $B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu)$ objectives

In this section we study the solution for $B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu)$ objectives. The next lemma shows that if the winning set for two-dimensional mean-payoff games along with Büchi condition is empty, then the solution for sum of two mean-payoff objectives along with Büchi condition is also empty.

Lemma 11. *For every arena and every threshold ν we have $\text{Win}_1(B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu)) = \emptyset$ if and only if for all $x, y \in \mathbb{R}$ such that $x + y \geq \nu$ we have $\text{Win}_1(B_1 \wedge \Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y)) = \emptyset$.*

Proof. To prove the implication from left to right observe that for all $x, y \in \mathbb{R}$ with $x + y \geq \nu$ we have $\text{Win}_1(B_1 \wedge \Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y))$ is contained in $\text{Win}_1(B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu))$.

To prove the converse direction we assume that for every $x, y \in \mathbb{R}$ with $x + y \geq \nu$ the region $\text{Win}_1(B_1 \wedge \Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y)) = \emptyset$. By Proposition 10 applied to every position of the arena, we have that there exists $\epsilon > 0$ such that for all $x, y \in \mathbb{R}$ with $x + y \geq \nu - \epsilon$ it holds that $\text{Win}_1(B_1 \wedge \Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y)) = \emptyset$. We define a strategy σ_2 , which is a winning strategy for Player 2 against the objective $B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu)$.

1. Let AV_1 (resp., AV_2) be the average weight of the play until the current position in the first (resp., second) dimension.
2. If $\text{AV}_1 + \text{AV}_2 \leq \nu - \epsilon$, take an arbitrary move and return to Step 1.
3. Otherwise, if $\text{AV}_1 + \text{AV}_2 > \nu - \epsilon$, then let $\delta = \frac{\text{AV}_1 + \text{AV}_2 - (\nu - \epsilon)}{2}$. Then, play according to a strategy that ensures violation of the objective $B_1 \wedge \Phi_1^{\text{inf}}(\text{AV}_1 - \delta) \wedge \Phi_2^{\text{inf}}(\text{AV}_2 - \delta)$, until either the average value of the first dimension is less than $\text{AV}_1 - \frac{\delta}{2}$ or the average value of the second dimension is less than $\text{AV}_2 - \frac{\delta}{2}$. Then return to Step 1.

In Step 3 we have $AV_1 - \delta + AV_2 - \delta \geq \nu - \epsilon$, therefore a strategy for Player 2 to ensure violation of $B_1 \wedge \Phi_1^{\text{inf}}(AV_1 - \delta) \wedge \Phi_2^{\text{inf}}(AV_2 - \delta)$ exists.

We now prove that the strategy σ_2 is winning for Player 2. Towards contradiction suppose that ρ is a play that is consistent with σ_2 and satisfies $B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu)$. Hence, by definition, there is a round $\ell \in \mathbb{N}$ after which the average value of the first (resp., second) dimension is always at least MAXAV_1 (resp., MAXAV_2) and $\text{MAXAV}_1 + \text{MAXAV}_2 > \nu - \epsilon$. Then, either the strategy executes Step 2 infinitely often or after some time it remains in Step 3 forever.

- Consider that the strategy executes Step 2 infinitely often. Then, every time the strategy executes Step 2, it proceeds to Step 3, hence Step 3 is executed infinitely often. After the execution of Step 3, we have that either the average weight of the first dimension is less than $\text{MAXAV}_1 - \frac{\delta}{2}$ or that the average weight of the second dimension is less than $\text{MAXAV}_2 - \frac{\delta}{2}$ and the contradiction follows. Since Step 3 is executed infinitely often, we get a contradiction with the assumption that after the round ℓ , the average value of the first (resp., second) dimension is always at least MAXAV_1 (resp., MAXAV_2).
- Now, consider the case where after some time the strategy remains in Step 3 forever. Since Player 2 plays a strategy to violate $B_1 \wedge \Phi_1^{\text{inf}}(AV_1 - \delta) \wedge \Phi_2^{\text{inf}}(AV_2 - \delta)$ and average of the first (resp., the second) dimension is above $AV_1 - \frac{\delta}{2}$ (resp., $AV_2 - \frac{\delta}{2}$), it follows that the B_1 is violated.

The desired result follows. \square

Hence to solve games with Büchi and $\Phi_{\text{sum}}^{\text{inf}}(\nu)$ objectives, we first present solution of conjunction of a Büchi and two mean-payoff objectives, extending the results of [18].

Algorithm ALGOTWOMP. We present an algorithm (which we refer to as ALGOTWOMP) to compute $\text{Win}_1(B_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$ recursively in the following way:

1. Compute $W_1 = \text{Win}_1(B_1)$ and $W_2 = \text{Win}_1(\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$.
2. Let $X = V \setminus (W_1 \cup W_2)$.
3. If $X \neq \emptyset$, recursively call the algorithm on the arena $\mathcal{G} \upharpoonright (V \setminus \text{Attr}_2(X))$.
4. Otherwise, if $X = \emptyset$, then $\text{Win}_1(B_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$ is the whole arena.

We now present the correctness and complexity analysis of the algorithm.

Correctness. Observe that if $X \neq \emptyset$, then from every position in X , Player 2 can falsify $B_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$, and hence all positions in $\text{Attr}_2(X)$ belong to the winning set of Player 2 (i.e., do not belong to $\text{Win}_1(B_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$). Hence the computation correctly continues recursively over $\mathcal{G} \upharpoonright (V \setminus \text{Attr}_2(X))$. We now show that if $\text{Win}_1(B_1) = \text{Win}_1(\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)) = V$, then Player 1 has a winning strategy for $B_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$.

Lemma 12. *Let \mathcal{G} be an arena with the set of positions V such that $\text{Win}_1(B_1) = \text{Win}_1(\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)) = V$. Then, $\text{Win}_1(B_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)) = V$.*

Proof. We first prove that for $B_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$ objective, if Player 2 can violate the objective, then he can do it by a memoryless strategy. In order to prove that memoryless strategies suffice for Player 2, we observe that if two plays π_1 and π_2 satisfy both the Büchi and the mean-payoff objectives, then any arbitrary *mix* between the weight sequences of π_1 and π_2 and between the accepting positions in sequences of π_1 and π_2 will still satisfy both the Büchi and the mean-payoff objectives. Hence, this objective is a *convex objective*, as defined in [25], and moreover, it follows from [25] that against convex objectives, if Player 2 has a violating strategy, then he has a memoryless one. Thus it is enough to show that Player 1 wins in the whole arena against every Player 2 memoryless strategy.

Let σ_2 be an arbitrary Player-2 memoryless strategy. We define \mathcal{G}^{σ_2} as the graph resulting from fixing choices of Player 2 in the arena \mathcal{G} according to the strategy σ_2 . Let S be an arbitrary bottom strongly connected component (an SCC with no edges out of the SCC) in \mathcal{G}^{σ_2} and let u be an arbitrary position in S . As $u \in \text{Win}_1(B_1)$ it follows that S has a position from B_1 (otherwise, σ_2 is a Player 2 winning strategy from u). Since $u \in \text{Win}_1(\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$, then by Lemma 12 in [26], there exist k simple cycles C_1, \dots, C_k in S such that for some k positive integers N_1, \dots, N_k we have (i) $\sum_{i=1}^k N_i \cdot \text{wt}_1(C_i) \geq q$; and (ii) $\sum_{i=1}^k N_i \cdot \text{wt}_2(C_i) \geq p$. Let c_i be an arbitrary position in the cycle C_i and let π_i be the shortest path from c_i to c_{i+1} for $i = 1, \dots, k-1$ and let π_k be the shortest path from c_k to c_1 that goes through a position in B_1 . Clearly, for $i = 1, \dots, k-1$ we have $|\pi_i| \leq n$ and $|\pi_k| \leq 2n$. We denote by ρ_i the finite cyclic path

$$(C_1)^{N_1 \cdot i} \pi_1 (C_2)^{N_2 \cdot i} \pi_2 \dots \pi_{k-1} (C_k)^{N_k \cdot i} \pi_k$$

and by ρ the infinite path $\rho_1 \cdot \rho_2 \cdot \rho_3 \cdot \dots$. Clearly, ρ satisfies the Büchi condition B_1 and the fact that the path also satisfies the mean-payoff condition $\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$ follows from [26, Lemma 11]. The desired result follows. \square

Complexity. Algorithm ALGOTWOMP makes at most $|V|$ recursive calls. The set $\text{Win}_1(B_1)$ can be computed in quadratic time [12, 27, 28], and due to [18], $\text{Win}_1(\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$ can be computed in polynomial time in nW if $p, q \in \mathbb{Q}$ are represented by $O(\log(nW))$ bits. Therefore, each call takes polynomial time in nW . Therefore, the whole algorithm works in polynomial time in nW .

Lemma 13. *Let \mathcal{G} be a game arena with n positions and weights from $[-W, W]$. Given $p, q \in \mathbb{Q}$ represented by $O(\log(nW))$ bits, algorithm ALGOTWOMP correctly computes $\text{Win}_1(B_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$ in polynomial time in nW .*

Proposition 14. *Let \mathcal{G} be a game arena with n positions and weights from $[-W, W]$. The following assertions hold: (i) For every initial position v_0 , the maximum threshold ν , for which Player 1 can satisfy $B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu)$ is either $-\infty$ or it can be encoded by $8 \log(nW)$ bits and it is computable in polynomial time in nW . (ii) For a fixed threshold ν , the region $\text{Win}_1(B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu))$ can be computed in polynomial time in nW .*

Proof. By Lemma 9 and Lemma 11 it is enough to compute $\text{Win}_1(B_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$ for all rationals p, q whose numerators and denominators belong to $\{-(nW)^4, \dots, (nW)^4\}$, i.e., rationals that can be encoded by $8 \log(nW)$ bits. Then, for (i) we return the maximum $p + q$ over the winning regions containing v_0 . For (ii) we return the union of winning regions with $p + q \geq \nu$. Since there are only polynomially many threshold couples p, q to consider, and since for every p, q the computation of the winning region can be done in polynomial time (Lemma 13), the total complexity is polynomial in nW . \square

5.3.3 Solving games with $B_1 \wedge (coB_2 \vee \Phi_{\text{sum}}^{\text{inf}}(\nu))$ objectives

The next lemma suggests that we can replace the original winning objective with two simpler objectives.

Lemma 15. *For every arena, we have $\text{Win}_1(B_1 \wedge (coB_2 \vee \Phi_{\text{sum}}^{\text{inf}}(\nu))) = \emptyset$ if and only if $\text{Win}_1(B_1 \wedge coB_2) = \emptyset$ and $\text{Win}_1(B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu)) = \emptyset$.*

Proof. To prove the implication from right to left observe that $\text{Win}_1(B_1 \wedge coB_2)$ and $\text{Win}_1(B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu))$ are contained in $\text{Win}_1(B_1 \wedge (coB_2 \vee \Phi_{\text{sum}}^{\text{inf}}(\nu)))$.

For the implication from left to right assume that $\text{Win}_1(B_1 \wedge coB_2) = \emptyset$ and $\text{Win}_1(B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu)) = \emptyset$, i.e., Player 2 wins in the entire arena against both $B_1 \wedge coB_2$ and $B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu)$ objectives for Player 1. Also, we consider that $\text{Win}_1(B_1) \neq \emptyset$, since $\text{Win}_1(B_1) = \emptyset$ trivially implies $\text{Win}_1(B_1 \wedge (coB_2 \vee \Phi_{\text{sum}}^{\text{inf}}(\nu))) = \emptyset$. Then, there exists the maximal threshold ν_{\max} for which Player 1 can satisfy $B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu_{\max})$ and ν_{\max} is smaller than ν . Thus, by Proposition 14 it follows that there is $\epsilon > 0$ such that Player 2 wins in the entire arena also against the $B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu - \epsilon)$ objective. Let σ_2^1 (resp., σ_2^2) be a Player 2 violating strategy for $B_1 \wedge coB_2$ objective (resp., $B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu - \epsilon)$ objective). Note that we can choose σ_2^1 to be a memoryless strategy (since $B_1 \wedge coB_2$ objective is a parity objective with three priorities). We construct a winning strategy σ_2^{win} for Player 2 violating $B_1 \wedge (coB_2 \vee \Phi_{\text{sum}}^{\text{inf}}(\nu))$ objective. The strategy σ_2^{win} is to alternate the following steps.

Step 1. Play according to σ_2^1 until a position from B_2 is visited.

Step 2. Starting at the current round N , play according to σ_2^2 until the first round K for which there exist two rounds $N \leq i_1, i_2 \leq K$ such that the average weight of the first (respectively, the second) dimension in round i_1 (resp., i_2) is p (resp., q) and $p + q \leq \nu - \frac{\epsilon}{2}$. Note that in this step, σ_2^2 is played considering the history from the start of this step (i.e., it ignores the history that is before the start of this current step), but for the partial average we consider the weights of the entire history. Goto Step 1.

Consider a counter strategy σ_1 and the play π according to the strategy σ_1 and σ_2^{win} . If σ_2^{win} plays σ_2^1 infinitely long and cannot visit a position from B_2 , then since σ_2^1 is a strategy that ensures violation of $B_1 \wedge coB_2$, the play π violates condition B_1 . Similarly, if σ_2^{win} plays infinitely long σ_2^2 and does not achieve the specified condition on the partial averages, then the play π satisfies $\Phi_{\text{sum}}^{\text{inf}}(\nu - \frac{\epsilon}{2})$. This follows from the finite-history independent winning property that even with the finite history up to the beginning of Step 2, it still ensures winning against $B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu - \frac{\epsilon}{2})$. Since σ_2^2 is winning against $B_1 \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu - \epsilon)$ objective, π must violate condition B_1 . Finally, if the strategy σ_2^{win} alternates between Step 1 and Step 2 infinitely often, then the play π violates coB_2 and it violates $\Phi_{\text{sum}}^{\text{inf}}(\nu)$. \square

The next theorem is the main result of this section.

Theorem 16. *Let \mathcal{G} be a game arena with n positions and weights from $[-W, W]$. For the implication game with (Büchi, Büchi, LIMAVERAGE + LIMAVERAGE) objective, the value of every position can be computed in polynomial time in nW .*

Proof. For a fixed ν , we can compute the winning region by computing $R_1 = \text{Win}_1(B_1 \wedge \text{co}B_2)$ and $R_2 = \text{Win}(B_1 \wedge \Phi_{sum}^{\text{inf}}(\nu))$. If both R_1 and R_2 are empty, then by Lemma 15, Player 1 does not win for threshold ν from anywhere in the arena. Otherwise, $\text{Attr}_1(R_1 \cup R_2)$ is a part of Player 1 winning region, and we continue the computation recursively for $\mathcal{G} \upharpoonright (V \setminus \text{Attr}_1(R_1 \cup R_2))$. The computation is polynomial thanks to Proposition 14.

To compute the maximal threshold, we first compute the winning region for $\nu = 2W + 1$, and clearly if Player 1 wins for that threshold, then he wins for every threshold (so the value is ∞). In addition, for every pair of rationals p, q whose numerators and denominators belong to $\{-(nW)^4, \dots, (nW)^4\}$, we compute the winning region for threshold $\nu = p + q$, and for every position we assign the maximal threshold he can satisfy. We note that by Proposition 14, it is enough to check only thresholds that are the sum of rationals whose numerators and denominators belong to $\{-(nW)^4, \dots, (nW)^4\}$. Finally, for every position that cannot satisfy the threshold $\nu = -2W - 1$, we assign the value $-\infty$.

The complexity is polynomial in nW due to the fact that we need to consider only polynomial number of thresholds, and because the computation for a fixed threshold is polynomial. \square

5.4 Implication games with objectives (Büchi, Büchi, LIMAVGSUP + LIMAVGSUP)

In this section, we consider implication games defined by (Büchi, Büchi, LIMAVGSUP₁ + LIMAVGSUP₂) objective, where LIMAVGSUP₁, LIMAVGSUP₂ are defined as follows. Given weight functions wt_1 and wt_2 we have LIMAVGSUP₁ = LIMAVGSUP^{wt₁} and LIMAVGSUP₂ = LIMAVGSUP^{wt₂}. Similar to the previous section, given a threshold ν we will use the following notation:

$$\Phi_{sum}^{\text{sup}}(\nu) = \{\pi : \text{LIMAVGSUP}_1(\pi) + \text{LIMAVGSUP}_2(\pi) \geq \nu\};$$

i.e., the objective requires the sum of the mean-payoff objectives is at least ν .

Basic idea. As in the previous subsection, we solve the game by considering two simpler winning conditions, namely, B_1 and $(B_1 \wedge \text{co}B_2) \vee \Phi_{sum}^{\text{sup}}(\nu)$. We show that if Player 1 wins for both the objectives, then he also wins the original objective (and the converse direction is trivial). We first prove the above assertion and then we analyze games with $(B_1 \wedge \text{co}B_2) \vee \Phi_{sum}^{\text{sup}}(\nu)$ objective. We show that if we exchange the role of the players (namely, when Player 1 now wishes to violate the condition), then the new winning condition becomes a conjunction of a parity condition and a condition of the form $\Phi_{sum}^{\text{inf}}(\nu)$. We rely on the results of the previous section and extend the analysis from a conjunction of Büchi condition and $\Phi_{sum}^{\text{inf}}(\nu)$ to a conjunction of parity condition and $\Phi_{sum}^{\text{inf}}(\nu)$. We rely on the fact that the required parity condition consists of only three priorities and obtain a polynomial time solution in nW .

Single-dimensional vs. two-dimensional objectives. Recall Example 8 that shows objectives $\Phi_{sum}^{\text{inf}}(\nu) = \{\pi : \text{LIMAVGINF}^{\text{wt}_1}(\pi) + \text{LIMAVGINF}^{\text{wt}_2}(\pi) \geq \nu\}$ and $\Phi^*(\nu) = \{\pi : \text{LIMAVGINF}^{\text{wt}_1 + \text{wt}_2}(\pi) \geq \nu\}$ are different. Consider $\mathcal{G}'_1, \mathcal{G}'_2$ obtained from $\mathcal{G}_1, \mathcal{G}_2$ by multiplying all weights by -1 in Example 8 and all positions belong to Player 1. Then, every play π on \mathcal{G}'_2 has value 0, i.e., $\text{LIMAVGSUP}^{\text{wt}_1 + \text{wt}_2}(\pi) = 0$, while Player 1, using strategy σ_1 from Example 8, can achieve $\text{LIMAVGSUP}^{\text{wt}_1}(\pi) = \text{LIMAVGSUP}^{\text{wt}_2}(\pi) = 1$. Therefore, Player 1 can satisfy $\Phi_{sum}^{\text{sup}}(2)$, while the sum of the dimensions is zero at every position.

5.4.1 Reduction to $(\text{co}B_1 \vee B_2) \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$

In this section we will show that the crux of the analysis of implication games with (Büchi, Büchi, LIMAVGSUP₁ + LIMAVGSUP₂) objective can be reduced to analysis of games with $(\text{co}B_1 \vee B_2) \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$ objective. The first step to achieve the reduction is as follows: in the following lemma we show that if Player 1 wins in the entire arena for the B_1 and $(B_1 \wedge \text{co}B_2) \vee \Phi_{sum}^{\text{sup}}(\nu)$ objectives, then he can win the implication game for the entire arena.

Lemma 17. *Let \mathcal{G} be an arena and let V be the set of positions of \mathcal{G} . If $\text{Win}_1(B_1) = V$ and $\text{Win}_1((B_1 \wedge \text{co}B_2) \vee \Phi_{sum}^{\text{sup}}(\nu)) = V$, then $\text{Win}_1(B_1 \wedge (\text{co}B_2 \vee \Phi_{sum}^{\text{sup}}(\nu))) = V$.*

Proof. We construct a Player 1 winning strategy for $B_1 \wedge (\text{co}B_2 \vee \Phi_{sum}^{\text{sup}}(\nu))$ in the following way. Let σ_1 be a winning strategy for $(B_1 \wedge \text{co}B_2) \vee \Phi_{sum}^{\text{sup}}(\nu)$ for all positions in V . Initially set $\epsilon = 1$ and $N = 0$.

Step 1. Starting at the current round N , play according to σ_1 until the first round K for which there are two rounds $N \leq i_1, i_2 \leq K$ such that in round i_1 (resp., i_2) the average weight (from the beginning of the play) of the first (resp., second) dimension is at least r_1 (resp., r_2) and $r_1 + r_2 \geq \nu - \epsilon$. Note that in this step, σ_1 is played considering the history from the start of this step (i.e., it ignores the history that is before the start of this current step), but for the partial average we consider the weights of the entire history.

Step 2. Play according to a winning strategy for B_1 until a position from B_1 is visited (i.e., an attractor strategy to B_1 is played).

Step 3. Set $N = K + 1$, $\epsilon = \frac{\epsilon}{2}$ and return to Step 1.

In order to prove that the strategy is winning for $B_1 \wedge (coB_2 \vee \Phi_{sum}^{sup}(\nu))$ we consider two distinct cases.

1. In the first case, Steps 2 and 3 are executed only finitely many times. Hence, the strategy gets stuck forever from some point in Step 1. In other words, there is an infinite suffix π in which Player 1 plays only according to a winning strategy for $(B_1 \wedge coB_2) \vee \Phi_{sum}^{sup}(\nu)$, and for some $\epsilon > 0$ we get that $LIMAVGSUP_1(\pi) + LIMAVGSUP_2(\pi) \leq \nu - \epsilon$. Hence, it must be the case that $(B_1 \wedge coB_2)$ is satisfied, and thus $B_1 \wedge (coB_2 \vee \Phi_{sum}^{sup}(\nu))$ is also satisfied and Player 1 wins.
2. In the second case, Steps 2 and 3 are executed infinitely often. Note that in Step 2 the attractor strategy is only played for $|V|$ steps, and the strategy never gets stuck forever in Step 2. Let π be a play according to the strategy above. Since Step 2 is executed infinitely often, we get that π satisfies B_1 . Let $\nu_1 = LIMAVGSUP_1(\pi)$ and $\nu_2 = LIMAVGSUP_2(\pi)$. Towards a contradiction, let us assume that there exists $\delta > 0$ such that $\nu_1 + \nu_2 \leq \nu - \delta$. By the definition of $LIMAVGSUP$, there exists a round N such that in every round after round N the average weight of dimension i is at most $\nu_i + \frac{\delta}{4}$ (for $i = 1, 2$). Hence, for ϵ that is smaller than $\frac{\delta}{4}$, the execution of Step 1 will never terminate and the contradiction follows. Hence, we get that $LIMAVGSUP_1(\pi) + LIMAVGSUP_2(\pi) \geq \nu$ and π satisfies B_1 . Thus, the π belongs to $B_1 \wedge (coB_2 \vee \Phi_{sum}^{sup}(\nu))$.

Note that (as in Lemma 15) we also use the finite-history independent winning property above. To conclude, we get that in both cases the strategy is winning for $B_1 \wedge (coB_2 \vee \Phi_{sum}^{sup}(\nu))$ and the result follows. \square

Remark 18. (Implication of the above lemma). Note that it is trivially true that from

$$X = (V \setminus \text{Win}_1(B_1)) \cup (V \setminus \text{Win}_1((B_1 \wedge coB_2) \vee \Phi_{sum}^{sup}(\nu))),$$

Player 2 has a strategy to violate the objective of the implication game (as violation of any of B_1 or $(B_1 \wedge coB_2) \vee \Phi_{sum}^{sup}(\nu)$ implies the violation of the condition for implication game). Hence if X is non-empty, then the Player-2 attractor of X can be removed from the arena, and the computation can continue on the sub-arena. Hence solving the implication game reduces to solving games with Büchi condition (that is well-known); and games with $(B_1 \wedge coB_2) \vee \Phi_{sum}^{sup}(\nu)$ condition, which we study below.

Games with $(B_1 \wedge coB_2) \vee \Phi_{sum}^{sup}(\nu)$ conditions. In order to compute the winning region for the objective $(B_1 \wedge coB_2) \vee \Phi_{sum}^{sup}(\nu)$, we compute Player 2 winning region for the complementary objective, namely, for $(coB_1 \vee B_2) \wedge \neg \Phi_{sum}^{sup}(\nu)$, where $\neg \Phi_{sum}^{sup}(\nu) = \{\pi : LIMAVGSUP_1(\pi) + LIMAVGSUP_2(\pi) < \nu\}$. This objective for the arena \mathcal{G} is equivalent to the objective $(coB_1 \vee B_2) \wedge \Phi_{sum}^{inf}(\nu)$ in an arena \mathcal{G}' that is obtained by multiplying all the weights in \mathcal{G} by -1 . In addition, we observe that the condition $(coB_1 \vee B_2)$ can be encoded by a parity condition with three priorities. Hence, in some cases we formulate the objective as a conjunction of a parity condition P and a mean-payoff sum. In the rest of this section (until Theorem 27), we switch the roles for the players and consider that Player 1 wishes to satisfy $(coB_1 \vee B_2) \wedge \Phi_{sum}^{inf}(\nu)$. The next lemma extends Lemma 11 to parity condition.

Lemma 19. Let P be a parity condition. Then $\text{Win}_1(P \wedge \Phi_{sum}^{inf}(\nu)) = \emptyset$ if and only if for all $x, y \in \mathbb{R}$ with $x + y \geq \nu$ we have $\text{Win}_1(P \wedge \Phi_1^{inf}(x) \wedge \Phi_2^{inf}(y)) = \emptyset$.

Proof. The proof follows exactly by the same arguments as the proof of Lemma 11 (simply by replacing the Büchi condition with a parity condition). \square

The above lemma provides a reduction of games with $(coB_1 \vee B_2) \wedge \Phi_{sum}^{inf}(\nu)$ objective to the conjunction of the parity and the existential version of two-dimensional mean-payoff objectives $\Phi_{exi}^{inf}(\nu)$. In order to complete the reduction of this section, we show that the quantification of x, y in the $\Phi_{exi}^{inf}(\nu)$ can be restricted to rationals represented by at most $O(\log(nW))$ bits. This allows to reduce $\text{Win}_1(P \wedge \Phi_{sum}^{inf}(\nu))$ to polynomial number of invocations of a procedure computing, given p, q , the region $\text{Win}_1(P \wedge \Phi_1^{inf}(q) \wedge \Phi_2^{inf}(p))$.

Lemma 20. For every threshold ν and an arena \mathcal{G} with n positions and weights from $[-W, W]$: Player 1 can ensure $(coB_1 \vee B_2) \wedge \Phi_1^{inf}(x) \wedge \Phi_2^{inf}(y)$ for some $x, y \in \mathbb{R}$ such that $x + y \geq \nu$ iff there is a threshold couple $p, q \in \mathbb{Q}$ whose numerator and denominator are bounded by $\pm(nW)^4$ with $p + q \geq \nu$ and Player 1 wins in $(coB_1 \vee B_2) \wedge \Phi_1^{inf}(q) \wedge \Phi_2^{inf}(p)$.

Proof. We first prove that for such objectives, if Player 2 can violate the objective, then he can do it by a memoryless strategy, and the rest of the proof is similar to the proof of Lemma 9. In order to prove that memoryless strategies suffice, we view this objective as a conjunction of parity and two-dimensional mean-payoff objectives, and we observe that if two plays π_1 and π_2 satisfy the parity and mean-payoff objectives, then any arbitrary *mix* between the weight sequences of π_1 and π_2 and between the priority sequences of π_1 and π_2 will still satisfy both the parity and the mean-payoff objectives. Hence, this objective is a *convex objective*, as defined in [25], and it follows from [25] that against convex objectives, if Player 2 has a violating strategy, then he has a memoryless one.

Recall that for a finite path $\pi = e_1 e_2 \dots e_n$ and $i = 1, 2$ we denote $\text{wt}_i(\pi) = \sum_{j=1}^n \text{wt}_i(e_j)$. Then, we define $\text{Avg}(\pi) = (\frac{\text{wt}_1(\pi)}{n}, \frac{\text{wt}_2(\pi)}{n})$. In one-player games, Player 1 can satisfy the mean-payoff objectives for thresholds x, y if there is an SCC C with the set of simple cycles of C being \mathbb{C} such that $(x, y) \in \text{CONV}\{\text{Avg}(c) : c \in \mathbb{C}\}$ (e.g., see [24]). Hence, when adding a parity condition, we require an SCC \mathbb{C} , with minimal priority even,¹ where the set of simple cycles of C is \mathbb{C} , and $(x, y) \in \text{CONV}\{\text{Avg}(c) : c \in \mathbb{C}\}$. Hence, we can repeat the same analysis as in the proof of Lemma 9 with a change that we consider only SCCs with even minimal priority. The priorities over the positions do not change the size of the extreme points in the constructed polygon. Hence, the analysis yields the same results as in Lemma 9. \square

The following proposition concludes this section.

Proposition 21. *Let \mathcal{G} be an arena with n positions and weights from $[-W, W]$. The game with $B_1 \wedge (coB_2 \vee \Phi_{sum}^{\text{sup}}(\nu))$ objective Turing-reduces in polynomial time in nW to games with $(coB_1 \vee B_2) \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$ objectives on arenas with at most n positions and weights from $[-W, W]$ and $p, q \in \mathbb{Q}$ represented by $O(\log(nW))$ bits.*

5.4.2 Solving games with $(coB_1 \vee B_2) \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$ objectives

We present the following lemmas to obtain an algorithm for computing the winning region for the condition $\varphi = (coB_1 \vee B_2) \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$. The first lemma gives a general scheme for computing the winning regions of a tail objective, and the other two lemmas show how to compute the winning region for sub-formulas of φ .

Lemma 22. *Let ϕ be a tail objective and let \mathcal{G} be an arena with n positions and weights from $[-W, W]$. Assume that for every game sub-arena \mathcal{G}' of \mathcal{G} , there is an algorithm (oracle \textcircled{O}) that in polynomial time in nW returns some $v \in \mathcal{G}'$ with an information $v \in \text{Win}_1(\phi)$ or $v \in \text{Win}_2(\phi)$. Then the winning regions of the game can be computed in polynomial time in nW .*

Proof. The following recursive scheme computes the winning regions with at most n recursive calls:

- If $n = 0$, then we are done.
- Run \textcircled{O} and find a position v that belongs either to $\text{Win}_1(\varphi)$ or to $\text{Win}_2(\varphi)$.
- If $v \in \text{Win}_i(\varphi)$, then $\text{Attr}_i(v)$ is part of Player i winning region.
- Continue the computation recursively over $\mathcal{G} \upharpoonright (V \setminus \text{Attr}_i(v))$, where V is the set of positions of \mathcal{G} .

In every recursive call the size of the arena is reduced by at least 1. As the computation of the attractor can be done in linear time, we get a polynomial algorithm with at most n calls to \textcircled{O} . \square

Lemma 23. *Let \mathcal{G} be a game arena with n positions and weights from $[-W, W]$. There is an algorithm that, given $p, q \in \mathbb{Q}$ represented by $O(\log(nW))$ bits, computes the region $\text{Win}_1(coB_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$ in \mathcal{G} in polynomial time in nW .*

Proof. Let V be the set of positions of \mathcal{G} . By Lemma 22 it is enough to present a polynomial time algorithm in nW that returns a position and says whether it belongs to the winning region of Player 1 or Player 2. The algorithm is as follows: We first compute $\text{Attr}_2(B_1)$ and consider the following cases.

- If $\text{Attr}_2(B_1) = V$, then Player 1 loses everywhere since Player 2 can ensure that B_1 is visited infinitely often. Thus we return an arbitrary position and mark it as part of Player-2 winning region.
- Otherwise, we compute $\text{Win}_1(\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$ over $\mathcal{G} \upharpoonright (V \setminus \text{Attr}_2(B_1))$. We now consider two sub-cases.
 - If $v \in \text{Win}_1(\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$, then it is also in $\text{Win}_1(coB_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$, and we can return v as a Player-1 winning position.

¹note that in case of adding a parity condition the SCC need not be a maximal SCC

- If $\text{Win}_1(\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)) = \emptyset$ in $\mathcal{G} \upharpoonright (V \setminus \text{Attr}_2(B_1))$, then we claim that $\text{Win}_1(\text{co}B_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$ over \mathcal{G} is also empty, and thus, we can return an arbitrary position and mark it as part of Player-2 winning region. Indeed, the following strategy is winning for Player 2:

1. If in $\text{Attr}_2(B_1)$, play an attractor strategy to reach B_1 .
2. If not in $\text{Attr}_2(B_1)$, play to violate the mean-payoff condition $\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$, as long as the play is not in $\text{Attr}_2(B_1)$. Return to 1.

Clearly, this strategy either violates $\text{co}B_1$ or the mean-payoff condition (or both).

The desired result follows. \square

Lemma 24. *Let \mathcal{G} be an arena with the set of positions V , two Büchi conditions B_1 and B_2 , and two-dimensional weight function. Let $x, y \in \mathbb{R}$. Assume that (i) Player-1 winning region for $\Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y)$ is V ; and (ii) in the arena $\mathcal{G} \upharpoonright (V \setminus \text{Attr}_1(B_2))$, Player 1 wins the condition $\text{co}B_1 \wedge \Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y)$ from every initial position. Then Player 1 winning region for $(\text{co}B_1 \vee B_2) \wedge \Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y)$ is the entire arena \mathcal{G} .*

Proof. In order to prove the lemma we describe a Player-1 winning strategy. We present the key ideas of the construction as the details are similar to previous results. Informally, the following is a Player-1 winning strategy: Initially set $\epsilon = 1$

1. Play according to a winning strategy that satisfies the mean-payoff conditions (recall that we assume that $\text{Win}_1(\Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y)) = V$), until the average weight is at least $x - \epsilon$ in the first dimension and $y - \epsilon$ in the second dimension.
2. If the play is not in $\text{Win}_1(\text{co}B_1 \wedge \Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y))$:
 - Then the play is in $\text{Attr}_1(B_2)$. Play attractor strategy until a position from B_2 is reached.
 - Else, play to satisfy $\text{co}B_1 \wedge \Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y)$. If at some point Player 2 advances the token to $\text{Attr}_1(B_2)$, then play attractor strategy until a position from B_2 is reached.
3. Set $\epsilon = \frac{\epsilon}{2}$ and return to 1.

Intuitively, the correctness of the construction of the winning strategy follows from the fact that the mean-payoff conditions are satisfied, and either B_2 is visited infinitely often, or there is an infinite suffix that is played according to a winning strategy for $\text{co}B_1 \wedge \Phi_1^{\text{inf}}(x) \wedge \Phi_2^{\text{inf}}(y)$. The desired result follows. \square

Given the above lemmas we now present the main result of this section.

Lemma 25. *Let \mathcal{G} be a game arena with n positions and weights from $[-W, W]$. There is an algorithm to compute $\text{Win}_1((\text{co}B_1 \vee B_2) \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$ in polynomial time in nW for all p, q represented by $O(\log(nW))$ bits.*

Proof. Let V be the set of positions of \mathcal{G} . We recall that by Lemma 22 it is enough to present an algorithm that returns a position and mark it either as a Player-1 winning position or as a Player-2 winning position. The algorithm is as follows:

- We first compute $\text{Win}_1(\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$ (can be done in polynomial time in nW [18]). Clearly, if $v \notin \text{Win}_1(\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$, then $v \notin \text{Win}_1((\text{co}B_1 \vee B_2) \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$ and we return v and mark it as a Player-2 winning position.
- Otherwise, Player 1 wins $\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$ everywhere in the arena. Next, we compute $\text{Win}_1(B_2 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$ (polynomial due to Lemma 13). If $v \in \text{Win}_1(B_2 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$, then surely $v \in \text{Win}_1((\text{co}B_1 \vee B_2) \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$, and we return v as a Player-1 winning position.
- Otherwise, $\text{Win}_1(B_2 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)) = \emptyset$. Next, we compute $\text{Attr}_1(B_2)$ and compute $\text{Win}_1(\text{co}B_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$ over $\mathcal{G} \upharpoonright (V \setminus \text{Attr}_1(B_2))$ (the computation is polynomial by Lemma 23). If $v \notin \text{Win}_1(\text{co}B_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$, then $v \notin \text{Win}_1((\text{co}B_1 \vee B_2) \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p))$, since a violating strategy for Player 2 over $\mathcal{G} \upharpoonright (V \setminus \text{Attr}_1(B_2))$ can avoid reaching $\text{Attr}_1(B_2)$ in \mathcal{G} , and thus Player 2 violates the condition. Thus, if such v exists, then we return v and mark it as a Player-2 winning position.
- Otherwise, (i) Player 1 wins $\Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$ everywhere; and (ii) Player 1 wins $\text{co}B_1 \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$ everywhere in $\mathcal{G} \upharpoonright (V \setminus \text{Attr}_1(B_2))$. Hence, by Lemma 24, Player 1 wins the condition $(\text{co}B_1 \vee B_2) \wedge \Phi_1^{\text{inf}}(q) \wedge \Phi_2^{\text{inf}}(p)$ everywhere in \mathcal{G} and we return an arbitrary position marked as a Player-1 winning position.

The desired result follows. \square

5.4.3 The main result.

Proposition 26. *Let \mathcal{G} be a game arena with n positions and weights from $[-W, W]$. The maximum threshold ν , for which Player 1 can satisfy $(\text{co}B_1 \vee B_2) \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu)$ from an initial position v_0 , is either $-\infty$ or it can be encoded by $8 \log(nW)$ bits and it is computable in polynomial time in nW . Moreover, for a fixed ν , it is possible to compute $\text{Win}_1((\text{co}B_1 \vee B_2) \wedge \Phi_{\text{sum}}^{\text{inf}}(\nu))$ in polynomial time in nW .*

Proof. The proof is similar to the proof of Proposition 14, we simply replace Lemma 9 and Lemma 11 by Lemma 20 and Lemma 25, respectively. \square

We use Remark 18 and similar to the proof of Theorem 16 (by using Proposition 26 instead of Proposition 14) we obtain the following result.

Theorem 27. *Let \mathcal{G} be a game arena with n positions and weights from $[-W, W]$. For the implication game with (Büchi, Büchi, LIMAVGSUP + LIMAVGSUP) objective, the value of every position can be computed in polynomial time in nW .*

6 Discounted-sum Implication Games

In this section we will present solution of implication discounted-sum games where the Boolean objectives are Büchi objectives. However, before we present the solution of discounted-sum parity games. While mean-payoff parity games have been studied before, discounted-sum parity games have not been studied, and we present their solution as they may be of independent interest.

6.1 Discounted-sum parity games

We consider *discounted-sum parity games* that are implication games $(\Phi_1, \Phi_2, \text{DISC}_\lambda)$ where Φ_1 is a parity objective and Φ_2 is tautology (satisfied for all plays). In other words, the objective is to ensure that the parity objective satisfied and minimize the discounted-sum objective.

Theorem 28. *The following assertions hold for discounted-sum parity games:*

1. *There is a polynomial-time reduction (Turing reduction) to parity games and discounted-sum games.*
2. *For every $\epsilon > 0$, Player 1 has a finite-memory ϵ -optimal strategy.*
3. *Player 2 has an optimal memoryless strategy.*

Proof. Consider an arena \mathcal{G} , a priority function p , a labeling function wt , and a discount factor $\lambda \in (0, 1)$. Let the parity objective defined by the function p be Φ . Let $\mathcal{W}_1 = \text{Win}_1(\Phi)$ be the set of winning positions for the parity objective. We define an arena \mathcal{G}' which is \mathcal{G} restricted to \mathcal{W}_1 , i.e., $\mathcal{G}' = \mathcal{G} \upharpoonright \mathcal{W}_1$. Note that for every Player 2 position v in \mathcal{W}_1 all its outgoing edges must remain in \mathcal{W}_1 (as otherwise v would not belong to \mathcal{W}_1). We establish the following: (A) For every position not in \mathcal{W}_1 the value is ∞ ; and (B) for every position in \mathcal{W}_1 the value for the discounted-sum parity objective in \mathcal{G} is the value for the discounted-sum objective in \mathcal{G}' . Since for every position v not in \mathcal{W}_1 Player 2 has a winning strategy to satisfy $\bar{\Phi}$ (i.e., to falsify Φ) from v it follows that the value for the discounted-sum parity objective at v is ∞ . We now establish (B). Consider a position $v \in \mathcal{W}_1$ and let t be the value at v for the discounted-sum game played on \mathcal{G}' . We prove that (a) the value for the discounted-sum parity objective at v in \mathcal{G} is at least t , and (b) for every $\epsilon > 0$, there is a strategy for Player 1 in \mathcal{G} to achieve the value at most $t + \epsilon$ for the discounted-sum parity objective. Clearly, (a) and (b) imply (B) for v . This implies that the decision problem for discounted-sum parity games (whether the value of the game is below a given threshold) can be solved in polynomial time with oracles for parity games and discounted-sum games. We now establish (a) and (b).

(a) *The value of \mathcal{G} is at least t .* Consider a strategy σ_1 for Player 1 in \mathcal{G} . If against this strategy Player 2 can ensure to leave \mathcal{W}_1 , then Player 2 can ensure that the parity objective Φ is violated. Hence for such a strategy the value is ∞ . Otherwise, if the strategy σ_1 ensures that \mathcal{W}_1 is never left, then the strategy is also a strategy for the discounted-sum game on \mathcal{G}' . Thus, its value cannot be lower than the value of the discounted-sum game on \mathcal{G}' .

(b) *For every $\epsilon > 0$, there is a strategy for Player 1 to achieve the value at most $t + \epsilon$.* Let W be the maximal absolute weight in \mathcal{G} . Given $\epsilon > 0$, let $k > 0$ satisfy the following inequality $W \cdot \lambda^k \cdot \frac{1}{1-\lambda} < \frac{\epsilon}{2}$. Since λ^k decreases exponentially, clearly for every $\epsilon > 0$ such a k exists. Consider a strategy for Player 1 in which he plays according to a memoryless optimal strategy in the discounted-sum game on \mathcal{G}' for k rounds, and then plays according to a memoryless winning strategy in the parity game on \mathcal{G} . Note that since \mathcal{W}_1 is not left for the k


rounds, Player 1 can switch to a winning strategy for the parity objective. First, observe that since Player 1 plays an optimal strategy for the discounted-sum objective for k rounds, the discounted-sum payoff for the first k rounds is at most $t + W \cdot \lambda^k \cdot \frac{1}{1-\lambda}$ (otherwise the discounted-sum value would have exceeded t), and the discounted-sum payoff after k rounds is at most $W \cdot \lambda^k \cdot \frac{1}{1-\lambda}$. Moreover, since a winning strategy for the parity objective is played after a finite number of steps, it follows that the parity objective is satisfied. Hence the value of the resulting play does not exceed $t + 2 \cdot W \cdot \lambda^k \cdot \frac{1}{1-\lambda} < t + \epsilon$. Also note that the strategy for Player 1 is a finite-memory ϵ -optimal strategy as it plays a memoryless strategy for k rounds (that can be implemented by a bounded counter), which is a memoryless optimal strategy for the discounted-sum games, and then switches to another memoryless strategy, which is a memoryless winning strategy for the parity objective.

The above concludes the proof for the first two items, and now we prove the existence of memoryless optimal strategies for Player 2.

Memoryless strategies for Player 2: Both, discounted-sum and parity games admit memoryless optimal/winning strategies. Consider the following strategy σ_2^* for Player 2. On positions in \mathcal{W}_1 , Player 2 plays according to an optimal memoryless strategy in the discounted-sum game on \mathcal{G}' and in positions from $V \setminus \mathcal{W}_1$ he plays according to a memoryless winning strategy (to falsify Φ) in the parity game on \mathcal{G} . We now show that σ_2^* , which is memoryless, is an optimal strategy for Player 2 in the discounted-sum parity game. Consider a strategy σ_1 for Player 1. If a play defined by σ_1 and σ_2^* reaches $V \setminus \mathcal{W}_1$, then the strategy σ_2^* ensures that the parity objective is violated and the value is infinite. Otherwise, if the play stays in \mathcal{W}_1 , then σ_2^* coincides with an optimal memoryless strategy in the discounted-sum game on \mathcal{G}' . Thus, its value is not smaller than that value of the discounted-sum parity game on \mathcal{G} . \square

In the previous theorem we established the existence of finite-memory ϵ -optimal strategies for Player 1 for all $\epsilon > 0$. We now show that optimal strategies need not exist in general for Player 1 (in contrast to mean-payoff parity games where optimal strategies exist for Player 1, Theorem 5 item (ii)).


Example 29. (*Non-existence of optimal strategies*). Consider an arena depicted below:



There are two positions A, B with priorities 1 and 0, both owned by Player 1. There are two loops, at A of weight 0 and at B of weight 1. Once Player 1 leaves A to B , he cannot return to A . The value of the game is 0 since Player 1 can take arbitrarily many times the loop at A of weight 0, and then move to B . But, to satisfy the parity condition, which is a Büchi condition in our case that requires to visit B infinitely often, Player 1 has to leave A eventually. Therefore, the value of every play that satisfies the parity condition is strictly positive. Also note that in the example, for discount factor $\frac{1}{2}$ and sufficiently small $\epsilon > 0$, every ϵ -optimal strategy requires memory.

In mean-payoff parity games, Player 1 has optimal strategies, but they require infinite memory. However, in mean-payoff parity games, if there is a finite-memory optimal strategy, then there is a memoryless optimal strategy (see Theorem 5, item (iv)). We now show that in contrast, there exist discounted-sum parity games, where finite-memory optimal strategies exist, but no memoryless strategy is an optimal one.

Example 30. (*Finite-memory stronger than memoryless for optimal strategies*). Consider the game with positions A, B, C, D, E , where position E is an absorbing position (only a self-loop as the transition) with weight 0, and priority 0. Position A belongs to Player 1 where the outgoing edges are to B with weight 0 and to E with weight 2. Position B belongs to Player 2, and the outgoing edges are to C and D with weight 0. Positions C and D have A as the unique successor and the edge weights are 0 for C to A and -1 for D to A . Positions A, B, C have priority 2 and D has priority 1. The following figure shows the game pictorially.



Consider the game with discount factor $\frac{1}{2}$. We first claim that the value of the game is 0. First observe that Player 2 by always choosing the edge from B to C ensures that the value is at least 0. We now present a finite-memory optimal strategy for Player 1 to ensure value 0, and the strategy for Player 1 is as follows: it chooses A to B as long as Player 2 chooses B to C ; however, if Player 2 chooses B to D , then Player 1 chooses A to E . Given the strategy for Player 1, either A, B, C is visited infinitely often, and the discounted-sum payoff is 0, and the parity objective is satisfied. Otherwise, the edge weight -1 is followed by edge weight 2, and all other edge weights are 0. The effective discounted-sum contribution for edge weight 2 is 1 as compared to edge weight -1 as it appears one step after and the discount factor is $\frac{1}{2}$. Hence the discounted-sum is 0, and the minimum priority visited infinitely often is 0. Hence the value of the game is 0 and there is a finite-memory optimal strategy. However, consider the two memoryless strategies for Player 1. If Player 1 chooses A to E , then the discounted-sum parity payoff is 1. If Player 1 always chooses A to B , then by choosing B to D Player 2 ensures that the parity objective is violated, and the payoff is ∞ . It follows that in this game though finite-memory optimal strategies exist, there is no memoryless optimal strategy.

6.2 Implication games with objectives (Büchi, Büchi, DISC_λ).

In the section we study implication games with objective (Büchi, Büchi, DISC_λ), which correspond to the (quantitative) simulation between DISC_λ -automata with Büchi acceptance conditions. Note that for discounted-sum objectives, the sum of two discounted-sum objectives $\text{DISC}_\lambda^{\text{wt}_1} + \text{DISC}_\lambda^{\text{wt}_2}$ is equivalent to one discounted-sum objective with the sum of the weight functions, i.e., $\text{DISC}_\lambda^{(\text{wt}_1 + \text{wt}_2)}$, where the function $(\text{wt}_1 + \text{wt}_2)$ for every edge assigns the sum of wt_1 and wt_2 . Hence we only consider implication games where the quantitative objective is a discounted-sum objective.

Theorem 31. *Implication games with objective (Büchi, Büchi, DISC_λ) reduce in polynomial time to discounted-sum games.*

Proof. Consider an arena \mathcal{G} , Büchi conditions F_1, F_2 , a labelling function wt , and a discount factor $\lambda \in (0, 1)$. Consider the Boolean objective Φ_1 that requires F_1 to be visited infinitely often and F_2 finitely often, and the Boolean objective Φ_2 that requires F_1 to be visited finitely often. Let $\mathcal{W}_1 = \text{Win}_1(\Phi_1)$ and $\mathcal{W}_2 = \text{Win}_2(\Phi_2)$. Observe that for a position $v \in \mathcal{W}_1$, Player 1 can always force the value of the implication game to be $-\infty$, by ensuring its own Büchi condition and falsifying the Büchi condition of the opponent. Similarly, if $v \in \mathcal{W}_2$, then Player 2 can force the value of the implication game to be ∞ by violating the Büchi condition for Player 1. The sets $\mathcal{W}_1, \mathcal{W}_2$ can be computed in polynomial time [20, 12, 27, 28].

Let $X = V \setminus (\mathcal{W}_1 \cup \mathcal{W}_2)$, and we now consider positions in X . We consider the arena $\mathcal{G}' = \mathcal{G} \upharpoonright X$. Observe that Player-1 positions in X have outgoing edges only to positions in X and \mathcal{W}_2 , but not to \mathcal{W}_1 ; whereas Player-2 positions in X have outgoing edges only to positions in X and \mathcal{W}_1 , but not to \mathcal{W}_2 . For example, if from a position $v \in X$ for Player 2 there is a move to \mathcal{W}_2 , then v would be in \mathcal{W}_2 , and similarly for positions for Player 1. We claim that the value at v of the implication game on \mathcal{G} is equal to the value of the discounted-sum game at v on \mathcal{G}' . Intuitively, the first player who leaves \mathcal{G}' loses in the implication game, i.e., the value of the play is ∞ if Player 1 left \mathcal{G}' and $-\infty$ if Player 2 left \mathcal{G}' first. Otherwise, in \mathcal{G}' Player 1 can ensure that his Büchi condition is satisfied (since the game is outside \mathcal{W}_2), and Player 2 can ensure that either his Büchi condition is satisfied or the Büchi condition of Player 1 is violated (since the game is outside \mathcal{W}_1). Thus for each $\epsilon > 0$, each player can play the discounted-sum game long enough to ensure that value is within ϵ of the discounted-sum game, and then switch to satisfying their respective Boolean objective (own Büchi condition for Player 1, and $\bar{\Phi}_1$ for Player 2). We present one case of the formal argument and the other case is similar. Recall that the value of the play in the implication game is given by the following function f : let Ψ_1 be the objective to visit F_1 infinitely often, and Ψ_2 the objective to visit F_2 infinitely often; then

$$f(\pi) = \begin{cases} \text{DISC}_\lambda^{\text{wt}}(\pi) & \text{if } \Psi_1(\pi) = \Psi_2(\pi) = 1, \\ \infty & \text{if } \Psi_1(\pi) = 0, \\ -\infty & \text{if } \Psi_1(\pi) = 1 \text{ and } \Psi_2(\pi) = 0. \end{cases}$$

Let $\sigma'_1 \in \mathcal{S}_1[\mathcal{G}', M]$ be a memoryless optimal strategy for Player 1 for the discounted-sum objective on \mathcal{G}' and let $\epsilon > 0$. Let N be such that the discounted-sum after N steps does not exceed $\frac{\epsilon}{2}$ (for every $\epsilon > 0$ such a N exists as in the proof of Theorem 28). Consider σ_1 that plays according to σ'_1 for histories contained in \mathcal{G}' for N steps, after that it plays to satisfy the Büchi condition F_1 . On histories ending in positions from \mathcal{W}_1 , the strategy σ_1 plays according to a winning strategy that satisfies Φ_1 . Finally, just for completeness, on the remaining histories strategy σ_1 takes some moves. Consider a strategy σ_2 of Player 2. If the play $\pi(\sigma_1, \sigma_2, v)$ is contained in \mathcal{G}' , it satisfies the Büchi condition F_1 . Moreover, there is a strategy $\sigma'_2 \in \mathcal{S}_2[\mathcal{G}']$ that follows σ_2 on \mathcal{G}' . Both

the play $\pi(\sigma_1, \sigma_2, v)$ and the play $\pi(\sigma'_1, \sigma'_2, v)$ coincide on first N steps, therefore either (a) $\pi(\sigma_1, \sigma_2, v)$ satisfies the Büchi condition F_2 , and $|f(\pi(\sigma_1, \sigma_2, v)) - \text{DISC}_\lambda^{\text{wt}}(\pi(\sigma'_1, \sigma'_2, v))| < \epsilon$; or (b) it violates the Büchi condition F_2 , and $-\infty = f(\pi(\sigma_1, \sigma_2, v)) < \text{DISC}_\lambda^{\text{wt}}(\pi(\sigma'_1, \sigma'_2, v))$. If $\pi(\sigma_1, \sigma_2, v)$ leaves \mathcal{G}' , it is due to a move of Player 2, and it must lead to \mathcal{W}_1 as σ_1 on histories contained in \mathcal{G}' stays within \mathcal{G}' . Once Player 2 moves to \mathcal{W}_1 , then σ_1 coincides with a winning strategy that satisfies Φ_1 and thus, the value of the play in $-\infty \leq \text{DISC}_\lambda^{\text{wt}}(\pi(\sigma'_1, \sigma'_2, v))$. This shows the argument for Player 1, and the argument for Player 2 is similar. \square

Note that the above result implies that for implication games with Büchi conditions and discounted-sum objectives there exist pseudo-polynomial time algorithms.

Remark 32. (Non-existence of optimal strategies). Note that implication games with Büchi conditions and discounted-sum objectives generalize discounted-sum Büchi games, and hence it follows from Example 29 that optimal strategies do not exist for Player 1 in general. However, whereas for discounted-sum parity games memoryless optimal strategies exist for Player 2, in implication games where the first objective is tautology, the second objective is a Büchi objective, we obtain a discounted-sum Büchi game for Player 2. It follows from Example 29 that optimal strategies do not exist for Player 2 in general, i.e., in implication games with Büchi conditions and discounted-sum objectives optimal strategies do not exist for both players in general.

7 Conclusion and Future Works

In this paper, we defined a new kind of games called *implication games*. In the framework of implication games we studied Quantitative Fair Simulation Games (QFSGs), which define simulation between weighted automata with Büchi acceptance conditions. We obtained polynomial-time algorithms for these games under the assumption that weights are encoded in unary. Note that since we consider non-deterministic automata it suffices to consider Büchi acceptance conditions for ω -regular conditions. Our framework of implication games can also be applied to obtain solution of computing simulation distances between two non-deterministic (non-weighted) automata where each of the automaton is equipped with a Büchi acceptance condition over the infinite runs. Since we consider only Büchi acceptance conditions for the automata problem, for our simulation games we only need to consider implication games where the Boolean objectives are Büchi conditions. However, the solution problems for implication games with more general Boolean conditions, such as parity, Streett, Rabin objectives, are interesting open problems. Moreover, for quantitative objectives we consider the sum of LIMAVGINF and the sum of LIMAVGSUP objectives, whereas mixing the sum of LIMAVGINF and LIMAVGSUP appears to be another interesting problem. Finally, implication games with other quantitative objectives, such as energy objectives, Boolean combination of mean-payoff or discounted-sum objectives, are other interesting direction of future works.

The following table shows the results of our paper in the context of previously studied notions of simulation.

Simulation type		Game type	Strategy complexity		Algorithmic complexity
			Player 1	Player 2	
Simulation		Safety games	memoryless		polynomial
Fair Simulation		Parity-3 games			
Quantitative Simulation	LIMAVG-autom.	Mean-payoff games			
	DISC $_\lambda$ -autom.	Discounted-sum games			
Quantitative Fair Simulation	LIMAVGINF-autom.	Implication games	infinite memory		pseudo-poly.
	LIMAVGSUP-autom.				
	DISC $_\lambda$ -autom.		may not exist		

In the first column we have the type of the simulation notion. The simulation of weighted automata without Büchi acceptance conditions is separated into simulation of LIMAVG- and DISC $_\lambda$ -automata, as the reduction to mean-payoff games for LIMAVGSUP-automata as well as LIMAVGINF-automata coincide, while it is different for DISC $_\lambda$ -automata. In contrast, simulation of weighted automata with Büchi acceptance conditions is separately considered for LIMAVGINF-, LIMAVGSUP- and DISC $_\lambda$ -automata. In the second column we have the game corresponding to the simulation notion of the first column. In the following two columns, we present the strategy complexity of the optimal strategies for the players, and the algorithmic complexity of the decision problem for the corresponding game.

Acknowledgement. This research was funded in part by the European Research Council (ERC) under grant agreement 267989 (QUAREM), by the Austrian Science Fund (FWF) project S11402-N23 (RiSE), FWF Grant No P23499- N23, FWF NFN Grant No S11407-N23 (RiSE), ERC Start grant (279307: Graph Games), and Microsoft faculty fellows award.

References

References

- [1] R. Milner, An algebraic definition of simulation between programs, in: IJCAI, 1971, pp. 481–489.
- [2] R. Alur, T. A. Henzinger, Computer-aided verification, unpublished, available at <http://www.cis.upenn.edu/cis673/> (2004).
- [3] R. Alur, T. A. Henzinger, O. Kupferman, M. Y. Vardi, Alternating refinement relations, in: D. Sangiorgi, R. de Simone (Eds.), CONCUR '98: Concurrency Theory, 9th International Conference, Nice, France, September 8-11, 1998, Proceedings, Vol. 1466 of Lecture Notes in Computer Science, Springer, 1998, pp. 163–178. doi:10.1007/BFb0055622.
URL <http://dx.doi.org/10.1007/BFb0055622>
- [4] T. A. Henzinger, O. Kupferman, S. K. Rajamani, Fair simulation, *Inf. Comput.* 173 (1) (2002) 64–81. doi:10.1006/inco.2001.3085.
URL <http://dx.doi.org/10.1006/inco.2001.3085>
- [5] K. Chatterjee, L. Doyen, T. A. Henzinger, Quantitative languages, *ACM Trans. Comput. Log.* 11 (4). doi:10.1145/1805950.1805953.
URL <http://doi.acm.org/10.1145/1805950.1805953>
- [6] K. Chatterjee, L. Doyen, T. A. Henzinger, Expressiveness and closure properties for quantitative languages, in: Proceedings of the 24th Annual IEEE Symposium on Logic in Computer Science, LICS 2009, 11-14 August 2009, Los Angeles, CA, USA, 2009, pp. 199–208. doi:10.1109/LICS.2009.16.
URL <http://dx.doi.org/10.1109/LICS.2009.16>
- [7] U. Boker, K. Chatterjee, T. A. Henzinger, O. Kupferman, Temporal specifications with accumulative values, in: Proceedings of the 26th Annual IEEE Symposium on Logic in Computer Science, LICS 2011, June 21-24, 2011, Toronto, Ontario, Canada, 2011, pp. 43–52. doi:10.1109/LICS.2011.33.
URL <http://dx.doi.org/10.1109/LICS.2011.33>
- [8] R. Bloem, K. Chatterjee, T. A. Henzinger, B. Jobstmann, Better quality in synthesis through quantitative objectives, in: A. Bouajjani, O. Maler (Eds.), Computer Aided Verification, 21st International Conference, CAV 2009, Grenoble, France, June 26 - July 2, 2009. Proceedings, Vol. 5643 of Lecture Notes in Computer Science, Springer, 2009, pp. 140–156. doi:10.1007/978-3-642-02658-4.
URL <http://dx.doi.org/10.1007/978-3-642-02658-4>
- [9] A. Bohy, V. Bruyère, E. Filiot, J. Raskin, Synthesis from LTL specifications with mean-payoff objectives, in: Tools and Algorithms for the Construction and Analysis of Systems - 19th International Conference, TACAS 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings, 2013, pp. 169–184.
- [10] N. Immerman, Number of quantifiers is better than number of tape cells, *Journal of Computer and System Sciences* 22 (1981) 384–406.
- [11] C. Beeri, On the membership problem for functional and multivalued dependencies in relational databases, *ACM Trans. on Database Systems* 5 (1980) 241–259.
- [12] M. Jurdzinski, Small progress measures for solving parity games, in: STACS'00, LNCS 1770, Springer, 2000, pp. 290–301.
- [13] K. Chatterjee, S. Chaudhary, P. Kamath, Faster algorithms for alternating refinement relations, in: CSL, LIPIcs 16, Schloss Dagstuhl, 2012, pp. 167–182.
- [14] U. Zwick, M. Paterson, The complexity of mean payoff games on graphs, *Theoretical Computer Science* 158 (1) (1996) 343–359.
- [15] A. Condon, The complexity of stochastic games, *Inf. Comput.* 96 (2) (1992) 203–224. doi:10.1016/0890-5401(92)90048-K.
URL [http://dx.doi.org/10.1016/0890-5401\(92\)90048-K](http://dx.doi.org/10.1016/0890-5401(92)90048-K)

- [16] K. Chatterjee, L. Doyen, H. Edelsbrunner, T. A. Henzinger, P. Rannou, Mean-payoff automaton expressions, in: CONCUR 2010 - Concurrency Theory, 21th International Conference, CONCUR 2010, Paris, France, August 31-September 3, 2010. Proceedings, 2010, pp. 269–283.
- [17] P. Cerný, T. A. Henzinger, A. Radhakrishna, Simulation distances, in: P. Gastin, F. Laroussinie (Eds.), CONCUR 2010 - Concurrency Theory, 21th International Conference, CONCUR 2010, Paris, France, August 31-September 3, 2010. Proceedings, Vol. 6269 of Lecture Notes in Computer Science, Springer, 2010, pp. 253–268. doi:10.1007/978-3-642-15375-4-18.
URL <http://dx.doi.org/10.1007/978-3-642-15375-4-18>
- [18] K. Chatterjee, Y. Velner, Hyperplane separation technique for multidimensional mean-payoff games, in: Proceedings of the 24th International Conference on Concurrency Theory, CONCUR'13, Springer-Verlag, Berlin, Heidelberg, 2013, pp. 500–515.
- [19] D. A. Martin, Borel determinacy, *Annals of Mathematics* (1975) 363–371.
- [20] E. A. Emerson, C. S. Jutla, Tree automata, mu-calculus and determinacy (extended abstract), in: 32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991, IEEE Computer Society, 1991, pp. 368–377. doi:10.1109/SFCS.1991.185392.
URL <http://doi.ieeecomputersociety.org/10.1109/SFCS.1991.185392>
- [21] A. Ehrenfeucht, J. Mycielski, Positional strategies for mean payoff games, *International Journal of Game Theory* 8 (2) (1979) 109–113.
- [22] K. Chatterjee, T. A. Henzinger, M. Jurdzinski, Mean-payoff parity games, in: 20th IEEE Symposium on Logic in Computer Science (LICS 2005), 26-29 June 2005, Chicago, IL, USA, Proceedings, 2005, pp. 178–187.
- [23] Y. Velner, Robust multidimensional mean-payoff games are undecidable, CoRR abs/1410.5703.
URL <http://arxiv.org/abs/1410.5703>
- [24] R. Brenguier, J.-F. Raskin, Optimal Values of Multidimensional Mean-Payoff Games, Research report, Université Libre de Bruxelles (U.L.B.), Belgium, 28 pages (Oct. 2014).
URL <https://hal.archives-ouvertes.fr/hal-00977352>
- [25] E. Kopczynski, Half-positional determinacy of infinite games, in: Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II, 2006, pp. 336–347.
- [26] Y. Velner, K. Chatterjee, L. Doyen, T. A. Henzinger, A. Rabinovich, J. Raskin, The complexity of multi-mean-payoff and multi-energy games, CoRR abs/1209.3234.
URL <http://arxiv.org/abs/1209.3234>
- [27] K. Chatterjee, M. Henzinger, An $O(n^2)$ time algorithm for alternating Büchi games, in: SODA, 2012, pp. 1386–1399.
- [28] K. Chatterjee, M. Henzinger, Efficient and dynamic algorithms for alternating büchi games and maximal end-component decomposition, *J. ACM* 61 (3) (2014) 15. doi:10.1145/2597631.
URL <http://doi.acm.org/10.1145/2597631>