

Infinite-Duration Bidding Games*

Guy Avni¹, Thomas A. Henzinger¹, and Ventsislav Chonev²

1 IST Austria

2 Max Planck Institute for Software Systems (MPI-SWS)

Abstract

Two-player games on graphs are widely studied in formal methods as they model the interaction between a system and its environment. The game is played by moving a token throughout a graph to produce an infinite path. There are several common modes to determine how the players move the token through the graph; e.g., in turn-based games the players alternate turns in moving the token. We study the *bidding* mode of moving the token, which, to the best of our knowledge, has never been studied in infinite-duration games. Both players have separate *budgets*, which sum up to 1. In each turn, a bidding takes place. Both players submit bids simultaneously, and a bid is legal if it does not exceed the available budget. The winner of the bidding pays his bid to the other player and moves the token. For reachability objectives, repeated bidding games have been studied and are called *Richman games* [36, 35]. There, a central question is the existence and computation of *threshold budgets*; namely, a value $t \in [0, 1]$ such that if Player 1's budget exceeds t , he can win the game, and if Player 2's budget exceeds $1 - t$, he can win the game. We focus on parity games and mean-payoff games. We show the existence of threshold budgets in these games, and reduce the problem of finding them to Richman games. We also determine the strategy-complexity of an optimal strategy. Our most interesting result shows that memoryless strategies suffice for mean-payoff bidding games.

1998 ACM Subject Classification J.4 Social and Behavioral Sciences, F.1.2 Modes of Computation

Keywords and phrases Bidding games, Richman games, Parity games, Mean-payoff games

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2017.17

1 Introduction

Two-player infinite-duration games on graphs are an important class of games as they model the interaction of a system and its environment. Questions about automatic synthesis of a reactive system from its specification [40] are reduced to finding a winning strategy for the “system” player in a two-player game. The game is played by placing a token on a vertex in the graph and allowing the players to move it throughout the graph, thus producing an infinite trace. The winner or value of the game is determined according to the trace. There are several common modes to determine how the players move the token that are used to model different types of systems (c.f., [4]). The most well-studied mode is *turn-based*, where the vertices are partitioned between the players and the player who controls the vertex on which the token is placed, moves it. Other modes include *probabilistic* and *concurrent* moves.

We study a different mode of moving, which we refer to as *bidding*, and to the best of our knowledge, has never been studied for infinite-duration games. Both players have *budgets*, where for convenience, we have $B_1 + B_2 = 1$. In each turn a bidding takes place for the right to move the token. The players submit bids simultaneously, where a bid is legal if it does not exceed the available budget. Thus, a bid is a real number in $[0, B_i]$, for $i \in \{1, 2\}$. The player who bids higher

* This research was supported in part by the Austrian Science Fund (FWF) under grants S11402-N23 (RiSE/SHiNE) and Z211-N23 (Wittgenstein Award).



pays the other player, and decides where the token moves. Draws can occur and one needs to devise a mechanism for resolving them (e.g., giving advantage to Player 1), and our results do not depend on a specific mechanism.

Bidding arises in many settings and we list several examples below. The players in a two-player game often model concurrent processes. Bidding for moving can model an interaction with a scheduler. The process that wins the bidding gets scheduled and proceeds with its computation. Thus, moving has a cost and processes are interested in moving only when it is critical. When and how much to bid can be seen as quantifying the resources that are needed for a system to achieve its objective, which is an interesting question. Other takes on this problem include reasoning about which input signals need to be read by the system at its different states [20, 2] as well as allowing the system to read chunks of input signals before producing an output signal [28, 27, 33]. Also, our bidding game can model *scrip systems* that use internal currencies for bidding in order to prevent “free riding” [31]. Such systems are successfully used in various settings such as databases [43], group decision making [42], resource allocation, and peer-to-peer networks (see [29] and references therein). Finally, repeated bidding is a form of a sequential auction [37], which is used in many settings including online advertising.

Recall that the winner or value of the game is determined according to the outcome, which is an infinite trace. There are several well-studied objectives in games. The simplest objective is *reachability*, where Player 1 has a target vertex and a trace is winning for him iff it visits the target. Bidding reachability games are equivalent to *Richman games* [36, 35], named after David Richman. Richman games are the first to study the bidding mode of moving. The central question that is studied on Richman games regards a *threshold budget*, which is a function $\text{THRESH} : V \rightarrow [0, 1]$ such that if Player 1’s budget exceeds $\text{THRESH}(v)$ at a vertex v , then he has a strategy to win the game. On the other hand, if Player 2’s budget exceeds $1 - \text{THRESH}(v)$, he can win the game (recall that the budgets add up to 1). In [36, 35], the authors show that threshold budgets exist, are unique, and that finding them is in NP. We slightly improve their result by showing that the problem is in NP and coNP. We illustrate the bidding model and the threshold problem in the following example.

► **Example 1.** Consider for example, the bidding reachability game that is depicted in Figure 1. Player 1’s goal is to reach t , and Player 2’s goal is to prevent this from happening. How much budget suffices for Player 1 to guarantee winning? Clearly, even if Player 1 has all the budget, he cannot win in v_1 , thus $\text{THRESH}(v_1) = 1$. Similarly, even if Player 2 has all the budget in t , Player 1 has already won, thus $\text{THRESH}(t) = 0$. We show a naive solution in which Player 1 wins when his budget exceeds 0.75. Indeed, if Player 1’s budget is $0.75 + \epsilon$, for $\epsilon > 0$, then since the budgets add up to 1, Player 2’s budget is $0.25 - \epsilon$. In the first turn, Player 1 bids $0.25 + \frac{\epsilon}{2}$ and wins the bidding since Player 2 cannot bid above 0.25. He pays his bid to Player 2 and moves the token to v_2 . Thus, at the end of the round, the budgets are $0.5 + \frac{\epsilon}{2}$ and $0.5 - \frac{\epsilon}{2}$ and the token is on v_2 . In the second bidding, Player 1 bids all his budget, wins the bidding since Player 2 cannot bid above 0.5, moves the token to t , and wins the game. It turns out that the threshold budgets are lower: it follows from Theorem 3 that they are $\text{THRESH}(v_0) = 2/3$ and $\text{THRESH}(v_2) = 1/3$. ◀

We introduce and study infinite duration bidding games with richer qualitative objectives as well as quantitative objectives. *Parity games* are an important class of qualitative games as the problem of reactive synthesis from LTL specifications is reduced to a parity game. The vertices in a parity game are labeled by an index in $\{0, \dots, d\}$, for some $d \in \mathbb{N}$, and an infinite trace is winning for Player 1 iff the parity of the maximal index that is visited infinitely often is odd. The quantitative games we focus on are *mean-payoff* games. An infinite outcome has a value, which can be thought of as the amount of money that Player 1 pays Player 2. Accordingly, we refer to the players in a mean-payoff game as *Maximizer* (Max, for short) and *Minimizer* (Min, for short). The vertices of a mean-payoff game are labeled by values in \mathbb{Z} . Consider an infinite trace π . The *energy* of a prefix

π^n of length n of π , denoted $E(\pi^n)$, is the sum of the values it traverses. The mean-payoff value of π is $\liminf_{n \rightarrow \infty} E(\pi^n)/n$. We are interested in cases where Min can guarantee a non-positive mean-payoff value. It suffices to show that he can guarantee that an infinite outcome π either has infinitely many prefixes with $E(\pi^n) = 0$, or that the energy is bounded, thus there is $N \in \mathbb{N}$ such that for every $n \in \mathbb{N}$, we have $E(\pi^n) \leq N$. We stress the point that there are two “currencies” in the game: a “monopoly money” that is used to determine who moves the token and which the players do not care about once the game ends, and the values on the vertices, which is the value that Min and Max seek to minimize and maximize, respectively. We illustrate mean-payoff games with the following example.



■ **Figure 1** On the left, a bidding reachability game. On the right, a bidding mean-payoff game where the weights are depicted on the edges.

► **Example 2.** Consider the mean-payoff bidding game that is depicted in Figure 1, where for convenience the values are placed on the edges and not on the vertices. We claim that Min has a strategy that guarantees a non-positive mean-payoff value. Without loss of generality, Max always chooses the 1-valued edge. Min’s strategy is a *tit-for-tat*-like strategy, and he always takes the (-1) -valued edge. The difficulty is in finding the right bids. Initially, Min bids 0. Assume Max wins a bidding with $b > 0$. Min will try and *match* this win: he bids b until he wins with it. Let b_1, \dots, b_n be Max’s winning bids before Min wins with b . We call these *un-matched* bids. The next bid Min attempts to match is $b' = \min_{1 \leq i \leq n} b_i$; he bids b' until he wins with it, and continues similarly until all bids are matched.

We claim that the tit-for-tat strategy guarantees a non-positive mean-payoff value. Observe first that if a prefix of the outcome has k unmatched bids, then the energy is k . In particular, if all bids are matched, the energy is 0. Suppose Min bids b . We claim that the number of un-matched bids is at most $\lceil 1/b \rceil$. Otherwise, since b is less than all other un-matched bids, Max would need to invest more than a budget of 1. It follows that an infinite outcome that never reaches energy level 0 has bounded energy, thus the mean-payoff value is non-positive. ◀

We study the existence and computation of threshold budgets in parity and mean-payoff bidding games. Also, we determine the strategy complexity that is necessary for winning. Recall that a winning strategy in a game typically corresponds to an implementation of a system. A strategy that uses an unbounded memory, like the tit-for-tat strategy above, is not useful for implementing. Thus, our goal is to find strategies that use little or no memory, which are known as *memoryless* strategies.

We show that parity bidding games are linearly-reducible to Richman games allowing us to obtain all the positive results from these games; threshold budgets exist, are unique, and computing them is no harder than for Richman games, i.e., the problem is in NP and coNP. We find this result quite surprising since for most other modes of moving, parity games are considerably harder than reachability games. The crux of the proof considers bottom strongly-connected components (BSCCs, for short) in the arena, i.e., SCCs with no exiting edges. We show that in a strongly connected bidding parity game, exactly one of the players wins with every initial budget, thus the threshold budgets of the vertices of a BSCC are in $\{0, 1\}$. If the vertex with highest parity in a BSCC is odd, then Player 1 wins, i.e., the threshold budgets are all 0, and otherwise Player 2 wins, i.e., the threshold budgets are all 1. We can thus construct a Richman game by setting the target of Player 1 to the BSCCs that are winning for him and the target of Player 2 to the ones that are winning for him. Moreover, we show that memoryless strategies are sufficient for winning in these games.

We proceed to study mean-payoff bidding games. We adapt the definition of threshold values; we say that $t \in [0, 1]$ is a threshold value for Min if with a budget that exceeds t , Min can guarantee a non-positive mean-payoff value. On the other hand, if Max's budget exceeds $1 - t$, he can guarantee a positive mean-payoff value. We show that threshold values exist and are unique in mean-payoff bidding games. The crux of the existence proof again considers the BSCCs of the game. We show that in a strongly-connected mean-payoff bidding game, the threshold budgets are in $\{0, 1\}$, thus again either Min "wins" or Max "wins" the game. Moreover, this classification can be determined in NP and coNP, thus the complexity of solving bidding mean-payoff games coincides with Richman games. Our results for strongly-connected games are obtained by developing the connection that was observed in [36, 35] between the threshold budget and the reachability probability in a probabilistic model on the same structure as the game. We show a connection between bidding mean-payoff games and *one-counter 2.5-player games* [14, 13] to prove the classification of BSCCs. In turn, these games are equivalent to discrete *quasi-birth-death processes* [24] and generalize *solvency games* [11], which can be thought of as a rewarded Markov decision process with a single vertex.

The classification above is existential in nature and does not provide any insight on how a player guarantees a mean-payoff value. Our most technically challenging results concern the constructions strategies for Min and Max. The challenging part of the construction is reasoning about strongly-connected bidding mean-payoff games. Consider a strongly-connected game in which Min can guarantee a non-positive mean-payoff value. The idea of our construction is to tie between changes in Min's budget with changes in the energy; *investing* one unit of budget (with the appropriate normalization) implies a decrease of a unit of energy, and on the other hand, an increase of a unit of energy implies a *gain* of one unit of budget. Since the budgets are bounded by 1, the value cannot increase arbitrarily. Finding the right bids in a general SCC is not trivial, and we find our solution to be surprisingly elegant. The case where Max can guarantee a positive mean-payoff value, is more challenging. Unlike a memoryless strategy for Min, the normalization factor must decrease as the value increases so that Max does not exhaust his budget. We show constant memory strategies in general and identify a fragment in which we show memoryless strategies.

Further bidding games Variants of bidding games where studied in the past. Already in [35] several variants are studied including a *poorman* version in which the winner of the bidding pays the bank, thus the amount of money in the game decreases as the game proceeds. Motivated by recreational games, e.g., bidding chess, *discrete bidding games* are studied in [23], where the money is divided into chips, so a bid cannot be arbitrarily small as in the bidding games we study. In *all-pay* bidding games [38], the players all pay their bids to the bank. Non-zero-sum two-player games were recently studied in [30]. They consider a bidding game on a directed acyclic graph. Moving the token throughout the graph is done by means of bidding. The game ends once the token reaches a sink, and each sink is labeled with a pair of payoffs for the two players that do not necessarily sum up to 0. They show existence of *subgame perfect equilibrium* for every initial budget and a polynomial algorithm to compute it.

Due to lack of space, most of the proofs appear in the full version [5].

2 Preliminaries

An *arena* is a pair $\langle G, \alpha \rangle$, where G is a directed graph and α is an objective. A game is played on an arena as follows. A token is placed on a vertex in the arena and the players move it throughout the graph. The *outcome* is an infinite path π . The winner or value is determined according to π and α as we elaborate below. There are several common modes in which the players move the token. In *turn-based games* the vertices are partitioned between the players and the player who controls the vertex on which the token is placed, moves it. Another mode is *probabilistic* choices, where the

game can be thought of as a *Markov chain*, thus the edges are labeled with probabilities, and the edge on which the token proceeds is chosen randomly. A combination of these two modes is called *2.5-player games*, where the vertices are partitioned into three sets: Player 1 vertices, Player 2 vertices, and probabilistic vertices. Finally, in *concurrent games*, each player has a possible (typically finite) set of actions he can choose from in a vertex. The players select an action simultaneously, and the choice of actions dictates to which vertex the token moves.

We study a different mode of moving, which we call *bidding*. Both players have budgets, where for convenience, we have $B_1 + B_2 = 1$. In each turn, a bidding takes place to determine who moves the token. Both players submit bids simultaneously, where a bid is a real number in $[0, B_i]$, for $i \in \{1, 2\}$. The player who bids higher pays the other player and decides where the token moves. Note that the sum of budgets always remains 1. While draws can occur, in the questions we study we try avoid the issue of draws.

A *strategy* prescribes to a player which *action* to take in a game, given a finite *history* of the game, where we define these two notions below. In 2.5-player games, histories are paths and actions are vertices. Thus, a strategy for Player i , for $i \in \{1, 2\}$, takes a finite path that ends in a Player i vertex, and prescribes to which vertex the token moves to next. In bidding games, histories and strategies are more complicated as they maintain the information about the bids and winners of the bids. A history is a sequence of the form $v_0, \langle v_1, b_1, i_1 \rangle, \langle v_2, b_2, i_2 \rangle, \dots, \langle v_k, b_k, i_k \rangle \in V \cdot (V \times [0, 1] \times \{1, 2\})^*$, where, for $j \geq 1$, in the j -th round, the token is placed on vertex v_{j-1} , the winning bid is b_j , and the winner is Player i_j , and Player i_j moves the token to vertex v_j . An action for a player is $\langle b, v \rangle \in ([0, 1] \times V)$, where b is the bid and v is the vertex to move to upon winning. An initial vertex v_0 and strategies f_1 and f_2 for Players 1 and 2, respectively, determine a unique *outcome* π for the game, denoted $out(v_0, f_1, f_2)$, which is an infinite sequence in $V \cdot (V \times [0, 1] \times \{1, 2\})^\omega$. We sometimes abuse notation and refer to $out(v_0, f_1, f_2)$ as a finite prefix of the infinite outcome. We drop v_0 when it is clear from the context. We define the outcome inductively. The first element of the outcome is v_0 . Suppose π_1, \dots, π_j is defined. The players bids are given by $\langle b_1, v_1 \rangle = f_1(\pi_1, \dots, \pi_j)$ and $\langle b_2, v_2 \rangle = f_2(\pi_1, \dots, \pi_j)$. If $b_1 > b_2$, then $\pi_{j+1} = \langle v_1, b_1, 1 \rangle$, and dually when $b_1 < b_2$, we have $\pi_{j+1} = \langle v_2, b_2, 2 \rangle$. We assume there is some tie-breaking mechanism that determines who the winner is when $b_1 = b_2$, and our results are not affected by what the tie-breaking mechanism is. Consider a finite outcome π . The *payment* of Player 1 in π , denoted $\mathcal{B}_1(\pi)$, is $\sum_{1 \leq j \leq |\pi|} (-1)^{3-i_j} b_j$, and Player 2's payment, denoted $\mathcal{B}_2(\pi)$ is defined similarly. For $i \in \{1, 2\}$, consider an initial budget $B_i^{init} \in [0, 1]$ for Player i . A strategy f is *legal* for Player i with respect to B_i^{init} if for every $v_0 \in V$ and strategy g for the other player, Player i 's bid in a finite outcome $\pi = out(v_0, f, g)$ does not surpass his budget. Thus, for $\langle b, v \rangle = f(\pi)$, we have $b \leq B_i^{init} - \mathcal{B}_i(\pi)$.

Richman games and threshold budgets The simplest qualitative objective is reachability: Player 1 has a target vertex v_R and an infinite outcome is winning for him if it visits v_R . Reachability bidding games are known as Richman games [36, 35]. In Richman games both players have a target, which we denote by v_R and v_S . The game ends once one of the targets is reached. Note that this definition is slightly different from standard reachability games since there, Player 2 has no target and his goal is to keep the game from v_R . Though, we show that for our purposes, since Richman games have no ties, reachability games are equivalent to Richman games (see Lemma 4).

The central question that is studied on bidding games regards a *threshold budget*. A threshold budget is a function $THRESH : V \rightarrow [0, 1]$ such that if Player 1's budget exceeds $THRESH(v)$ at a vertex v , then he has a strategy to win the game. On the other hand, if Player 2's budget exceeds $1 - THRESH(v)$, he can win the game. We sometimes use $THRESH_1(v)$ to refer to $THRESH(v)$ and $THRESH_2(v)$ to refer to $1 - THRESH(v)$. We formalize the problem of finding threshold budgets as a decision problem. We define the THRESH-BUDG problem, which takes as input a bidding game

\mathcal{G} , a vertex v , and a value $t \in [0, 1]$, and the goal is to decide whether $\text{THRESH}(v) = t$.

Threshold values are shown to exist in [36] as well as how to compute them. We review briefly their results. Consider a Richman game $\mathcal{G} = \langle V, E, v_R, v_S \rangle$. We define the *Richman function* as follows. We first define $R(v, i)$, for $i \in \mathbb{N} \cup \{0\}$, where the intuition is that if Player 1's budget exceeds $R(v, i)$, he can win in at most i steps. We define $R(v_R, 0) = 0$ and $R(v, 0) = 1$ for every other vertex $v \in V$. Indeed, Player 1 can win in 0 steps from v_R no matter what his initial budget is, and even if he has all the budget, he cannot win in 0 steps from anywhere else. Consider $i \in \mathbb{N}$ and $v \in V$. We denote by $\text{adj}(v) \subseteq V$, the adjacent vertices to v , so $u \in \text{adj}(v)$ iff $E(v, u)$. Let v^+ be the vertex that maximizes the expression $\max_{u \in \text{adj}(v)} R(u, i - 1)$, and let v^- be the vertex that minimizes the expression $\min_{u \in \text{adj}(v)} R(u, i - 1)$. We define $R(v, i) = \frac{1}{2}(R(v^+, i - 1) + R(v^-, i - 1))$. We define $R(v) = \lim_{i \rightarrow \infty} R(v, i)$. The following theorem shows that $R(v)$ equals $\text{THRESH}(v)$, and throughout the paper we use them interchangeably. We give the proof of the theorem for completeness.

► **Theorem 3.** [36] *For every $v \in V$, we have $\text{THRESH}(v) = R(v)$, thus if Player 1's budget at v exceeds $R(v)$, he can win from v , and if Player 2's budget exceeds $1 - R(v)$, he can win from v .*

Proof. We prove for Player 1 and the proof for Player 2 is dual. Let $t \in \mathbb{N}$ be an index such that $B_1^{\text{init}} > R(v, t)$. We prove by induction on t that Player 1 wins in at most t steps. The base case is easy. For the inductive step, assume Player 1 has a budget of $R(v, i) + \epsilon$. He bids $b_1 = \frac{1}{2}(R(v^+, i - 1) - R(v^-, i - 1))$. If he wins the bidding, he proceeds to v^- with a budget of $R(v^-, i - 1) + \epsilon$. If he loses, then Player 2's bid exceeds b_1 and the worst he can do is move to v^+ . But then Player 1's budget is at least $R(v^+, i - 1) + \epsilon$. By the induction hypothesis, Player 1 wins in at most $i - 1$ steps from both positions. ◀

We make precise the equivalence between reachability and Richman games.

► **Lemma 4.** *Consider a bidding reachability game $\mathcal{G} = \langle V, E, T \rangle$, where $T \subseteq V$ is a target set of vertices for Player 1. Let $S \subseteq V$ be the vertices with no path to T . Consider the Richman game $\mathcal{G}' = \langle V \cup \{v_R, v_S\}, E', v_R, v_S \rangle$, where $E' = E \cup \{\langle v, v_R \rangle : v \in T\} \cup \{\langle v, v_S \rangle : v \in S\}$. For every $v \in V$, the threshold budget of v in \mathcal{G} equals the threshold budget of v in \mathcal{G}' .*

Finding threshold budgets The authors in [35] study the complexity of threshold-budget problem and show that is in NP. They guess, for each vertex v its neighbors v^- and v^+ , and devise a linear program with the constraints $R(v) = \frac{1}{2}(R(v^-) + R(v^+))$ and, for every neighbor v' of v , we have $R(v^-) \leq R(v') \leq R(v^+)$. The program has a solution iff the guess is correct. They leave open the problem of determining the exact complexity of finding the threshold budgets, and they explicitly state that it is not known whether the problem is in P or NP-hard.

We improve on their result by showing that THRESH-BUDG is in NP and coNP. Our reduction uses an important observation that is made in [36], which will be useful later on. They connect between threshold budgets and reachability probabilities in Markov chains.

► **Observation 5.** Consider a Richman game $\mathcal{G} = \langle V, E, v_R, v_S \rangle$. Let $M(\mathcal{G})$ be a Markov chain in which for each vertex $v \in V$, the probability of the edges $\langle v, v^+ \rangle$ and $\langle v, v^- \rangle$ is $\frac{1}{2}$ and the other outgoing edges from v have probability 0. Then, since $R(v) = \frac{1}{2}(R(v^+) + R(v^-))$, in $M(\mathcal{G})$, the probability of reaching v_R from v is $\text{THRESH}(v)$.

We reduce THRESH-BUDG to the problem of “solving” a *simple stochastic game* (SSG, for short) [22]. An SSG has two players; one tries to minimize the probability that the target is reached, and the second player tries to minimize it. It is well-known that the game has a *value*, which is the probability of reaching the target when both players play optimally. The problem of finding the value of an SSG is known to be in $\text{NP} \cap \text{coNP}$. The SSG we construct can be seen as a turn-based

game in which the player whose turn it is to move is chosen uniformly at random. The details of the proof can be found in the full version.

► **Theorem 6.** *THRESH-BUDG for Richman games is in $NP \cap coNP$.*

We stress the fact that the strategies in SSGs are very different from bidding games. As mentioned above, there, the strategies only prescribe which vertex to move the token to, whereas in bidding games, a strategy also prescribes what the next bid should be. So, a solution of a Richman game by reducing it to an SSG is existential in nature and does not give insight on the bids a player uses in his winning strategy. We will return to this point later on.

Objectives We study zero-sum games. The qualitative games we focus on are parity games. A parity game is a triple $\langle V, E, p \rangle$, where $p : V \rightarrow \{0, \dots, d\}$ is a parity function that assigns to each vertex a *parity index*. An infinite outcome is winning for Player 1 iff the maximal index that is visited infinitely often is odd. The quantitative games we focus on are mean-payoff games. A mean-payoff game is $\langle V, E, w \rangle$, where $w : V \rightarrow \mathbb{Z}$ is a weight function on the vertices. We often refer to the sum of weights in a path as its *energy*. Consider an infinite outcome $\pi = v_0, \langle v_1, b_1, i_1 \rangle, \dots$. For $n \geq 0$, we use π^n to refer to the prefix of length n of π . The energy of π^n , denoted $E(\pi^n)$, is $\sum_{0 \leq i \leq n-1} w(v_i)$. We define the mean-payoff value of π to be $\liminf_{n \rightarrow \infty} \frac{E(\pi^n)}{n}$. The value of π can be thought of as the amount of money Player 1 pays Player 2. Note that the mean-payoff values do not affect the budgets of the players. That is, the game has two currencies: a “monopoly money” that is used to determine who moves the token and which the players do not care about once the game ends, and the mean-payoff value that is determined according to the weights of the vertices, which is the value that Min and Max seek to minimize and maximize, respectively. Consider a finite outcome π . We use $\mathcal{B}_m(\pi)$ and $\mathcal{B}_M(\pi)$ to denote the sum of payments of Min and Max in the bids. Throughout the paper we use m and M to refer to Min and Max, respectively.

Strategy complexity Recall that a winning strategy in a two-player game often corresponds to a system implementation. Thus, we often search for strategies that use limited or no memory. That is, we ask whether a player can win even with a *memoryless* strategy, which is a strategy in which the action depends only on the position of the game and not on the history. For example, in turn-based games, for $i \in \{1, 2\}$, a memoryless strategy for Player i prescribes, for each vertex $v \in V_i$, a successor vertex u . It is well known that memoryless strategies are sufficient for winning in a wide variety of games, including turn-based parity games and turn-based mean-payoff games. In Richman games, the threshold budgets tell us who the winner of the game is. But, they do not give insight on how the game is won game, namely what are the bids the winning player bids in order to win. Particularly, when the threshold budgets are 0 as we shall see in Lemmas 7 and 12.

We extend the definition of memoryless strategies to bidding games, though the right definition is not immediate. One can define a memoryless strategy as a function from vertex and budget to action (i.e., bid and vertex) similar to the definition in other games. However, this definition does not preserve the philosophy of implementation with no additional memory. Indeed, recall the proof of Theorem 3. One can define a strategy that, given a vertex $v \in V$ and a budget B , bids according to $R_t(v)$, where t is the minimal index such that $R_t(v) < B$. Clearly, the memory that is needed to implement such a strategy is infinite.

To overcome this issue, we use a different definition. We define a memoryless strategy in a vertex $v \in V$ with initial budget $B \in [0, 1]$ as a pair $\langle u, f_v^B \rangle$, where $u \in \text{adj}(v)$ is the vertex to proceed to upon winning and $f_v^B : [0, 1] \rightarrow [0, 1]$ is a function that takes the current budget and, in mean-payoff games, also the energy, and returns a bid. We require that f_v^B is *simple*, namely a polynomial or a selection between a constant number of polynomials. For simplicity, we assume a memoryless

strategy is generated for an initial vertex with an initial budget, thus there can be different strategies depending where the game starts and with what budget. Also, we call a concatenation of memoryless strategies, a memoryless strategy.

3 Parity Bidding Games

We study threshold budgets in bidding parity games. We first study strongly-connected parity games and show a classification for them; either Player 1 wins with every initial budget or Player 2 wins with every initial budget.

► **Lemma 7.** *Consider a strongly-connected parity game $\mathcal{G} = \langle V, E, p \rangle$. There exists $\tau \in \{0, 1\}$ such that for every $v \in V$, we have $R(v) = \tau$. Moreover, we have $\tau = 0$ iff $\max_{v \in V} p(v)$ is odd.*

Proof. The proof relies on the following claim: Player 1 wins a Richman game in which only his target is reachable, with every initial budget. The claim clearly implies the lemma as we view a strongly-connected bidding parity game as a Richman game in which Player 1 tries to force the game to the vertex with the highest parity index, and Player 2 has no target, thus Player 1 wins with every initial budget. The claim is similar for Player 2. The proof of the claim follows from the fact that the threshold budget of a vertex $v \in V$ is some average between $\text{THRESH}(v_R)$ and $\text{THRESH}(v_S)$, and the average depends on the distances of v to the two targets. When only Player 1's target is reachable, we have $\text{THRESH}(v) = 0$. The details of the proof can be found in the full version. ◀

Consider a bidding parity game $\mathcal{G} = \langle V, E, p \rangle$. Let R and S be the set of vertices in the BSCCs that are winning for Player 1 and Player 2, respectively. Let \mathcal{G}' be the Richman game that is obtained from \mathcal{G} by setting the target of Player 1 to be the vertices in R and the target of Player 2 to be the vertices in S . The following lemma follows from Lemma 7.

► **Lemma 8.** *For every $v \in V$, we have that $\text{THRESH}(v)$ in \mathcal{G} equals $\text{THRESH}(v)$ in \mathcal{G}' .*

Lemma 8 allows us to obtain the positive results of Richman games in parity bidding games. In the full version, we construct memoryless strategies in Richman games. The idea is to show that a bid at a vertex v of the form $\frac{R(v^+) - R(v^-)}{2} + \epsilon$ guarantees that either v_R is reached within $|V|$ steps, or Player 1's budget increases by a constant. Thus, we have the following.

► **Theorem 9.** *The threshold budgets in parity bidding games exist, are unique, THRESH-BUDG is in $NP \cap coNP$, and memoryless strategies suffice for winning.*

4 Mean-Payoff Bidding Games

We proceed to study mean-payoff games. We adjust the definition of threshold budgets to the quantitative setting.

► **Definition 10.** Consider a mean-payoff bidding game $\mathcal{G} = \langle V, E, w \rangle$. The threshold budget in a vertex $v \in V$, denoted $\text{THRESH}(v)$, is a value $t \in [0, 1]$ such that

1. If Min's budget exceeds t at v , then he can guarantee a non-positive mean-payoff value, and
2. if Max's budget exceeds $1 - t$, then he can guarantee a strictly positive value.

4.1 Solving Bidding Mean-Payoff Games

In this section we solve the problem of finding threshold values in bidding mean-payoff games. Our solution relies on work on probabilistic models, namely *one-counter simple stochastic games* [14, 13], and it is existential in nature. Namely, knowing what the threshold budget is in v does not

give much insight on how Min guarantees a non-negative value even if he has sufficient budget, and similarly for Max. Constructing concrete memoryless strategies for the two players is much more challenging and we show constructions in the following sections.

Recall that in bidding parity games, we showed a classification for strongly-connected games; namely, the threshold budgets in all vertices are in $\{0, 1\}$, thus either Player 1 wins with every initial budget or Player 2 wins with every initial budget. We show a similar classification for strongly-connected bidding mean-payoff games: the threshold budgets in all vertices of a strongly-connected bidding mean-payoff game are in $\{0, 1\}$, thus in a strongly-connected bidding mean-payoff game, for every initial energy and every initial budget, either Min can guarantee a non-positive mean-payoff value or Max can guarantee a positive mean-payoff value. The classification uses a generalization of the Richman function to weighted graphs. Consider a strongly-connected bidding mean-payoff game $\mathcal{G} = \langle V, E, w \rangle$ and a vertex $u \in V$. We construct a graph $\mathcal{G}^u = \langle V^u, E^u, w^u \rangle$ by making two copies u_s and u_t of u , where u_s has no incoming edges and u_t has no outgoing edges. Thus, a path from u_s to u_t in \mathcal{G}^u corresponds to a loop in \mathcal{G} . Recall that we denote by $w(v)$ the weight of the vertex v .

► **Definition 11.** Consider a strongly-connected bidding mean-payoff game $\mathcal{G} = \langle V, E, w \rangle$, a vertex $u \in V$. We define the *weighted Richman function* $W : V \rightarrow \mathbb{Q}$ first on \mathcal{G}^u . We define $W(u_t) = 0$ and for every $v \in (S \setminus \{u_t\})$, we define $W(v) = \frac{1}{2}(W(v^+) + W(v^-)) + w(v)$. In order to define W on \mathcal{G} , we define $W(u)$ to be $W(u_s)$ in \mathcal{G}^u .

We use the connection with probabilistic models as in Observation 5 in order to show that W is well defined. We view \mathcal{G}^u as a rewarded Markov chain, in which, for $v \in V$, the outgoing edges from v with positive probability probabilities are $\langle v, v^+ \rangle$ and $\langle v, v^- \rangle$, and their probability is $1/2$. The function W coincides with the expected reward of a run that starts and returns to u , which in turn is well-defined since the probability of returning to u is 1.

Similarly to the connection we show in Theorem 6 between Richman values and reachability probabilities in a simple-stochastic game, we prove Lemma 12 by connecting the threshold value in bidding mean-payoff games to the probability that a counter in a one-counter simple-stochastic games reaches value 0. We then use results from [14, 13] on this model to prove the lemma. The proof can be found in the full version.

► **Lemma 12.** Consider a strongly-connected bidding mean-payoff game $\mathcal{G} = \langle V, E, w \rangle$. There is $\tau \in \{0, 1\}$ such that for every $v \in V$, we have $\text{THRESH}(v) = \tau$. Moreover, we have $\tau = 0$ iff there exists $u \in V$ with $W(u) \leq 0$.

Lemma 13, which is also helpful in the following sections, shows how to connect the mean-payoff value with the objective of reaching energy 0 or maintaining non-negative energy. Its proof can be found in the full version.

- **Lemma 13.** Consider a strongly-connected bidding mean-payoff game \mathcal{G} and a vertex u in \mathcal{G} .
- Suppose that for every initial budget and initial energy, Min has a strategy f_m and there is a constant $N \in \mathbb{N}$ such that for every Max strategy f_M , a finite outcome $\pi = \text{out}(u, f_m, f_M)$ either reaches energy 0 or the energy is bounded by N throughout π . Then, Min can guarantee a non-positive mean-payoff value in \mathcal{G} .
 - If for every initial budget $B_M^{\text{init}} \in [0, 1]$ for Max there exists an initial energy level $n \in \mathbb{N}$ such that Max can guarantee a non-negative energy level in \mathcal{G} , then Max can guarantee a positive mean-payoff value in \mathcal{G} .

The proof of the following theorem can be found in the full version. Deciding the classification in Lemma 12 can be done in NP and coNP by guessing the neighbors the vertices and using

linear programming, similarly to Richman games. Then, we reduce bidding mean-payoff games to Richman games in a similar way to the proof of Lemma 8 for parity games.

► **Theorem 14.** *Threshold budgets exist in bidding mean-payoff games, they are unique, and THRESH-BUDG for bidding mean-payoff games is in $NP \cap coNP$.*

4.2 A Memoryless Optimal Strategy for Min

We turn to the more challenging task of finding memoryless strategies for the players, and in this section we focus on constructing a strategy for Min. Theorem 9 and Lemma 12 allow us to focus on strongly-connected bidding mean-payoff games. Consider a strongly-connected bidding mean-payoff game $\mathcal{G} = \langle V, E, w \rangle$ that has a vertex $u \in V$ with $W(u) \leq 0$. We construct a Min memoryless strategy that guarantees that for every initial energy and every initial budget, either the energy level reaches 0 or it is bounded. By Lemma 13, this suffices for Min to guarantee a non-positive mean-payoff value in \mathcal{G} .

The idea behind our construction is to tie between changes in the energy level and changes of the budget. That is, in order to decrease the energy by one unit, Min needs to *invest* at most one unit of budget (with an appropriate normalization), and when Max increases the energy by one unit, Min's *gain* is at least one unit of budget. Our solution builds on an alternative solution to the two-loop game in Figure 1. This solution is inspired by a similar solution in [35].

► **Example 15.** Consider the bidding mean-payoff game that is depicted in Figure 1. We show a Min strategy that guarantees a non-positive mean-payoff value. Consider an initial Min budget of $B_m^{init} \in [0, 1]$ and an initial energy level of $k_I \in \mathbb{N}$. Let $N \in \mathbb{N}$ be such that $B_m^{init} > \frac{k_I}{N}$. Min bids $\frac{1}{N}$ and takes the (-1) -weighted edge upon winning. Intuitively, Min invests $\frac{1}{N}$ for every decrease of unit of energy and, since by losing a bidding he gains at least $\frac{1}{N}$, this is also the amount he gains when the energy increases. Formally, it is not hard to show that the following invariant is maintained: if the energy level reaches $k \in \mathbb{N}$, Min's budget is at least $\frac{k}{N}$. Note that the invariant implies that either an energy level of 0 is reached infinitely often, or the energy is bounded by N . Indeed, in order to cross an energy of N , Max would need to invest a budget of more than 1. Lemma 13 implies that the mean-payoff value is non-positive, and we are done. ◀

Extending this result to general strongly connected games is not immediate. Consider a strongly-connected game $\mathcal{G} = \langle V, E, w \rangle$ and a vertex $u \in V$. We would like to maintain the invariant that upon reaching u with energy k , the budget of Min exceeds k/N , for a carefully chosen N . The game in the simple example above has two favorable properties that general SCCs do not necessarily have. First, unlike the game in the example, there can be infinite paths that avoid u , thus Min might need to invest budget in drawing the game back to u . Moreover, different paths from u to itself may have different energy levels, so bidding a uniform value (like the $\frac{1}{N}$ above) is not possible. The solution to these problems is surprisingly elegant and uses the weighted Richman function in Definition 11.

Consider an initial budget of $B_m^{init} \in [0, 1]$ for Min and an initial energy $k_I \in \mathbb{N}$. We describe Min's strategy f_m . At a vertex $v \in V$, Min's bid is $\frac{W(v^+) - W(v^-)}{2} \cdot \frac{1}{N}$ and he proceeds to v^- upon winning, where we choose $N \in \mathbb{N}$ in the following. Let w_M be the maximal absolute weighted Richman value in \mathcal{G} , thus $w_M = \max_{v \in V} |W(v)|$. Let b_M be the maximal absolute "bid", thus $b_M = \max_{v \in V} \left| \frac{W(v^+) - W(v^-)}{2} \right|$. We choose $N \in \mathbb{N}$ such that $B_m^{init} > \frac{k_I + b_M + w_M}{N}$.

In the following lemmas we prove that f_m guarantees that an outcome either reaches energy level 0 or that the energy is bounded, as well as showing that f_m is legal, i.e., that Min always bids less than his budget. The following lemma is the crux of the construction as it connects the weighted Richman function with the change in energy and in budget. Recall that for a finite outcome π the accumulated energy in π is $E(\pi)$ and the payments of Min throughout π is $\mathcal{B}_m(\pi)$.

► **Lemma 16.** *Consider a Max strategy f_M , and let $\pi = \text{out}(f_m, f_M)$ be a finite outcome that starts in a vertex v and ends in v' . Then, we have $W(v) - W(v') \geq E(\pi) + N \cdot \mathcal{B}_m(\pi)$.*

Proof. We prove by induction on the length of π . In the base case $v = v'$, thus $E(\pi) = \mathcal{B}_m(\pi) = 0$ and the claim is trivial. For the induction step, let b be the winning bid in the first round and let π' be the suffix of π after the first bidding. We distinguish between two cases. In the first case, Min wins the bidding, pays $b = \frac{W(v^+) - W(v^-)}{2} \cdot \frac{1}{N}$, and proceeds to v^- . Thus, we have $E(\pi) + N \cdot \mathcal{B}_m(\pi) = w(v) + E(\pi') + N(b + \mathcal{B}_m(\pi'))$. By the induction hypothesis, we have $E(\pi') + N \cdot \mathcal{B}_m(\pi') \leq W(v^-) - W(v')$, thus $E(\pi) + N \cdot \mathcal{B}_m(\pi) \leq w(v) + \frac{W(v^+) - W(v^-)}{2} + W(v^-) - W(v') = W(v) - W(v')$, and we are done. For the second case, suppose Max wins the bidding. Min's gain is $-b < -\frac{W(v^+) - W(v^-)}{2} \cdot \frac{1}{N}$, and Max proceeds to v'' having $W(v'') \leq W(v^+)$. Similar to the previous case, we have $E(\pi) + N \cdot \mathcal{B}_m(\pi) = w(v) + E(\pi') + N(-b + \mathcal{B}_m(\pi')) \leq w(v) - \frac{W(v^-) - W(v^+)}{2} + W(v^+) - W(v') = W(v) - W(v')$, and we are done. ◀

The following corollary of Lemma 16 explains why we refer to our technique as “tying energy and budget”. Its proof follows from the fact that $W(u_s) \leq 0$ and $W(u_t) = 0$.

► **Corollary 17.** *Consider a Max strategy f_M , and let $\pi = \text{out}(f_m, f_M)$ be a finite outcome from u to u . Then, we have $-N \cdot \mathcal{B}_m(\pi) \leq E(\pi)$.*

We formalize the intuition above by means of an invariant that is maintained throughout the outcome. Recall that the game starts from a vertex $u \in V$ with $W(u) \leq 0$, the initial energy is $k_I \in \mathbb{N}$, Min's initial budget is $B_m^{\text{init}} \in [0, 1]$, and N is such that $B_m^{\text{init}} > \frac{k_I + b_M + w_M}{N}$.

► **Lemma 18.** *Consider a Max strategy f_M , and let $\pi = \text{out}(f_m, f_M)$ be a finite outcome. Then, when the energy level reaches k , Min's budget is at least $\frac{k + b_M}{N}$.*

Proof. The invariant clearly holds initially. Consider a partition $\pi = \pi_1 \cdot \pi_2$, where π_1 is a maximal prefix of π that ends in u and π_2 starts in u and ends in a vertex $v \in V$. The energy level at the end of π is $k = k_I + E(\pi)$. Recall that $\mathcal{B}_m(\pi)$ is the sum of Min's payments in π , thus his budget at the end of π is $B_m^{\text{init}} - (\mathcal{B}_m(\pi_1) + \mathcal{B}_m(\pi_2))$. By Corollary 17, we have $-\mathcal{B}_m(\pi_1) \geq \frac{1}{N}E(\pi_1)$ and by Lemma 16, we have $-\mathcal{B}_m(\pi_2) \geq \frac{1}{N}(E(\pi_2) - W(u) + W(v)) \geq \frac{1}{N}(E(\pi_2) - 0 - w_M)$. Combining with $B_m^{\text{init}} \geq \frac{k_I + b_M + w_M}{N}$, we have that the new budget is at least $\frac{k_I + b_M + w_M}{N} + \frac{E(\pi_1)}{N} + \frac{E(\pi_2) - w_M}{N} = \frac{k + b_M}{N}$, and we are done. ◀

Lemma 18 implies that Min always has sufficient budget to bid according to f_m , thus the strategy is legal. Moreover, since Min's budget cannot exceed 1, Lemma 18 implies that if the energy does not reach 0, then it is bounded by $N - b_M$. Thus, Lemma 13 implies that Min has a memoryless strategy that guarantees a non-positive mean-payoff value in a strongly-connected bidding mean-payoff game having a vertex u with $W(u) \leq 0$. Combining with the memoryless strategy in parity games, we have the following.

► **Theorem 19.** *Consider a bidding mean-payoff game $\mathcal{G} = \langle V, E, w \rangle$ and a vertex $v \in V$. If Min's initial budget exceeds $\text{THRESH}(v)$, he has a memoryless strategy that guarantees a non-positive mean-payoff value.*

4.3 A Memoryless Optimal Strategy for Max

The complementary result of the previous section is more involved. Consider a strongly-connected bidding mean-payoff game \mathcal{G} with a vertex u that has $W(u) > 0$. We devise a Max strategy that guarantees a positive mean-payoff value in \mathcal{G} . We start with a fragment of the general case called

recurrent SCCs, and we generalize our solution later. We say that an SCC $G = \langle V, E \rangle$ is a recurrent, if there is a vertex $u \in V$ such that every cycle in G includes u . We refer to u as the *root* of G .

Intuitively, the construction has two ingredients. First, we develop the idea of tying energy and budget. We construct a Max strategy f_M that guarantees the following: when Max invests a unit of budget (with an appropriate normalization), then the energy increases by at least one unit, and when the energy decreases by one unit, Max's gain is at least $z > 1$ units of budget, where z arises from the game graph. The second ingredient concerns the normalization factor. Recall that in the previous section it was a constant $\frac{1}{N}$. Here on the other hand, it cannot be constant. Indeed, if the normalization does not decrease as the energy increases, Max's budget will eventually run out, which is problematic since with a budget of 1, Min can guarantee reaching energy level 0, no matter how high the energy is. The challenge is to decide when and how to decrease the normalization factor. We split \mathbb{N} into *energy blocks* of size M , for a carefully chosen $M \in \mathbb{N}$. The normalization factor of the bids depends on the block in which the energy is in, and we refer to it as the *currency* of the block. The currency of the n -th block is z^{-n} . Note that the currency of the $(n-1)$ -th block is higher by a factor of z from the currency of the n -th block. This is where the first ingredient comes in: investing in the n -th block is done in the currency of the n -th block, whereas gaining in the n -th block is in the higher currency of the $(n-1)$ -th block. We switch between the currencies when the energy moves between energy blocks only at the root u of \mathcal{G} . This is possible since \mathcal{G} is a recurrent SCC. The mismatch between gaining and investing is handy when switching between currencies as we cannot guarantee that when we reach u the energy is exactly in the boundary of an energy block.

We formalize this intuition. We start by finding an alternative definition for the weighted Richman function. Recall that in order to define W , we constructed a new graph \mathcal{G}^u by splitting u into u_s and u_t . We define the *contribution* of a vertex $v \in V$ to $W(u_s)$, denoted $cont(v)$, as follows. We have $cont(u_s) = 1$. For a vertex $v \in V$, we define $pre(v) = \{v' \in V : v = v'^- \text{ or } v = v'^+\}$. For $v \in V$, we define $cont(v) = \sum_{v' \in pre(v)} \frac{1}{2} \cdot cont(v')$. The proof of the following lemma uses the connection with probabilistic models, and follows from standard arguments there.

► **Lemma 20.** *We have $W(u) = \sum_{v \in V} (cont(v) \cdot w(v))$.*

Let $z = \left(\sum_{v:w(v) \geq 0} cont(v) \cdot w(v) \right) \cdot \left(\sum_{v:w(v) < 0} cont(v) \cdot |w(v)| \right)^{-1}$. Since $W(u) > 0$, we have $z > 1$. Let \mathcal{G}^z be the game that is obtained from \mathcal{G} by multiplying the negative-weighted vertices by z , thus $\mathcal{G}^z = \langle V, E, w^z \rangle$, where $w^z(v) = w(v)$ if $w(v) \geq 0$ and otherwise $w^z(v) = z \cdot w(v)$. We denote by W^z the weighted threshold budgets in \mathcal{G}^z . The following lemma follows immediately from Lemma 20.

► **Lemma 21.** *We have $W^z(u) = 0$.*

We define the partition into energy blocks. Let $cycles(u)$ be the set of simple cycles from u to itself and $w_M = \max_{\pi \in cycles(u)} |E(\pi)|$. We choose M such that $M \geq (b_M + 3w_M)/(1 - z^{-1})$, where b_M is the maximal bid as in the previous section. We partition \mathbb{N} into blocks of size M . For $n \geq 1$, we refer to the n -th block as M_n , and we have $M_n = \{M(n-1), M(n-1)+1, \dots, Mn-1\}$. We use β_n^\downarrow and β_n^\uparrow to mark the upper and lower boundaries of M_n , respectively. We use a $M_{\geq n}$ to denote the set $\{M_n, M_{n+1}, \dots\}$. Consider a finite outcome π that ends in u and let $visit_u(\pi)$ be the set of indices in which π visits u . Let $k_I \in \mathbb{N}$ be an initial energy. We say that π *visits* M_n if $k_I + E(\pi) \in M_n$. We say that π *stays in* M_n starting from an index $1 \leq i \leq |\pi|$ if for all $j \in visit_u(\pi)$ such that $j \geq i$, we have $k_I + E(\pi_1, \dots, \pi_j) \in M_n$.

We are ready to describe Max's strategy f_M . Suppose the game reaches a vertex v and the energy in the last visit to u was in M_n , for $n \geq 1$. Then, Max bids $z^{-n} \cdot \frac{1}{2} (W^z(v^+) - W^z(v^-))$ and proceeds to v^+ upon winning. Note that currency changes occur only in u . Recall that for an outcome π , the sum of payments of Max is $\mathcal{B}_M(\pi)$ and let $E^z(\pi)$ be the change in energy in \mathcal{G}^z . The proof of Lemma 16 can easily be adjusted to this setting.

► **Lemma 22.** Consider a Min strategy f_m , and let $\pi = \text{out}(f_m, f_M)$ be a finite outcome that starts in v , ends in v' , and stays within a block M_n , for $n \geq 1$. We have $W^z(v) - W^z(v') \leq E^z(\pi) - z^n \cdot \mathcal{B}_M(\pi)$. In particular, for $\pi \in \text{cycles}(u)$, we have $E^z(\pi) \leq z^n \cdot \mathcal{B}_M(\pi)$.

We relate between the changes in energy in the two structures. The proof of the following lemma can be found in the full version.

► **Lemma 23.** Consider an outcome $\pi \in \text{cycles}(u)$. Then, $E(\pi) \geq E^z(\pi)$ and $E(\pi) \geq zE^z(\pi)$.

A corollary of Lemmas 22 and 23 is the following. Recall that $\mathcal{B}_M(\pi)$ is the amount that Max pays, thus it is negative when Max gains budget. Intuitively, the corollary states that if the energy increases in M_n , then Max invests in the currency of M_n , and if the energy decreases, he gains in the currency of M_{n-1} .

► **Corollary 24.** Consider a Min strategy f_m , and let $\pi = \text{out}(f_m, f_M)$ be a finite outcome such that $\pi \in \text{cycles}(u)$. Then, we have $E(\pi) \geq z^n \cdot \mathcal{B}_M(\pi)$ and $zE(\pi) \geq z^n \cdot \mathcal{B}_M(\pi)$.

Consider an initial Max budget $B_M^{\text{init}} \in [0, 1]$. We choose an initial energy $k_I \in \mathbb{N}$ with which f_M guarantees that energy level 0 is never reached. Recall the intuition that increasing the energy by a unit requires an investment of a unit of budget in the right currency. Thus, increasing the energy from the lower boundary β_n^\downarrow of M_n to its upper boundary β_n^\uparrow , costs $M \cdot z^{-n}$. We use $\text{cost}(M_n)$ to refer to $M \cdot z^{-n}$ and $\text{cost}(M_{\geq n}) = \sum_{i=n}^{\infty} \text{cost}(M_i)$. A first guess for k_I would be β_n^\downarrow such that $B_M^{\text{init}} > \text{cost}(M_{\geq n})$. This is almost correct. We need some *wiggle room* to allow for changes in the currency. Let $\text{wiggle} = 2w_M + b_M$, where recall that $w_M = \max_{\pi \in \text{cycles}(u)} E(\pi)$ and that b_M is the maximal bid. We define k_I to be β_n^\downarrow such that $B_M^{\text{init}} > \text{wiggle} \cdot z^{-(n-1)} + \text{cost}(M_{\geq n})$ and $\sum_{i=1}^n \text{cost}(M_i) > 1$, thus Min cannot decrease the energy to 0.

Consider a Min strategy f_m , and let $\pi = \text{out}(f_m, f_M)$ be a finite outcome. We partition π into subsequences in which the same currency is used. Let $\pi = \pi_1 \cdot \pi_2 \cdot \dots \cdot \pi_\ell$ be a partition of π . For $1 \leq i \leq \ell$, we use π^i to refer to the prefix $\pi_1 \cdot \dots \cdot \pi_i$ of π , and we use $e^i = k_I + E(\pi^i)$ to refer to the energy at the end of π^i . Consider the partition in which, for $1 \leq i \leq \ell$, the prefix π^i visits u and π_i is a maximal subsequence that stays in some energy block.

Suppose π_i stays in M_n . There can be two options; either the energy decreases in π_i , thus the energy before it e^{i-1} is in M_{n+1} and the energy after it e^i is in M_n , or it increases, thus $e^{i-1} \in M_{n-1}$ and $e^i \in M_n$. We then call π^i *decreasing* and *increasing*, respectively. The definition of w_M and the fact that \mathcal{G} is recurrent imply that upon entering M_n , the energy is within w_M of the boundary. Thus, in the case that π^i is decreasing, the energy at the end of π^i is $e^i \geq \beta_n^\uparrow - w_M$ and in the case it is increasing, we have $e^i \leq \beta_{n+1}^\downarrow + w_M$. Let $\ell_0 = 0$, and for $i \geq 1$, let $\ell_i = (\beta_{n+1}^\downarrow - w_M) - e^i$ in the first case and $\ell_i = (\beta_n^\uparrow + w_M) - e^i$ in the second case. Note that $\ell_i \in \{0, \dots, 2w_M\}$. In the full version, we prove the following invariant on Max's budget when changing between energy blocks.

► **Lemma 25.** For every $i \geq 0$, suppose π^i ends in M_n . Then, Max's budget is at least $(\text{wiggle} + \ell_i) \cdot z^{-(\hat{n}-1)} + \text{cost}(M_{\geq \hat{n}})$, where $\hat{n} = n + 1$ if π^i is decreasing and $\hat{n} = n$ if π^i is increasing.

It is not hard to show that Lemma 25 implies that f_M is legal. That is, consider a finite outcome π that starts immediately after a change in currency. Using Lemma 22, we can prove by induction on the length of π that Max has sufficient budget for bidding. The harder case is when π decreases, and the proof follows from the fact that *wiggle* is in the higher currency of the lower block. Combining Lemma 25 with our choice of the initial energy, we get that the energy never reaches 0 as otherwise Min invests a budget of more than 1. Lemma 13 implies that Max guarantees a positive mean-payoff value in a strongly-connected game.

► **Theorem 26.** Consider a bidding mean-payoff game $\mathcal{G} = \langle V, E, w \rangle$ in which all BSCCs are recurrent. For a vertex $v \in V$, if Max has an initial budget that is greater than $1 - \text{THRESH}(v)$, he has a memoryless strategy that guarantees a positive mean-payoff value.

General strongly-connected games In the full version, we develop a constant-memory strategy for Max that guarantees a positive mean-payoff value. The difficulty lies in coping with outcomes in which the energy forms a sine-like wave on the boundary of an energy block. In recurrent SCCs, we can change currency every time the wave changes block, which does not work in general SCCs as we show in an example. We develop further the two ingredients that are used in f_M . First, recall that investing in an energy block M_n is in the currency of the n -th block, whereas gaining is in the higher currency of the $(n - 1)$ -th block. In general games, we need a stronger property; investing in M_n is in the lower currency of the $(n + 1)$ -th block while gaining is still in the higher currency of the $(n - 1)$ -th block. Next, we differentiate between even blocks, i.e., M_{2n} , and odd blocks, i.e., M_{2n+1} , for some $n \in \mathbb{N}$. When the energy level reaches an even block M_{2n} , the currency used is z^{-n} . In order to determine the currency in the odd blocks, we take the history of the play into account; the currency matches the currency in the last energy block that was visited before entering M_{2n+1} . Hence, we call our strategy a *constant-memory* strategy. The odd blocks serve as “buffers” so that when we change currency, there is a sufficiently large change in energy that in turn implies that Max’s budget sufficiently increases compared with the change in energy. Combining with the memoryless strategy in parity games of Theorem 9, we have the following.

► **Theorem 27.** *Consider a bidding mean-payoff game $\mathcal{G} = \langle V, E, w \rangle$. For a vertex $v \in V$, if Max has an initial budget that is greater than $1 - \text{THRESH}(v)$, he has a constant-memoryless strategy that guarantees a positive mean-payoff value.*

5 Discussion and Future Directions

We introduce and study infinite-duration bidding games in which the players bid for the right to move the token. This work belongs to a line of works that transfer concepts and ideas between the areas of formal methods and algorithmic game theory (AGT, for short). Richman games originated in the game theory community in the 90s and recently gained interest by the AGT community [30]. We combine them with the study of infinite-duration games, which is well-studied in the formal methods community. Prior to this work, a series of works focused on applying concepts and ideas from formal methods to *resource-allocation games* [10, 8, 9, 6, 7, 34], which constitutes a well-studied class of games in AGT. More to the formal methods side, there are many works on games that share similar concepts to these that are studied in AGT. For example, logics for reasoning about multi-agent systems [3, 19, 39], studies of equilibria in games related to synthesis and repair problems [18, 25, 1, 15], and studies of infinite-duration non-zero-sum games [21, 16, 17, 12].

There are several problems we left open as well as plenty of future research directions. We list a handful of them below. We showed that the complexity of THRESH-BUDG is in NP and coNP. We leave open the problem of determining its exact complexity. We conjecture that it is reducible from solving simple stochastic games, which will show that it is as hard as several other problems whose exact complexity is unknown. In this work we focused on parity and mean-payoff games. *Energy games* are games that are played on a weighted graph, where one of the players tries to reach negative energy and the second player tries to prevent it. Note that unlike parity and mean-payoff, the energy objective is not *prefix independent*. We can show that threshold budgets exist in energy games. The complexity of THRESH-BUDG in energy games is interesting and is tied with recent work on optimizing the probability of reaching a destination in a weighted MDP [26, 41]. For acyclic energy bidding games, the problem is PP-hard using a result in [26], and for a single-vertex games the problem is in P using the direct formula of [32]. For general games the problem is open.

Acknowledgments

We thank Petr Novotný for helpful discussions and pointers.

References

- 1 S. Almagor, G. Avni, and O. Kupferman. Repairing multi-player games. In *Proc. 26th CONCUR*, pages 325–339, 2015.
- 2 S. Almagor, D. Kuperberg, and O. Kupferman. The sensing cost of monitoring and synthesis. In *Proc. 35th FSTTCS*, pages 380–393, 2015.
- 3 R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002.
- 4 K.R. Apt and E. Grädel. *Lectures in Game Theory for Computer Scientists*. Cambridge University Press, 2011.
- 5 G. Avni, S. Guha, and O. Kupferman. An abstraction-refinement methodology for reasoning about network games. In *Proc. 26th IJCAI*, 2017.
- 6 G. Avni, S. Guha, and O. Kupferman. Timed network games. In *Proc. 42nd MFCS*, 2017.
- 7 G. Avni, T. A. Henzinger, and V. Chonev. Infinite-duration bidding games. *CoRR*, abs/1705.01433, 2017. <https://arxiv.org/abs/1705.01433>.
- 8 G. Avni, T. A. Henzinger, and O. Kupferman. Dynamic resource allocation games. In *Proc. 9th SAGT*, pages 153–166, 2016.
- 9 G. Avni, O. Kupferman, and T. Tamir. Congestion games with multisets of resources and applications in synthesis. In *Proc. 35th FSTTCS*, pages 365–379, 2015.
- 10 G. Avni, O. Kupferman, and T. Tamir. Network-formation games with regular objectives. *Inf. Comput.*, 251:165–178, 2016.
- 11 N. Berger, N. Kapur, L. J. Schulman, and V. V. Vazirani. Solvency games. In *Proc. 28th FSTTCS*, pages 61–72, 2008.
- 12 P. Bouyer, R. Brenguier, N. Markey, and M. Ummels. Nash equilibria in concurrent games with büchi objectives. In *Proc. 31st FSTTCS*, pages 375–386, 2011.
- 13 T. Brázdil, V. Brozek, K. Etessami, and A. Kucera. Approximating the termination value of one-counter mdps and stochastic games. In *Proc. 38th ICALP*, pages 332–343, 2011.
- 14 T. Brázdil, V. Brozek, K. Etessami, A. Kucera, and D. Wojtczak. One-counter markov decision processes. In *Proc. 21st SODA*, pages 863–874, 2010.
- 15 R. Brenguier, L. Clemente, P. Hunter, G. A. Pérez, M. Randour, J.-F. Raskin, O. Sankur, and M. Sassolas. Non-zero sum games for reactive synthesis. In *Proc. 10th LATA*, pages 3–23, 2016.
- 16 T. Brihaye, V. Bruyère, J. De Pril, and H. Gimbert. On subgame perfection in quantitative reachability games. *Logical Methods in Computer Science*, 9(1), 2012.
- 17 V. Bruyère, N. Meunier, and J.-F. Raskin. Secure equilibria in weighted games. In *Proc. CSL-LICS*, pages 26:1–26:26, 2014.
- 18 K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Games with secure equilibria. *Theor. Comput. Sci.*, 365(1-2):67–82, 2006.
- 19 K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. *Inf. Comput.*, 208(6):677–693, 2010.
- 20 K. Chatterjee, R. Majumdar, and T. A. Henzinger. Controller synthesis with budget constraints. In *Proc. 11th HSCC*, pages 72–86, 2008.
- 21 K. Chatterjee, R. Majumdar, and M. Jurdzinski. On nash equilibria in stochastic games. In *Proc. 13th CSL*, pages 26–40, 2004.
- 22 A. Condon. On algorithms for simple stochastic games. In *Proc. DIMACS*, pages 51–72, 1990.

- 23 M. Develin and S. Payne. Discrete bidding games. *The Electronic Journal of Combinatorics*, 17(1):R85, 2010.
- 24 K. Etessami, D. Wojtczak, and M. Yannakakis. Quasi-birth-death processes, tree-like qbds, probabilistic 1-counter automata, and pushdown systems. *Perform. Eval.*, 67(9):837–857, 2010.
- 25 D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *Proc. 16th TACAS*, pages 190–204, 2010.
- 26 C. Haase and S. Kiefer. The odds of staying on budget. In *Proc. 42nd ICALP*, pages 234–246, 2015.
- 27 M. Holtmann, L. Kaiser, and W. Thomas. Degrees of lookahead in regular infinite games. *Logical Methods in Computer Science*, 8(3), 2012.
- 28 F. A. Hosch and L. H. Landweber. Finite delay solutions for sequential conditions. In *ICALP*, pages 45–60, 1972.
- 29 K. Johnson, D. Simchi-Levi, and P. Sun. Analyzing scrip systems. *Operations Research*, 62(3):524–534, 2014.
- 30 G. Kalai, R. Meir, and M. Tennenholtz. Bidding games and efficient allocations. In *Proc. 16th EC*, pages 113–130, 2015.
- 31 I. A. Kash, E. J. Friedman, and J. Y. Halpern. Optimizing scrip systems: crashes, altruists, hoarders, sybils and collusion. *Distributed Computing*, 25(5):335–357, 2012.
- 32 G. Katriel. Gambler’s ruin probability – a general formula. *Statistics and Probability Letters*, 83:2205–2210, 2013.
- 33 F. Klein and M. Zimmermann. How much lookahead is needed to win infinite games? *Logical Methods in Computer Science*, 12(3), 2016.
- 34 O. Kupferman and T. Tamir. Hierarchical network formation games. In *Proc. 23rd TACAS*, pages 229–246, 2017.
- 35 A. J. Lazarus, D. E. Loeb, J. G. Propp, W. R. Stromquist, and D. H. Ullman. Combinatorial games under auction play. *Games and Economic Behavior*, 27(2):229–264, 1999.
- 36 A. J. Lazarus, D. E. Loeb, J. G. Propp, and D. Ullman. Richman games. 29:439–449, 1996.
- 37 R. Paes Leme, V. Syrgkanis, and É. Tardos. Sequential auctions and externalities. In *Proc. 23rd SODA*, pages 869–886, 2012.
- 38 M. Menz, J. Wang, and J. Xie. Discrete all-pay bidding games. *CoRR*, abs/1504.02799, 2015.
- 39 F. Mogavero, A. Murano, and M. Y. Vardi. Reasoning about strategies. In *Proc. 30th FSTTCS*, pages 133–144, 2010.
- 40 A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. 16th POPL*, pages 179–190, 1989.
- 41 M. Randour, J.-F. Raskin, and O. Sankur. Variations on the stochastic shortest path problem. In *Proc. 16th VMCAI*, pages 1–18, 2015.
- 42 D. M. Reeves, B. M. Soule, and T. Kasturi. Yootopia! *SIGecom Exchanges*, 6(2):1–26, 2007.
- 43 M. Stonebraker, P. M. Aoki, W. Litwin, A. Pfeffer, A. Sah, J. Sidell, C. Staelin, and A. Yu. Mari-
posa: A wide-area distributed database system. *VLDB J.*, 5(1):48–63, 1996.