Infinite-Duration Richman Bidding Games*

Guy Avni
IST Austria
Klosterneuburg, Austria
guy.avni@ist.ac.at

Thomas A. Henzinger
IST Austria
Klosterneuburg, Austria
tah@ist.ac.at

Ventsislav Chonev
Google
v.chonev@gmail.com

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. Reachability bidding games, called *Richman games*, have been studied in [35, 34]. 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 and mean-payoff games. We show the existence of threshold budgets and show that the complexity of finding them coincides with the $NP \cap coNP$ complexity of reachability bidding games. The solution for mean-payoff consists of our most technically challenging contribution, where we construct optimal strategies for the players while extending and generalizing the probabilistic connection that was known for reachability bidding games.

1 Introduction

Two-player infinite-duration games on graphs are a central class of games in formal verification [4] and have deep connections to foundations of logic [44]. Questions about automatic synthesis of a reactive system from its specification [42] 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 payoff 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 new mode of moving in infinite-duration games, which is called *bidding*, and in which the players *bid* for the right to move the token. The bidding mode of moving was introduced in [34, 35] for reachability games, where two bidding rules were defined. The first bidding rule, which we focus on in this paper and is called the *Richman* rule (named after David Richman), is as follows: Each player has a budget, and before each move, the players submit bids simultaneously, where a bid is legal if it does not exceed the available budget. The player who bids higher wins the bidding, pays the bid to other player, and moves the token. A second bidding rule is called *poorman* bidding in [34], is similar except that the winner of the bidding pays the "bank" rather than the other player.

*The work described in this paper was published in the proceedings of the 28th International Conference on Concurrency Theory, pages 21:1–21:18, LIPIcs 85, 2017.

Submitted to:

Bidding for moving is a general concept that is relevant in any setting in which a scheduler needs to decide the order in which selfish agents perform actions. For example, 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 [19, 2] as well as allowing the system to read chunks of input signals before producing an output signal [27, 26, 31]. Also, our bidding game can model *scrip systems* that use internal currencies for bidding in order to prevent "free riding" [30]. Such systems are successfully used in various settings such as databases [46], group decision making [45], resource allocation, and peer-to-peer networks (see [28] and references therein). Finally, repeated bidding is a form of a sequential auction [36], which is used in many settings including online advertising.

Recall that the winner or payoff of the game is determined according to the play, which is an infinite trace. 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 were studied in [35, 34], where the bidding mode of moving was introduced. The central question that is studied is a necessary and sufficient budget for winning, which we call a *threshold budget*. Formally, the threshold budget of a vertex v, denoted $\text{TH}(v)$ is a value in $[0,1]$ such that if the game starts at v and Player 1's budget exceeds $\text{TH}(v)$, then he has a strategy to win the game. On the other hand, if Player 2's budget exceeds $1-\text{TH}(v)$, he can win the game (recall that the budgets add up to 1). This is a central property of the game, which is a form of *determinacy*, and shows that no ties can occur.[1] In [35, 34], 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.

More interesting, from the synthesis and logic perspective, are infinite winning conditions. We introduce and study infinite duration bidding games with richer qualitative objectives as well as quantitative objectives. We start with qualitative games and show that *parity* bidding games are linearly-reducible to bidding reachability 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 bidding reachability games, i.e., the problem is in NP and coNP. We show that in a strongly-connected game, one of the players can win with any positive initial budget. In a general game, we first classify the *bottom strongly-connected components* (BSCCs, for short) of the graph to the ones that are winning for each player, and construct a reachability game in which each player tries to force the game to a BSCC that is winning for him.

Our most interesting results concern *mean-payoff* bidding games, which are quantitative games; an infinite play π of the game is associated with a value $c \in \mathbb{R}$. Player 1's *payoff* in π is c and Player 2's payoff is $-c$. Accordingly, we refer to the players in a mean-payoff game as Max and Min. The payoff of π is determined according to the weights it traverses and, as in the previous games, the bids are only used to determine whose turn it is to move. The central question in these games is: Given a value $c \in \mathbb{Q}$, what is the initial budget that is necessary and sufficient for Max to guarantee a payoff of c? More formally, we say that c is the *value* with respect to an initial budget $B \in [0,1]$ if for every $\varepsilon > 0$, we have (1) when Max's initial budget is $B+\varepsilon$, he can guarantee a payoff of at least c, and (2) intuitively, Max cannot hope for more: if Max's initial ratio is $B-\varepsilon$, then Min can guarantee a payoff of at most c.

The crux of the solution again concerns the BSCCs of the game. We extend the known, somewhat

[1] When the initial budget of Player 1 is exactly $\text{TH}(v)$, the winner of the game depends on how we resolve draws in biddings, and our results hold for any tie-breaking mechanism.

unexpected, probabilistic connection for bidding reachability games to bidding mean-payoff games. A *random-turn based game* (RTB game, for short) is a graph game in which in each round, the player who chooses the next move is selected uniformly at random. More formally, an RTB game is a special case of a 2.5-*player* game [21]. For qualitative RTB games, the *value* of the game is the maximal probability with which Player 1 can guarantee winning. It was shown in [35] that the threshold budget in a Richman reachability game coincides with the value of a RTB game played on the same structure.

We extend this probabilistic connection to mean-payoff games. The *mean-payoff value* of a mean-payoff RTB game is the maximal expected payoff Max can guarantee. We show that in a strongly-connected Richman mean-payoff game, no matter what the initial budgets are, the optimal payoff Max can guarantee coincides with the mean-payoff value in a mean-payoff RTB game played on the same structure. We provide two proofs of this claim. The first relies on a connection between bidding mean-payoff games and *one-counter* 2.5-*player games* [13, 14]. In turn, these games are equivalent to discrete *quasi-birth-death processes* [23] and generalize *solvency games* [10], which can be thought of as a rewarded Markov decision process with a single vertex.

The drawback of the approach above is that it is existential in nature and does not give any insight on how to construct optimal strategies. The second proof technique, which constitutes our most technically challenging results, is direct; we construct optimal strategies for Min and Max. Beyond the importance of constructing strategies, the ideas developed in the construction were later used to solve mean-payoff games with poorman bidding rules [7] for which a probabilistic connection does not exist (see more details in Section 5). 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. One of the technical challenges is that in general strongly-connected graphs the bids must differ between the vertices, and we need to decide in each vertex how "important" it is to move. Our technique relies on the concept of *potential*, which was developed in the context of the *strategy improvement* algorithm to solve graph games (c.f., [25]).

Further related work Beyond the works that are directly relevant to us, which we have compared to above, we list previous work on Richman games. Motivated by recreational games, e.g., bidding chess [11, 33], *discrete bidding games* are studied in [22], where the money is divided into chips, so a bid cannot be arbitrarily small unlike the bidding games we study. In *all-pay* bidding games [37], both the winner and loser of a bidding pay their bids to the bank. Non-zero-sum two-player Richman games were recently studied in [29]. In addition to these works, Richman games where studied in the context of the equivalent model of RTB games (c.f., [41]) and the infinity Laplacian (c.f., [40]).

Due to lack of space, most of the proofs appear in the full version [6].

2 Preliminaries

A graph game is played on a directed graph $G = \langle V, E \rangle$, where V is a finite set of vertices and $E \subseteq V \times V$ is a set of edges. The *neighbors* of a vertex $v \in V$, denoted $N(v)$, is the set of vertices $\{u \in V : \langle v, u \rangle \in E\}$, and we say that G has out-degree 2 if for every $v \in V$, we have $|N(v)| = 2$. A *path* in G is a finite or infinite sequence of vertices $v_1, v_2, \ldots$ such that for every $i \geq 1$, we have $\langle v_i, v_{i+1} \rangle \in E$.

Objectives An objective O is a set of infinite paths. In reachability games, Player 1 has a target vertex v_R and an infinite path is winning for him if it visits v_R. In *parity* games each vertex has a parity index in $\{1, \ldots, d\}$, and an infinite path is winning for Player 1 iff the maximal parity index that is visited infinitely often is odd. We also consider games that are played on a weighted graph $\langle V,E,w\rangle$, where $w: V \to \mathbb{Q}$. Consider an infinite path $\pi = v_1, v_2, \ldots$. For $n \in \mathbb{N}$, we use π^n to denote the prefix of length n of π. We call the sum of weights that π^n traverses the *energy* of the game, denoted $E(\pi^n)$. Thus, $E(\pi^n) = \sum_{1\leq j<n} w(v_j)$. In *energy* games, the goal of Player 1 is to keep the energy level positive, thus he wins an infinite path iff for every $n \in \mathbb{N}$, we have $E(\pi^n) > 0$. Unlike the previous objectives, a path in a *mean-payoff* game is associated with a payoff, which can be thought of as a reward for Player 1 and a cost for Player 2. Thus, Player 1's goal is to maximize the payoff and Player 2's goal is to minimize it. Accordingly, in mean-payoff games, we refer to Player 1 as Max and Player 2 as Min. We define the payoff of π to be $\liminf_{n\to\infty} \frac{1}{n}E(\pi^n)$. We say that Max wins an infinite path of a mean-payoff game if the payoff is non-negative.

Strategies and plays 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. For example, in turn-based 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. Intuitively, a strategy prescribes an action $\langle b, v\rangle$, where b is a bid that does not exceed the available budget and v is a vertex to move to upon winning. Formally, a history in a bidding game is a sequence $\tau = v_0, \langle v_1, b_1, \ell_1\rangle, \langle v_2, b_2, \ell_2\rangle, \ldots, \langle v_k, b_k, \ell_k\rangle \in V \cdot (V \times \mathbb{R} \times \{1,2\})^*$, 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 ℓ_j, and Player ℓ_j moves the token to vertex v_j. For $i \in \{1,2\}$, let $W_i(\tau)$ denote the indices in which Player i wins the bidding, thus $W_i(\tau) = \{1 \leq j \leq k : \ell_j = i\}$. The *payment* of Player i in τ, denoted $pay_i(\tau)$, is $\sum_{j\in W_i(\tau)} b_j$. Let B and $1-B$ be Player 1 and 2's initial budgets, respectively. Player 1's budget following τ is $B - pay_1(\tau) + pay_2(\tau)$ and Player 2's budget is $(1-B) + pay_1(\tau) - pay_2(\tau)$. We restrict attention to *legal strategy* in which $\langle v_k, v\rangle$ must be an edge and b must not exceed the available budget.

An initial vertex v_0 and strategies f_1 and f_2 for Players 1 and 2, respectively, determine a unique *play* π for the game, denoted $play(v_0, f_1, f_2)$, which is an infinite sequence in $V \cdot (V \times \mathbb{R} \times \{1,2\})^\omega$. We sometimes abuse notation and refer to $play(v_0, f_1, f_2)$ as a finite prefix of the infinite play. We drop v_0 when it is clear from the context. We define the play inductively. The first element is v_0. Suppose $\pi_1, \ldots, \pi_j$ is defined. The players bids are given by $\langle b_1, v_1\rangle = f_1(\pi_1, \ldots, \pi_j)$ and $\langle b_2, v_2\rangle = f_2(\pi_1, \ldots, \pi_j)$. If $b_1 > b_2$, then Player 1 wins the bidding and decides where to move the token to, thus $\pi_{j+1} = \langle v_1, b_1, 1\rangle$. Dually when $b_1 < b_2$ Player 2 wins the bidding and 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 an objective O. An infinite play $v_0, \langle v_1, b_1, \ell_1\rangle, \langle v_2, b_2, \ell_2\rangle, \ldots$ satisfies O iff the infinite path $v_0, v_1, v_2, \ldots$ satisfies O. We call a strategy f_1 *winning* for Player 1 if for every strategy f_2 of the other player $play(v_0, f_1, f_2)$ satisfies O. Winning strategies for Player 2 are defined dually.

The first question that arises in the context of bidding games asks what is the necessary and sufficient initial ratio to guarantee an objective. We generalize the definition in [34, 35]:

Definition 1. *(***Threshold budgets***) Consider a bidding game $\mathcal{G}$, a vertex v, and an initial budget $B \in [0,1]$ for objective O for Player 1. The threshold budget in v, denoted $\text{TH}(v)$, is a number in $[0,1]$ such that*

- *if $B > \text{TH}(v)$, then Player 1 has a winning strategy that guarantees O is satisfied, and*
- *if $B < \text{TH}(v)$, then Player 2 has a winning strategy that violates O.*

Random-turn based games A 2.5-player mean-payoff game is played on an arena $\langle V_1, V_2, V_N, E, \Pr\rangle$, where the sets of vertices V_1, V_2, and V_N are disjoint and respectively represent the set of states that are controlled by Player 1, Player 2, and *nature*, $E \subseteq (V_1 \cup V_2) \times V_N$ is a set of deterministic edges, and $\Pr : V_N \times (V_1 \cup V_2) \to [0,1]$ assigns probabilities to outgoing edges from nature vertices. Whenever the token reaches a vertex in $v \in V_1$, Player 1 chooses how to move it, where a legal move is $u \in V_N$ such that $E(v,u)$, and similarly for Player 2. When the token reaches $u \in V_N$, it continues to $v \in (V_1 \cup V_2)$ with probability $\Pr(u,v)$. We denote $V = (V_1 \cup V_2 \cup V_N)$.

A 2.5-player game with reachability objectives is called a *simple stochastic game* (SSG, for short) [21]. Intuitively, the *value* of an SSG $\mathscr{G}$, denoted $val(\mathscr{G})$, is the maximal probability with which Player 1 can guarantee reaching the target. Formally, let $\mathscr{G} = \langle V_1, V_2, V_N, E, \Pr, v_R\rangle$, where $v_R \in V$ is the target for Player 1. It is well known that optimal positional strategies exist in SSGs. Given two positional strategies we construct a Markov chain $\mathscr{M}^{f,g}$ by trimming away edges that do not comply with the strategies. That is, we leave edges of the form $\langle v,u\rangle \in E$, where $v \in V_1$ and $u = f(v)$, or $v \in V_2$ and $u = g(v)$.[2] Then, the value of $v \in V$ with respect to f and g, denote $val^{f,g}_{\mathscr{G}}(v)$, is the probability of reaching v_R in $\mathscr{M}^{f,g}$. The value of v is $val_{\mathscr{G}}(v) = \max_f \min_g val^{f,g}_{\mathscr{G}}(v)$. It is well known that $val_{\mathscr{G}}(v) = \min_g \max_f val^{f,g}_{\mathscr{G}}(v)$.

A mean-payoff 2.5-player game is $\mathscr{G} = \langle V_1, V_2, V_N, E, \Pr, w\rangle$, where $w : V \to \mathbb{Q}$. We again restrict attention to *ergodic* 2.5-player games in which each player has a strategy $f_{u,v}$, for every $u, v \in V$, such that $f_{u,v}$ guarantees reaching v starting from u with probability 1. Intuitively, the *mean-payoff value* of a mean-payoff 2.5-player game $\mathscr{G}$, denoted $\texttt{MP}(\mathscr{G})$, is the maximal expected payoff Max can guarantee, and it is easy to see that the mean-payoff value does not depend on the initial vertex. Formally, a *rewarded Markov chain* is a tuple $\mathscr{M} = \langle V, \Pr, w\rangle$, where V is a set of vertices, $\Pr : (V \times V) \to [0,1]$ is a probability function, and $w : V \to \mathbb{Q}$ assigns weights to vertices. The *stationary distribution* in a vertex $v \in V$ intuitively states what precent of the time a random infinite walk on $\mathscr{M}$ stays in v. In order to compute it, we construct a linear program with a variable x_v, for every vertex $v \in V$, a constraint $x_v = \sum_{v' \in V} \Pr[v', v] \cdot x_{v'}$, and a normalizing constraint $\sum_{v \in V} x_v = 1$. The mean-payoff value in $\mathscr{M}$, denoted $\texttt{MP}(\mathscr{M})$, is $\sum_{v \in V} x_v \cdot w(v)$. It is well known that optimal positional strategies exist in mean-payoff 2.5-player games. As in the qualitative case, given such an ergodic game $\mathscr{G}$ and two strategies f and g, we construct a rewarded Markov chain $\mathscr{M}^{f,g}$, and define $\texttt{MP}^{f,g}(\mathscr{G}) = \texttt{MP}(\mathscr{M}^{f,g})$. We define $\texttt{MP}(\mathscr{G}) = \max_f \min_g \texttt{MP}^{f,g}(\mathscr{G})$.

Consider a game $\mathscr{G} = \langle V, E\rangle$. The *random-turn based game* that is associated with $\mathscr{G}$ is a 2.5-player game that intuitively simulates the fact that the player to choose the next move is chosen uniformly at random. Formally, we define $\texttt{RTB}(\mathscr{G}) = \langle V_1, V_2, V_N, E, \Pr\rangle$, where each vertex in V is split into three vertices, each controlled by a different player, thus for $\alpha \in \{1, 2, N\}$, we have $V_\alpha = \{v_\alpha : v \in V\}$, nature vertices simulate the random choice, thus $\Pr[v_N, v_1] = \Pr[v_N, v_2] = 0.5$, and reaching a vertex that is controlled by one of the two players means that he chooses the next move, thus $E = \{\langle v_\alpha, u_N\rangle : \langle v,u\rangle \in E \text{ and } \alpha \in \{1,2\}\}$. When $\mathscr{G}$ is weighted, then the weights of v_1, v_2, and v_N equal that of v.

3 Qualitative objectives

Bidding games with reachability objectives where studied in [35, 34]. They study a slightly different model, which we call *double-reachability games*, and in which each player has a target and the game ends once one of the targets is reached. They show the following.

Theorem 2. *[35] Threshold budgets exist in double-reachability games. Moreover, threshold budgets have the following property. Consider a double-reachability bidding game* $\mathscr{G} = \langle V, E, v_R, v_S\rangle$. *We have*

[2] A similar construction can be obtained for arbitrary strategies, and these simple definitions suffice for our needs.

$\textsc{Th}(v_R) = 0$ *and* $\textsc{Th}(v_S) = 1$, *and for every* $v \in V \setminus \{v_R, v_S\}$, *we have* $\textsc{Th}(v) = \frac{1}{2}\big(\textsc{Th}(v^+) + \textsc{Th}(v^-)\big)$, *where* $v^-, v^+ \in N(v)$ *are such that for every* $v' \in N(v)$, *we have* $\textsc{Th}(v^-) \le \textsc{Th}(v') \le \textsc{Th}(v^+)$.

We make precise the equivalence between the two types of reachability objectives.

Lemma 3. *Consider a bidding reachability game* $\mathscr{G} = \langle V, E, T \rangle$, *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* $\mathscr{G}' = \langle V \cup \{v_R, v_S\}, E', v_R, v_S \rangle$, *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* $\mathscr{G}$ *equals the threshold budget of* v *in* $\mathscr{G}'$.

An important probabilistic connection was recognized in [35], which we formalize below.

Corollary 4. *[35] Consider a double-reachability bidding game* $\mathscr{G} = \langle V, E, v_R, v_S \rangle$. *For* $v \in V$, *we have* $\textsc{Th}(v) = val_{RTB(\mathscr{G})}(v)$.

We turn to study the problem of finding threshold budgets. Formally, the THRESH-BUDG problem gets as input a bidding game and a vertex v in it and the goal is to determine whether $\textsc{Th}(v) \ge 0.5$. It is shown in [34] that THRESH-BUDG is in NP, and it is explicitly stated that it is not known whether the problem is in P or NP-hard. Since finding the value in an SSG is in NP and coNP, using the corollary above we can slightly improve their result.

Theorem 5. *THRESH-BUDG for double-reachability games is in NP* $\cap$ *coNP.*

We continue to 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 6. *Consider a strongly-connected parity game* $\mathscr{G} = \langle V, E, p \rangle$. *There exists* $\rho \in \{0, 1\}$ *such that for every* $v \in V$, *we have* $\textsc{Th}(v) = \rho$. *Moreover, we have* $\rho = 0$ *iff* $\max_{v \in V} p(v)$ *is odd.*

Proof. The proof relies on the following claim: in a reachability bidding game with a target that is reachable from every vertex, Player 1 wins with every positive initial budget. The claim clearly implies the lemma as we view a strongly-connected bidding parity game as a reachability bidding game in which Player 1 tries to force the game to the vertex with the highest parity index. The proof of the claim follows from the fact that the threshold budget of a vertex $v \in V$ is some average between $\textsc{Th}(v_R)$ and $\textsc{Th}(v_S)$, and the average depends on the distances of v to the two targets. When only Player 1's target is reachable, we have $\textsc{Th}(v) = 0$. The details of the proof can be found in the full version. □

Lemma 6 allows us to reduce a bidding parity game $\mathscr{G} = \langle V, E, p \rangle$ to a double reachability game (and in turn to a reachability game). Consider a BSCC $S \subseteq V$. We call S *winning* for Player 1 if for every $v \in S$, we have $\textsc{Th}(v) = 0$. Otherwise, we have $\textsc{Th}(v) = 1$, and we call S *losing*. Then, in the double-reachability game, the target for Player 1 is the set of winning BSCCs and the target for Player 2 is the set of losing BSCCs. Formally, we have the following.

Lemma 7. *Parity bidding games are linearly reducible to reachability bidding games. Consider a bidding parity game* $\mathscr{G} = \langle V, E, p \rangle$, *and let* $W, L \subseteq V$ *be the set of vertices in BSCCs that are winning and losing for Player* 1, *respectively. For* $v \in V \setminus (L \cup W)$, $\textsc{Th}(v)$ *in* $\mathscr{G}$ *equals* $\textsc{Th}(v)$ *in the double-reachability game* $\langle V \setminus (L \cup W), E, L, W \rangle$.

Lemma 7 allows us to obtain the positive results of reachability bidding games in parity bidding games.

Theorem 8. *Threshold budgets exist in parity bidding games and THRESH-BUDG is in NP* $\cap$ *coNP.*

4 Mean-Payoff Bidding Games

This section consists of our most technically challenging contribution. We show that threshold budgets exist in mean-payoff bidding games and construct optimal strategies for the players. The crux of the proof considers the BSCCs of the game.

Consider a strongly-connected mean-payoff bidding game $\mathscr{G}$. For an initial budget $B \in [0,1]$, we denote by $\texttt{MP}^B(\mathscr{G})$ the *mean-payoff value* of $\mathscr{G}$, which is intuitively the optimal payoff Min can guarantee. More formally, suppose Min starts with a budget of B. Then, he can guarantee a payoff of at most $\texttt{MP}^B(\mathscr{G})$, and he cannot hope for more: for every $\varepsilon > 0$, Max can guarantee a payoff of at least $\texttt{MP}^B(\mathscr{G}) - \varepsilon$. Lemma 7 implies that the optimal payoff does not depend on the initial vertex.

Recall that $\texttt{RTB}(\mathscr{G})$ is a 2.5-player game in which the player who moves next is chosen uniformly at random, and that $\texttt{MP}(\texttt{RTB}(\mathscr{G}))$ is the optimal expected payoff both players can guarantee. We show the following probabilistic connection.

Theorem 9. *Consider a strongly-connected bidding mean-payoff game $\mathscr{G}$. The mean-payoff value of $\mathscr{G}$ exists, does not depend on the initial budget, and equals the mean-payoff value of the random-turn based mean-payoff game $\texttt{RTB}(\mathscr{G})$ in which the player who chooses the next move is selected uniformly at random, thus for every $B \in [0,1]$, we have $\texttt{MP}^B(\mathscr{G}) = \texttt{MP}(\texttt{RTB}(\mathscr{G}))$.*

We illustrate an existential proof for Theorem 9 the details of which can be found in the full version. The proof relies on the probabilistic connection and extensive work on *one-counter SSG* [13, 14]. In the following sections we show an alternative constructive proof. The draw-back of the existential proof is that it does not give any insight on how to construct optimal strategies. Indeed, a strategy in a 2.5-player game only prescribes which edges a player should choose and does not give any insight on how much to bid, which is the difficult part of constructing strategies in bidding games.

Reasoning about the payoff of a play is complicated. Instead, we reason about the energy of finite plays, which, recall, is the sum of weights that are traversed by the play. Consider a strongly-connected mean-payoff game $\mathscr{G}$. We view $\mathscr{G}$ as an energy game, thus Max wins an infinite play iff the energy in every finite prefix is positive. We construct a RTB game $\texttt{RTB}(\mathscr{G})$, which is a one-counter SSG. We show in the full version that the threshold budget of the energy game $\mathscr{G}$ that starts in a vertex v with energy $k \in \mathbb{N}$ is equivalent to the value of $\texttt{RTB}(\mathscr{G})$ that starts in v, and with a counter value of k. It is shown in [13, 14] that when $\texttt{MP}(\texttt{RTB}(\mathscr{G})) = 0$, then the value of the game is 1 for every initial counter value, thus Min has a strategy that wins with probability 1 no matter what the initial counter value is. Also, when $\texttt{MP}(\texttt{RTB}(\mathscr{G})) > 0$, then the value tends to 0 as the initial counter value increases. Lemma 10 shows that these properties suffice for showing that $\texttt{MP}(\mathscr{G}) = 0$. Its proof is simple and can be found in the full version.

Lemma 10. *Consider a strongly-connected bidding mean-payoff game $\mathscr{G}$ and a vertex u in $\mathscr{G}$.*

- *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 play $\pi = play(u, f_m, f_M)$ either reaches energy 0 or the energy is bounded by N throughout π. Then, Min can guarantee a non-positive payoff in $\mathscr{G}$.*
- *If for every initial budget for Max there exists an initial energy level $n \in \mathbb{N}$ such that Max can guarantee a non-negative energy level in $\mathscr{G}$, then Max can guarantee a positive payoff in $\mathscr{G}$.*

Similar to the qualitative case, we reduce mean-payoff bidding games to reachability games by first reasoning about the BSCCs of the game. Solving one-counter SSGs can be done in NP and coNP, thus we are obtain the same complexity upper bound as for reachability games. The details of the following theorem can be found in the full version.

Theorem 11. *Threshold budgets exist in mean-payoff bidding games and THRESH-BUDG is in NP $\cap$ coNP.*

4.1 Constructing an optimal strategy for Min

Consider a strongly-connected game $\mathscr{G}$ with $\texttt{MP}(\texttt{RTB}(\mathscr{G})) = 0$. We show the first direction in Theorem 9 and construct a strategy for Min that guarantees a non-positive payoff no matter what the initial budget is. We construct a Min positional strategy that guarantees that for every initial energy and every initial budget, either the energy level reaches 0 or it is bounded. By Lemma 10, this suffices for Min to guarantee a non-positive mean-payoff value in $\mathscr{G}$. We develop intuition for the construction in the following example.

Example 12. Consider the bidding mean-payoff game that is depicted in Figure 1. We show a Min strategy that guarantees a non-positive payoff no matter what the initial energy is. 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 10 implies that the payoff is non-positive, and we are done. □

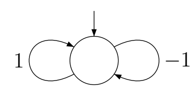

Figure 1: A bidding mean-payoff game where the weights are depicted on the edges.

Extending this result to general strongly connected games is not immediate. Consider a strongly-connected game $\mathscr{G} = \langle V, E, w\rangle$ 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.

Consider a strongly-connected mean-payoff bidding game $\mathscr{G} = \langle V, E, w\rangle$. We need a way to determine how "important" it is to move in a vertex in V. Our solution relies on the concept of *potential*, which was defined in the context of the strategy improvement algorithm to solve probabilistic games. Consider two optimal positional strategies f and g in $\texttt{RTB}(\mathscr{G})$, for Min and Max, respectively. Recall that when constructing $\texttt{RTB}(\mathscr{G})$, for every vertex $v \in V$, we add two copies v_{Min} and v_{Max}, that are controlled by Min and Max, respectively. For $v \in V$, let $v^-, v^+ \in V$ be such that $f(v_{Min}) = v^-$ and $g(v_{Max}) = v^+$. The *potential* of v, denoted $\text{Po}(v)$, is a known concept in probabilistic models and its existence is guaranteed [43]. We use the potential to define the *strength* of v, denoted $\text{St}^r(v)$, which intuitively measures how much the potentials of the neighbors of v differ. We assume w.l.o.g. that $\texttt{MP}(\texttt{RTB}(\mathscr{G})) = 0$ as otherwise we can decrease all weights by this value. The potential of v is a function that satisfies the following:

$$\text{Po}(v) = \frac{1}{2}\big(\text{Po}(v^+) + \text{Po}(v^-)\big) + w(v) \text{ and } \text{St}(v) = \frac{1}{2}\big(\text{Po}(v^+) - \text{Po}(v^-)\big)$$

There are optimal strategies for which $\mathrm{Po}(v^-) \leq \mathrm{Po}(v') \leq \mathrm{Po}(v^+)$, for every $v' \in N(v)$, which can be found for example using the strategy iteration algorithm.

Consider a finite path $\pi = v_1, \ldots, v_n$ in $\mathscr{G}$. We intuitively think of π as a play, where for every $1 \leq i < n$, the bid of Min in v_i is $\mathrm{St}(v_i)$ and he moves to v_i^- upon winning. Thus, if $v_{i+1} = v_i^-$, we say that Min won in v_i, and if $v_{i+1} \neq v_i^-$, we say that Min lost in v_i. Let $W(\pi)$ and $L(\pi)$ respectively be the indices in which Min wins and loses in π. We call Min wins *investments* and Min loses *gains*, where intuitively he *invests* in increasing the energy and *gains* a higher ratio of the budget whenever the energy decreases. Let $G(\pi)$ and $I(\pi)$ be the sum of gains and investments in π, respectively, thus $G(\pi) = \sum_{i \in L(\pi)} \mathrm{St}(v_i)$ and $I(\pi) = \sum_{i \in W(\pi)} \mathrm{St}(v_i)$. Recall that the energy of π is $E(\pi) = \sum_{1 \leq i < n} w(v_i)$. The following lemma connects the strength, potential, and accumulated energy.

Lemma 13. *Consider a strongly-connected game $\mathscr{G}$ such that $\mathtt{MP}(\mathtt{RTB}(\mathscr{G})) = 0$, and a finite path π in $\mathscr{G}$ from v to u. Then, $Po(v) - Po(u) \leq E(\pi) - G(\pi) + I(\pi)$.*

Proof. We prove by induction on the length of π. For $n = 1$, the claim is trivial since both sides of the equation are 0. Suppose the claim is true for paths of length n and we prove for paths of length $n+1$. We distinguish between two cases. In the first case, Min wins in v, thus the second vertex in π is v^-. Let π' be the prefix of π starting from v^-. Note that since Min wins the first bidding, we have $G(\pi) = G(\pi')$ and $I(\pi) = \mathrm{St}(v) + I(\pi')$. Also, we have $E(\pi) = E(\pi') + w(v)$. Combining these, we have $E(\pi) - G(\pi) + I(\pi) = E'(\pi) + w(v) - G(\pi') + I(\pi') + \mathrm{St}(v)$. By the induction hypothesis, we have $\mathrm{Po}(v^-) - \mathrm{Po}(u) \leq E(\pi') - G(\pi') + I(\pi')$. Combining these with the definition of $\mathrm{St}(v)$, we have the following.

$$E(\pi) - G(\pi) + I(\pi) \geq \mathrm{St}(v) + \mathrm{Po}(v^-) + w(v) - \mathrm{Po}(u) =$$

$$= \frac{1}{2}\big(- \mathrm{Po}(v^-) + \mathrm{Po}(v^+)\big) + \mathrm{Po}(v^-) + w(v) - \mathrm{Po}(u) = \mathrm{Po}(v) - \mathrm{Po}(u)$$

We continue to the second case in which Max wins in v and let v' be the second vertex in π. Recall that we have $\mathrm{Po}(v^+) \geq \mathrm{Po}(v')$. Dually to the first case, we have $G(\pi) = \mathrm{St}(v) + G(\pi')$ and $I(\pi) = I(\pi')$. Combining with the induction hypothesis, we have $E(\pi) - G(\pi) + I(\pi) \geq w(v) + \mathrm{St}(v) + \mathrm{Po}(v^+) - \mathrm{Po}(u)$. Plugging in $\mathrm{St}(v) = \frac{1}{2}\big(\mathrm{Po}(v^+) - \mathrm{Po}(v^-)\big)$, we have $E(\pi) - G(\pi) + I(\pi) \geq \frac{1}{2}\big(\mathrm{Po}(v^-) + \mathrm{Po}(v^+)\big) - \mathrm{Po}(u) = \mathrm{Po}(v) - \mathrm{Po}(u)$, and we are done. □

We are ready to describe a strategy f_m for Min. Consider a positive initial budget $B \in (0,1]$ for Min and an initial energy $k_I \in \mathbb{N}$. Let $\mathrm{Po}_M = \max_{v \in V} |\mathrm{Po}(v)|$ and $\mathrm{St}_M = \max_{v \in V} |\mathrm{St}(v)|$. We choose $N \in \mathbb{N}$ such that $B > \frac{k_I + \mathrm{St}_M + 2\mathrm{Po}_M}{N}$. When the game reaches $v \in V$, Min bids $\mathrm{St}(v)/N$ and moves to v^- upon winning. We formalize the intuition of tying energy and budget by means of an invariant that is maintained throughout a play.

Lemma 14. *Consider a Max strategy f_M, and let $\pi = play(f_m, f_M)$ be a finite play after which the energy is k, i.e., $k = k_I + E(\pi)$. Then, Min's budget following π is at least $\frac{k + St_M}{N}$.*

Proof. The invariant clearly holds initially. Let B be Min's initial budget and B' his budget following π, thus $B' = B + \big(G(\pi) - I(\pi)\big)/N$. From Lemma 13, we have $2\mathrm{Po}_M \leq E(\pi) - \big(G(\pi) - I(\pi)\big)$. Recall that that $B > \frac{k_I + \mathrm{St}_M + 2\mathrm{Po}_M}{N}$. We combining with $k = k_I + E(\pi)$ and re-arranging:

$$2\mathrm{Po}_M \leq \frac{k - k_I}{N} - B' + B \leq \frac{k - k_I}{N} - B' + B \leq \frac{k - k_I}{N} - B' + \frac{k_I + \mathrm{St}_M + 2\mathrm{Po}_M}{N}$$

□

Lemma 14 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 14 implies that if the energy does not reach 0, then it is bounded by $N - \mathrm{St}_M$. Combining with Lemma 10, we have the following.

Theorem 15. *Given a positive initial budget for Min and any initial initial energy, the strategy f_m guarantees that either an energy level* 0 *is reached or that the energy is bounded from above.*

4.2 Constructing an optimal strategy for Max

Constructing an optimal strategy for Max is more involved since our definition of payoff gives Min the advantage. Consider a strongly-connected bidding mean-payoff game $\mathscr{G}$ with $\mathtt{MP}(\mathtt{RTB}(\mathscr{G})) > 0$. Given an initial budget B for Max, we provide an initial energy level k_I such that when the game starts at k_I, Max keeps the energy level non-negative. By Lemma 10, this suffices to construct a strategy that guarantees a positive payoff in $\mathscr{G}$. In this section we describe the construction for a fragment of the general case called *recurrent SCCs*. The difficulties of the general case and the main ideas in its solution already appear in this simpler case, and the extension to general games can be found in the full version. An SCC $G = \langle V, E \rangle$ is 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. We consider mean-payoff games played on recurrent SCCs.

We describe the intuition of the construction. A first attempt for constructing a Max strategy would be to tie energy and budget as in the strategy for Min in the previous section. This attempt fails since Min can allow Max to win for a while, draw the energy to N, where Max's budget runs out. When Min has all (or most) of the budget, he can win an arbitrary number of biddings in a row, which means that he can draw the energy arbitrary low, causing Max to lose. Avoiding the exhaustion of the budget is the key difficulty in constructing Max's strategy. The moral of this attempt is that the "normalization factor", which was $1/N$ in the previous section, must decrease as the energy increases.

We split the natural numbers into blocks, and define a different normalization factor in each block, which we call the *currency* of the block. For $n \in \mathbb{N}$, when the energy is in the n-th block, Max bids in the currency of the n-th block. Inside the n-th block, we tie between energy and budget, thus if the energy stays within a block, we are done. The difficulty is handling plays that switch between blocks. Indeed, consider a play with a sinusoidal energy behavior: Min loses for a bit and causes the energy to increase by, say, c units from the top of Block n to the bottom of Block $n+1$ and then wins for a bit causing the energy to decrease by c units into the top of Block n. Since energy is tied with budget, every time the energy increases, Max "invests" c units in the higher currency of the n-th block, and every time the energy decreases, he "gains" c units in the lower currency of the $(n+1)$-th block. His budget will eventually run out. To overcome this issue, we develop the idea of tying energy and budget to give an advantage to Max: investing is done in the currency of the current block while gaining is done in the higher currency of the lower block.

We formalize this intuition. Consider a strongly-connected recurrent mean-payoff bidding game $\mathscr{G} = \langle V, E, w \rangle$ with $\mathtt{MP}(\mathtt{RTB}(\mathscr{G})) > 0$. We alter the weights to give advantage to Min. For $z > 1$, let $\mathscr{G}^z = \langle V, E, w^z \rangle$, where $w^z(v) = w(v)$ if $w(v) \geq 0$ and otherwise $w^z(v) = z \cdot w(v)$. Clearly, $\mathtt{MP}(\mathtt{RTB}(\mathscr{G})) \geq \mathtt{MP}(\mathtt{RTB}(\mathscr{G}^z))$. We select $z > 1$ such that $\mathtt{MP}(\mathtt{RTB}(\mathscr{G}^z)) \geq 0$. We respectively denote by Po^z and St^z, the potential and strength functions of $\mathscr{G}^z$ as in the previous section. We define the partition into energy blocks. Let $cycles(u)$ be the set of simple cycles from u to itself and $E_M = \max_{\pi \in cycles(u)} |E(\pi)|$. We choose $M \in \mathbb{N}$ such that $M \geq (\mathrm{St}^z_M + 3E_M)/(1 - z^{-1})$, where St^z_M is the maximal strength 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 $M_n = \{M(n-1), M(n-1)+1, \ldots, Mn-1\}$. We use $\beta_n^{\downarrow}$ and $\beta_n^{\uparrow}$ to mark the upper and lower boundaries of M_n, respectively. We use a $M_{\geq n}$ to denote the set $\{M_n, M_{n+1}, \ldots\}$. Consider a finite play 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 $k_I + E(\pi_1, \ldots, \pi_j) \in M_n$.

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 $z^{-n} \cdot \mathrm{St}^z(v)$ and proceeds to v^+ upon winning. Thus, the currency in M_n is z^{-n}. Note that currency changes occur only in u. We formalize the asymmetry between "gaining" and "investing". Let $E^z(\pi)$ be the change in energy in $\mathscr{G}^z$.

Lemma 16. *Consider an play $\pi \in cycles(u)$. Then, $E(\pi) \geq E^z(\pi)$ and $zE(\pi) \geq E^z(\pi)$.*

Proof. Let $E^{\geq 0}(\pi)$ and $E^{<0}(\pi)$ be the sum of non-negative weights and negative weights in π, respectively. We have $E(\pi) = E^{\geq 0}(\pi) + E^{<0}(\pi)$ and $E^z(\pi) = E^{\geq 0}(\pi) + zE^{<0}(\pi)$. The inequality $E(\pi) \geq E^z(\pi)$ is immediate. For the second inequality, we multiply the first equality by z and subtract it from the first to get $E^z(\pi) - zE(\pi) = E^{\geq 0}(\pi) - zE^{\geq 0}(\pi) \leq 0$, and we are done. □

Let $pay_M(\pi)$ denote the amount that Max pays in a finite play π, thus it is negative when Max gains budget. Adjusting Lemma 13 to our setting, for $\pi \in cycles(u)$, we have $E^z(\pi) \geq z^n \cdot pay_M(\pi)$. Intuitively, the corollary states that in M_n, investing in the increase of energy is done in the currency z^{-n} of M_n while gaining due to decrease of energy is done in the higher currency $z^{-(n-1)}$ of M_{n-1}.

Corollary 17. *Consider a Min strategy f_m, and let $\pi = play(f_m, f_M)$ be a finite play such that $\pi \in cycles(u)$. Then, we have $E(\pi) \geq z^n \cdot pay_M(\pi)$ and $zE(\pi) \geq z^n \cdot pay_M(\pi)$.*

Consider an initial Max budget $B_M^{init} \in [0,1]$. 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 $\beta_n^{\downarrow}$ of M_n to its upper boundary $\beta_n^{\uparrow}$, costs $M \cdot z^{-n}$. We define $cost(M_n) = M \cdot z^{-n}$ and $cost(M_{\geq n}) = \sum_{i=n}^{\infty} cost(M_n)$. Finally, we need some *wiggle room* to allow for changes in the currency. Let $wiggle = 2E_M + \mathrm{St}_M^z$.

Definition 18. *Let k_I be $\beta_n^{\downarrow}$ such that $B_M^{init} > wiggle \cdot z^{-(n-1)} + cost(M_{\geq n})$ and $\sum_{i=1}^{n} cost(M_i) > 1$.*

Consider a Min strategy f_m, and let $\pi = play(f_m, f_M)$ be a finite play. We partition π into subsequences in which the same currency is used. Let $\pi = \pi_1 \cdot \pi_2 \cdot \ldots \cdot \pi_\ell$ be a partition of π. For $1 \leq i \leq \ell$, we use π^i to refer to the prefix $\pi_1 \cdot \ldots \cdot \pi_i$ 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 $\mathscr{G}$ is recurrent imply that upon entering M_n, the energy is within E_M of the boundary. Thus, in the case that π^i is decreasing, the energy at the end of π^i is $e^i \geq \beta_n^{\uparrow} - E_M$ and in the case it is increasing, we have $e^i \leq \beta_n^{\downarrow} + E_M$. Let $\ell_0 = 0$, and for $i \geq 1$, let $\ell_i = (\beta_{n+1}^{\downarrow} - E_M) - e^i$ in the first case and $\ell_i = (\beta_n^{\downarrow} + E_M) - e^i$ in the second case. Note that $\ell_i \in \{0, \ldots, 2E_M\}$. In the full version, we prove the following invariant on Max's budget when changing between energy blocks.

Lemma 19. *For every $i \geq 0$, suppose π^i ends in M_n. Then, Max's budget is at least $(wiggle + \ell_i) \cdot z^{-(\hat{n}-1)} + cost(M_{\geq \hat{n}})$, 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 19 implies that f_M is legal. That is, consider a finite play π that starts immediately after a change in currency. Using Lemma 13, 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 19 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 10 implies that Max guarantees a positive mean-payoff value in a strongly-connected game.

Theorem 20. *Consider a strongly-connected recurrent mean-payoff bidding game $\mathcal{G}$ with $\mathtt{MP}(\mathtt{RTB}(\mathcal{G})) > 0$. Suppose the game starts with a positive initial budget for Max and the corresponding initial energy as in Definition 18. The strategy f_M guarantees that the energy stays non-negative with these initial values.*

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. We showed existence of threshold budgets in parity and mean-payoff bidding games and constructed optimal strategies for the players. This work belongs to a line of works that transfer concepts and ideas between the areas of formal verification and algorithmic game theory [39], two fields with a different take on game theory and with complementary needs. For example, formally reasoning about multi-agent safety critical systems, e.g., components of an autonomous car, requires insights on rationality. On the other side, formally verifying the correctness of auctions or reasoning about ongoing auctions, are both challenges that can benefit from the experience of the formal methods community. Examples of works in the intersection of the two fields include logics for specifying multi-agent systems [3, 18, 38], studies of equilibria in games related to synthesis and repair problems [17, 16, 24, 1], non-zero-sum games in formal verification [12, 15, 20], and applying concepts from formal methods to *resource allocation games* such as rich specifications [9], efficient reasoning about very large games [5, 32], and a dynamic selection of resources [8].

The concept of bidding for moving is general and there are endless ways to take it. We list several directions for future studies. It is interesting to study further bidding rules and add concepts like multi-player games or partial information. A step towards the later is a better understanding of poorman bidding in which the higher bidder pays the "bank" rather than the other player. Since there is no constraint that the sum of budgets is 1, poorman bidding is more amendable for extensions. Infinite-duration poorman bidding games were studied in [7]. We describe some of their results. The central quantity that is studied in poorman bidding is the ratio of the total budget. For reachability games, the good news that are shown in [34] is that "threshold ratios" exist, similar to threshold budgets in Richman games. The bad news, however, are that poorman games are more complicated in that threshold ratios need not be rational and a probabilistic connection does not exist. Extending the good news to richer qualitative objectives is similar to the Richman case. Things get interesting in mean-payoff poorman games. There, in strongly-connected games, unlike the Richman case, the value of the game depends on the initial ratio. Moreover and quite surprisingly, the probabilistic connection pops up: Suppose Max's initial ratio is $r \in [0,1]$, then the value of a strongly-connected poorman mean-payoff game $\mathcal{G}$ equals the mean-payoff value of a RTB that is played on $\mathcal{G}$ in which Max chooses a move with probability r and Min with probability $1-r$. Interestingly, in a strongly-connected mean-payoff game $\mathcal{G}$, the value of $\mathcal{G}$ when viewed as a poorman game with initial ratio 0.5 equals the value of $\mathcal{G}$ when viewed as a Richman game. The constructions for poorman games use the ideas developed in this paper.

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*, pages 70–76, 2017.

[6] G. Avni, T. A. Henzinger, and V. Chonev. Infinite-duration bidding games. *CoRR*, abs/1705.01433, 2017. `https://arxiv.org/abs/1705.01433`.

[7] G. Avni, T. A. Henzinger, and R. Ibsen-Jensen. Infinite-duration poorman-bidding games. *CoRR*, abs/1804.04372, 2018. `https://arxiv.org/abs/1804.04372`.

[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. Network-formation games with regular objectives. *Inf. Comput.*, 251:165–178, 2016.

[10] N. Berger, N. Kapur, L. J. Schulman, and V. V. Vazirani. Solvency games. In *Proc. 28th FSTTCS*, pages 61–72, 2008.

[11] J. Bhatt and S. Payne. Bidding chess. *Math. Intelligencer*, 31:37–39, 2009.

[12] P. Bouyer, R. Brenguier, N. Markey, and M. Ummels. Pure Nash equilibria in concurrent deterministic games. *Logical Methods in Computer Science*, 11(2), 2015.

[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] 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.

[16] K. Chatterjee. Nash equilibrium for upward-closed objectives. In *Proc. 15th Annual Conf. of the European Association for Computer Science Logic*, volume 4207 of *Lecture Notes in Computer Science*, pages 271–286. Springer, 2006.

[17] K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Games with secure equilibria. *Theor. Comput. Sci.*, 365(1-2):67–82, 2006.

[18] K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. *Inf. Comput.*, 208(6):677–693, 2010.

[19] K. Chatterjee, R. Majumdar, and T. A. Henzinger. Controller synthesis with budget constraints. In *Proc. 11th HSCC*, pages 72–86, 2008.

[20] K. Chatterjee, R. Majumdar, and M. Jurdzinski. On Nash equilibria in stochastic games. In *Proc. 13th CSL*, pages 26–40, 2004.

[21] A. Condon. On algorithms for simple stochastic games. In *Proc. DIMACS*, pages 51–72, 1990.

[22] M. Develin and S. Payne. Discrete bidding games. *The Electronic Journal of Combinatorics*, 17(1):R85, 2010.

[23] 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.

[24] D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *Proc. 16th TACAS*, pages 190–204, 2010.

[25] A. J. Hoffman and R. M. Karp. On nonterminating stochastic games. *Management Science*, 12(5):359—370, 1966.

[26] M. Holtmann, L. Kaiser, and W. Thomas. Degrees of lookahead in regular infinite games. *Logical Methods in Computer Science*, 8(3), 2012.

[27] F. A. Hosch and L. H. Landweber. Finite delay solutions for sequential conditions. In *ICALP*, pages 45–60, 1972.

[28] K. Johnson, D. Simchi-Levi, and P. Sun. Analyzing scrip systems. *Operations Research*, 62(3):524–534, 2014.

[29] G. Kalai, R. Meir, and M. Tennenholtz. Bidding games and efficient allocations. In *Proc. 16th EC*, pages 113–130, 2015.

[30] 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.

[31] F. Klein and M. Zimmermann. How much lookahead is needed to win infinite games? *Logical Methods in Computer Science*, 12(3), 2016.

[32] O. Kupferman and T. Tamir. Hierarchical network formation games. In *Proc. 23rd TACAS*, pages 229–246, 2017.

[33] U. Larsson and J. Wästlund. Endgames in bidding chess. *Games of No Chance 5*, 70, 2018.

[34] 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.

[35] A. J. Lazarus, D. E. Loeb, J. G. Propp, and D. Ullman. Richman games. *Games of No Chance*, 29:439–449, 1996.

[36] R. Paes Leme, V. Syrgkanis, and É. Tardos. Sequential auctions and externalities. In *Proc. 23rd SODA*, pages 869–886, 2012.

[37] M. Menz, J. Wang, and J. Xie. Discrete all-pay bidding games. *CoRR*, abs/1504.02799, 2015.

[38] F. Mogavero, A. Murano, G. Perelli, and M. Y. Vardi. Reasoning about strategies: On the model-checking problem. *ACM Trans. Comput. Log.*, 15(4):34:1–34:47, 2014.

[39] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[40] Y. Peres, O. Schramm, S. Sheffield, and D. B. Wilson. *Tug-of-War and the Infinity Laplacian*, pages 595–638. Springer New York, 2011.

[41] Y. Peres, O. Schramm, S. Sheffield, and D. Bruce Wilson. Random-turn hex and other selection games. *The American Mathematical Monthly*, 114(5):373–387, 2007.

[42] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. 16th POPL*, pages 179–190, 1989.

[43] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, USA, 2005.

[44] M.O. Rabin. Decidability of second order theories and automata on infinite trees. *Transaction of the AMS*, 141:1–35, 1969.

[45] D. M. Reeves, B. M. Soule, and T. Kasturi. Yootopia! *SIGecom Exchanges*, 6(2):1–26, 2007.

[46] M. Stonebraker, P. M. Aoki, W. Litwin, A. Pfeffer, A. Sah, J. Sidell, C. Staelin, and A. Yu. Mariposa: A wide-area distributed database system. *VLDB J.*, 5(1):48–63, 1996.