The Cost of Exactness in Quantitative Reachability*

Krishnendu Chatterjee[†] Laurent Doyen[§] Thomas A. Henzinger[†]

[†] IST Austria

[§] CNRS & LSV, ENS Paris-Saclay, U. Paris-Saclay

Abstract. In the analysis of reactive systems a quantitative objective assigns a real value to every trace of the system. The value decision problem for a quantitative objective requires a trace whose value is at least a given threshold, and the exact value decision problem requires a trace whose value is exactly the threshold. We compare the computational complexity of the value and exact value decision problems for classical quantitative objectives, such as sum, discounted sum, energy, and mean-payoff for two standard models of reactive systems, namely, graphs and graph games.

1 Introduction

The formal analysis of reactive systems is a fundamental problem in computer science. Traditionally the analysis focuses on correctness properties, where a Boolean objective classifies the traces of the reactive system as either correct or incorrect. Recently there has been significant interest in the performance analysis of reactive systems as well as the analysis of reactive systems in resource-constrained environments such as embedded systems. In such scenarios quantitative objectives are necessary. A quantitative objective assigns a real value to every trace of the system which measures how desirable the trace is.

Given a reactive system and a quantitative objective, we consider two variants of the decision problem. First, the value decision problem for a quantitative objective requires a trace whose value is at least a given threshold. Second, the exact value decision problem requires a trace whose value is exactly the threshold.

Based on the length of the traces to be analyzed, quantitative objectives can be classified into three categories as follows: (a) infinite-horizon objectives where traces of infinite length are considered; (b) finite-horizon objectives where traces of a given bounded length are considered; (c) indefinite-horizon objectives where, given source and target vertices of the system, traces starting at the source and ending at the target are considered. While infinite-horizon and finite-horizon objectives have been traditionally studied, indefinite-horizon objectives

* This research was supported in part by the Austrian Science Fund (FWF) under grants S11402-N23 and S11407-N23 (RiSE/SHiNE), and Z211-N23 (Wittgenstein Award), ERC Start grant (279307: Graph Games), Vienna Science and Technology Fund (WWTF) through project ICT15-003.

are natural in many applications, such as robotics, where the robot must reach a goal state while optimizing the cost of the path [4].

In this work, we focus on two finite-state models of reactive systems, namely, graphs and graph games. Every transition of the system is assigned an integer-valued weight representing a reward (or cost). We consider three classical quantitative objectives, which are variants of the sum of the weights: (i) the standard sum of the weights, (ii) the discounted sum of the weights, and (iii) the energy objective, which is the sum but requires that all partial sums along the trace are non-negative. We study the computational complexity of the value and exact-value decision problems for the indefinite-horizon case of the above three quantitative objectives, both for graphs and games. We also distinguish whether the numbers are represented in unary and binary. We show how to extend and adapt existing results from the literature to obtain a comprehensive picture about the computational complexity of the problems we study. The results are summarized in Table 1 for graphs and Table 2 for graph games.

Related works. The value decision problem for quantitative objectives has been extensively studied for graphs and games. For the finite-horizon case the standard solution is the value iteration (or dynamic programming) approach [17, 27]. For the infinite-horizon case there is a rich literature: for mean-payoff objectives in graphs [25] and games [16, 19, 30, 8], for energy objectives in graphs and games [9, 5, 8], for discounted-sum objectives in graphs [1] and games [30, 21]. The exact value decision problem represents an important special case of the problem where there are multiple objectives. The multiple objectives problem has been studied for mean-payoff and energy objectives [29, 14, 24]. For discounted-sum objectives the problem has been studied in other contexts (such as for randomized selection of paths) [13, 12]. The special case of multiple objectives defined using a single quantitative function leads to interval objectives [23]. While finite-horizon and infinite-horizon problems have been studied for graphs and games, the indefinite-horizon problem has been studied mainly in artificial intelligence and robotics for different models (such as partially-observable MDPs) [4, 11, 10]. In this work we present a comprehensive study for indefinite-horizon objectives in graphs and games.

2 Preliminaries

A *weighted graph* $G = \langle V, E, w \rangle$ consists of a finite set V of vertices, a set $E \subseteq V \times V$ of edges, and a function $w : E \to \mathbb{Z}$ that assigns an integer weight to each edge of the graph. In the sequel, we consider weights encoded in unary, as well as in binary.

A *path* in G is a sequence $\rho = v_0 v_1 \dots v_k$ such that $(v_i, v_{i+1}) \in E$ for all $0 \le i < k$. We say that ρ is a path from v_0 to v_k. Given two vertices $s, t \in V$, we denote by $\mathsf{Paths}(s, t)$ the set of all paths from s to t in G (we assume that the graph G is clear from the context). A *prefix* of ρ is a sequence $\rho[0 \dots j] = v_0 v_1 \dots v_j$ where $j \le k$. We denote by $\mathsf{Pref}(\rho)$ the set of all prefixes of ρ.

	≥ 0		$= 0$	
	unary	binary	unary	binary
Sum	PTIME		PTIME	NP-c
Disc_λ	PTIME		Decidability is open Finite-path hard	
Energy	PTIME		PTIME	NP-c

Table 1. The complexity of the quantitative (s, t)-reachability problem for graphs, for threshold and exact value, with weights encoded in unary or in binary.

The total weight of ρ is defined by $\mathsf{Sum}(\rho) = \sum_{i=0}^{k-1} w(v_i, v_{i+1})$, and given a discount factor $0 < \lambda < 1$, the discounted sum of ρ is $\mathsf{Disc}_\lambda(\rho) = \sum_{i=0}^{k-1} \lambda^i \cdot w(v_i, v_{i+1})$. Note that for $\lambda = 1$, we have $\mathsf{Disc}_1(\rho) = \mathsf{Sum}(\rho)$. In the sequel, we consider a rational discount factor represented by two integers encoded like the weights in the graph (in unary or in binary). A winning condition is a set of paths. We consider the following winning conditions, which contain paths from s to t satisfying quantitative constraints. For $\sim \in \{=, \geq\}$, define

- $\mathsf{Sum}^{\sim 0}(s, t) = \{\rho \in \mathsf{Paths}(s, t) \mid \mathsf{Sum}(\rho) \sim 0\}$,
- $\mathsf{Disc}_\lambda^{\sim 0}(s, t) = \{\rho \in \mathsf{Paths}(s, t) \mid \mathsf{Disc}_\lambda(\rho) \sim 0\}$,
- $\mathsf{Energy}^{\sim 0}(s, t) = \{\rho \in \mathsf{Sum}^{\sim 0}(s, t) \mid \mathsf{Sum}(\rho') \geq 0 \text{ for all } \rho' \in \mathsf{Pref}(\rho)\}$.

Note that the energy condition is a variant of the sum requiring that all partial sums are nonnegative. For example, $\mathsf{Sum}^{=0}(s, t)$ is the set of all paths from s to t with a total weight equal to 0, and $\mathsf{Energy}^{=0}(s, t)$ are all such paths that maintain the total weight nonnegative along all their prefixes. Note that the energy condition is the same as the requirement that counters remain nonnegative used in VASS (Vector Addition Systems with States) and counter automata [2].

Definition 1 (Quantitative (s, t)-reachability problem for graphs). *The quantitative (s, t)-reachability problem for graphs asks, given a graph G and a winning condition $\varphi \in \{\mathsf{Sum}^{\sim 0}, \mathsf{Disc}_\lambda^{\sim 0}, \mathsf{Energy}^{\sim 0}\}$, whether the set $\varphi(s, t)$ is nonempty.*

3 Graphs

In this section we assume without loss of generality that there is no incoming edge in vertex s, and no outgoing edge from vertex t. We discuss the details of the complexity results for graphs.

Theorem 1. *The complexity bounds for the quantitative (s, t)-reachability problem for graphs are shown in Table 1.*

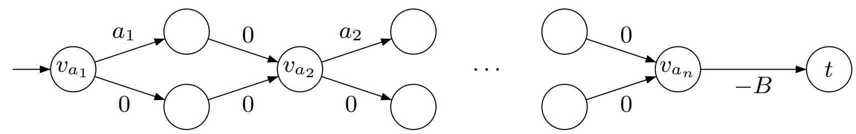

Fig. 1. Reduction from the subset-sum problem with $A = \{a_1, \ldots, a_n\}$ for the NP-hardness result of $\mathsf{Sum}^{=0}$ in graphs ($s = v_{a_1}$).

3.1 Results for Sum

Results for $\mathsf{Sum}^{\geq 0}$ The problem asks whether there exists a path from s to t of total weight at least 0. We can compute the longest path between s and t using Bellman-Ford algorithm, which detects positive cycles (the algorithm is the same as for finding a shortest path with opposite sign of the weights). Hence the problem is in PTIME for weights encoded in binary (and thus for weights encoded in unary as well).

Results for $\mathsf{Sum}^{=0}$ The problem asks whether there exists a path from s to t of total weight exactly 0. A pseudo-polynomial algorithm is known for this problem [26, Theorem 6]. Therefore, the problem is in PTIME for weights encoded in unary. It is also known that the problem is NP-complete for weights encoded in binary [26, Theroem 1,Theorem 9]. The NP upper bound is obtained by a reduction to integer linear programming (ILP) over variables x_e ($e \in E$) that represent the number of times edge e is used in a path from s to t, and where the ILP constraints require that for every vertex v, the number of incoming edges in v is equal to the number of outgoing edges from v (except for the source and target nodes s and t). The solution of the ILP should form a strongly connected component when a back-edge is added from t to s, which can be checked in polynomial time. The NP lower bound is obtained by a reduction from the subset sum problem, which asks, given a finite set $A \subseteq \mathbb{N}$ and a number $B \in \mathbb{N}$, whether there exists a subset $Z \subseteq A$ such that $\Sigma_{z \in Z} z = B$ (the sum of the elements of Z is B). The reduction, illustrated in Fig. 1, consists in constructing a graph in which there is a vertex v_a for each $a \in A$, and from v_a there are two outgoing edges, one with weight a, the other with weight 0. The two edges lead to intermediate vertices from which there is one edge with weight 0 to the vertex $v_{a'}$ (where a' is the successor of a in some total order over A). From the last vertex v_a, there is an edge to t with weight $-B$. The answer to the subset sum problem is Yes if and only if there is a path of total weight 0 from the first vertex v_a to t.

3.2 Results for Disc_λ

Results for $\mathsf{Disc}_\lambda^{\geq 0}$ We present a polynomial-time algorithm for weights and discount factor encoded in binary (thus also for weights and discount factor encoded in unary). First we compute the co-reachable vertices in the graph, namely the set $\mathsf{coReach}(t) = \{v \in V \mid \mathsf{Paths}(v,t) \neq \varnothing\}$ of vertices from which there exists a path to t, and we consider the graph $G' = \langle V \cap \mathsf{coReach}(t), E \cap (\mathsf{coReach}(t) \times \mathsf{coReach}(t))\rangle$ in which the vertex t has a self-loop with weight 0.

Then, we compute for each vertex $v \in \mathsf{coReach}(t)$, the largest value of the discounted sum of an infinite path from v in G'. The discounted sum of an infinite path $v_0 v_1 \dots \in V^\omega$ is $\sum_{i=0}^{\infty} \lambda^i \cdot w(v_i, v_{i+1})$. Note that the series converges because $\lambda < 1$ and the weights are bounded. The largest discounted sum of a path from a given vertex can be computed in polynomial time using linear programming [1, Section 3.1].

We consider the following cases:

- If the value $val(s)$ in the source vertex s is strictly greater than 0, then the answer to the (s,t)-reachability problem is Yes. Indeed, consider a prefix ρ' of length n of an optimal path ρ from s, where n is such that $\frac{2\lambda^n}{1-\lambda} \cdot W < val(s)$ (where W is the largest weight of G' in absolute value). Then it is easy to show that ρ' can be continued to a path that reaches t with positive weight.
- If $val(s) < 0$ then the answer to the (s,t)-reachability problem is No, as all finite paths from s to t have negative value (otherwise, there would be an infinite path with value at least 0, by prolonging the path through the self-loop on t).
- If $val(s) = 0$ then consider the graph G'' obtained from G' by keeping only the optimal edges, where an edge $e = (v, v')$ is *optimal* if $val(v) = w(v, v') + \lambda \cdot val(v')$. The answer to the (s,t)-reachability problem is Yes if and only if there is a path from s to t in G'', which can be computed in polynomial time. Indeed, if there exists a path from s to t with discounted sum equal to 0, then this path is optimal for the infinite-path problem since $val(s) = 0$, and therefore it uses only optimal edges. Moreover, if an infinite path from s uses only optimal edges, then it has value $val(s) = 0$, thus if such a path reaches t, then it gives a solution to the problem since from t the only outgoing edge is a self-loop with weight 0.

Results for $\mathsf{Disc}_\lambda^{=0}$ The decidability of the problem is open. Note that the decidability of the problem of finding an infinite path with exact discounted sum 0 is also open [3].

3.3 Results for Energy

Results for $\mathsf{Energy}^{\geq 0}$ The problem asks whether there exists a path from s to t that maintains the total weight (of all its prefixes) at least 0. We present a polynomial-time algorithm for weights encoded in binary (thus also for weights encoded in unary).

The algorithm relies on the fact that if there exists a path from s to t, then either the path is acyclic, or it contains a cycle, and that cycle needs to be positive (otherwise, we can remove the cycle and get an equally good path). Accordingly, the algorithm has two steps. First, we compute for each vertex $v \in V$ the largest

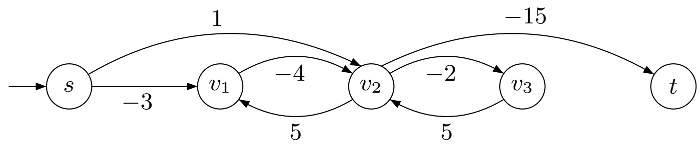

0	$-\infty$	$-\infty$	$-\infty$	$-\infty \; = \mu_0$
$\mathsf{Post}(\mu_0) = 0$	$-\infty$	1	$-\infty$	$-\infty \; = \mu_1$
$\mathsf{Post}(\mu_1) = 0$	6	1	$-\infty$	$-\infty \; = \mu_2$
$\mathsf{Post}(\mu_2) = 0$	6	2	$-\infty$	$-\infty \; = \mu_3$
$\mathsf{Post}(\mu_3) = 0$	7	2	0	$-\infty \; = \mu_4$
$\mathsf{Post}(\mu_4) = 0$	7	5	0	$-\infty \; = \mu_5$

$-\infty$	$-\infty$	5	$-\infty$	$-\infty \; = \tilde{\mu}_0$ (at v_2)
$\mathsf{Post}(\tilde{\mu}_0) = -\infty$	10	5	3	$-\infty \; = \tilde{\mu}_1$
$\mathsf{Post}(\tilde{\mu}_1) = -\infty$	10	8	3	$-\infty \; = \tilde{\mu}_2$
$\mathsf{Post}(\tilde{\mu}_2) = -\infty$	13	8	6	$-\infty \; = \tilde{\mu}_3$
$\mathsf{Post}(\tilde{\mu}_3) = -\infty$	13	11	6	$-\infty \; = \tilde{\mu}_4$
$\mathsf{Post}(\tilde{\mu}_4) = -\infty$	16	11	9	$-\infty \; = \tilde{\mu}_5$

Fig. 2. Sample of the fixpoint iterations to decide if there exists a path from s to t with energy (sum of weights) always at least 0.

total weight of a path from s to v (where the path must have all its prefixes nonnegative). To do that, we start from a function $\mu_0 : V \to \mathbb{N} \cup \{-\infty\}$ such that $\mu_0(s) = 0$ and $\mu_0(v) = -\infty$ for all $v \in V \setminus \{s\}$ and we iterate the operator $\mathsf{Post} : (V \to \mathbb{N} \cup \{-\infty\}) \to (V \to \mathbb{N} \cup \{-\infty\})$ defined as follows:

$$\mathsf{Post}(\mu)(v) = \max\{\mu(u) + w(u,v) \mid (u,v) \in E \wedge \mu(u) + w(u,v) \geq 0\} \cup \{\mu(v)\}$$

where $\max \varnothing = -\infty$. Consider $\mathsf{Post}^n(\mu_0)$, the nth iterate of Post on μ_0 where $n = |V|$. Intuitively, the value $\mathsf{Post}^n(\mu_0)(v)$ is the largest credit of energy (i.e., total weight) with which it is possible to reach v from s with a path of length at most n while maintaining the energy always nonnegative along the way. If $\mathsf{Post}^n(\mu_0)(t) \geq 0$, then the answer to the (s,t)-reachability problem is Yes. Otherwise, it means that there is no acyclic path from s to t that satisfies the energy constraint, and thus a positive cycle is necessary to reach t.

The second step of the algorithm is to check, for each vertex v with initial energy $\tilde{\mu}_0$ defined by $\tilde{\mu}_0(v) = \mathsf{Post}^n(\mu_0)(v)$ and $\tilde{\mu}_0(v') = -\infty$ for all $v' \neq v$, whether a positive cycle can be executed from v, that is whether $\mathsf{Post}^n(\tilde{\mu}_0)(v) > \tilde{\mu}_0(v)$. Note that only the vertices v such that $\mathsf{Post}^n(\mu_0)(v) \neq -\infty$ need to be considered. If there exists such a vertex from which t is reachable (without any constraint on the path from v to t), then the answer to the (s,t)-reachability problem is Yes. Otherwise, the answer to the problem is No because if there existed a path from s to t satisfying the energy constraint, it could not be an acyclic path (by the result of the first step of the algorithm), and it could not contain a cycle because (i) there is no positive cycle that can be reached from s and executed (by the result of the second step of the algorithm), and (ii) all negative cycles can be removed from the path to obtain a simpler, eventually acyclic, path that satisfies the energy constraint, which is impossible, as shown above. It is easy to see that the above computation can be done in polynomial time.

Consider the weighted graph with five vertices in Fig. 2. The graph has two cycles around v_2, both are positive. The vertex t is reachable from s, but in order to maintain the energy level nonnegative, we need to go through a cycle around v_2 which increases the energy level and allows to take the transition from v_2 to t with weight -15.

The algorithm first computes for each vertex the largest energy level that can be obtained by a path of length 5 from s (while maintaining the energy level always nonnegative). The result is shown in Fig. 2 as $\mu_5 = \mathsf{Post}^5(\mu_0)$. Note that $\mu_5(t) = -\infty$, thus there is no acyclic path from s to t with nonnegative energy level. The second step of the algorithm is a positive cycle detection, from each vertex of the graph. The computation from v_2 is illustrated in Fig. 2. Since the value at v_2 has strictly increased, a positive cycle is detected, and since t is reachable from v_2 (even if it is by a negative path), we can reach t from s (through v_2) with energy always at least 0.

Results for $\mathsf{Energy}^{=0}$ The problem is in PTIME for weights encoded in unary, as this a reachability problem in VASS (Vector Addition Systems with States) of dimension one, which is known to be NL-complete [2, Section 1]. The problem is NP-complete for weights encoded in binary, as this is exactly the reachability problem for one-counter automata, which is NP-complete [20, Proposition 1, Theorem 1]. The NP upper bound follows since one-counter automata allow zero-tests along the execution, and the NP lower bound holds even without zero-tests using a reduction from the subset-sum problem (see also Fig. 1).

4 Games

A *game* consists of a weighted graph $G = \langle V, E, w\rangle$ where V is partitioned into the sets V_1 of player-1 vertices, and the set V_2 of player-2 vertices. We assume that player-1 vertices and player-2 vertices alternate, i.e., $E \subseteq (V_1 \times V_2) \cup (V_2 \times V_1)$. This incurs no loss of generality as we can insert intermediate vertices along every transition that does not 'alternate'. We also assume that every vertex has a successor, that is for all $v \in V$ there exists $v' \in V$ such that $(v, v') \in E$.

A *strategy* of player 1 is a function $\sigma : V^*V_1 \to V$ such that $(v, \sigma(\rho \cdot v)) \in E$ for all $\rho \in V^*$ and all $v \in V_1$. A strategy σ is *memoryless* if it depends on the last vertex only, that is $\sigma(\rho \cdot v) = \sigma(\rho' \cdot v)$ for all $\rho, \rho' \in V^*$ and all $v \in V_1$. Given

	≥ 0		$= 0$	
	unary	binary	unary	binary
Sum	PTIME	NP $\cap$ coNP	PSPACE-c	EXPSPACE-c
Disc_λ	PTIME	NP $\cap$ coNP[1]	Decidability is open Finite-path hard	
Energy	PTIME	NP $\cap$ coNP	PSPACE-c	EXPSPACE-c

Table 2. The complexity of the quantitative (s,t)-reachability problem for games, for threshold and exact value, with weights encoded in unary or in binary.

an initial vertex v, and a strategy σ of player 1, we say that an infinite path $\rho = v_0 v_1 \ldots$ is an *outcome* of σ from v if $v_0 = v$ and $\sigma(v_0 \ldots v_j) = v_{j+1}$ for all $j \geq 0$ such that $v_j \in V_1$. We denote by $\mathsf{Outcome}_v^\omega(\sigma)$ the set of all outcomes of strategy σ from vertex v.

Definition 2 (Quantitative (s,t)-reachability problem for games). *The quantitative (s,t)-reachability problem for games asks, given a game G and a winning condition $\varphi \in \{\mathsf{Sum}^{\sim 0}, \mathsf{Disc}_\lambda^{\sim 0}, \mathsf{Energy}^{\sim 0}\}$, whether there exists a strategy σ of player 1 such that for all outcomes $\rho \in \mathsf{Outcome}_s^\omega(\sigma)$ there exists a prefix of ρ that belongs to the set $\varphi(s,t)$.*

Theorem 2. *The complexity bounds for the quantitative (s,t)-reachability problem for games are shown in Table 2.*

We now discuss the details of the results.

4.1 Results for Sum

Results for $\mathsf{Sum}^{\geq 0}$ The game problem for $\mathsf{Sum}^{\geq 0}$ is also known as the max-cost reachability problem. The problem is in NP $\cap$ coNP for weights encoded in binary [18, Theorem 5.2]. The result of [18, Theorem 5.2] holds for the winning conditions defined as a strict threshold, namely $\mathsf{Sum}^{>0}(s,t) = \{\rho \in \mathsf{Paths}(s,t) \mid \mathsf{Sum}(\rho) > 0\}$, and the same proof idea works for non-strict threshold. The result is obtained by a reduction to mean-payoff games [16, 30], which can be viewed as games where one player 1 wins if he can ensure that all cycles formed along a play are positive. Such games can be solved in NP $\cap$ coNP. The reduction constructs a mean-payoff game as a copy of the original game over the set of states from which player 1 can ensure to reach t, and adds an edge from t back to s with weight 0. Then player 1 can ensure a positive cycle if and only if he can ensure the objective $\mathsf{Sum}^{>0}(s,t)$: either he can reach t and loop through it (with a positive total weight), or he can ensure positive cycles that can be repeated until the total weight is sufficiently high to let him reach t while the

[1] The problem can be solved in PTIME if the weights in the graph are in binary, and the discount factor is in unary [21].

total weight remains positive. Conversely, if he can reach t with positive total weight, then he can win the mean-payoff game by repeatedly reaching t. Note that memoryless strategies are sufficient for player 2, but not for player 1 as he may need to accumulate weights along positive cycles before switching to the strategy that ensures reaching t.

It is not known whether mean-payoff games can be solved in polynomial time. The game problem for $\mathsf{Sum}^{\geq 0}$ is at least as hard as mean-payoff games, thus in the same status as mean-payoff games. This result is analogous to [7, Theorem 1(2)]. The idea of the reduction is, given a mean-payoff game G with initial vertex v, to construct a game G' from G by adding a transition with weight $-nW - 1$ from every player-1 vertex to a new vertex t where n is the number of vertices in G and W is the largest absolute weight in G. The reduction works because player-1 vertices and player-2 vertices alternate. The reduction is correct because if player 1 wins the mean-payoff game (with strict threshold), then he has a memoryless strategy to ensure that all reachable cycles are positive. Then, in G' player 1 can play the mean-payoff winning strategy long enough to accumulate total weight $nW+1$, and then use the transition with weight $-nW-1$ to reach t, and thus win in G'. In the other direction, if player 1 does not win the mean-payoff game, then player 2 can fix a memoryless strategy to ensure that all cycles are non-positive. Hence, the total weight of all finite prefixes of all outcomes is at most nW (the largest possible weight of an acyclic path), which is not sufficient to reach t, thus player 2 wins in G'.

For weights encoded in unary, the game problem for $\mathsf{Sum}^{\geq 0}$ can be solved in polynomial time using the algorithm of [7, Theorem 1(4)], a fixpoint iteration that relies on backward induction to compute the optimal cost for $i + 1$ rounds of the game, knowing the optimal cost for i rounds, similar to the pseudo-polynomial algorithm for solving mean-payoff games [30, 8]. The fixpoint iteration stops when the cost stabilizes to a finite value, or exceeds nW indicating that an arbitrary large cost can be achieved.

Results for $\mathsf{Sum}^{=0}$ The game problem for $\mathsf{Sum}^{=0}$ is a reachability problem where the target consists in both the vertex and the weight value. The problem was shown to be PSPACE-complete for weights encoded in unary [28, Theorem 5], and EXPSPACE-complete for weights encoded in binary [22, Theorem 1].

4.2 Results for Disc_λ

Results for $\mathsf{Disc}_\lambda^{\geq 0}$ The game problem for $\mathsf{Disc}_\lambda^{\geq 0}$ is in NP $\cap$ coNP for weights encoded in binary, by an argument similar to [18, Theorem 5.2] which shows the result for strict threshold (where the winning condition is the set of paths ρ from s to t such that $\mathsf{Disc}_\lambda(\rho) > 0$). The solution for strict threshold can be modified for non-strict threshold along the same idea as for the graph problem, thus by a reduction to infinite-horizon discounted sum games, which are solvable in NP $\cap$ coNP [30], and even in PTIME for unary encoding of the discount factor (even if the weights are encoded in binary) [21].

It is not known whether discounted sum games can be solved in polynomial time, and we show that discounted sum games reduce to the (s,t)-reachability problem for $\mathsf{Disc}_\lambda^{\geq 0}$. Given an infinite-horizon discounted sum game G, consider the game G' obtained from G by adding a vertex t, and edges (v,t) for all vertices v of player 1 in G (with weight 0). The reduction works because player-1 vertices and player-2 vertices alternate. Given the rational threshold ν for the game G, due to the separation of values in discounted sum games (which means that the optimal value in discounted sum games is the value of a play consisting of an acyclic prefix followed by a simple cycle, thus a rational number with denominator bounded by b^n, where b is the denominator of the discount factor λ, and n is the number of vertices in G) we can construct a number $\nu' < \nu$ such that if the optimal value in G is smaller than ν, then it is also smaller than ν'. The reduction produces the game G' with vertices s,t and threshold ν' (which can be replaced by threshold 0, by subtracting $(1-\lambda)\cdot\nu'$ to all weights). It is easy to see that (i) if player 1 can ensure discounted sum at least ν from an initial vertex s in G, then by playing sufficient long the optimal strategy from s, player 1 can ensure a value sufficiently close to ν to ensure reaching t with value at least ν'. Conversely, (ii) if player 1 does not win the discounted sum game G from s with threshold ν, then player 1 cannot win for threshold ν' and thus he cannot win in G' for (s,t)-reachability, which establishes the correctness of the reduction.

Results for $\mathsf{Disc}_\lambda^{=0}$ For $\mathsf{Disc}_\lambda^{=0}$, the decidability of the problem is open, as it is already open for graphs.

4.3 Results for Energy

Results for $\mathsf{Energy}^{\geq 0}$ For $\mathsf{Energy}^{\geq 0}$, the problem is inter-reducible with energy games: we consider infinite-horizon energy games where the winning condition for player 1 requires to maintain the total payoff (i.e., the energy) at least 0 along all prefixes of the (infinite) play, starting with initial energy 0. Memoryless strategies are sufficient for player 1 in energy games, and after fixing a memoryless strategy, all finite outcomes have nonnegative total weight thus all reachable simple cycles are nonnegative.

The reductions follow the same general ideas as between $\mathsf{Sum}^{\geq 0}$ and mean-payoff games, with some additional care. While nonnegative cycles are sufficient for player 1 in energy games, the reduction works only for the slightly stronger winning condition that asks player 1 to form only strictly positive cycles (while maintaining the energy condition on all acyclic outcomes as well). This stronger winning condition is equivalent to an energy condition in a modified graph where the weights are decreased by a value $\epsilon > 0$ where ϵ is sufficiently small to ensure that negative simple cycles remain negative (thus $n\epsilon < 1$). Moreover, since the initial energy 0 may now no longer be sufficient to survive the acyclic paths, we need to give a slightly positive initial energy value (by an initial transition of weight $n\epsilon$). Note that this initial energy does not allow player 1 to survive a

	≥ 0		$= 0$	
	unary	binary	unary	binary
Disc_λ	PTIME		Decidability is open Infinite-path hard	
$\overline{\mathsf{MP}}$, $\underline{\mathsf{MP}}$	PTIME		PTIME	

Table 3. The complexity of the infinite-horizon quantitative problem for graphs, for threshold and exact value, with weights encoded in unary or in binary.

negative finite prefix as $n\epsilon < 1$. We can take $\epsilon = \frac{1}{n+1}$ and scale up the weights by a factor $n+1$ to get integer weights. From this game graph with modified weights, we can use the same reductions as between $\mathsf{Sum}^{\geq 0}$ and mean-payoff games.

It follows that the problem has the same status as energy games with fixed initial credit, namely it is in NP $\cap$ coNP for weights encoded in binary, and in PTIME for weights encoded in unary [5, Proposition 12,Theorem 13].

Results for $\mathsf{Energy}^{=0}$ For $\mathsf{Energy}^{=0}$, the game problem is PSPACE-complete for weights encoded in unary [6, Theorem 11], and EXPSPACE-complete for weights encoded in binary [22, Theorem 1].

5 Survey of Infinite-horizon Quantitative Objectives

We present a survey of the computational complexity for the problem of satisfying a quantitative objective over an infinite duration, that requires an infinite trace with value either at least, or exactly a given threshold.

We consider winning conditions defined by the following quantitative measures over infinite paths (we denote by Paths^ω the set of all infinite paths in the graph G, where G is clear from the context), for $\sim \in \{=, \geq\}$:

- $\mathsf{Disc}_\lambda^{\sim 0} = \{\rho \in \mathsf{Paths}^\omega \mid \mathsf{Disc}_\lambda(\rho) \sim 0\}$,
- $\overline{\mathsf{MP}}^{\sim 0} = \{\rho \in \mathsf{Paths}^\omega \mid \limsup_{n\to\infty} \frac{1}{n} \cdot \mathsf{Sum}(\rho[0 \dots n]) \sim 0\}$,
- $\underline{\mathsf{MP}}^{\sim 0} = \{\rho \in \mathsf{Paths}^\omega \mid \liminf_{n\to\infty} \frac{1}{n} \cdot \mathsf{Sum}(\rho[0 \dots n]) \sim 0\}$,

The discounted sum is well defined for infinite paths (the infinite sum always exists). The $\overline{\mathsf{MP}}$ and $\underline{\mathsf{MP}}$ conditions are the mean-payoff objectives (see also Section 4.1), which are well defined as the limsup and liminf always exist, although the limit itself may not exist.

5.1 Results for graphs

We consider the infinite-horizon quantitative problem for graphs, which is to decide, given a graph G and a winning condition $\varphi \in \{\mathsf{Disc}^{\sim 0}, \overline{\mathsf{MP}}^{\sim 0}, \underline{\mathsf{MP}}^{\sim 0}\}$, whether the set φ is nonempty.

	≥ 0		$= 0$	
	unary	binary	unary	binary
Disc_λ	PTIME	NP $\cap$ coNP[2]	Decidability is open Infinite-path hard	
$\overline{\mathsf{MP}}$, $\underline{\mathsf{MP}}$	PTIME	NP $\cap$ coNP	PTIME	NP $\cap$ coNP

Table 4. The complexity of the infinite-horizon quantitative problem for games, for threshold and exact value, with weights encoded in unary or in binary.

Theorem 3. *The complexity bounds for the infinite-horizon quantitative problem for graphs are shown in Table 3.*

The results of Table 3 for discounted sum follow from the linear programming approach for computing the largest discounted sum of an infinite path [1, Section 3.1], and the decidability of the exact-value problem is open [3].

For mean-payoff, the infinite-horizon quantitative problem can be solved in polynomial time using Karp's algorithm to compute the reachable cycle with largest mean value, which runs in polynomial time [25]. For the exact-value problem, it is easy to see that the answer is Yes if and only if there exists a strongly connected component (scc) that contains both a nonnegative cycle and a nonpositive cycle. The path that reaches such an scc and then alternates between the two cycles (essentially repeating the nonnegative cycle until the partial sum of weights becomes positive, then switching to the nonpositive cycle until the partial sum of weights becomes negative, and so on) has mean-payoff value 0 (for both limsup and liminf) because the partial sum of the acyclic parts of the path (obtained by removing all cycles) is bounded by nW, where n is the number of vertices in G and W is the largest absolute weight in G. The scc decomposition and cycle with largest (resp., least) mean value can be computed in polynomial time.

5.2 Results for games

We consider the infinite-horizon quantitative problem for games, which is to decide, given a graph G, an initial vertex v, and a winning condition $\varphi \in \{\mathsf{Disc}^{\sim 0}, \overline{\mathsf{MP}}^{\sim 0}, \underline{\mathsf{MP}}^{\sim 0}\}$, whether there exists a strategy σ of player 1 such that $\mathsf{Outcome}^{\omega}_v(\sigma) \subseteq \varphi$.

Theorem 4. *The complexity bounds for the infinite-horizon quantitative problem for games are shown in Table 4.*

The results of Table 4 for discounted sum follow from the results of [30, 21, 3] (see also Section 4.2).

[2] The problem can be solved in PTIME if the weights in the graph are in binary, and the discount factor is in unary [21].

For mean-payoff, the results for the threshold problem follow from [16, 30] in particular there is a pseudo-polynomial algorithm for solving mean-payoff games [30, 8]. The NP $\cap$ coNP result for the exact-value problem follows from [23, Corollary 6], and the set Z of initial vertices from which player 1 has a winning strategy has the following characterization: from every vertex in Z, player 1 has a strategy to ensure nonnegative mean-payoff value, and player 1 has a (possibly different) strategy to ensure nonpositive mean-payoff value. Moreover if from some vertex player 1 does not have a strategy to ensure nonnegative (or nonpositive) mean-payoff value, then player 1 does not have a winning strategy from that vertex for the exact-value objective. By an argument analogous to the case of graphs, we can show that player 1 wins from every vertex in Z by switching between the strategies to ensure nonpositive and nonnegative mean-payoff value, because the partial sums will remain bounded by nW, thus the mean-payoff value is 0 (both for limsup and liminf).

We can compute the set Z by removing from the set V of vertices the vertices that are losing for player 1, iteratively as follows, until a fixpoint is obtained: at each iteration, remove the vertices where player 1 does not win either the non-positive or the nonnegative mean-payoff objective, and remove the vertices from which player 2 can ensure to reach an already removed vertex (this amounts to solving a reachability game, thus in polynomial time). The number of iterations is at most n, thus the algorithm is polynomial for weights in unary. Note that player 2 has a memoryless strategy from all removed vertices, to ensure that the mean-payoff value is not 0.

It follows that for weights in binary, the exact-value problem can be solved in NP by guessing the set Z and checking that from every vertex in Z player 1 wins the nonpositive mean-payoff objective as well as the nonnegative mean-payoff objective (possibly with a different strategy), and in coNP by guessing a memoryless winning strategy for player 2 in $V \setminus Z$ and solving in PTIME the exact-value problem for mean-payoff in graphs.

Note that the exact-value problem can be reduced to a two-dimensional mean-payoff objective, which is known to be solvable in NP $\cap$ coNP for $\overline{\mathsf{MP}}$, but only in coNP for $\underline{\mathsf{MP}}$ [29]. In contrast, the exact-value problem is solvable in NP as well for $\underline{\mathsf{MP}}$.

6 Conclusion

In this work we studied the complexity of the value decision problem and the exact-value decision problem for sum, discounted sum, and energy objectives for the indefinite-horizon case. We studied them for graphs and graph games, and also distinguished the representation of numbers in unary and binary. In several cases the exact decision problem is computationally harder as compared to the non-exact counterpart. An interesting direction of future work is to consider the problems we studied in other related models, such as stochastic games (extending the work of [15]), Markov decision processes, timed games, etc.

References

1. D. Andersson. An improved algorithm for discounted payoff games. In *Proc. 11th ESSLLI Student Session*, pages 91–98, 2006.
2. M. Blondin, A. Finkel, S. Göller, C. Haase, and P. McKenzie. Reachability in two-dimensional vector addition systems with states is PSPACE-complete. In *Proc. of LICS: Symposium on Logic in Computer Science*, pages 32–43. IEEE Computer Society, 2015.
3. U. Boker, T. A. Henzinger, and J. Otop. The target discounted-sum problem. In *Proc. of LICS: Symposium on Logic in Computer Science*, pages 750–761. IEEE Computer Society, 2015.
4. B. Bonet and H. Geffner. Solving POMDPs: RTDP-Bel vs. Point-based Algorithms. In *Proc. of IJCAI: International Joint Conference on Artificial Intelligence*, pages 1641–1646, 2009.
5. P. Bouyer, U. Fahrenberg, K. G. Larsen, N. Markey, and J. Srba. Infinite runs in weighted timed automata with energy constraints. In *Proc. of FORMATS: Formal Modeling and Analysis of Timed Systems*, LNCS 5215, pages 33–47. Springer, 2008.
6. T. Brázdil, P. Jancar, and A. Kucera. Reachability games on extended vector addition systems with states. In *Proc. of ICALP: Automata, Languages, and Programming*, LNCS 6199, pages 478–489. Springer, 2010.
7. T. Brihaye, G. Geeraerts, A. Haddad, and B. Monmege. To reach or not to reach? efficient algorithms for total-payoff games. In *Proc. of CONCUR: Concurrency Theory*, volume 42 of *LIPIcs*, pages 297–310. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.
8. L. Brim, J. Chaloupka, L. Doyen, R. Gentilini, and J.-F. Raskin. Faster algorithms for mean-payoff games. *Formal Methods in System Design*, 38(2):97–118, 2011.
9. A. Chakrabarti, L. de Alfaro, T. A. Henzinger, and M. Stoelinga. Resource interfaces. In *Proc. of EMSOFT: Embedded Software*, LNCS 2855, pages 117–133. Springer, 2003.
10. K. Chatterjee and M. Chmelik. Indefinite-horizon reachability in Goal-DEC-POMDPs. In *Proc. of ICAPS: International Conference on Automated Planning and Scheduling*, pages 88–96. AAAI Press, 2016.
11. K. Chatterjee, M. Chmelik, R. Gupta, and A. Kanodia. Optimal cost almost-sure reachability in POMDPs. *Artif. Intell.*, 234:26–48, 2016.
12. K. Chatterjee, V. Forejt, and D. Wojtczak. Multi-objective discounted reward verification in graphs and MDPs. In *Proc. of LPAR: Logic for Programming, Artificial Intelligence, and Reasoning*, LNCS 8312, pages 228–242. Springer, 2013.
13. K. Chatterjee, R. Majumdar, and T. A. Henzinger. Markov decision processes with multiple objectives. In *Proc. of STACS: Symposium on Theoretical Aspects of Computer Science*, LNCS 3884, pages 325–336. Springer, 2006.
14. K. Chatterjee and Y. Velner. Hyperplane separation technique for multidimensional mean-payoff games. In *Proc. of CONCUR: Concurrency Theory*, LNCS 8052, pages 500–515. Springer, 2013.
15. T. Chen, V. Forejt, M. Z. Kwiatkowska, A. Simaitis, A. Trivedi, and M. Ummels. Playing stochastic games precisely. In *Proc. of CONCUR: Concurrency Theory*, LNCS 7454, pages 348–363. Springer, 2012.
16. A. Ehrenfeucht and J. Mycielski. Positional strategies for mean payoff games. *Int. Journal of Game Theory*, 8(2):109–113, 1979.
17. J. Filar and K. Vrieze. *Competitive Markov Decision Processes.* Springer-Verlag, 1997.

18. E. Filiot, R. Gentilini, and J.-F. Raskin. Quantitative languages defined by functional automata. *Logical Methods in Computer Science*, 11(3), 2015.
19. V. A. Gurvich, A. V. Karzanov, and L. G. Khachiyan. Cyclic games and an algorithm to find minimax cycle means in directed graphs. *USSR Computational Mathematics and Mathematical Physics*, 28(5):85–91, 1988.
20. C. Haase, S. Kreutzer, J. Ouaknine, and J. Worrell. Reachability in succinct and parametric one-counter automata. In *Proc. of CONCUR: Concurrency Theory*, LNCS 5710, pages 369–383. Springer, 2009.
21. T. D. Hansen, P. B. Miltersen, and U. Zwick. Strategy iteration is strongly polynomial for 2-player turn-based stochastic games with a constant discount factor. *J. ACM*, 60(1):1:1–1:16, 2013.
22. P. Hunter. Reachability in succinct one-counter games. In *Proc. of RP: Reachability Problems*, LNCS 9328, pages 37–49. Springer, 2015.
23. P. Hunter and J.-F. Raskin. Quantitative games with interval objectives. In *Proc. of FSTTCS*, volume 29 of *LIPIcs*, pages 365–377. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2014.
24. M. Jurdzinski, R. Lazic, and S. Schmitz. Fixed-dimensional energy games are in pseudo-polynomial time. In *Proc. of ICALP (II): Automata, Languages, and Programming*, pages 260–272, 2015.
25. R. M. Karp. A characterization of the minimum cycle mean in a digraph. *Discrete Mathematics*, 23:309–311, 1978.
26. M. Nykänen and E. Ukkonen. The exact path length problem. *J. Algorithms*, 42(1):41–53, 2002.
27. M. L. Puterman. *Markov Decision Processes*. John Wiley and Sons, 1994.
28. J. Reichert. On the complexity of counter reachability games. In *Proc. of RP: Reachability Problems*, LNCS 8169, pages 196–208. Springer, 2013.
29. Y. Velner, K. Chatterjee, L. Doyen, T. A. Henzinger, A. M. Rabinovich, and J.-F. Raskin. The complexity of multi-mean-payoff and multi-energy games. *Inf. Comput.*, 241:177–196, 2015.
30. U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theor. Comput. Sci.*, 158(1&2):343–359, 1996.