

Computing Threshold Budgets in Discrete-Bidding Games

Received Jan 4, 2024
Revised Sep 27, 2024
Accepted Nov 23, 2024
Published Jan 22, 2025
Updated Jul 2, 2025

Key words and phrases

Discrete bidding games, Richman games, parity games, reachability games

Guy Avni^a ✉ 

Suman Sadhukhan^a ✉ 

^a Department of Computer Science, University of Haifa, Israel

ABSTRACT. In a two-player zero-sum graph game, the players move a token throughout a graph to produce an infinite play, which determines the winner of the game. *Bidding games* are graph games in which in each turn, an auction (bidding) determines which player moves the token: the players have budgets, and in each turn, both players simultaneously submit bids that do not exceed their available budgets, the higher bidder moves the token, and pays the bid to the lower bidder (called *Richman* bidding). We focus on *discrete*-bidding games, in which, motivated by practical applications, the granularity of the players' bids is restricted, e.g., bids must be given in cents.

A central quantity in bidding games is *threshold budgets*: a necessary and sufficient initial budget for winning the game. Previously, thresholds were shown to exist in parity games, but their structure was only understood for reachability games. Moreover, the previously-known algorithms have a worst-case exponential running time for both reachability and parity objectives, and output strategies that use exponential memory. We describe two algorithms for finding threshold budgets in parity discrete-bidding games. The first is a fixed-point algorithm. It reveals, for the first time, the structure of threshold budgets in parity discrete-bidding games. Based on this structure, we develop a second algorithm that shows that the problem of finding threshold budgets is in NP and coNP for both reachability and parity objectives. Moreover, our algorithm constructs strategies that use only linear memory.

This is a corrected version of the paper (<https://arxiv.org/abs/2210.02773v4>) published originally on Jan 22, 2025.

This research was supported in part by ISF grant no. 1679/21. A preliminary version of this article appeared at FSTTCS 22 [14].

1. Introduction

Two-player zero-sum *graph games* are a central class of games. A graph game proceeds as follows. A token is placed on a vertex and the players move it throughout the graph to produce an infinite play, which determines the winner of the game. The central algorithmic problem in graph games is to identify the winner and to construct winning strategies. One key application of graph games is *reactive synthesis* [29], in which the goal is to synthesize a *reactive system*, namely a policy for interacting with an adversarial environment, that satisfies a given specification no matter how the environment behaves.

Two orthogonal classifications of graph games are according to the *mode* of moving the token and according to the players' *objectives*. For the latter, we focus on two canonical qualitative objectives. In *reachability* games, there is a set of target vertices and Player 1 wins if a target vertex is reached. In *parity* games, each vertex is labeled with a parity index and an infinite path is winning for Player 1 iff the highest parity index that is visited infinitely often is odd. The simplest and most studied mode of moving is *turn-based*: the players alternate turns in moving the token. We note that reactive synthesis reduces to solving a turn-based parity game. Examples of other modes of moving are *concurrent* and *probabilistic* moves (see [4]).

We study *bidding graph games* [25, 24], which apply the following mode of moving: both players have budgets, and in each turn, an auction (bidding) determines which player moves the token. Concretely, we focus on *Richman bidding* (named after David Richman): in each turn, both players simultaneously submit bids that do not exceed their available budget, the higher bidder moves the token, and pays his bid to the lower bidder. Note that the sum of budgets stays constant throughout the game. We distinguish between *continuous*- and *discrete*-bidding, where in the latter, the granularity of the players' bids is restricted. The central questions in bidding games revolve around the *threshold budgets*, namely a necessary and sufficient initial budget for winning the game.

Continuous-bidding games. This paper focuses on discrete-bidding. We briefly survey the relevant literature on continuous-bidding games, which have been more extensively studied than their discrete-bidding counterparts. Bidding games were introduced in [25, 24]. The objective that was considered is a variant of reachability, which we call *double reachability*: each player has a target and a player wins if his target is reached (unlike reachability games in which Player 2's goal is to prevent Player 1 from reaching his target). The targets are assumed to be distinct, thus the game is zero sum. Note that a priori, it is possible that a play results in a tie if no target is visited. It was shown, however, that in continuous-bidding games, a target is necessarily reached, thus double-reachability games essentially coincide with reachability games continuous-bidding games.

Threshold budgets were shown to exist; namely, each vertex v has a value $\text{Th}(v)$ such that if Player 1's budget is strictly greater than $\text{Th}(v)$, he wins the game from v , and if his budget is strictly less than $\text{Th}(v)$, Player 2 wins the game. Moreover, it was shown that the threshold function Th is the *unique* function that satisfies the following property, which we call the *average property*. Suppose that the sum of budgets is 1, and t_i is Player i 's target, for $i \in \{1, 2\}$. Then, Th assigns a value in $[0, 1]$ to each vertex such that at the “end points”, we have $\text{Th}(t_1) = 0$ and $\text{Th}(t_2) = 1$, and the threshold at every other vertex is the average of two of its neighbors.

Uniqueness implies that the problem of finding threshold budgets¹ is in NP and coNP. Moreover, an intriguing equivalence was observed between reachability continuous-bidding games and a class of stochastic game [19] called *random-turn games* [28]. Intricate equivalences between *mean-payoff* continuous-bidding games and random-turn games have been shown in [8, 9, 10, 11] (see also [7]).

Parity continuous-bidding games were studied in [8]. The following key property was identified in games played on strongly-connected graphs. With every positive initial budget, a player can force visiting all vertices in the graph infinitely often. Consider a strongly-connected parity continuous-bidding game \mathcal{G} . It follows that if the maximal parity index in \mathcal{G} is odd, then Player 1 wins with any positive initial budget, i.e., the thresholds in \mathcal{G} are all 0. Dually, if the maximal parity index in \mathcal{G} is even, then the thresholds are all 1. This property gives rise to a simple reduction from parity bidding games to double-reachability bidding games: roughly, a player's goal is to reach a bottom strongly-connected component in which he can win with any positive initial budget.

Discrete-bidding games. This paper studies *discrete-bidding games*, which are similar to continuous-bidding games except that the granularity of the bids is restricted: the sum of the budgets in the game is fixed to $k \in \mathbb{N}$ and bids are restricted to be integers. A key difference between continuous- and discrete-bidding games is bidding ties, which now need to be handled explicitly. We focus on the tie-breaking mechanism that was defined in [20]: one of the players has the *advantage* and when a tie occurs, the player with the advantage chooses between (1) use the advantage to win the bidding and pass it to the other player, or (2) keep the advantage and let the other player win. Other tie-breaking mechanisms and the properties that they lead to were studied in [1]. For example, the tie-breaking mechanism “alternate tie breaking” in which the players alternate turns in winning ties, is not determined, i.e., there is a game with an initial position such that neither player has a (pure) winning strategy. On the other hand, it was shown that any tie-breaking mechanism that breaks ties without considering the history of ties, admits determinacy.

The motivation to study discrete-bidding games is practical: in most applications, the assumption that bids can have arbitrary granularity is unrealistic. We point out that the results

¹ Stated as a decision problem: given a game and a vertex v , decide whether $\text{Th}(v) \geq 0.5$.

in continuous-bidding games, particularly those on infinite-duration games, do in fact develop strategies that bid arbitrarily small bids. It is highly questionable whether such strategies are applicable in practical applications.

Bidding games model ongoing and stateful auctions. We list examples of domains in which such auctions arise. An immediate example is auctions for online advertisements [27]. Bidding games were applied in [12] as a scheduling mechanism in a “decoupled” synthesis procedure: given an objective of the form $\psi_1 \wedge \psi_2$, the idea is to find, independently, two bidding strategies f_1 for ψ_1 and f_2 for ψ_2 and an initial budget allocation such that the strategies guarantee that the outcome of playing f_1 against f_2 satisfies $\psi_1 \wedge \psi_2$. For example, the strongest guarantees are obtained when each strategy f_i , for $i = 1, 2$, is winning for ψ_i , i.e., it guarantees ψ_i even against an adversary. Bidding as a mechanism for scheduling arises in *blockchain* technology, where *miners* accept transaction fees, which can be thought of as bids, and prioritize transactions based on them. Verification against attacks is a well-studied problem [18, 5]. Attacks based on manipulations of these fees are hard to detect, can cause significant losses, and thus call for verification of the protocols [18, 5]. Another example is applying bidding games as a mechanism for fair allocation of resources. Non-zero-sum Richman-bidding games were studied and applied to resource allocation in [26] and *poorman*-bidding games (in which winning bids are paid to the “bank”) were applied in [15]. Poorman discrete-bidding were studied in [13]. In addition, researchers have studied training of agents that accept “advice” from a “teacher”, where the advice is equipped with a “bid” that represents its importance [3]. Finally, recreation bidding games have been studied, e.g., bidding chess [16], as well as combinatorial games that apply bidding instead of alternating turns [30].

We reiterate that practical applications of bidding games require some granularity on the bids. At the same time, we seek a high granularity to enable flexibility. A high granularity translates to choosing a large sum of budgets k .

Previous results. For reachability objectives, the theory of continuous bidding games was largely adapted to discrete-bidding in [20]: threshold budgets were shown to exist and satisfy a discrete version of the average property and winning strategies are derived from the threshold budgets. However, the only known algorithm to compute thresholds is a value-iteration algorithm whose worst-case running time is exponential when k is given in binary.

For parity discrete-bidding games, there were large gaps in our understanding. In [1], *determinacy* was shown, namely from each configuration of the game, one of the players has a (pure) winning strategy. Determinacy is achieved by showing that the game satisfies a local property, which implies “global” determinacy. Using the observation that an additional budget will not harm a player, we obtain existence of thresholds. Importantly, this technique does not show that the thresholds satisfy the average property, and it was left open whether they indeed satisfy the average property. In terms of complexity, the algorithm to decide the winner from a

configuration is naive: construct and solve the explicit concurrent game that corresponds to a bidding game. The running time of the algorithm is exponential when k is given in binary. Another disadvantage of the algorithm is that the strategies that it produces use exponential memory and do not connect between bids and thresholds, as is done in reachability discrete-bidding games. To make matters worse, it was observed that unlike the properties of thresholds in reachability discrete-bidding games, which are conceptually similar to those in continuous-bidding, thresholds in parity discrete- and continuous-bidding games are inherently different: a strongly-connected discrete-bidding game \mathcal{G} was shown in which a player cannot force visiting a vertex v infinitely often even if he is initially allocated all of the budget. That is, when \mathcal{G} is a Büchi game and v is the only accepting vertex, then under continuous-bidding semantics, the threshold are 0 whereas under discrete-bidding, the thresholds are $k + 1$ (meaning that even a budget of k does not suffice for winning).

Our results. We develop two complementary algorithms for computing threshold budgets in parity discrete-bidding games. Our first algorithm is a fixed-point algorithm. It repeatedly solves (i.e., finds threshold budgets) in *frugal-reachability* bidding games, which is an objective that we introduce in which on top of a reachability objective, in order to win, a player must reach its target with a sufficient budget. Our algorithm is inspired by algorithms to solve turn-based games such as Zielonka’s [31] and Kupferman and Vardi’s [23] algorithms, whereas continuous-bidding games reduce to stochastic games. Recently, the fixed-point algorithm was adapted to bidding games *with charging* [6]. This algorithm shows, for the first time, that threshold budgets in parity discrete-bidding games satisfy the average property. Moreover, the strategies that it produces are derived from the thresholds, as in reachability discrete-bidding games. On the downside, the algorithm runs in exponential time when k is given in binary.

Second, we show that the problem of finding threshold budgets in parity discrete-bidding games² is in NP and coNP. The bound applies also to reachability discrete-bidding games for which only an exponential-time algorithm was known. We briefly describe the idea of our proof. A first attempt to find thresholds is to guess thresholds (this is possible since the budgets are discrete) and verify that the guess satisfies the average property (recall that in continuous-bidding games, functions that satisfy the average property are unique). We show, however, that functions that satisfy the discrete average property are not unique. That is, even if a function satisfies the average property, it might not represent the thresholds in the game. We point out that this observation holds already in reachability discrete-bidding games, and to the best of our knowledge, was never made before. We overcome this challenge as follows. Our algorithm first guesses a function, checks whether it satisfies the average property, then verifies that it coincides with the thresholds. This last step is done via a reduction to turn-based parity games and is based on the structure of the thresholds and strategies that our first algorithm establishes.

2 Formally, given a discrete-bidding game \mathcal{G} , a vertex v , and a threshold ℓ , decide whether $\text{Th}(v) \geq \ell$.

Another advantage of this algorithm is that it outputs a strategy that can be implemented using linear memory (previously, only construction of exponential-size strategies was known).

The new version of this paper corrects typos that previously appeared in crucial places and addresses the corner case of losing vertices in Lem. 5.2 and Lem. 5.4.

2. Preliminaries

2.1 Concurrent games

We define the formal semantics of bidding games via two-player *concurrent games* [2]. Intuitively, a concurrent game proceeds as follows. A token is placed on a vertex of a graph. In each turn, both players concurrently select actions, and their joint actions determine the next position of the token. The outcome of a game is an infinite path. A game is accompanied by an objective, which specifies which plays are winning for Player 1. In this paper, we will consider *reachability* and *parity* objectives.

Formally, a concurrent game is played on an *arena* $\langle A, Q, \lambda, \delta \rangle$, where A is a finite non-empty set of actions, Q is a finite non-empty set of states (in order to differentiate, we use “states” or “configurations” in concurrent games and “vertices” in bidding games), the function $\lambda : Q \times \{1, 2\} \rightarrow 2^A \setminus \{\emptyset\}$ specifies the allowed actions for Player i in vertex v , and the transition function is $\delta : Q \times A \times A \rightarrow Q$. Suppose that the token is placed on a state $q \in Q$ and, for $i \in \{1, 2\}$, Player i chooses action $a_i \in \lambda(q, i)$. Then, the token moves to $\delta(q, a_1, a_2)$. For $q, q' \in Q$, we call q' a *neighbor* of q if there is a pair of actions $\langle a_1, a_2 \rangle \in \lambda(q, 1) \times \lambda(q, 2)$ with $q' = \delta(q, a_1, a_2)$. We denote the neighbors of q by $N(q) \subseteq Q$.

A (finite) *history* is a sequence $\langle q_0, a_0^1, a_0^2 \rangle, \dots, \langle q_{n-1}, a_{n-1}^1, a_{n-1}^2 \rangle, q_n \in (Q \times A \times A)^* \cdot Q$ such that, for each $0 \leq i < n$, we have $q_{i+1} = \delta(q_i, a_i^1, a_i^2)$. A *strategy* is a “recipe” for playing the game. Formally it is a function $\sigma : (Q \times A \times A)^* \cdot Q \rightarrow A$. We restrict attention to *legal* strategies; namely, strategies that for each history $\pi \in (Q \times A \times A)^* \cdot Q$ that ends in $q \in Q$, choose an action in $\lambda(q, i)$, for $i \in \{1, 2\}$. A *memoryless* strategy is a strategy that, for every state $q \in Q$, assigns the same action to every history that ends in q .

Two strategies σ_1 and σ_2 for the two players and an initial state q_0 , give rise to a unique *play*, denoted $\text{play}(q_0, \sigma_1, \sigma_2)$, which is a sequence in $(Q \times A \times A)^\omega$ and is defined inductively as follows. The first element of $\text{play}(q_0, \sigma_1, \sigma_2)$ is q_0 . Suppose that the prefix of length $j \geq 1$ of $\text{play}(q_0, \sigma_1, \sigma_2)$ is defined to be $\pi^j \cdot q_j$, where $\pi^j \in (Q \times A \times A)^*$. Then, at turn j , for $i \in \{1, 2\}$, Player i takes action $a_i^j = \sigma_i(\pi^j \cdot q_j)$, the next state is $q_{j+1} = \delta(q_j, a_1^j, a_2^j)$, and we define $\pi^{j+1} = \pi^j \cdot \langle v_j, a_1^j, a_2^j \rangle \cdot q_{j+1}$. The *path* that corresponds to $\text{play}(q_0, \sigma_1, \sigma_2)$ is $q_0, q_1, \dots \in Q^\omega$.

For $i \in \{1, 2\}$, we say that Player i *controls* a state $q \in Q$ if, intuitively, the next state is determined solely according to their chosen action. Formally, q is controlled by Player 1 if for every action $a_1 \in A$, there is a state q' such that no matter which action $a_2 \in A$ Player 2 takes, we have $q' = \delta(q, a_1, a_2)$, and the definition is dual for Player 2. *Turn-based games* are a special

case of concurrent games in which all states are controlled by one of the players. Note that a concurrent game that is not turn-based might still contain some vertices that are controlled by one of the players.

2.2 Bidding games

A discrete-bidding game is played on an arena $G = \langle V, E, k \rangle$, where V is a set of vertices, $E \subseteq V \times V$ is a set of directed edges, and $k \in \mathbb{N}$ is the sum of the players' budgets. For a vertex $v \in V$, we slightly abuse notation and use $N(v)$ to denote the neighbors of v in G , namely $N(v) = \{u : E(v, u)\}$. We will consider decision problems in which G is given as input. We then assume that k is encoded in binary, thus the size of G is $O(|V| + |E| + \log(k))$.

Intuitively, in each turn, both players simultaneously choose a bid that does not exceed their available budgets. The higher bidder moves the token and pays the other player. Note that the sum of budgets is constant throughout the game. Tie-breaking needs to be handled explicitly in discrete-bidding games as it can affect the properties of the game [1]. In this paper, we focus on *advantage-based* tie-breaking mechanism [20]: exactly one of the players holds the *advantage* at a turn, and when a tie occurs, the player with the advantage chooses between (1) win the bidding and pass the advantage to the other player, or (2) let the other player win the bidding and keep the advantage. We describe the semantics of bidding games formally below.

We will describe the formal semantics of a bidding game by constructing the explicit concurrent game that it corresponds to. We introduce the required notation. Following [20], we denote the advantage with $*$. Let \mathbb{N} denote the non-negative integers, \mathbb{N}^* the set $\{0, 0^*, 1, 1^*, 2, 2^*, \dots\}$, and $[k]$ the set $\{0, 0^*, \dots, k, k^*\}$. We define an order $<$ on \mathbb{N}^* by $0 < 0^* < 1 < 1^* < \dots$. Let $m \in \mathbb{N}^*$. When saying that Player 1 has a budget of $m^* \in [k]$, we mean that Player 1 has the advantage, and implicitly, we mean that Player 2's budget is $k - m$, and she does not have the advantage. We use $|m|$ to denote the integer part of m , i.e., if $m = x^*$ for some $x \in \mathbb{N}$, we denote $|m| = x$. Specifically, for $m \in \mathbb{N}$, we have $|m| = m$. We define operators \oplus and \ominus over \mathbb{N}^* . Intuitively, we use \oplus as follows: suppose that Player 1's budget is m^* and Player 2 wins a bidding with a bid of b_2 , then Player 1's budget is updated to $m^* \oplus b_2$. Similarly, for $\ell \leq m$, a bid of $b_1 = \ell^*$ means that Player 1 will use the advantage if a tie occurs and $b_1 = \ell$ means that he will not use it. Upon winning the bidding, his budget is updated to $m^* \ominus b_1$.

DEFINITION 2.1 (\oplus and \ominus operators). For $x, y \in \mathbb{N}$, define $x^* \oplus y = x \oplus y^* = (x + y)^*$, $x \oplus y = x + y$. For $x, y \in \mathbb{N}$, define $x \ominus y = x - y$, $x^* \ominus y = (x - y)^*$, and in particular $x^* \ominus y^* = x - y$. For notational consistency, for $x, y \in \mathbb{N}$, we define $x^* \oplus y^* = (x + y + 1)^*$, and $x \ominus y^* = (x - y - 1)^*$. Recall that \ominus is intuitively used to deduct the winning bid from the winner's budget and \oplus is intuitively used to add the winning bid to the losing player's budget, hence the latter two cases do not follow this intuitive meaning.

Next, we highlight two special cases, which are used frequently throughout the paper.

DEFINITION 2.2 (Successor and predecessor). For $B \in \mathbb{N}^*$, we denote by $B \oplus 0^*$ and $B \ominus 0^*$ respectively the *successor* and *predecessor* of B in \mathbb{N}^* according to $<$, defined as $B \oplus 0^* = \min\{x > B\}$ and $B \ominus 0^* = \max\{x < B\}$. We note that this notation is convenient since it applies both for budgets that include the advantage and those that do not. When the status of the advantage is known we use the following notation. When $B \in \mathbb{N}$ does not include the advantage, we use B^* as shorthand for $B \oplus 0^*$. When $B = x^*$ for some $x \in \mathbb{N}$, we use $|B| + 1 = x + 1$ as shorthand for $B \oplus 0^*$.

2.3 Bidding games as concurrent games

Consider an arena $\langle V, E, k \rangle$ of a bidding game. The corresponding *configurations* are $C = \{\langle v, B \rangle \in V \times ([k] \cup \{k + 1\})\}$, where a configuration $c = \langle v, B \rangle \in C$ means that the token is placed on vertex $v \in V$ and Player 1's budget is B . Implicitly, Player 2's budget is $k^* \ominus B$. Note that, vertices of the form $\langle v, k + 1 \rangle \in C$ are symbolic, namely they do not represent configurations of the game since Player 1's budget cannot exceed k . The arena of the explicit concurrent game is $\langle A, C, \lambda, \delta \rangle$, where $A = [k] \times V$, and we define the allowed actions in each configuration and transitions next. An action $\langle b, v \rangle \in A$ means that the player bids b and proceeds to v upon winning the bidding. We require the player with the advantage to decide prior to the bidding whether they will use the advantage or not. Thus, when Player 1's budget is B^* , Player 1's legal bids are $[B]$ and Player 2's legal bids are $\{0, \dots, k - B\}$, and when Player 1's budget is B , Player 1's legal bids are $\{0, 1, \dots, B\}$ and Player 2's legal bids are $[k \ominus B]$. Next, we describe the transitions. Suppose that the token is placed on a configuration $c = \langle v, B \rangle$ and Player i chooses action $\langle b_i, u_i \rangle$, for $i \in \{1, 2\}$. If $b_1 > b_2$, Player 1 wins the bidding and the game proceeds to $\langle u_1, B_1 \ominus b_1 \rangle$. The definition for $b_2 > b_1$ is dual. The remaining case is a tie, i.e., $b_1 = b_2$. Since only one of the players has the advantage, a tie can occur only when the player who has the advantage does not use it. Suppose that $c = \langle v, B^* \rangle$, i.e., Player 1 has the advantage, and the definition when Player 2 has the advantage is dual. Player 2 wins the bidding, Player 1 keeps the advantage, and we proceed to $\langle u_2, B^* \oplus b_2 \rangle$. Note that the size of the arena is $O(|V| \times k)$, which is exponential in the size of \mathcal{G} since k is given in binary.

Consider two strategies f and g and an initial configuration $c = \langle v, B \rangle$. We sometimes abuse notation and treat $\text{play}(v, f, g) = \langle v_0, B_0 \rangle, \langle v_1, B_1 \rangle, \dots$ as the infinite path v_0, v_1, \dots in the bidding game.

2.4 Objectives and threshold budgets

A bidding game is $\mathcal{G} = \langle V, E, k, O \rangle$, where $\langle V, E, k \rangle$ is an arena and $O \subseteq V^\omega$ is an *objective*, which specifies the infinite paths that are winning for Player 1.

We introduce notations on paths before defining the objectives that we consider. Consider a path $\pi = v_0, v_1, \dots$ and consider a subset of vertices $A \subseteq V$. We say that π *visits* A if there is

$j \geq 0$ such that $v_j \in A$. We denote by $\text{inf}(\pi) \subseteq V$, the set of vertices that π visits infinitely often. We say that π *enters* A at time $j \geq 1$ if $v_j \in A$ and $v_{j-1} \notin A$, and it is *exited* at time j if $v_j \notin A$ and $v_{j-1} \in A$.

We consider the following two canonical objectives:

- **Reachability:** A reachability bidding game is $\langle V, E, k, S \rangle$, where $S \subseteq V$ is a set of sinks. Player 1, the reachability player, wins an infinite play π iff it visits S , and we then say that π *ends* in S . *Safety* objectives are dual to reachability objectives; the safety player wins a play iff it never visits S .
- **Parity:** A parity bidding game is $\langle V, E, k, p \rangle$, where $p : V \rightarrow \{1, \dots, d\}$ assigns to each vertex a *parity index*, for $d \in \mathbb{N}$. A play τ is winning for Player 1 iff $\max_{v \in \text{inf}(\tau)} p(v)$ is odd. The special case in which p assigns parities in $\{2, 3\}$ is called *Büchi* objective; Player 1 wins a play iff it visits the set $\{v \in V : p(v) = 3\}$ infinitely often.

We introduce *frugal* objectives in bidding games in which, roughly, Player 1 wins by reaching a target with a sufficient budget.

DEFINITION 2.3 (Frugal objectives).

- A *frugal-reachability* bidding game is $\langle V, E, k, S, \text{fr} \rangle$, where V, E , and k are as in bidding games, $S \subseteq V$ is a set of target vertices, and $\text{fr} : S \rightarrow [k]$ assigns a *frugal-target budget* to each target. Consider a play π . Player 1 wins π iff π reaches S , i.e., π ends in a configuration $\langle s, B \rangle$ with $s \in S$, and Player 1's budget at S exceeds the frugal-target budget, i.e., π ends in $\langle s, B \rangle$ with $B \geq \text{fr}(s)$. Note that a reachability bidding game is a special case of a frugal-reachability bidding game in which $\text{fr} \equiv 0$.
- The *frugal-safety* objective is dual to frugal-reachability. We describe the winning condition explicitly. A frugal-safety bidding game is $\langle V, E, k, S, \text{fr} \rangle$, where V, E , and k are as in bidding games, $S \subseteq V$ is a set of sinks, and $\text{fr} : S \rightarrow [k]$ assigns a frugal-target budget to each sink. Player 1, the safety player, wins a play π if: (1) π never reaches S , or (2) π reaches a configuration $\langle s, B \rangle$ with $s \in S$ and $B \geq \text{fr}(s)$. Note that a safety bidding game is a special case of a frugal-safety bidding game in which $\text{fr} \equiv k + 1$.
- A *frugal-parity* bidding game is $\langle V, E, k, p, S, \text{fr} \rangle$, where $p : (V \setminus S) \rightarrow \{0, \dots, d\}$ and the other components are as in the above. Player 1 wins a play π if (1) π does not reach S and satisfies the parity objective, or (2) π satisfies a frugal-reachability objective: it ends in a configuration $\langle s, B \rangle$ with $s \in S$ and $B \geq \text{fr}(s)$.

REMARK 2.4. We point out that an ω -regular objective (e.g., reachability and parity) is a collection of paths in the graph, i.e., a subset of V^ω . On the other hand, a frugal objective is a subset of configurations, i.e., a subset of $(V \times [k])^\omega$. We do not know of a reduction from a frugal objective game to a non-frugal objective game that preserves the thresholds. At the

same time, as we show in the next section, it is not hard to extend algorithms for reachability bidding games to frugal-reachability games. Moreover, we are not aware of applications of frugal objectives, and define them as a means to solve parity bidding games.

Next, we define winning strategies.

DEFINITION 2.5 (Winning strategies). Consider a configuration $c = \langle v, B \rangle$ and an objective O .

- A Player 1 strategy f is winning from c if for every strategy g , $\text{play}(c, f, g)$ satisfies O .
- A Player 2 strategy g is winning from c if for every strategy f , $\text{play}(c, f, g)$ does not satisfy O .

A player wins from c if they have a winning strategy from c .

The central quantity in bidding games is the *threshold budget* at a vertex, which is the necessary and sufficient initial budget at that vertex for Player 1 to guarantee winning the game. It is formally defined as follows.

DEFINITION 2.6 (Threshold budgets). Consider a bidding game \mathcal{G} . The *threshold budget* at a vertex v in \mathcal{G} , denoted $\text{Th}_{\mathcal{G}}(v)$, is such that

- Player 1 wins from every configuration $\langle v, B \rangle$ with $B \geq \text{Th}_{\mathcal{G}}(v)$, and
- Player 2 wins from every configuration $\langle v, B \rangle$ with $B < \text{Th}_{\mathcal{G}}(v)$.

We refer to the function $\text{Th}_{\mathcal{G}}$ as the *threshold budgets*.

REMARK 2.7 (Thresholds in losing vertices). Note that a game might contain vertices that are losing for Player 1 with every initial budget. For example, in a reachability game, a sink with no path to the target is always losing for Player 1. Following [20], in such a losing vertex v , we set the threshold to $\text{Th}(v) = k + 1$. Since the highest possible budget for a player in the game is k^* , setting $\text{Th}(v) = k + 1$ can intuitively be understood as Player 1 requires more budget for winning than he can possibly obtain.

REMARK 2.8. We point out that existence of threshold budgets is not trivial. Indeed, existence of threshold budgets implies *determinacy*: from each configuration, there is a player who has a winning strategy. By observing that an additional budget will not harm a player in a bidding game, we obtain that determinacy implies existence of thresholds. Recall that bidding games are succinctly-represented concurrent games, and concurrent games are often not determined, for example, the simple concurrent game “matching pennies” is not determined since neither player can win the game. Typically, the vertices of a concurrent game can be partitioned into *surely winning* vertices from which Player 1 has a winning strategy, *surely losing* from which Player 2 has a winning strategy, and in the rest, neither player has a winning strategy. An optimal strategy from a vertex in the last set is *mixed*; it assigns a probability distribution over actions. Interestingly, in bidding games, all vertices are either surely winning or surely losing.

Previously, existence of thresholds in reachability games with advantage-based tie-breaking was shown in [20] by identifying the structure of the thresholds. An alternative technique was developed in [1] in which bidding games were shown to have a local property (called “local determinacy”) that implies “global” determinacy, and was used to show that Muller bidding games with various tie-breaking mechanisms are determined. Other subclasses of concurrent games (beyond bidding games) were shown to have a local property that implies determinacy in [17].

3. Frugal-Reachability Discrete-Bidding Games

The study in [20] focuses primarily on reachability discrete-bidding games played on DAGs. We revisit their results, provide explicit and elaborate proofs for games played on general graphs, and extend the results to frugal-reachability games. Specifically, Theorem 3.7 points to an issue in bidding games played on general graphs that was not explicitly addressed in [20].

3.1 Background: reachability continuous-bidding games

Many of the techniques used in reachability discrete-bidding games are adaptations of techniques developed for reachability continuous-bidding games [25, 24]. In order to develop intuition and ease presentation of discrete-bidding games, in this section, we illustrate the ideas and techniques of continuous-bidding games.

Recall that in continuous-bidding games there is no restriction on the granularity of bids, i.e., bids can be arbitrarily small. Throughout this section we assume that the sum of the players’ budgets is 1. Note that since winning bids are paid to the opponent, the sum of budgets stays constant throughout the game.

DEFINITION 3.1 (Continuous threshold budgets). The *continuous threshold budget* at a vertex v is a budget $\text{Th}(v) \in [0, 1]$ such that for every $\epsilon > 0$:

- if Player 1’s budget is $\text{Th}(v) + \epsilon$, he wins the game from v , and
- if Player 1’s budget is $\text{Th}(v) - \epsilon$, Player 2 wins the game from v .

REMARK 3.2 (Losing vertices in continuous-bidding games). Recall that in discrete bidding, the threshold budget in a losing vertex is $k + 1$, which is higher than the highest budget a player can obtain in a discrete bidding game (Remark 2.7). The treatment of losing vertices in continuous bidding is similar, though implicit in Def. 3.1: the continuous threshold budget in a losing vertex is 1, thus by Def. 3.1, Player 1 wins if his budget is $1 + \epsilon$, which is never the case since it is higher than the total budget.

REMARK 3.3. We point out that the issue of tie breaking is avoided in continuous-bidding games by considering initial budgets that differ from the threshold. That is, the guarantee is

that if a player's budget is strictly above the threshold, he wins the game no matter which tie-breaking mechanism is used, e.g., even if the opponent wins all bidding ties.

A *double-reachability* continuous-bidding game is $\langle V, E, t_1, t_2 \rangle$, where for $i \in \{1, 2\}$, the vertex t_i is the target of Player i and every vertex $v \neq t_1, t_2$ has a path to both. The game ends once one of the targets is reached, and the player whose target is reached is the winner. The careful reader might notice that the definition does not define a winner when no target is reached. We will show below that this case is avoided.

DEFINITION 3.4 (Continuous average property). Consider a double-reachability continuous-bidding game $\mathcal{G} = \langle V, E, t_1, t_2 \rangle$ and a function $T : V \rightarrow [0, 1]$. For $v \in V$, denote $v_T^+ := \arg \max_{u \in N(v)} T(u)$ and $v_T^- := \arg \min_{u \in N(v)} T(u)$. We say that T has the *continuous average property* if for every vertex $v \in V$:

$$T(v) = \begin{cases} 1 & \text{if } v = t_2 \\ 0 & \text{if } v = t_1 \\ \frac{T(v_T^-) + T(v_T^+)}{2} & \text{otherwise} \end{cases}$$

when the function T is clear from the context, we simply refer v_T^+ and v_T^- as v^+ and v^- respectively. Note that, there could be more than one vertex u (similarly, w) such that $T(u) = T(v^-)$ (respectively, $T(w) = T(v^+)$), but for the sake of convenience, we collectively denote any of them as v^- and v^+ respectively.

The next theorem presents the main results on reachability continuous-bidding games: a function that satisfies the continuous average property is unique, and it coincides with the continuous threshold budgets. We illustrate the proof techniques, in particular how to construct a winning bidding strategy given the thresholds in the game.

THEOREM 3.5. [25, 24] Consider a double-reachability continuous-bidding game $\langle V, E, t_1, t_2 \rangle$. Continuous threshold budgets exist, and the threshold budgets $Th : V \rightarrow [0, 1]$ is the unique function that has the continuous average property. Moreover, the problem of deciding whether $Th(v) \leq 0.5$ is in NP and coNP.

PROOF. (SKETCH) Let T be a function that satisfies the continuous average property, where we omit the proof of existence of such a function. We prove that for every vertex v , the continuous threshold budget at v is $T(v)$. Uniqueness follows immediately.

The complexity bound is obtained by guessing, for every vertex v , two neighbors v^- and v^+ , and constructing and solving a linear program based on the constraints in Def. 3.4 (for more details, see [8], which shows a reduction to stochastic games [19]).

From a function with the continuous-average property to a strategy. Suppose that Player 1's budget at v is $T(v) + \varepsilon$, for $\varepsilon > 0$. We describe a winning Player 1 strategy. Recall that

$v^+, v^- \in N(v)$ are respectively the neighbors of v that attain the maximal and minimal values according to T . Let

$$b(v) := \frac{T(v^+) - T(v^-)}{2}.$$

The key observation is that $T(v) + b(v) = T(v^+)$ and $T(v) - b(v) = T(v^-)$.

Consider the following Player 1 strategy. At vertex $v \in V$, bid $b(v)$ and proceed to v^- upon winning. We show that the strategy maintains the following invariant:

Invariant: When the game reaches a configuration $\langle u, B \rangle$, then $B > T(u)$.

We list two implications of the invariant. First, it implies that the strategy is legal, namely Player 1's budget at v suffices to bid $b(v)$. Second, it implies that Player 1 does not lose, namely no matter how Player 2 plays, the game will not reach t_2 . Indeed, assume towards contradiction that t_2 is reached. Then, the invariant implies that Player 1's budget is strictly greater than 1, which violates the assumption that the sum of budgets is 1.

Note that “not losing” does not suffice for winning, namely that Player 1 forces the game to t_1 . These details, however, are not relevant to this paper. For completeness, we describe the rough idea. Suppose that the game reaches configuration $\langle u, B \rangle$. The invariant implies $B > T(u)$. We call $B - T(u)$ Player 1's *spare change*. The idea is to choose Player 1's bids carefully in a way that ensures that as the game continues, his spare change strictly increases so that eventually his budget suffices to win $|V|$ times in a row. We point out that this idea can be extended to show that in a strongly-connected game, a player can force infinitely many visits to a vertex with any positive initial budget, which is at the core of solving parity continuous-bidding games.

We prove that Player 1's strategy maintains the invariant against any Player 2 strategy. Note that the invariant holds initially. Suppose that the game reaches configuration $\langle u, B \rangle$ with $B > T(u)$. We claim that the invariant is maintained in the next turn. Indeed, if Player 1 wins the bidding, the next configuration is $\langle v^-, B - b(v) \rangle$, and the claim follows from $T(v) - b(v) = T(v^-)$. If Player 2 wins the bidding, she bids at least $b(v)$, thus Player 1's updated budget is at least $B + b(v)$, and the worst that Player 2 can do for Player 1 is to move to v^+ . The claim follows from $T(v) + b(v) = T(v^+)$.

Reasoning about the flipped game. Finally, we show that Player 2 wins when Player 1's budget is $T(v) - \varepsilon$. We intuitively “flip” the game and associate Player 1 with Player 2. More formally, let \mathcal{G}' be the same as \mathcal{G} except that Player 1's goal is to reach t_2 and Player 2's goal is to reach t_1 . For every $u \in V$, define T' by $T'(u) = 1 - T(u)$. A key observation is that T' satisfies the continuous average property in \mathcal{G}' . In particular, note that $T'(t_1) = 1$ and $T'(t_2) = 0$. Now, in order to win from v in \mathcal{G} when Player 1's budget is $T(v) - \varepsilon$, Player 2 follows a winning Player 1 strategy in \mathcal{G}' with an initial budget of $1 - T(v) + \varepsilon$. ■




Figure 1. A discrete-bidding reachability game with two functions that satisfy the average property.

3.2 Frugal-Reachability discrete-bidding games

We turn to study discrete-bidding games.

3.2.1 The discrete average property

In this section, we adapt the definition of the continuous average property (Def. 3.4) to the discrete setting and analyze its properties.

DEFINITION 3.6 (Average property). Consider a frugal-reachability discrete-bidding game $\mathcal{G} = \langle V, E, k, S, \text{fr} \rangle$. We say that a function $T : V \rightarrow [k] \cup \{k+1\}$ has the *average property* if for every $s \in S$, we have $T(s) = \text{fr}(s)$, and for every $v \in V \setminus S$,

$$T(v) = \left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor + \varepsilon \text{ where } \varepsilon = \begin{cases} 0 & \text{if } |T(v^+)| + |T(v^-)| \text{ is even and } T(v^-) \in \mathbb{N} \\ 1 & \text{if } |T(v^+)| + |T(v^-)| \text{ is odd and } T(v^-) \in \mathbb{N}^* \setminus \mathbb{N} \\ * & \text{otherwise} \end{cases}$$

where $v^+ := \arg \max_{u \in N(v)} T(u)$ and $v^- := \arg \min_{u \in N(v)} T(u)$

The following theorem shows, somewhat surprisingly and for the first time, that functions that satisfy the discrete average property are not unique. That is, there are functions satisfying the discrete average property but not coinciding with the threshold budgets. This is in stark contrast to continuous-bidding games in which there is a unique function that satisfies the average property.

THEOREM 3.7. *The reachability discrete-bidding game \mathcal{G}_1 that is depicted in Fig. 1 with target t for Player 1 has more than one function that satisfies the average property.*

PROOF. Assume a total budget of $k = 5$. We represent a function $T : V \rightarrow [k]$ as a vector $\langle T(v_0), T(v_1), T(v_2), T(t) \rangle$. It is not hard to verify that both $\langle 4, 3^*, 3, 2 \rangle$ and $\langle 5, 4^*, 3^*, 2 \rangle$ satisfy the average property. (The latter represents the threshold budgets). ■

The following lemma intuitively shows that the “complement” of T , denoted below T' , satisfies the average property. For a vertex v , the value $T'(v)$ should be thought of as Player 2’s

budget when Player 1's budget is strictly less than $T(v)$. We will later show that since T' satisfies the average property, Player 2 can win with a budget of $T'(v)$. A similar idea is used in continuous-bidding in the last point in the proof of Theorem 3.5. It follows that a function T that satisfies the average property satisfies $T \leq \text{Th}_{\mathcal{G}}$. Indeed, on the one hand, if Player 1's budget is less than $T(v)$, then Player 2, the safety player, wins by avoiding Player 1's targets. However, it could be the case that $T < \text{Th}_{\mathcal{G}}$; namely, Player 1 cannot win (force a visit to a target) with a budget of $T(v)$.

DEFINITION 3.8 (Complement of T). Let $\mathcal{G} = \langle V, E, k, S, \text{fr} \rangle$ be a discrete-bidding game with a frugal objective. Let $T : V \rightarrow [k] \cup \{k+1\}$ be a function. We define “complement” of T , denoted by $T' : V \rightarrow [k] \cup \{k+1\}$ as: $T'(v) = (k+1) \ominus T(v)$ for all $v \in V$.

OBSERVATION 3.9 (Relation between T value and T' value of a vertex). *We make the following observations about such a function T and its complement T' , which can be directly derived using the definition of \ominus , in the following.*

- I. For every $v \in V$, we have that $T(v)$ and $T'(v)$ agree on which player has the advantage, formally $T(v) \in \mathbb{N}$ iff $T'(v) \in \mathbb{N}$.
- II. For $v \in V$, a neighbouring vertex with maximum T -value is the neighbouring vertex with minimum T' -value, and vice-versa. Notationally, $v_T^+ = v_{T'}^-$ and $v_T^- = v_{T'}^+$.
- III For any $v \in V$, if $T(v) \in \mathbb{N}$ then $T'(v) = |T'(v)| = (k+1) - |T(v)|$, otherwise $|T'(v)| = k - |T(v)|$.

LEMMA 3.10. *Let $\mathcal{G} = \langle V, E, k, S, \text{fr} \rangle$ be a discrete-bidding game with a frugal objective. Let $T : V \rightarrow [k] \cup \{k+1\}$ be a function that satisfies the average property. Then, the complement of T , denoted by $T' : V \rightarrow [k] \cup \{k+1\}$, also satisfies the average property.*

PROOF. Let us first fix the notation for v^- and v^+ as to denote the neighbouring vertices of v with minimum and maximum T -value, respectively for this proof. We need to show that $T'(v) = \left\lfloor \frac{|T'(v^-)| + |T'(v^+)|}{2} \right\rfloor + \varepsilon'$, where ε' comes from the definition of the average property. We already have $T(v) = \left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor + \varepsilon$. Note that $T'(v^-) = T'(v_{T'}^+)$ and $T'(v^+) = T'(v_{T'}^-)$, as per Observation 3.9.

We proceed to prove the lemma.

First, we consider the scenario when $T(v) = 0$. This necessarily means $\varepsilon = 0$, $T(v^-) = 0$, $|T(v^+)| = 0$, and $T'(v) = (k+1) \ominus T(v) = k+1$. This implies $T'(v_{T'}^+) = T'(v^-) = k+1$ and $T'(v_{T'}^-) = T'(v^+) \in \{k^*, k+1\}$. If $T'(v_{T'}^-)$ is k^* , we have $|T'(v^+)| + |T'(v^-)| = 2k+1$ and $\varepsilon' = 1$. On the other hand, if $T'(v_{T'}^-) = k+1$, then we have $|T'(v^+)| + |T'(v^-)| = 2(k+1)$, and $\varepsilon' = 0$. Thus, we have $T'(v) = k+1 = \left\lfloor \frac{|T'(v^-)| + |T'(v^+)|}{2} \right\rfloor + \varepsilon'$

Now, we assume that $T(v) > 0$. We divide the analysis in four exhaustive cases in the following. Before delving into the case analysis, for better reading comprehension, we explain

the structure that each of these analyses follows. We divide the analysis into four exhaustive cases according to the advantage statuses of $T(v^+)$ and $T(v^-)$, the first being where both of them are in \mathbb{N} , in the second one, both of them are in $\mathbb{N}^* \setminus \mathbb{N}$, and the final two cases are when one of them is in \mathbb{N} and the other one is in $\mathbb{N}^* \setminus \mathbb{N}$.

Then, in each case, we look into the relation between the parity of $|T(v^+)| + |T(v^-)|$ and $|T'(v^+)| + |T'(v^-)|$, and what the $\varepsilon, \varepsilon'$ values are. Once those are settled, we start with the corresponding expression of $\left\lfloor \frac{|T'(v^+)| + |T'(v^-)|}{2} \right\rfloor + \varepsilon'$, and show that it is indeed $T'(v)$.

We now delve into the technical part of this analysis below:

— $T(v^+), T(v^-) \in \mathbb{N}$.

First, from Observation 3.9, we have $T'(v^-), T'(v^+) \in \mathbb{N}$ as well.

Thus, it is enough to discuss the two cases corresponding to when the sum is even or odd, respectively. Note that $|T(v^+)| + |T(v^-)|$ is even iff $|T'(v^-)| + |T'(v^+)|$ is even, simply because in this case $T'(v^+) = (k+1) - T(v^+)$ and $T'(v^-) = (k+1) - T(v^-)$. It follows that irrespective of whether $|T(v^-)| + |T(v^+)|$ is odd or even, we have $\varepsilon' = \varepsilon$.

When $\varepsilon = \varepsilon' = 0$ (i.e., the sum is even), we have

$$\begin{aligned} \frac{|T'(v^+)| + |T'(v^-)|}{2} &= \frac{(k+1 - |T(v^-)|) + (k+1 - |T(v^+)|)}{2} \\ &= k+1 - \frac{|T(v^+)| + |T(v^-)|}{2} \\ &= (k+1) - T(v) = T'(v) \end{aligned}$$

When $\varepsilon = \varepsilon' = 0^*$ (i.e., the sum is odd), we have

$$\begin{aligned} \left(\left\lfloor \frac{|T'(v^+)| + |T'(v^-)|}{2} \right\rfloor \right)^* &= \left(\left\lfloor k+1 - \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor \right)^* \\ &= \left(k+1 - \left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor - 1 \right)^* \\ &= (k+1) \ominus \left(\left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor \right)^* \\ &= (k+1) \ominus T(v) = T'(v) \end{aligned}$$

— $T(v^+), T(v^-) \in \mathbb{N}^* \setminus \mathbb{N}$.

In this case too, $|T(v^+)| + |T(v^-)|$ is even iff $|T'(v^-)| + |T'(v^+)|$ is even. Therefore, we have $\varepsilon' = \varepsilon$.

When $\varepsilon = \varepsilon' = *$ (i.e., the sum is even and the minimum is in $\mathbb{N}^* \setminus \mathbb{N}$), we have

$$\begin{aligned}
 \left(\frac{|T'(v^+)| + |T'(v^-)|}{2} \right)^* &= \left(\frac{(k - |T(v^-)|) + (k - |T(v^+)|)}{2} \right)^* \\
 &= \left(k - \frac{|T(v^-)| + |T(v^+)|}{2} \right)^* \\
 &= \left((k+1) - \frac{|T(v^-)| + |T(v^+)|}{2} - 1 \right)^* \\
 &= (k+1) \ominus \left(\frac{|T(v^-)| + |T(v^+)|}{2} \right)^* \\
 &= (k+1) \ominus T(v) = T'(v)
 \end{aligned}$$

When $\varepsilon = \varepsilon' = 1$ (i.e., the sum is odd and the minimum is in $\mathbb{N}^* \setminus \mathbb{N}$), we have

$$\begin{aligned}
 \left\lfloor \frac{|T'(v^+)| + |T'(v^-)|}{2} \right\rfloor + 1 &= \left\lfloor \frac{(k - |T(v^+)|) + (k - |T(v^-)|)}{2} \right\rfloor + 1 \\
 &= \left\lfloor k - \frac{|T(v^-)| + |T(v^+)|}{2} \right\rfloor + 1 \\
 &= k - \left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor - 1 + 1 \\
 &= (k+1) - \left(\left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor + 1 \right) \\
 &= (k+1) \ominus T(v) = T'(v)
 \end{aligned}$$

— $T(v^+) \in \mathbb{N}$ and $T(v^-) \in \mathbb{N}^* \setminus \mathbb{N}$

Note that, in this case, $|T(v^+)| + |T(v^-)|$ is even iff $|T'(v^-)| + |T'(v^+)|$ is odd.

We can also see that when $|T(v^+)| + |T(v^-)|$ is even, we have $\varepsilon = \varepsilon' = 0^*$. Therefore, we have

$$\begin{aligned}
 \left(\left\lfloor \frac{|T'(v^+)| + |T'(v^-)|}{2} \right\rfloor \right)^* &= \left(\left\lfloor \frac{(k+1) - |T(v^+)| + k - |T(v^-)|}{2} \right\rfloor \right)^* \\
 &= \left(\left\lfloor \frac{2k+1 - (|T(v^+)| + |T(v^-)|)}{2} \right\rfloor \right)^* \\
 &= \left(k - \frac{|T(v^+)| + |T(v^-)|}{2} \right)^* \\
 &= \left(k+1 - \frac{|T(v^+)| + |T(v^-)|}{2} - 1 \right)^* \\
 &= (k+1) \ominus \left(\frac{|T(v^+)| + |T(v^-)|}{2} \right)^* \\
 &= (k+1) \ominus T(v) = T'(v)
 \end{aligned}$$

On the other hand, when $|T(v^+)| + |T(v^-)|$ is odd, we have $\varepsilon = 1, \varepsilon' = 0$. In this case, we proceed as follows:

$$\begin{aligned}
 \frac{|T'(v^+)| + |T'(v^-)|}{2} &= \frac{k - |T(v^+)| + (k+1) - |T(v^-)|}{2} \\
 &= \frac{2k+1 - (|T(v^+)| + |T(v^-)|)}{2} \\
 &= k - \frac{|T(v^+)| + |T(v^-)| - 1}{2} \\
 &= k - \left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor \\
 &= (k+1) - \left(\left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor + 1 \right) \\
 &= (k+1) \ominus T(v) = T'(v)
 \end{aligned}$$

— $T(v^+) \in \mathbb{N}^*$ and $T(v^-) \in \mathbb{N}$

In this case, like the earlier one, $|T(v^+)| + |T(v^-)|$ is even iff $|T'(v^-)| + |T'(v^+)|$ is odd. But unlike the earlier case, here when $|T(v^+)| + |T(v^-)|$ is even, we have $\varepsilon = 0$ and $\varepsilon' = 1$. On the other hand, when $|T(v^+)| + |T(v^-)|$ is odd, then we have $\varepsilon = \varepsilon' = *$.

We first consider the case when $|T(v^+)| + |T(v^-)|$ is even, in the following:

$$\begin{aligned}
 \left\lfloor \frac{|T'(v^+)| + |T'(v^-)|}{2} \right\rfloor + 1 &= \left\lfloor \frac{k - |T(v^+)| + (k+1) - |T(v^-)|}{2} \right\rfloor \\
 &= \left\lfloor \frac{2k+1 - (|T(v^+)| + |T(v^-)|)}{2} \right\rfloor \\
 &= \left\lfloor k - \frac{|T(v^+)| + |T(v^-)| - 1}{2} \right\rfloor \\
 &= k - \left\lfloor \frac{|T(v^+)| + |T(v^-)| - 1}{2} \right\rfloor - 1 \\
 &= k - \left(\frac{|T(v^+)| + |T(v^-)|}{2} - 1 \right) - 1 \\
 &= (k+1) - \left(\frac{|T(v^+)| + |T(v^-)|}{2} + 1 \right) \\
 &= (k+1) \ominus T(v) = T'(v)
 \end{aligned}$$

On the other hand, when $|T(v^+)| + |T(v^-)|$ is odd, we have:

$$\begin{aligned}
 \left(\left\lfloor \frac{|T'(v^+)| + |T'(v^-)|}{2} \right\rfloor \right)^* &= \left(\left\lfloor k - \frac{|T(v^+)| + |T(v^-)| - 1}{2} \right\rfloor \right)^* \\
 &= \left(k - \frac{|T(v^+)| + |T(v^-)| - 1}{2} \right)^* \\
 &= \left(k - \left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor \right)^* \\
 &= \left((k+1) - \left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor - 1 \right)^* \\
 &= (k+1) \ominus \left(\left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor \right)^* \\
 &= (k+1) \ominus T(v) = T'(v)
 \end{aligned}$$

As we establish exhaustively that $\left\lfloor \frac{|T'(v^+)| + |T'(v^-)|}{2} \right\rfloor$ is indeed $T'(v)$, we conclude with the fact that T' satisfies the average property when T does. ■

3.2.2 From a function that satisfies the average to a bidding strategy

Throughout this section, fix a function $T : V \rightarrow [k] \cup \{k+1\}$ that satisfies the average property. We show how to construct a bidding strategy from T . We will use this construction also for parity games.

DEFINITION 3.11 (Partial strategy). A *partial strategy* based on T is a function $f_T : C \rightarrow [k] \times 2^V$ that chooses, at each configuration, a bid and a set of *allowed* vertices.

Note that f_T is not a strategy since it does not assign a unique vertex to proceed to upon winning the bidding.

DEFINITION 3.12 (f' agrees with f_T). Consider a strategy f' and a partial strategy $f_T : C \rightarrow [k] \times 2^V$. Consider a configuration c . Let $\langle b, A \rangle = f_T(c)$ and $\langle b', u' \rangle = f'(c)$. We say that f' agrees with f_T at c if $b = b'$ and $u' \in A$. We say that f' agrees with f_T if f' agrees with f_T in all configurations.

We describe the intuition behind the construction of f_T . We construct f_T so that every strategy f' that agrees with f_T maintains the following invariant. Suppose that the game starts from configuration $\langle v, B \rangle$ with $B \geq T(v)$. Then, against any opponent strategy, when the game reaches a configuration $\langle u, B' \rangle$, we have $B' \geq T(u)$. Since for every sink $s \in S$, we have $T(s) \geq \text{fr}(s)$, the invariant implies that f' guarantees that the frugal objective is not violated. Technically, the construction of f_T is similar in spirit to the construction in continuous bidding (Theorem 3.5). There, a non-losing strategy maintains the invariant that Player 1's budget

exceeds $\text{Th}(v)$ by bidding $b = \frac{\text{Th}(v^+) - \text{Th}(v^-)}{2}$ at a vertex v , and proceeding to v^- upon winning the bidding. Recall that the invariant is maintained since $\text{Th}(v) - b = \text{Th}(v^-)$ and $\text{Th}(v) + b = \text{Th}(v^+)$.

We describe f_T formally. Consider $v \in V$ and a budget $B \in [k]$ with $B \geq T(v)$. Let $v^+ = \arg \max_{u \in N(v)} T(u)$ and $v^- = \arg \min_{u \in N(v)} T(u)$. We define $f_T(\langle v, B \rangle) = \langle b^T(v, B), A(v) \rangle$ as follows. First, we define the allowed vertices

$$A(v) = \begin{cases} \{u \in N(v) : T(u) = T(v^-)\} & \text{if } T(v^-) \in \mathbb{N} \\ \{u \in N(v) : T(u) \leq T(v^-) \oplus 0^*\} & \text{if } T(v^-) \in \mathbb{N}^* \setminus \mathbb{N} \end{cases} \quad (1)$$

Second, the definition of the bid $b^T(v, B)$ is based on b_v^T , defined as follows:

$$b_v^T = \begin{cases} T(v) \ominus T(v^-) & \text{if } T(v^-) \in \mathbb{N} \\ T(v) \ominus (|T(v^-)| + 1) & \text{otherwise} \end{cases} \quad (2)$$

We define the bid chosen by f_T at a configuration $c = \langle v, B \rangle$. Intuitively, Player 1 “attempts” to bid b_v^T . This is not possible when b_v^T requires the advantage but Player 1 does not have it in c , i.e., $b_v^T \in \mathbb{N}^* \setminus \mathbb{N}$ and $B \in \mathbb{N}$. In such a case, Player 1 bids $|b_v^T| + 1 \in \mathbb{N}$. Formally, we define $b^T(v, B) = b_v^T$ when both b_v^T and B belong to either \mathbb{N} or $\mathbb{N}^* \setminus \mathbb{N}$, and $b_v^T \oplus 0^*$ otherwise.

3.2.3 Strategies that agree with f_T are not losing

Suppose that the game starts from a configuration $\langle v, B \rangle$ with $B \geq T(v)$ and Player 1 follows f' that agrees with f_T . In this section, we will show that f' maintains the invariant that whenever the play reaches a configuration $\langle u, B' \rangle$, we have $B' \geq T(u)$. This implies that f' is “not losing” since if the play reaches a sink $s \in S$, Player 1’s budget exceeds the frugal target budget. Note that “not losing” is not sufficient for winning; indeed, we require a winning strategy to draw the game to a sink. We will show in the next section that there is a winning strategy that agrees with f_T .

We start with the following technical observation.

OBSERVATION 3.13. *For every vertex v , we have $T(v)$ is in \mathbb{N} iff b_v^T is in \mathbb{N} .*

Next, assuming that Player 1 bids b_v^T from configuration $\langle v, T(v) \rangle$ (i.e., Player 1’s budget is $T(v)$), we establish a lower bound on his budget in the next round. The following observation takes care of the case that Player 1 wins the bidding at v .

OBSERVATION 3.14. *If $T(v^-) \in \mathbb{N}$, then $T(v) \ominus b_v^T = T(v^-)$, and if $T(v^-) \in \mathbb{N}^* \setminus \mathbb{N}$, then $T(v) \ominus b_v^T = |T(v^-)| + 1$*

The following lemma takes care of the case that Player 1 loses the bidding at v .

LEMMA 3.15. *Let T be a function that satisfies the average property and a vertex $v \in V$. Then $T(v) \oplus (b_v^T \oplus 0^*) = |T(v^+)|^*$*

PROOF. We establish this result by analysing four cases in the following, where each case corresponds to a parity of $|T(v^+)| + |T(v^-)|$ and an advantage status of $T(v^-)$:

— $|T(v^+)| + |T(v^-)|$ is even, and $T(v^-) \in \mathbb{N}$.

In this case, $b_v^T = T(v) \ominus T(v^-) = \frac{|T(v^+)| + |T(v^-)|}{2} \ominus T(v^-) = \frac{|T(v^+)| - |T(v^-)|}{2}$. Therefore, we have

$$T(v) \oplus (b_v^T \oplus 0^*) = \frac{|T(v^+)| + |T(v^-)|}{2} + \left(\frac{|T(v^+)| - |T(v^-)|}{2} \right)^* = |T(v^+)|^*$$

— $|T(v^+)| + |T(v^-)|$ is odd and $T(v^-) \in \mathbb{N}^* \setminus \mathbb{N}$.

In this case, $b_v^T = T(v) \ominus (|T(v^-)| + 1) = \left(\left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor + 1 \right) - (|T(v^-)| + 1) = \left\lfloor \frac{|T(v^+)| - |T(v^-)|}{2} \right\rfloor$.

Therefore, we have

$$T(v) \oplus (b_v^T \oplus 0^*) = \left(\left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor + 1 \right) + \left(\left\lfloor \frac{|T(v^+)| - |T(v^-)|}{2} \right\rfloor \right)^* = |T(v^+)|^*$$

— $|T(v^+)| + |T(v^-)|$ is even and $T(v^-) \in \mathbb{N}^* \setminus \mathbb{N}$.

In this case, $b_v^T = T(v) \ominus (|T(v^-)| + 1) = \left(\frac{|T(v^+)| + |T(v^-)|}{2} \right)^* \ominus (|T(v^-)| + 1) = \left(\frac{|T(v^+)| + |T(v^-)|}{2} - |T(v^-)| - 1 \right)^* = \left(\frac{|T(v^+)| - |T(v^-)|}{2} - 1 \right)^*$. Therefore, we have

$$T(v) \oplus (b_v^T \oplus 0^*) = \left(\frac{|T(v^+)| + |T(v^-)|}{2} \right)^* \oplus \frac{|T(v^+)| - |T(v^-)|}{2} = |T(v^+)|^*$$

— Finally, $|T(v^+)| + |T(v^-)|$ is odd and $T(v^-) \in \mathbb{N}$.

In this case, $b_v^T = T(v) \ominus T(v^-) = \left(\left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor \right)^* \ominus T(v^-) = \left(\left\lfloor \frac{|T(v^+)| - |T(v^-)|}{2} \right\rfloor \right)^*$. Therefore, we have

$$\begin{aligned} T(v) \oplus (b_v^T \oplus 0^*) &= \left(\left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor \right)^* \oplus \left(\left\lfloor \frac{|T(v^+)| - |T(v^-)|}{2} \right\rfloor + 1 \right) \\ &= \left(\frac{|T(v^+)| + |T(v^-)|}{2} - \frac{1}{2} + \frac{|T(v^+)| - |T(v^-)|}{2} - \frac{1}{2} + 1 \right)^* = |T(v^+)|^* \end{aligned}$$

■

Furthermore, we also make the following observation, which intuitively enlists at least how much Player 1's new budget would be after a bidding at v where he has a budget of $T(v) \oplus 0^*$ and he bids $b_v^T \oplus 0^*$.

OBSERVATION 3.16. *Let T be a function that satisfies the average property and a vertex $v \in V$, then*

- $(T(v) \oplus 0^*) \ominus (b_v^T \oplus 0^*) = T(v) \ominus b_v^T$, and
- $(T(v) \oplus 0^*) \oplus (b_v^T \oplus 1) = |T(v^+)|^* + 1$.

We proceed to prove that strategies that agree with f_T maintain an invariant on Player 1's budget.

LEMMA 3.17. *Suppose that Player 1 plays according to a strategy f' that agrees with f_T starting from configuration $\langle v, B \rangle$ satisfying $B \geq T(v)$. Then, against any Player 2 strategy, when the game reaches $u \in V$, Player 1's budget is at least $T(u)$.*

PROOF. The invariant holds initially by the assumption. Consider a history that ends in a configuration $\langle v, B \rangle$. Assume that $B \geq T(v)$. We claim that the invariant is maintained no matter the outcome of the bidding, namely $B \ominus b^T(v, B) \geq T(v^-)$ and $B \oplus (b^T(v, B) \oplus 0^*) \geq T(v^+)$.

We distinguish between two cases. First, when either both B and b_v^T are in \mathbb{N} or both B and b_v^T are in $\mathbb{N}^* \setminus \mathbb{N}$. In either case, Player 1 bids $b^T(v, B) = b_v^T$. It follows from Observation 3.14 and Lemma 3.15 that

$$B \ominus b^T(v, B) = B \ominus b_v^T \geq T(v) \ominus b_v^T \geq T(v^-), \text{ and}$$

$$B \oplus b^T(v, B) \oplus 0^* = B \oplus b_v^T \oplus 0^* \geq T(v) \oplus b_v^T \oplus 0^* = |T(v^+)|^*.$$

In the second case, Player 1 bids $b^T(v, B) = b_v^T \oplus 0^*$. Note that $B \geq T(v) \oplus 0^*$ because $T(v)$ and b_v^T have the same advantage status (Observation 3.13). It follows from Observation 3.16 that

$$B \ominus b^T(v, B) = B \ominus (b_v^T \oplus 0^*) \geq (T(v) \oplus 0^*) \ominus (b_v^T \oplus 0^*) = T(v) \ominus b_v^T \geq T(v^-), \text{ and}$$

$$B \oplus (b^T(v, B) \oplus 0^*) \geq (T(v) \oplus 0^*) \oplus (b_v^T \oplus 1) = |T(v^+)|^* + 1 > |T(v^+)|.$$

This concludes the proof. ■

The following proposition follows from Lemma 3.17.

PROPOSITION 3.18. *Suppose that Player 1 plays according to a strategy f' that agrees with f_T starting from configuration $\langle v, B \rangle$ satisfying $B \geq T(v)$, then:*

- f' is a legal strategy: the bid b prescribed by f_T does not exceed the available budget.
- f' does not underestimate the frugal-budget: if $s \in S$ is reached, Player 1's budget is at least $fr(s)$.

3.2.4 Existence of thresholds in frugal-reachability discrete-bidding games

We close this section by showing existence of threshold budgets in frugal-reachability discrete-bidding games. Recall that Theorem 3.7 shows that functions that satisfy the discrete average property are not unique. Let T be such a function. The following lemma shows that $T \leq \text{Th}_{\mathcal{G}}$. That is, if in a vertex v , Player 1 has a budget less than $T(v)$, then Player 2 has a winning strategy. This proves that the threshold budgets for Player 1 cannot be less than $T(v)$, when T is a function that satisfies the average property.

LEMMA 3.19. *Consider a frugal-reachability discrete-bidding game $\mathcal{G} = \langle V, E, k, S, fr \rangle$. If $T : V \rightarrow [k] \cup \{k+1\}$ is a function that satisfies the average property, then $T(v) \leq Th_{\mathcal{G}}(v)$ for every $v \in V$.*

PROOF. Given T that satisfies the average property, we construct T' as in Definition 3.8. Let $\langle v, B_1 \rangle$ be a configuration, where $v \in V$, Player 1's budget is B_1 , and implicitly, Player 2's budget is $B_2 = k^* \ominus B_1$. Note that $B_1 < T(v)$ iff $B_2 \geq T'(v)$. Moreover, for every $s \in S$, we have $T'(s) = (k+1) \ominus fr(s)$. We consider the “flipped” game; namely, we associate Player 2 with Player 1 (of Proposition 3.18), and construct a partial strategy $f_{T'}$ for Player 2. We construct a Player 2 strategy f' that agrees with $f_{T'}$: for each $v \in V$, we arbitrarily choose a neighbor u from the allowed vertices. By Lemma 3.17, no matter how Player 1 responds, whenever the game reaches $\langle u, B_1 \rangle$, we have $B_2 \geq T'(u)$. The invariant implies that f' is a winning strategy. Indeed, if the game does not reach a sink, Player 2 wins, and if it does, Player 1's frugal objective is not satisfied. ■

The following lemma shows the existence of a function that satisfies the average property and that coincides with threshold budgets.

LEMMA 3.20. *Consider a frugal-reachability discrete-bidding game $\mathcal{G} = \langle V, E, k, S, fr \rangle$. There is a function T that satisfies the average property with $T(v) \geq Th_{\mathcal{G}}(v)$, for every $v \in V$.*

PROOF. The proof is similar to the one in [20]. We illustrate the main ideas. For $n \in \mathbb{N}$, we consider the *truncated game* $\mathcal{G}[n]$, which is the same as \mathcal{G} except that Player 1 wins iff he wins in at most n steps. We find a sufficient budget for Player 1 to win in the vertices in $\mathcal{G}[n]$ in a backwards-inductive manner. For the base case, for every vertex $u \in V \setminus S$, since Player 1 cannot win from u in 0 steps, we have $T_0(u) = k+1$. For $s \in S$, we have $T_0(s) = fr(s)$. Clearly, $T_0 \neq Th_{\mathcal{G}[0]}$. For the inductive step, suppose that T_{n-1} is computed. For each vertex v , we define $T_n(v) = \left\lfloor \frac{|T_{n-1}(v^+)| + |T_{n-1}(v^-)|}{2} \right\rfloor + \varepsilon$ as in Def. 3.6. Following a similar argument to Theorem 3.5, it can be shown that if Player 1's budget is $T_n(v)$, he can bid a b so that if he wins the bidding, his budget is at least $T_{n-1}(v^-)$ and if he loses the bidding, his budget is at least $T_{n-1}(v^+)$. By induction, we get $Th_{\mathcal{G}[n]}(v) = T_n(v)$, for every $v \in V$, which also implies that T_n is a monotonically non-increasing function. Thus, for every vertex v , we let $T(v) = \lim_{n \rightarrow \infty} T_n(v)$, which is well-defined because of the monotonicity of T_n (which is coming from monotonicity of $Th_{\mathcal{G}[n]}$ and the fact that $Th_{\mathcal{G}[n]} = T_n$), and the fact that it only takes finitely many values, namely ranging over $[k] \cup \{k+1\}$. It is not hard to show that T satisfies the average property and that $T(v) \geq Th_{\mathcal{G}}(v)$, for every $v \in V$. ■

Let T be a function that results from the fixed-point computation from the proof of Lemma 3.20. Since it satisfies the average property, we apply Lemma 3.19 to show that Player 2 wins from v when Player 1's budget is $T(v) \ominus 0^*$. Since the values observed in a vertex during

an execution of the fixed-point algorithm are monotonically decreasing and since the number of values that a vertex can obtain is $2k + 1$, the running time is $O(|V| \cdot k)$. We thus conclude the following.

THEOREM 3.21. *Consider a frugal-reachability discrete-bidding game $\mathcal{G} = \langle V, E, k, S, fr \rangle$. Threshold budgets exist and satisfy the average property. Namely, there exists a function $Th : V \rightarrow [k] \cup \{k + 1\}$ such that for every vertex $v \in V$*

- *if Player 1's budget is $B \geq Th(v)$, then Player 1 wins the game, and*
- *if Player 1's budget is $B < Th(v)$, then Player 2 wins the game*

Moreover, there is an algorithm to compute Th that runs in time $O(|V| \cdot k)$, which is exponential in the size of the input when k is given in binary.

A frugal-safety objective is dual to a frugal-reachability objective. Thus, if $Th : V \rightarrow [k] \cup \{k + 1\}$ is the function providing the threshold budgets for the reachability player, then the complement of Th (Definition 3.8), denoted by Th' , provides the threshold budget for the safety player. Therefore, we conclude the following

COROLLARY 3.22. *Consider a frugal-safety discrete-bidding game $\mathcal{G} = \langle V, E, k, S, fr \rangle$. Threshold budgets exist and satisfy the average property. Namely, there exists a function $Th : V \rightarrow [k] \cup \{k + 1\}$ such that for every vertex $v \in V$*

- *if Player 1's budget is $B \geq Th(v)$, then Player 1 wins the game, and*
- *if Player 1's budget is $B < Th(v)$, then Player 2 wins the game*

Moreover, there is an algorithm to compute Th that runs in time $O(|V| \cdot k)$, which is exponential in the size of the input when k is given in binary.

REMARK 3.23. We point to a conceptual similarity with *concurrent stochastic games* [21], in which in each turn, both players concurrently choose actions, and the joint action gives rise to a probability distribution over next states. There, the value of the game is given by the least fixed point of what is referred to as a *value mapping* function [21, 22]. One can adapt their operator to our setting, thereby obtaining an operator that maps functions from states to budgets. Intuitively, each such function would be a candidate for the thresholds, and functions that satisfy the average property are the fixed points. As we show above, the maximal fixed point is the function that coincides with the thresholds.

4. A Fixed-Point Algorithm for Finding Threshold Budgets

In this section, we develop a fixed-point algorithm for finding threshold budgets in frugal-parity discrete-bidding games. While its worst-case running time is exponential in the input, the algorithm shows, for the first time, that threshold budgets in parity discrete-bidding games

satisfy the average property. This property is key in the development of the NP and coNP algorithm (Sec. 5).

4.1 Warm up: a fixed-point algorithm for Büchi bidding games

In this section, we illustrate the ideas of the fixed-point algorithm on the special case of Büchi games. The transition from Büchi to parity games involves an induction on the parity indices. A Büchi game is $\langle V, E, k, F \rangle$, where $F \subseteq V$ is a set of *accepting states*. Formally, Büchi games are a special case of parity games in which the vertices are labeled by $\{0, 1\}$. We stress that, for ease of presentation, we focus on games without a frugal objective.

Throughout this section we will take the perspective of Player 1, the co-Büchi player, whose objective is to visit F only finitely often. We present an algorithm that takes as input a Büchi game \mathcal{G} and outputs Player 1's thresholds, which we denote by coBü-Th . That is, for an initial configuration $\langle v, B \rangle$, we have:

- if $B \geq \text{coBü-Th}(v)$, Player 1 can guarantee that F is visited only *finitely* often, and
- if $B < \text{coBü-Th}(v)$, Player 2 can guarantee that F is visited *infinitely* often.

The fixed-point algorithm repeatedly finds thresholds in a sequence of increasingly easier objectives (from Player 1's perspective) whose limit is the co-Büchi objective. Roughly, recall that the co-Büchi objective requires that F is eventually not visited without specifying a bound on the number of visits to F . The objective Safe_i , for $i \geq 0$, introduces a restriction: F can be visited at most i times. In particular, Safe_0 is a safety game. Formally,

DEFINITION 4.1 (Bounded-eventual safety objectives). For $i \geq 0$, the objective Safe_i contains infinite paths that

- start in $V \setminus F$ and enter F at most i times before exiting F eventually, or
- start in F , exit F for the first time at some point, and then enter F at most $i - 1$ more times before eventually exiting F

The formal definition of when a path enters F can be found in Sec. 2.4. Note that we define Safe_0 so that every path that starts in F violates Safe_0 .

For $i \geq 0$, we denote by $\text{Th}_i : V \rightarrow [k] \cup \{k + 1\}$ the threshold for the objective Safe_i . We make two observations.

OBSERVATION 4.2. For $i \geq 0$ and $v \in V \setminus F$, we have:

- $\text{Th}_i(v) \geq \text{coBü-Th}(v)$, and
- $\text{Th}_i(v) \geq \text{Th}_{i+1}(v)$.

PROOF. First, ensuring the co-Büchi objective (i.e., entering F only finitely often) is easier than ensuring Safe_i (i.e., entering F at most i times), meaning that more budget is necessary for Safe_i

than for co-Büchi, thus $\text{Th}_i(v) \geq \text{coBü-Th}(v)$. Second, similarly, since the restriction imposed by Safe_i is harder than the restriction imposed by Safe_{i+1} , more budget is required for the latter, thus $\text{Th}_i(v) \geq \text{Th}_{i+1}(v)$. ■

It follows that the sequence $\text{Th}_0, \text{Th}_1, \dots$ of thresholds reaches a fixed point. We will show that the fixed point coincides with coBü-Th .

4.1.1 A recursive algorithm to compute thresholds for Safe_i

We describe a recursive algorithm to compute Th_i , for $i \geq 0$. The idea is to characterize Th_i as thresholds in two bidding games: a frugal-reachability game \mathcal{R}_i and a frugal-safety game \mathcal{S}_i . Throughout this section, we follow the convention of using v to denote a vertex in $V \setminus F$ and u to denote a vertex in F .

Base case. Recall that Safe_0 is a safety objective: Player 1 wins by ensuring that F is not visited at all. In particular, paths that start from F are losing for Player 1, thus as in Remark 2.7, we have the following.

LEMMA 4.3. *For $u \in F$, we have $\text{Th}_0(u) = k + 1$.*

Recursive step. Note that, in the base case, we have only computed $\text{Th}_0(u)$ for $u \in F$, and not $\text{Th}_0(v)$ for $v \in V \setminus F$. Therefore, in the general recursive step, we assume $\text{Th}_i(u)$ have already been computed for every $u \in F$ (recursive step hypothesis) and here we first compute $\text{Th}_i(v)$ for $v \in V \setminus F$, which is followed by the computation of $\text{Th}_{i+1}(u)$ for $u \in F$.

We first characterize the thresholds in vertices in $V \setminus F$ as thresholds in a frugal-safety game. Let $i \geq 0$ and suppose that Th_i has been computed for vertices in F . Recall that for $u \in F$, a budget of $\text{Th}_i(u)$ is the threshold to ensure the objective of exiting F and visiting F at most $i - 1$ more times. Suppose that the game starts in $v \in V \setminus F$. Player 1 wins by either not visiting F at all, or if $u \in F$ is reached, Player 1's budget should exceed $\text{Th}_{i-1}(u)$ since he can continue with a strategy that ensures Safe_{i-1} . Formally, we characterize Th_i in $V \setminus F$ as thresholds in a frugal-safety game played on the same arena as \mathcal{G} .

LEMMA 4.4. *Consider the frugal-safety game $\mathcal{S}_i = \langle V, E, k, F, \text{fr}_i \rangle$, where every vertex in F are sinks, and for each $u \in F$, we have $\text{fr}_i(u) = \text{Th}_i(u)$. Then, for every $v \in V \setminus F$, we have $\text{Th}_i(v) = \text{Th}_{\mathcal{S}_i}(v)$.*

REMARK 4.5. Note that, in \mathcal{S}_i , each vertex $u \in F$ are sinks. Thus, Player 1's threshold budget at those vertices in that game would be the same as the frugal budget. That is, for $u \in F$, $\text{Th}_{\mathcal{S}_i}(u) = \text{fr}_i(u) = \text{Th}_i(u)$. Therefore, we indeed have for every vertex $w \in V$ of \mathcal{S}_i , $\text{Th}_i(w) = \text{Th}_{\mathcal{S}_i}(w)$. Since $\text{Th}_{\mathcal{S}_i}$ satisfies the average property (Corollary 3.22), so does Th_i .

Next, we characterize the thresholds in vertices in F as thresholds in a frugal-reachability game. We assume that Th_i has been computed for vertices in $V \setminus F$. Recall that for a vertex $v \in V \setminus F$, a budget of $\text{Th}_i(v)$ is the threshold to ensure the objective of visiting F at most i times. Suppose that the game starts in $u \in F$. In order to ensure Safe_{i+1} , Player 1 must first force the game out of F and then ensure that F is visited at most i more times. This is achieved by ensuring that some $v \in V \setminus F$ is reached with a budget of $\text{Th}_i(v)$. Formally, we characterize Th_{i+1} in F as thresholds in a frugal-reachability game played on the same arena as \mathcal{G} .

LEMMA 4.6. *Consider the frugal-reachability game $\mathcal{R}_{i+1} = \langle V, E, k, V \setminus F, fr_{i+1} \rangle$, where every vertex in $V \setminus F$ are sinks and for each $v \in V \setminus F$, we define $fr_{i+1}(v) = \text{Th}_i(v)$. Then, for every $u \in F$, we have $\text{Th}_{i+1}(u) = \text{Th}_{\mathcal{R}_{i+1}}(u)$.*

Pseudocode We conclude this section with a pseudocode of the fixed-point algorithm. See Fig. 2 for a depiction of its operation. The algorithm calls two sub-routines FRUGAL-REACHABILITY and FRUGAL-SAFETY, which return the thresholds for all vertices that are not targets, respectively, in a frugal-reachability and frugal-safety game, e.g., by running the fixed-point algorithm described in Lemma 3.20.

Algorithm *co-Büchi-Thresholds(\mathcal{G})*

```

1:   $i := 0$ 
2:  Define the frugal-safety game  $\mathcal{S}_0 = \langle V, E, k, F, fr_0 \rangle$ , with  $fr_0 \equiv k + 1$ .
3:   $\text{Th}_{\mathcal{S}_0} = \text{FRUGAL-SAFETY}(\mathcal{S}_0)$ 
4:  Define  $\text{Th}_0(v) = \text{Th}_{\mathcal{S}_0}(v)$ , for  $v \in V \setminus F$ , and  $\text{Th}_0(u) = k + 1$ , for  $u \in F$ .
5:  do
6:     $i := i + 1$ 
7:    Define  $\mathcal{R}_i = \langle V, E, k, V \setminus F, \text{Th}_{\mathcal{S}_{i-1}} \rangle$ .
8:     $\text{Th}_{\mathcal{R}_i} := \text{FRUGAL-REACHABILITY}(\mathcal{R}_i)$  Thresholds for vertices in  $F$ .
9:    Define  $\mathcal{S}_i = \langle V, E, k, F, \text{Th}_{\mathcal{R}_i} \rangle$ 
10:    $\text{Th}_{\mathcal{S}_i} := \text{FRUGAL-SAFETY}(\mathcal{S}_i)$  Thresholds for vertices in  $V \setminus F$ .
11:   Define  $\text{Th}_i(v) = \text{Th}_{\mathcal{S}_i}(v)$ , for  $v \in V \setminus F$ , and  $\text{Th}_i(u) = \text{Th}_{\mathcal{R}_i}(u)$ , for  $u \in F$ .
12: while  $\text{Th}_{i-1} \neq \text{Th}_i$ 

```

Algorithm 1. A fixed-point algorithm to find threshold budgets in co-Büchi games.




Figure 2. A depiction of the fixed-point algorithm for the co-Büchi game on top of the figure. The goal of Player 1 (the co-Büchi player), is to visit t only finitely often. Set the total budget to $k = 5$. The lower part of the figure depicts the progress of thresholds, depicted in orange. For example, the game S_1 is a frugal-safety game in which the safety player wins either if the game never reaches t or if it reaches t , his budget is at least 3^* . The frugal-reachability games have a trivial state space. The threshold in t in \mathcal{R}_i coincides with v_1 in S_{i-1} since the reachability player guarantees reaching the target v_1 with a sufficient budget by bidding 0. The algorithm terminates once a fixed-point is reached. In this example, we see that the thresholds in \mathcal{G} are 0 in all vertices.

4.1.2 The fixed point coincides with coBü-Th

As mentioned above, the sequence $\text{Th}_0, \text{Th}_1, \dots$ of thresholds reaches a fixed point. We show that the fixed-point threshold coincides with the co-Büchi thresholds.

THEOREM 4.7. *Consider a Büchi bidding game \mathcal{G} . For $i \geq 0$, let Th_i be the threshold for satisfying the objective Safe_i , and $n \in \mathbb{N}$ be such that $\text{Th}_n = \text{Th}_{n+1}$. Then, Th_n coincides with the thresholds for the co-Büchi player, namely $\text{Th}_n = \text{coBü-Th}$. Moreover, coBü-Th satisfies the average property and computing it can be done in time $O(|V| \cdot k^2)$.*

PROOF. We show that $\text{Th}_n = \text{coBü-Th}$. First, $\text{Th}_n \geq \text{coBü-Th}$ follows immediately from Observation 4.2; recall that Th_n is the threshold for the objective Safe_n , which is harder for Player 1 to ensure than the co-Büchi objective.

Second, we show that $\text{Th}_n \leq \text{coBü-Th}$. Consider a vertex $v \in V \setminus F$. We show that Player 2, the Büchi player, wins from a configuration $\langle v, B \rangle$ with $B < \text{Th}_n(v) = \text{Th}_{n+1}(v)$. The case of an initial vertex in F is also captured in the following proof. Player 2 plays as follows from $\langle v, B \rangle$. Recall that $\text{Th}_{n+1}(v) = \text{Th}_{\mathcal{S}_{n+1}}(v)$. Thus, a budget of $B < \text{Th}_{\mathcal{S}_{n+1}}(v)$ means that the reachability player wins the frugal-safety game \mathcal{S}_i . Player 2 follows the reachability player's strategy to ensure that F is eventually reached with a budget below the frugal-target budget. Formally, a configuration $\langle u, B' \rangle$ is reached with $u \in F$ and $B' < \text{Th}_{\mathcal{R}_n}(u)$. Next, a budget of B' suffices for the safety player to win the frugal-reachability game \mathcal{R}_i and ensures that either: (1) F is never exited, or (2) $V \setminus F$ is visited with a budget that violates the frugal-target budget. Player 2 follows such a winning strategy to ensure either (1) F is never exited thus is clearly visited infinitely often or (2) a configuration $\langle v', B'' \rangle$ is reached with $v' \in V \setminus F$ and $B'' < \text{Th}_{\mathcal{S}_n}(v')$. Since $\text{Th}_{\mathcal{S}_n}(v') = \text{Th}_{\mathcal{S}_{n+1}}(v')$, Player 2 restarts her strategy from v' . Thus, in both cases Player 2 guarantees infinitely many visits to F , and we are done.

Finally, the thresholds satisfy the average property since so does each Th_i (see Remark 4.5). Regarding running time, note that the thresholds observed in a vertex are monotonically decreasing (Observation 4.2), thus the number of iterations until a fixed-point is reached is $O(|V| \cdot k)$. Each iteration includes two solutions of frugal-reachability games, each of running time $O(|V| \cdot k)$ (Theorem 3.21). ■

4.2 A fixed-point algorithm for frugal-parity bidding games

In this section, we extend the fixed-point algorithm developed in Sec. 4.1 to parity bidding games. The algorithm involves a recursion over the parity indices, which we carry out by strengthening the induction hypothesis and developing an algorithm for frugal-parity objectives instead of the special case of parity objectives.

For the remainder of this section, fix a frugal-parity game $\mathcal{G} = \langle V, E, k, p, S, \text{fr} \rangle$. Denote the maximal parity index by $d \in \mathbb{N}$. Recall that S is a set of sinks where the parity indices are

not defined and $\text{fr}(s)$ denotes Player 1's frugal target budget at $s \in S$. Thus, Player 1 wins a play π if

- π is infinite and satisfies the parity condition, or
- π is finite and ends in a configuration $\langle s, B \rangle$ with $s \in S$ and $B \geq \text{fr}(s)$.

We characterize the thresholds in \mathcal{G} by reasoning about games with a lower parity index. This characterization gives rise to a recursive algorithm to compute the thresholds.

LEMMA 4.8 (Base case). *Let $\mathcal{G} = \langle V, E, p, S, \text{fr} \rangle$ with only one parity index, i.e., $p(v) = d$, for all $v \in V$ and S is the set of sinks for which frugal-budgets are given by fr .*

- *Assume that d is odd. Let $\mathcal{S} = \langle V, E, S, \text{fr} \rangle$ be a frugal-safety game. Then, $\text{Th}_{\mathcal{G}} \equiv \text{Th}_{\mathcal{S}}$.*
- *Assume that d is even. Let $\mathcal{R} = \langle V, E, S, \text{fr} \rangle$ be a frugal-reachability game. Then, $\text{Th}_{\mathcal{G}} \equiv \text{Th}_{\mathcal{R}}$.*

PROOF. Clearly, in both cases, a finite play that ends in a sink is winning in \mathcal{G} iff it is winning in \mathcal{S} , and similarly for \mathcal{R} . When d is odd, any infinite play in \mathcal{G} is winning for Player 1, thus \mathcal{G} is a frugal-safety game. On the other hand, when d is even, any infinite play in \mathcal{G} is losing for Player 1, and the only way to win is by satisfying the frugal objective in a sink, thus \mathcal{G} is a frugal-reachability game. ■

COROLLARY 4.9. *When \mathcal{G} contains only one parity index, computing $\text{Th}_{\mathcal{G}}$ can be done by calling a sub-routine that finds the thresholds in a frugal-reachability (or a frugal-safety) bidding game. Moreover, by Theorem 3.21, $\text{Th}_{\mathcal{G}}$ satisfies the average property in this case.*

Recursive step Suppose that more than one parity index is used. Let $d \in \mathbb{N}$ denote the maximal parity index in \mathcal{G} . We assume access to a sub-routine that computes thresholds in frugal-parity games with a maximal parity index of $d - 1$, and we describe how to use it in order to compute thresholds in \mathcal{G} . We assume that d is even, and we describe the algorithm from Player 1's perspective. The definition for an odd d is dual from Player 2's perspective.

Let $F_d = \{v : p(v) = d\}$. Since d is even, a play that visits F_d infinitely often is losing for Player 1. Thus, a necessary (but not sufficient) requirement to win is to ensure that F_d is visited only finitely often. For example, a Büchi game can be modeled as follows: Player 1 is the co-Büchi player, the parity indices are 1 or 2, and the set F_2 denotes the accepting vertices, which Player 1 needs to visit only finitely often.

We define a bounded variant of the frugal-parity objective, similar to the definition of Safe_i in Sec. 4.1:

DEFINITION 4.10. For $i \geq 0$, a play π is in Fr-Parity_i if:

- π is finite and satisfies the frugal objective: ends in $\langle s, B \rangle$ with $s \in S$ and $B \geq \text{fr}(s)$, or
- π is infinite, satisfies the parity objective, and
 - starts from $V \setminus F_d$ and enters F_d at most i times before eventually exiting, or

- starts from F_d , exits F_d for the first time at some point, and then enters F_d at most $i - 1$ more times before exiting eventually.

In particular, every path that starts from F_d violates Fr-Parity_0 .

For $i \geq 0$, we denote by Th_i the thresholds for objective Fr-Parity_i . As in Observation 4.2, since the restriction monotonically decreases as i grows, the thresholds are monotonically non-increasing and they all lower-bound the thresholds in \mathcal{G} .

OBSERVATION 4.11. *For $i \geq 0$, we have $\text{Th}_{i+1} \leq \text{Th}_i$ and $\text{Th}_{\mathcal{G}} \leq \text{Th}_i$.*

It follows that the sequence of thresholds reaches a fixed-point, and we will show that the thresholds at the fixed point coincide with $\text{Th}_{\mathcal{G}}$.

We iteratively define and solve two sequences of games: a sequence of frugal-parity games $\mathcal{G}_0, \mathcal{G}_1, \dots$ each with maximal parity index $d - 1$ and a sequence of frugal-reachability games $\mathcal{R}_0, \mathcal{R}_1, \dots$. For $i \geq 0$, recall that $\text{Th}_{\mathcal{G}_i}$ and $\text{Th}_{\mathcal{R}_i}$ respectively denote the thresholds in \mathcal{G}_i and \mathcal{R}_i . We will show that Th_i can be characterized by $\text{Th}_{\mathcal{G}_i}$ and $\text{Th}_{\mathcal{R}_i}$: we will show that for $v \in F_d$ we have $\text{Th}_i(v) = \text{Th}_{\mathcal{G}_i}(v)$ and for $u \in V \setminus F_d$, we have $\text{Th}_i(u) = \text{Th}_{\mathcal{R}_i}(u)$.

We start with the frugal-parity games. The games share the same arena, which is obtained from \mathcal{G} by setting the vertices in F_d to be sinks. The games differ in the frugal target budgets. Formally, for $i \geq 0$, we define $\mathcal{G}_i = \langle V, E', p', S', \text{fr}_{\mathcal{G}_i} \rangle$, where the sinks are $S' = S \cup F_d$, the edges are restricted accordingly $E' = \{ \langle v, v' \rangle \in E : v \in V \setminus F_d \}$, the parity function p' coincides with p but is not defined over F_d , i.e., $p'(v) = p(v)$ for all $v \in V \setminus F_d$, and $\text{fr}_{\mathcal{G}_i}$ is the only component that changes as i changes, and it is defined below based on a solution to \mathcal{R}_i . Note that p' assigns at most $d - 1$ parity indices.

We construct the frugal-reachability games. Let $i \geq 0$. Intuitively, the game \mathcal{R}_i starts from F_d and Player 1's goal is to either satisfy the frugal objective in S or reach $V \setminus F_d$ with a budget that suffices to ensure that F_d is entered at most i more times. Formally, we construct the frugal-reachability game $\mathcal{R}_i = \langle V, E'', V \setminus F_d \cup S, \text{fr}_{\mathcal{R}_i} \rangle$, where $E'' = \{ \langle u, u' \rangle \in E : u \in F_d \}$ and

$$\text{fr}_{\mathcal{R}_i}(v) = \begin{cases} \text{fr}(v) & \text{if } v \in S \\ \text{Th}_{\mathcal{G}_i}(v) & \text{if } v \in V \setminus F_d. \end{cases}$$

LEMMA 4.12. *Let $i \geq 0$. Assume that for every $v \in V \setminus F_d$, we have $\text{Th}_{\mathcal{G}_i}(v) = \text{Th}_i(v)$. Then, for every $u \in F_d$, we have $\text{Th}_i(u) = \text{Th}_{\mathcal{R}_i}(u)$.*

PROOF. Suppose that \mathcal{G} starts from $\langle u, B \rangle$ with $u \in F_d$. We first show that when $B \geq \text{Th}_{\mathcal{R}_i}(u)$, Player 1 can ensure the objective Fr-Parity_i . Indeed, by following a winning strategy in \mathcal{R}_i , Player 1 guarantees that either (1) the frugal objective is satisfied in S , in which case the play is clearly winning in \mathcal{G} , or (2) the game reaches a configuration $\langle v, B' \rangle$ with $v \in V \setminus F_d$ and $B' \geq \text{Th}_{\mathcal{G}_i}(v)$, from which, by the assumption that $\text{Th}_{\mathcal{G}_i}(v) = \text{Th}_i(v)$, he can proceed with a winning strategy for Fr-Parity_i .

On the other hand, when $B < \text{Th}_{\mathcal{R}_i}(u)$, Player 2 violates Fr-Parity_i as follows. She first follows a winning strategy in \mathcal{R}_i , which ensures that no matter how Player 1 plays, the resulting play either (1) violates the frugal objective in \mathcal{S} , (2) stays in F_d , or (3) it reaches a configuration $\langle v, B' \rangle$ with $v \in V \setminus F_d$ and $B' < \text{Th}_{\mathcal{G}_i}(v) = \text{Th}_i(v)$. In Cases (1) and (2), the play is clearly winning for Player 2 for violating the objective Fr-Parity_i in \mathcal{G} , and in Case (3), the assumption on $\text{Th}_i(v)$ implies that Player 2 can continue with a strategy that violates Fr-Parity_i . ■

REMARK 4.13. Similar to Remark 4.5, here too, we obtain that Th_i satisfies the average property because it coincides with $\text{Th}_{\mathcal{R}_i}$ which is a function providing threshold budgets in a frugal-reachability game (Theorem 3.21).

We define the frugal target budgets $\text{fr}_{\mathcal{G}_i}$ of the frugal-parity game \mathcal{G}_i . Recall that we obtain \mathcal{G}_i from \mathcal{G} by setting F_d to be sinks. Thus, the sinks in \mathcal{G}_i consist of “old” sinks \mathcal{S} and “new” sinks F_d . The frugal target budgets of \mathcal{G} and \mathcal{G}_i agree on \mathcal{S} , thus for $s \in \mathcal{S}$ and $i \geq 0$, we have $\text{fr}_{\mathcal{G}_i}(s) = \text{fr}(s)$. For $u \in F_d$, we define $\text{fr}_{\mathcal{G}_0}(u) = k + 1$ and for $i > 0$, we define $\text{fr}_{\mathcal{G}_i}(u) = \text{Th}_{\mathcal{R}_{i-1}}(u)$.

LEMMA 4.14. *For $i \geq 0$ and $u \in F_d$, assume that a budget of $\text{fr}_{\mathcal{G}_i}(u)$ is the threshold to satisfy Fr-Parity_i . Then, for $v \in V \setminus F_d$, we have $\text{Th}_i(v) = \text{Th}_{\mathcal{G}_i}(v)$.*

PROOF. Recall that each \mathcal{G}_i agrees with \mathcal{G} on the parity indices in $V \setminus F_d$, thus an infinite path that satisfies the parity condition in \mathcal{G}_i satisfies it in \mathcal{G} , and that \mathcal{G}_i and \mathcal{G} agree on the frugal target budgets in \mathcal{S} .

Under the assumption in the statement, we prove that $\text{Th}_i(v) = \text{Th}_{\mathcal{G}_i}(v)$, for $v \in V \setminus F_d$. Suppose that \mathcal{G} starts from $\langle v, B \rangle$ with $v \in V \setminus F_d$. First, when $B \geq \text{Th}_{\mathcal{G}_i}(v)$, Player 1 ensures Fr-Parity_i by following a winning strategy in \mathcal{G}_i . Let π be the play that is obtained when Player 2 follows some strategy. Note that π is winning for Player 1 in \mathcal{G}_i , thus it satisfies one of the following:

1. π is finite and ends in $\langle s, B \rangle$ with $s \in \mathcal{S}$ and $B \geq \text{fr}_{\mathcal{G}_i}(s) = \text{fr}(s)$,
2. π is infinite (i.e., a sink is never reached) and satisfies the parity condition, or
3. π is finite and ends in $\langle u, B \rangle$ with $u \in F_d$ and $B \geq \text{fr}_{\mathcal{G}_i}(u)$.

Case (1) clearly satisfies the frugal objective of Fr-Parity_i , in Case (2) the parity condition is satisfied without visiting F_d once, thus again, Fr-Parity_i is satisfied. Finally, in Case (3), once the game reaches $\langle u, B \rangle$, the assumption on $\text{fr}_{\mathcal{G}_i}(u)$ implies that Player 1 can follow a strategy that ensures Fr-Parity_i . Second, if $B < \text{Th}_{\mathcal{G}_i}(v)$, Player 2 violates Fr-Parity_i by following a winning strategy in \mathcal{G}_i . The argument is dual to the above. ■

Note that since every path that starts from F_d violates Fr-Parity_0 , the threshold budget at every $u \in F_d$ is $k + 1$. This constitutes the proof of the base case of the following lemma, and the inductive step is obtained by combining Lemma 4.12 with Lemma 4.14.

LEMMA 4.15. *For $i \geq 0$, for $v \in V$ we have $Th_i(v) = Th_{\mathcal{G}_i}(v)$ and for $u \in V \setminus F_d$, we have $Th_i(u) = Th_{\mathcal{R}_i}(u)$.*

It follows from Observation 4.11 that the sequence Th_0, Th_1, \dots reaches a fixed point. We show that at the fixed point, the threshold coincides with $Th_{\mathcal{G}}$.

LEMMA 4.16. *Let $n \in \mathbb{N}$ such that $Th_n = Th_{n+1}$. Then, $Th_{\mathcal{G}} = Th_n$.*

PROOF. Lemma 4.15 and Observation 4.11 show that $Th_{\mathcal{G}} \leq Th_n$. To show equality, we show that Player 2 wins \mathcal{G} starting from a configuration $\langle v, B \rangle$ with $v \in V \setminus F_d$ and $B < Th_n(v)$. Player 2 proceeds by following a winning strategy in \mathcal{G}_{n+1} . Let π be a play that results from some Player 1 strategy. Since π is winning for Player 2 in \mathcal{G}_{n+1} , there are three cases:

1. π is finite and ends in $\langle s, B' \rangle$ with $s \in S$ and $B' < fr(s)$, thus it is winning also in \mathcal{G} ,
2. π is infinite and violates the parity objective, thus since \mathcal{G} and \mathcal{G}_{n+1} agree on the parity indices, π is winning for Player 2 in \mathcal{G} , or
3. π ends in $\langle u, B' \rangle$ with $B' < fr_{n+1}(u)$.

In Case (3), since $fr_{n+1}(u) = Th_{\mathcal{R}_n}(u)$, Player 2 continues by following a winning strategy for the safety player in \mathcal{R}_n . This guarantees that no matter how Player 1 plays, the play either stays within F_d , thus it necessarily violates the parity objective of \mathcal{G} , or it reaches $\langle v, B'' \rangle$ with $v \in V \setminus F_d$ and $B'' < fr_{\mathcal{R}_n}(v)$. In the latter case, since $fr_{\mathcal{R}_n}(v) = Th_n(v) = Th_{n+1}(v)$, Player 2 can restart her strategy. Note that Player 2's strategy guarantees that either F_d is eventually never reached, then she wins, or it is reached infinitely often, in which case she also wins since the play visits parity index d infinitely often. ■

Pseudocode The algorithm is described in Alg. 2 for an even d and from Player 1's perspective.

Note that since in a frugal-reachability game both the thresholds for the reachability and safety player satisfy the average property (Theorem 3.21) and the algorithm boils down to repeated calls to a solution of a frugal-reachability game, it outputs a function that satisfies the average property.

THEOREM 4.17. *Given a frugal-parity bidding game \mathcal{G} with maximal index d , Alg. 2 outputs the thresholds $Th_{\mathcal{G}}$. Moreover, $Th_{\mathcal{G}}$ satisfies the average property and Alg. 2 runs in time $O((|V| \cdot k)^d)$.*

REMARK 4.18. We point out that while we develop Alg. 2 for discrete-bidding games, it can be seen as a general “recipe” for extending a solution for frugal-reachability games to parity bidding games. While the algorithm that arises from this recipe might not be optimal complexity wise, it does provide a first upper bound, and importantly, it extends a proof that thresholds in frugal-reachability games have the average property to parity bidding games.

Algorithm *Frugal-Parity-Threshold*($\mathcal{G} = \langle V, E, k, p, S, fr \rangle$)

```

1:  if  $\mathcal{G}$  uses one parity index  $d$  then
2:      if  $d$  is odd then
3:          Return FRUGAL-SAFETY( $\mathcal{S} = \langle V, E, k, S, fr \rangle$ )
4:      else
5:          Return FRUGAL-REACHABILITY( $\mathcal{R} = \langle V, E, k, S, fr \rangle$ )
6:  Define  $E' = \{\langle v, v' \rangle \in E : v \in V \setminus F_d\}$  and  $E'' = \{\langle u, u' \rangle \in E : u \in F_d\}$ .
7:  Define  $\mathcal{G}_0 = \langle V, E', k, p|_{V \setminus F_d}, S \cup F_d, fr_{\mathcal{G}_0} \rangle$  with  $fr_{\mathcal{G}_0}(u) = \begin{cases} fr(u) & \text{if } u \in S \\ k+1 & \text{if } u \in F_d \end{cases}$ .
8:   $Th_{\mathcal{G}_0} = \text{FRUGAL-PARITY-THRESHOLD}(\mathcal{G}_0)$ 
9:  Define  $\mathcal{R}_0 = \langle V, E'', k, (V \setminus F_d) \cup S, fr_{\mathcal{R}_0} \rangle$ .
10:  $Th_{\mathcal{R}_0} = \text{FRUGAL-REACHABILITY-THRESHOLD}(\mathcal{R}_0)$ 
11: for  $i = 1, \dots$  do
12:      $Th_{\mathcal{G}_i} = \text{FRUGAL-PARITY-THRESHOLD}(\mathcal{G}_i = \langle V, E', k, p', S \cup F_d, fr_{\mathcal{G}_i} = fr \cup Th_{\mathcal{R}_{i-1}} \rangle)$ 
13:      $Th_{\mathcal{R}_i} = \text{FRUGAL-REACHABILITY-THRESHOLD}(\mathcal{R}_i = \langle V, E'', k, V \setminus F_d, Th_{\mathcal{G}_i} \rangle)$ 
14:     For each  $v \in F_d$ , define  $fr_{i+1}(v) = Th_{\mathcal{R}_i}(v)$ 
15:     if  $fr_i(v) = fr_{i+1}(v)$ , for all  $v \in F_d$  then
16:         Define  $Th_{\mathcal{G}}(v) = fr_i(v)$  for  $v \in F_d$ .
17:         Define  $Th_{\mathcal{G}}(u) = Th_{\mathcal{G}_i}(u)$  for  $u \in V \setminus F_d$ .
18:     Return  $Th_{\mathcal{G}}$ 

```

Algorithm 2. A fixed-point algorithm to find threshold budgets in frugal parity games.

5. Finding threshold budgets is in NP and coNP

We formalize the problem of finding threshold budgets as a decision problem:

Problem 1. (Finding Threshold Budgets). Given a frugal-parity bidding game $\mathcal{G} = \langle V, E, k, p, S, fr \rangle$, a vertex $v \in V$, and $\ell \in [k]$, decide whether $Th_{\mathcal{G}}(v) \geq \ell$.

We will show that Prob. 1 is in NP and coNP. Note that a function $T : V \rightarrow [k] \cup \{k+1\}$ can be represented using $O(|V| \cdot \log(k))$ bits, thus it is polynomial in the size of the input to Prob. 1. We describe a first attempt to show membership in NP and coNP. Guess T , verify that it satisfies the average property, and accept $\langle \mathcal{G}, v, \ell \rangle$ iff $T(v) \geq \ell$. Unfortunately, such an attempt fails. Even though by Theorem 4.17, the thresholds satisfy the average property, Theorem 3.7 shows that there can be other functions that satisfy it. That is, it could also be the case that T

satisfies the average property and $T \neq \text{Th}_{\mathcal{G}}$. We point out that in continuous-bidding games, if guessing such T would be possible, this scheme would have succeeded since there is a unique function that satisfies the continuous average property (Theorem 3.5).

In the remainder of this section, we will show that the following problem is in NP and coNP by reducing it to solving a turn-based parity game of size linear in the size of the graph (and not in the size of the encoding of k). Then an algorithm for Prob. 1 guesses both T and winning strategies in the turn-based game.

Problem 2. (Verifying a guess of T). *Given a frugal-parity discrete-bidding game \mathcal{G} with vertices V and a function $T : V \rightarrow [k] \cup \{k+1\}$ that satisfies the average property, decide whether $T \equiv \text{Th}_{\mathcal{G}}$.*

We describe the high-level idea. We find it instrumental to first recall an NP algorithm to decide whether Player 1 wins a turn-based parity game from an initial vertex v_0 . The algorithm first guesses a memoryless strategy f , which is a function that maps each vertex v that is controlled by Player 1 to an outgoing edge from v . The algorithm then verifies that f is winning for Player 1. The verification proceeds as follows. We solve the following problem: given a Player 1 strategy f , check whether Player 2 has a counter strategy g such that $\text{play}(v_0, f, g)$ violates Player 1's objective. If we find that Player 2 has a counter strategy g to f , then f is *not* winning, and we reject the guess. On the other hand, if Player 2 cannot counter f , then f is winning, and we accept the guess. Deciding whether Player 2 can counter f is done as follows. We trim every edge in the game that does not comply with f . This leaves a graph with only Player 2 choices, and we check if in every reachable SCC the highest priority index is odd. There exists a reachable SCC where the highest priority is even iff Player 2 can counter f iff f is *not* winning.³

Our algorithm for frugal-parity games follows conceptually similar steps. Let $T : V \rightarrow [k] \cup \{k+1\}$ that satisfies the average property. We verify whether $T \equiv \text{Th}_{\mathcal{G}}$ as follows. We construct a partial strategy f_T based on T . Recall that a partial strategy proposes a bid and a set of allowed vertices in each configuration. That is, guessing T is not quite like guessing a strategy as in the algorithm above, rather T gives rise to a partial strategy. We seek a Player 1 strategy that agrees with f_T and wins when the game starting from every configuration $\langle v, T(v) \rangle$ with $T(v) < k+1$. Note that if $T \equiv \text{Th}_{\mathcal{G}}$, then such a strategy exists. Given f_T , we describe an algorithm that decides whether Player 2 can counter every Player 1 strategy that agrees with f_T . Our algorithm constructs and solves a turn-based parity game.

³ An alternative description of the verification algorithm is the following. View the trimmed graph as an automaton with a singleton alphabet whose acceptance condition is Player 2's objective, and check whether the language of the automaton is empty. The language is empty iff Player 2 cannot counter f .

5.1 From bidding games to turn-based games

Let T be a function that satisfies the average property. Recall that the partial strategy f_T that is constructed in Sec. 3.2.2 is a function that, given a configuration $\langle v, B \rangle$, outputs $\langle b, A \rangle$, where $b \leq B$ is a bid and $A \subseteq V$ is a subset of neighbors of v that are called *allowed vertices*. A strategy f' agrees with f_T if from each configuration, it bids the same as f_T and chooses an allowed vertex upon winning the bidding.

We construct a parity turn-based game $G_{T,\mathcal{G}}$ such that, roughly, if Player 1 wins in every vertex in $G_{T,\mathcal{G}}$, then Player 1 has a strategy f' that agrees with f_T and wins from every configuration $\langle v, T(v) \rangle$ in \mathcal{G} , thus $T \geq \text{Th}_{\mathcal{G}}$.

We describe the intuition behind the construction of $G_{T,\mathcal{G}}$. Consider the following first attempt to construct $G_{T,\mathcal{G}}$. Recall the construction in Sec. 2.3 of the explicit concurrent game that corresponds to \mathcal{G} , and denote it by \mathcal{G}' . The vertices of \mathcal{G}' are the configurations C of \mathcal{G} . We construct a game \mathcal{G}'' on the configuration graph C . Recall that our goal is to check whether Player 2 can counter Player 1's strategy, which can be thought of as Player 2 responds to Player 1's actions in each turn. Thus, \mathcal{G}'' is turn-based: when the game is in configuration c , Player 1 first chooses $\langle b_1, v_1 \rangle$, and only then, Player 2 responds by choosing an action $\langle b_2, v_2 \rangle$. The next configuration is determined by these two actions in the same manner as the concurrent game. Next, we trim Player 1 actions in \mathcal{G}' that do not comply with f_T : in a configuration $c = \langle v, B \rangle$ in \mathcal{G}'' with $\langle b, A \rangle = f_T(c)$, Player 1 must bid b and choose a vertex in A . That is, an action $\langle b', v' \rangle$ is not allowed if $b' \neq b$ or if $v' \notin A$. Finally, we omit Player 2 actions that are dominated: observing a Player 1 bid of b , she chooses between bidding 0 and letting Player 1 win the bidding or bidding $b \oplus 0^*$ and winning the bidding. It is not hard to see that Player 1 wins \mathcal{G}'' from configuration $\langle v, B \rangle$ iff there is a strategy f' that agrees with f_T and wins \mathcal{G} from $\langle v, B \rangle$.

The first attempt fails since the size of \mathcal{G}'' is proportional to the number of configurations, which is exponential in \mathcal{G} . We overcome this key challenge as follows. Lemma 3.17 shows that when \mathcal{G} starts from configuration $\langle v, B \rangle$ with $B \geq T(v)$ a strategy f' that agrees with f_T maintains an invariant on Player 1's budget: the game only reaches configurations of the form $\langle v', B' \rangle$ with $B' \geq T(v')$. We shrink the size of the game by grouping all configurations in which Player 1's budget is greater than $T(v) \oplus 0^*$ into a vertex denoted $\langle v, \top \rangle$.

We describe the idea that allows keeping only three copies of each vertex (see details in Lemma 5.2). We refer to the distance from the invariant as *spare change*, formally at v with budget $B \geq T(v)$, the spare change is $|B \ominus T(v)|$. Recall from Sec. 3.2.2 that f_T chooses one of two bids in a vertex $v \in V$, and the choice depends on the advantage status and does not depend on the spare change. Thus, our winning strategy in \mathcal{G} emulates a winning strategy in $G_{T,\mathcal{G}}$: both bid according to f_T and the latter prescribes a vertex to move to upon winning a bidding. Thus, a play π in \mathcal{G} corresponds to a play $\tilde{\pi}$ in $G_{T,\mathcal{G}}$. There can be three outcomes: (1) $\tilde{\pi}$ is infinite, (2) $\tilde{\pi}$ ends in a sink S , or (3) $\tilde{\pi}$ ends in a sink $\langle v, \top \rangle$. The first two cases mean that π is winning in \mathcal{G} .

When Case (3) occurs and $G_{T,\mathcal{G}}$ reaches $\langle v, \top \rangle$, then \mathcal{G} reaches $\langle v, B \rangle$ with $B > T(v) \oplus 0^*$, and we restart $G_{T,\mathcal{G}}$ from either $\langle v, T(v) \rangle$ or $\langle v, T(v) \oplus 0^* \rangle$ depending on the advantage status. Note that, after restarting the game, Player 1 plays the same except that his spare change increased. This is the key idea. Since whenever Case (3) occurs, Player 1's spare change strictly increases and the spare change cannot exceed the total budget k , Case (3) can occur only finitely often.

Formally, we define the turn-based parity game $G_{T,\mathcal{G}} = \langle V_1, V_2, E, p \rangle$. The vertices controlled by Player i are V_i , for $i \in \{1, 2\}$, where


- $V_1 = \{\langle v, T(v) \rangle, \langle v, T(v) \oplus 0^* \rangle, \langle v, \top \rangle : v \in (V \cup S)\}$ and
- $V_2 = \{\langle v, c \rangle : v \in V, c \in C\}$.

Note that it is possible that $T(v) = k + 1$, which as in Remark 2.7, means that Player 1 loses from v in \mathcal{G} with every initial budget. We define each vertex $\langle v, k + 1 \rangle$ as losing for Player 1 in $G_{T,\mathcal{G}}$: it is a sink with even parity index. We define the edges in the game. A vertex $\langle v, B \rangle$ is a sink if $v \in S$ or if $B = \top$. Consider $c = \langle v, B \rangle \in V_1$ and let $\langle b, A \rangle = f_T(c)$. The neighbors of c are $\{\langle v', c \rangle : v' \in A\}$. Intuitively, $\langle v', c \rangle$ means that Player 1 chooses the action $\langle b, v' \rangle$ at configuration c ; the bid b is determined by f_T and v' is an allowed vertex. A vertex $\langle v', c \rangle$ is a Player 2 vertex. Intuitively, Player 2 makes two choices: who wins the bidding and where the token moves upon winning. Thus, a vertex $\langle v', c \rangle$ has two types of neighbors, depending on who wins the bidding at c :

- First, $\langle v', B \ominus b \rangle$ is a neighbor of $\langle v', c \rangle$, meaning Player 2 lets Player 1 win the bidding by bidding 0.
- Second, suppose that $k^* \ominus B \geq b \oplus 0^*$, i.e., Player 2 has sufficient budget to win the bidding. Let $B' = B \oplus (b \oplus 0^*)$ be Player 1's updated budget and $w \in N(v)$. If $c' = \langle w, B' \rangle \in V_1$, then c' is a neighbor of $\langle v', c \rangle$. We note that $c' \notin V_1$ when B' exceeds $T(w) \oplus 0^*$, then we trim the budget and set $\langle w, \top \rangle$ as a neighbor of $\langle v', c \rangle$.

For ease of presentation, we define parity indices only in Player 1 vertices. A non-sink vertex in $G_{T,\mathcal{G}}$ “inherits” its parity index from the vertex in \mathcal{G} ; namely, for $c = \langle v, B \rangle \in V_1$, we define $p'(c) = p(v)$. The parity index of a sink is odd so that Player 1 wins in sinks.

EXAMPLE 5.1. Fig. 1 depicts a frugal-reachability bidding game \mathcal{G}_1 with two functions that satisfy the average property: $T_1 = \langle 4, 3^*, 3, 2 \rangle$ and $T_2 = \langle 5, 4^*, 3^*, 2 \rangle$. Fig. 3 depicts the games G_{T_1,\mathcal{G}_1} and G_{T_2,\mathcal{G}_1} . For ease of presentation, Fig. 3 is slightly inconsistent with the construction of the games. The reason is that both f_{T_1} and f_{T_2} prescribe a singleton set of allowed vertices from all configurations, thus Player 1 makes no choices in the game. We thus skip his vertices and simplify Player 2's vertices: each vertex in Fig. 3 corresponds to a configuration, and all vertices are controlled by Player 2. Player 1's goal in both games is to reach a sink. An outgoing edge from vertex c labeled by $\langle b_1, b_2 \rangle$ represents the outcome of a bidding at configuration c in which Player i bids b_i , for $i \in \{1, 2\}$. Thus, each vertex c has two outgoing edges labeled by $\langle b_1, 0 \rangle$ and $\langle b_1, b_1 \oplus 0^* \rangle$, where b_1 is the bid that f_{T_1} or f_{T_2} prescribes at c . Note that some edges



(a) The turn-based game G_{T_1, \mathcal{G}_1} . Player 1 loses from some vertices, thus $T_1 \neq \text{Th}_{\mathcal{G}_1}$

(b) The turn-based game G_{T_2, \mathcal{G}_1} . Player 1 wins from all vertices, thus $T_2 \equiv \text{Th}_{\mathcal{G}_1}$

Figure 3. Turn-based reachability games that correspond to \mathcal{G}_1 from Fig. 1 for two functions that satisfy the average property. Player 1 vertices are omitted; that is, all depicted vertices are Player 2 vertices. The edge labeling depict bidding outcomes and are meant to ease presentation.

are disallowed. For example, in the configuration $\langle v_1, 5 \rangle$ in G_{T_2, \mathcal{G}_1} , the bid prescribed by f_{T_2} is $b_1 = 1$ and Player 2 cannot bid $b_1 \oplus 0^* = 1^*$ since it exceeds her available budget (indeed, $k = 5$, thus Player 2's budget in c is $k^* \ominus 5 = 0^*$).

Note that G_{T_1, \mathcal{G}_1} has a cycle. Thus, Player 1 does not win from every vertex and T_1 does not coincide with the threshold budgets. On the other hand, G_{T_2, \mathcal{G}_1} is a DAG. Thus, no matter how Player 2 plays, Player 1 wins from all vertices, and indeed $T_2 \equiv \text{Th}_{\mathcal{G}_1}$. \blacklozenge

5.2 Correctness

In this section, we prove soundness and completeness of the approach. We start with soundness.

LEMMA 5.2. *If Player 1 wins from every vertex $\langle v, B \rangle$ in $G_{T, \mathcal{G}}$ with $B < k + 1$, then $T \geq \text{Th}_{\mathcal{G}}$.*

PROOF. Suppose that Player 1 wins from every such vertex of $G_{T, \mathcal{G}}$ and let \tilde{f} be a Player 1 memoryless winning strategy. We construct a strategy f in \mathcal{G} based on \tilde{f} and show that it is winning from every configuration $\langle v, B \rangle$ where $B \geq T(v)$. This implies that $T \geq \text{Th}_{\mathcal{G}}$ since f witnesses that Player 1 can win with a budget of $T(v)$ from v . Note that we do not yet rule out that a different strategy wins with a lower budget, this will come later.

We introduce notation. Consider a configuration $c = \langle v, B \rangle$ in \mathcal{G} with $T(v) < k + 1$ and $B \geq T(v)$. The vertex in $G_{T, \mathcal{G}}$ that agrees with c , denoted by \tilde{c} , is the vertex in $\{\langle v, T(v) \rangle \langle v, T(v) \oplus 0^* \rangle\}$ that matches with c on the status of the advantage (and of course on the vertex of \mathcal{G}). Note the convention of calling c a *configuration* in \mathcal{G} and a *vertex* in $G_{T, \mathcal{G}}$. For example if $T(v) = 5^*$ for

some vertex v and $c = \langle v, 9 \rangle$ is a configuration in \mathcal{G} , then the vertex of $G_{T,\mathcal{G}}$ that *agrees* with c , denoted by \tilde{c} , is $\langle v, 6 \rangle$. Recall that even though the budget in c may be higher than that of \tilde{c} , the partial strategy f_T acts the same in both, i.e., $f_T(c) = f_T(\tilde{c})$. The *spare change* that is associated with c , denoted by $\text{Spare}(c)$ is $|B| - |T(v)|$.

In the following, we construct f based on f_T and \tilde{f} . Specifically, we define f to agree with f_T on the bid and choose the successor vertex according to \tilde{f} . Let $\langle b, A \rangle = f_T(c)$. Recall that \tilde{c} is a Player 1 vertex in $G_{T,\mathcal{G}}$ and its neighbours are of the form $\langle v', c \rangle$ such that v' is an allowed vertex, i.e., $v' \in A$. Intuitively, proceeding to vertex $\langle v', \tilde{c} \rangle$ in $G_{T,\mathcal{G}}$ is associated with moving to v' upon winning the bidding at \tilde{c} . Let $\langle v', \tilde{c} \rangle = \tilde{f}(\tilde{c})$. Then, we define $f(c) = \langle b, v' \rangle$.

We claim that f is winning from an initial configuration $c_0 = \langle v, B \rangle$ in \mathcal{G} with $T(v) < k + 1$ and $B \geq T(v)$. Let g be a Player 2 strategy in \mathcal{G} . The initial vertex \tilde{c}_0 in $G_{T,\mathcal{G}}$ is the vertex that *agrees* with c_0 . We construct a Player 2 strategy \tilde{g} in $G_{T,\mathcal{G}}$ so that $\text{play}(\tilde{c}_0, \tilde{f}, \tilde{g})$ in $G_{T,\mathcal{G}}$ simulates $\text{play}(c_0, f, g)$ in \mathcal{G} : when $\text{play}(c_0, f, g)$ is in a configuration c , $\text{play}(\tilde{c}_0, \tilde{f}, \tilde{g})$ is in a vertex \tilde{c} that *agrees* with c .

We define \tilde{g} inductively. Initially the invariant holds due to our choice of \tilde{c}_0 in $G_{T,\mathcal{G}}$. Suppose that \mathcal{G} is at configuration $c = \langle v, B \rangle$, then the $\text{play}(\tilde{c}_0, \tilde{f}, \tilde{g})$ in $G_{T,\mathcal{G}}$ is at the vertex c^* that *agrees* with c . Denote by $\hat{T}(v) \in \{T(v), T(v) \oplus 0^*\}$ such that $c^* = \langle v, \hat{T}(v) \rangle$. Let $\langle b_1, v_1 \rangle = f(c)$ as defined above, let $\langle b_2, v_2 \rangle = g(c)$ be Player 2's choice, and let d be the next configuration in \mathcal{G} . We extend the play in $G_{T,\mathcal{G}}$ as follows. We first register Player 1's move in $G_{T,\mathcal{G}}$ by proceeding to the Player 2 vertex $\langle v_1, c^* \rangle$. We distinguish between two cases. First, Player 1 wins the bidding in \mathcal{G} . We define \tilde{g} to choose $\langle v_1, \hat{T}(v) \ominus b_1 \rangle$ as the successor vertex from $\langle v_1, c^* \rangle$. Note that, in this case, the configuration d is $d^* = \langle v_1, B \ominus b_1 \rangle$. Since c^* *agrees* with c , then d^* *agrees* with d . Second, Player 2 wins the bidding in \mathcal{G} . We define \tilde{g} to proceed to $d^* = \langle v_2, \hat{T}(v) \oplus b_2 \rangle$, if it exists in $G_{T,\mathcal{G}}$. If d^* is in the graph, then again, d^* *agrees* with d . On the other hand, if d^* is not a vertex in $G_{T,\mathcal{G}}$, it intuitively means $\hat{T}(v) \oplus b_2 > T(v_2) \oplus 0^*$, and we define \tilde{g} to proceed to $\langle v_2, \top \rangle$. Let us denote \hat{d}^* be the vertex in $G_{T,\mathcal{G}}$ that *agrees* with d . We apply the same definition above starting from vertex \hat{d}^* in $G_{T,\mathcal{G}}$. That is, in the next turn, assuming that Player 1 proceeds in $G_{T,\mathcal{G}}$ to $\langle v_2, \hat{d}^* \rangle$, we define \tilde{g} according to $g(\langle v_2, B \oplus b_2 \rangle)$, as discussed above. We call this a *restart* in the simulation. Note that if the simulation is not a restart, then $\text{Spare}(c) = \text{Spare}(d)$, and if it is a restart, then $\text{Spare}(c) < \text{Spare}(d)$.

We claim that $\text{play}(c_0, f, g) = c_0, c_1, \dots$ is winning for Player 1 in \mathcal{G} . We slightly abuse notation and denote by $\text{play}(\tilde{c}_0, \tilde{f}, \tilde{g}) = \tilde{c}_0, c_1^*, \dots$ the sequence of Player 1 vertices that are traversed in $G_{T,\mathcal{G}}$, that is we skip Player 2 vertices. Since \tilde{f} is winning in $G_{T,\mathcal{G}}$, then $\text{play}(\tilde{c}_0, \tilde{f}, \tilde{g})$ is winning for Player 1. We distinguish between three cases. First, $\text{play}(\tilde{c}_0, \tilde{f}, \tilde{g})$ is infinite. Since for every $i \geq 0$, the vertex c_i^* *agrees* with c_i , the two plays agree on the parity indices that are visited, thus $\text{play}(c_0, f, g)$ satisfies the parity objective. Second, $\text{play}(\tilde{c}_0, \tilde{f}, \tilde{g})$ is finite and ends in a sink configuration $c_k = \langle s, B \rangle$. Note that $B \geq T(s)$, and the definition of T requires that $T(s) = \text{fr}(s)$. Since c_k^* *agrees* with c_k , it follows that c_k satisfies the frugal objective. Third,

$\text{play}(\tilde{c}_0, \tilde{f}, \tilde{g})$ is finite and ends in $c_k^* = \langle v, \top \rangle$. Let \hat{c}_k^* denote the vertex that agrees with c_k . We apply the reasoning above to $\text{play}(\hat{c}_k^*, \tilde{f}, \tilde{g})$.

Note that the third case can occur only finitely many times, since a restart causes the spare change to strictly increase, and the spare change is bounded by k . Thus, eventually, the play in $G_{T,\mathcal{G}}$ falls into one of the first two cases, which implies that $\text{play}(c_0, f, g)$ is winning for Player 1. ■

REMARK 5.3. The proof of Lemma 5.2 constructs a Player 1 winning strategy f in \mathcal{G} . Note that in order to implement f , we only need to keep track of a vertex in $G_{T,\mathcal{G}}$. Thus, its memory size equals the size of $G_{T,\mathcal{G}}$, which is linear in the size of \mathcal{G} . This is significantly smaller than previously known constructions in parity and reachability discrete-bidding games, where the strategy size is polynomial in k , and is thus exponential when k is given in binary.

The following lemma shows completeness; namely, that a correct guess of T implies that Player 1 wins from every vertex in $G_{T,\mathcal{G}}$.

LEMMA 5.4. *If $T \equiv \text{Th}_{\mathcal{G}}$, then Player 1 wins from every vertex $\langle v, B \rangle$ in $G_{T,\mathcal{G}}$ with $B \in \mathbb{N}^*$ and $B < k + 1$.*

PROOF. Assume towards a contradiction that $T \equiv \text{Th}_{\mathcal{G}}$ and there is a Player 1 vertex $\tilde{c}_0 = \langle v, B \rangle$ in $G_{T,\mathcal{G}}$ that is losing for Player 1. Let \tilde{g} be a Player 2 memoryless strategy that wins from vertex \tilde{c}_0 in $G_{T,\mathcal{G}}$. Recall that $B \in \{T(v), T(v) \oplus 0^*\}$. Note that $B \geq T(v)$. Since we assume $T \equiv \text{Th}_{\mathcal{G}}$, and $B < k + 1$ Player 1 wins from configuration $c_0 = \langle v, B \rangle$ in \mathcal{G} . Let f be a Player 1 winning strategy from c_0 in \mathcal{G} . Note that we follow the convention of referring to c in \mathcal{G} as a *configuration* and \tilde{c} in $G_{T,\mathcal{G}}$ as a *vertex*, even though both are $\langle v, B \rangle$. We will reach a contradiction by constructing a Player 2 strategy g in \mathcal{G} that counters f , thus showing that f is not winning. Recall that a winning Player 1 strategy can be thought of as a strategy that, in each turn, reveals Player 1's action first, and allows Player 2 to respond to Player 1's action.

We construct a Player 1 strategy \tilde{f} in $G_{T,\mathcal{G}}$ based on f as long as f agrees with f_T and a Player 2 strategy g in \mathcal{G} based on \tilde{g} . Both constructions are straightforward. First, for \tilde{f} , consider a Player 1 vertex \tilde{c} in $G_{T,\mathcal{G}}$. Recall that \tilde{c} is a configuration in \mathcal{G} , which we denote by c to avoid confusion. Suppose that $f(c)$ agrees with $f_T(c)$, that is denoting $\langle b, A \rangle = f_T(c)$, we have $\langle b, v \rangle = f(c)$ with $v \in A$. Then, in $G_{T,\mathcal{G}}$, from vertex \tilde{c} , the strategy \tilde{f} proceeds to $\langle v, \tilde{c} \rangle$. We stress that \tilde{f} is only defined when f agrees with f_T . Second, for g , recall that Player 2 vertices in $G_{T,\mathcal{G}}$ are of the form $\langle v, c \rangle$, and Player 2 chooses between, intuitively letting Player 1 win the bidding or bidding $b \oplus 0^*$, winning the bidding, and choosing the next vertex. Assume $\langle b, v \rangle = f(c)$ that agrees with f_T , then g responds by following \tilde{g} : if \tilde{g} lets Player 1 win from $\langle v, \tilde{c} \rangle$, then g bids 0 in c and lets Player 1 win the bidding, and if it wins the bidding by proceeding to vertex $\langle v', B' \rangle$, then g chooses $\langle b \oplus 0^*, v' \rangle$, i.e., it too wins the bidding in c and proceeds to vertex v' .

Let $\tilde{\pi}$ and π respectively denote the longest histories of $G_{T,\mathcal{G}}$ and \mathcal{G} that start from \tilde{c}_0 and c_0 that arise from applying \tilde{f} against \tilde{g} in $G_{T,\mathcal{G}}$ and f against g in \mathcal{G} , as long as f agrees with f_T . Note that, skipping Player 2 vertices in $G_{T,\mathcal{G}}$, the plays $\tilde{\pi}$ and π traverse the same sequence of configurations. We claim that the two plays cannot be infinite. Indeed, assume otherwise, then since we assume f is winning, π satisfies Player 1's objective, and since we assume \tilde{g} is winning, $\tilde{\pi}$ violates Player 1's objective, but both cannot hold at the same time. Also, $\tilde{\pi}$ cannot end in a sink. Indeed, visiting a sink that is winning for Player 1 is not possible since $\tilde{\pi}$ is consistent with a Player 2 winning strategy, and visiting a winning sink $\langle v, k+1 \rangle$ for Player 2 is not possible since it means that π , which is consistent with a Player 1 winning strategy in \mathcal{G} , visits a losing configuration for Player 1. We conclude that π and $\tilde{\pi}$ are finite and end in a configuration in which f does not agree with f_T .

Let $c = \langle v, B \rangle$, where $B \in \{T(v), T(v) \oplus 0^*\}$, be the last configuration in π . That is, c is the first configuration in which f chooses an action that does not agree with f_T . Let $\langle b, A \rangle = f_T(c)$ and $\langle b_1, v_1 \rangle = f(\pi)$. In the remainder of the proof, we consider the three ways in which f can disagree with f_T . In each of these cases, we subsequently define Player 2 response g , and show that she can win from the resulting configuration.

In particular, we show that in all but one subcases of these three cases, we have a “suitable” Player 2 response by g which results in a configuration of the form $c = \langle v', B' \rangle$, where $B' < T(v')$. Thus, the standard argument follows (in all but one subcase) from there as: because the budget “falls” below the threshold budget (by hypothesis, $T \equiv \text{Th}$), by definition, Player 2 has a winning strategy in \mathcal{G} from there onwards, and g simply follows that. In the remaining subcase (Case 2. (ii), in particular), we will see that even though this is also a way how f differs from f_T , it does not necessarily result Player 1's budget falling below $T(v')$ (assuming the resulting vertex is v) for any Player 2 response. But in this case, we argue that f eventually differs from f_T by other means. Hence, even though we may not have the desired “suitability” in Player 2's response g in this case, we will eventually encounter it when f eventually differs from f_T by other means.

We recall Observation 3.13, which intuitively states that when Player 1 has the advantage and the bids of f and f_T agree, then Player 1 uses the advantage.

We finally proceed to analyze the three ways in which f disagrees with f_T :

Case 1: f underbids; $b_1 < b$. Player 2 responds by bidding $b \ominus 0^*$. We show that she wins the bidding, but before that we show Player 2 can indeed bid $b \ominus 0^*$ at vertex v from her budget $k^* \ominus B$. Note that, here $b = b_v^T$ if $B = T(v)$, and $b = b_v^T \oplus 0^*$, if $B = T(v) \oplus 0^*$.

CLAIM. *When Player 1 has a budget $T(v)$ (alternatively, $T(v) \oplus 0^*$), Player 2 can bid $b_v^T \ominus 0^*$ (b_v^T respectively).*

Proof. We analyze the case when $B = T(v)$, as the other case is exactly similar. So, when Player 1's budget is $B = T(v)$, Player 2's budget is $k^* \ominus T(v)$. In order to establish that Player 2

can indeed bid $b \ominus 0^*$, we show the following:

$$(k^* \ominus T(v)) \ominus (b_v^T \ominus 0^*) \geq 0$$

We prove this by a case analysis akin to the proof of Lemma 3.15, i.e, we analyze four cases, each of which corresponds to a parity of $|T(v^+)| + |T(v^-)|$ and an advantage status of $T(v^-)$.

— $|T(v^+)| + |T(v^-)|$ is even and $T(v^-) \in \mathbb{N}$.

In this case,

$$\begin{aligned} (k^* \ominus T(v)) \ominus (b_v^T \ominus 0^*) &= \left(k^* \ominus \frac{|T(v^+)| + |T(v^-)|}{2} \right) \ominus \left(\frac{|T(v^+)| - |T(v^-)|}{2} - 1 \right)^* \\ &= \left(k - \frac{|T(v^+)| + |T(v^-)|}{2} \right)^* \ominus \left(\frac{|T(v^+)| - |T(v^-)|}{2} - 1 \right)^* \\ &= k - \frac{|T(v^+)| + |T(v^-)|}{2} - \frac{|T(v^+)| - |T(v^-)|}{2} + 1 \\ &= (k + 1) - |T(v^+)| \geq 0 \end{aligned}$$

— $|T(v^+)| + |T(v^-)|$ is odd and $T(v^-) \in \mathbb{N}^* \setminus \mathbb{N}$.

In this case, we have

$$\begin{aligned} (k^* \ominus T(v)) \ominus (b_v^T \ominus 0^*) &= \left(k^* \ominus \left(\left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor + 1 \right) \right) \ominus \left(\left\lfloor \frac{|T(v^+)| - |T(v^-)|}{2} \right\rfloor - 1 \right)^* \\ &= k - \left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor - 1 - \left\lfloor \frac{|T(v^+)| - |T(v^-)|}{2} \right\rfloor + 1 \\ &= k + 1 - \left(\frac{|T(v^+)| + |T(v^-)|}{2} + \frac{|T(v^+)| - |T(v^-)|}{2} \right) \\ &= (k + 1) - |T(v^+)| \geq 0 \end{aligned}$$

— $|T(v^+)| + |T(v^-)|$ is even and $T(v^-) \in \mathbb{N}^* \setminus \mathbb{N}$.

In this case, it goes as following:

$$\begin{aligned} (k^* \ominus T(v)) \ominus (b_v^T \ominus 0^*) &= \left(k^* \ominus \left(\frac{|T(v^+)| + |T(v^-)|}{2} \right)^* \right) \ominus \left(\frac{|T(v^+)| - |T(v^-)|}{2} - 1 \right) \\ &= (k + 1) - T(v^+) \geq 0 \end{aligned}$$

— Finally, $|T(v^+)| + |T(v^-)|$ is odd and $T(v^-) \in \mathbb{N}$.

Here, we have

$$\begin{aligned} (k^* \ominus T(v)) \ominus (b_v^T \ominus 0^*) &= \left(k^* \ominus \left(\left\lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \right\rfloor \right)^* \right) \ominus \left\lfloor \frac{|T(v^+)| - |T(v^-)|}{2} \right\rfloor \\ &= k - \frac{|T(v^+)| + |T(v^-)|}{2} + \frac{1}{2} - \frac{|T(v^+)| - |T(v^-)|}{2} + \frac{1}{2} \\ &= (k + 1) - T(v^+) \geq 0 \end{aligned}$$

Therefore, we conclude that when Player 1 has a budget of $T(v)$, Player 2 does have the enough budget to bid $b_v^T \ominus 0^*$. ◆

She then proceeds to a neighbour v' with T -value $T(v')$. Let $c' = \langle v', B' \rangle$ denote the resulting configuration. Intuitively, Player 2 pays less than she should for winning the bidding. Formally, we will show that $B' < T(v')$. This will conclude the proof. Indeed, since we assume $T \equiv \text{Th}_{\mathcal{G}}$, Player 2 has a winning strategy from c' , which she uses to counter f from c' onwards. We distinguish between two cases depending on whether Player 1 holds the advantage:

- (a) Player 1 holds the advantage, i.e., $B \in \mathbb{N}^* \setminus \mathbb{N}$. By Observation 3.13, he uses it according to f_T , thus $b \in \mathbb{N}^* \setminus \mathbb{N}$. Player 2 bids $b_2 = b \ominus 0^*$. First, note that the bid is legal. Indeed, since b contains the advantage, b_2 does not. Second, note that Player 2 wins the bidding. Indeed, if Player 1 bids less than b_2 , clearly Player 2 wins, and if he bids b_2 , then a tie occurs, and since he has the advantage and does not use it, Player 2 wins the bidding. As a result, Player 1's budget is updated to $B \oplus (b \ominus 0^*) < B \oplus b < B \oplus (b \oplus 0^*) = |T(v^+)|^*$, in particular, $B \oplus (b \ominus 0^*) < T(v^+)$.
- (b) Player 1 does not hold the advantage, i.e., $B \in \mathbb{N}$. Again, by Observation 3.13, b does not include the advantage and bidding $b_1 < b$ necessarily implies $b_1 < b \ominus 0^*$, simply because he does not hold the advantage. Player 2 bids $b \ominus 0^* \in \mathbb{N}^* \setminus \mathbb{N}$. It is not hard to see that she wins the bidding and showing that $B' < T(v')$ is done as in the previous case.

Case 2: f over bids; $b_1 > b$. We assume $B = T(v)$ and the case of $B = T(v) \oplus 0^*$ is similar. Note that Observation 3.13 implies that f_T proposes a bid of $b \oplus 0^*$ when Player 1's budget is $T(v) \oplus 0^*$. Intuitively, if Player 1 wins the bidding with his bid of b_1 , he will pay “too much”, and Player 2 indeed lets him win by bidding 0 (except for one case that we will explain later). The resulting configuration is $c' = \langle v_1, B \ominus b_1 \rangle$, and we will show that $B \ominus b_1 < T(v_1)$ (barring one case). As in the underbidding case, this concludes the proof: since we assume $T \equiv \text{Th}_{\mathcal{G}}$, Player 2 wins from c' .

We first consider the easier case when $b_1 > b \oplus 0^*$. Then, $B \ominus b_1 < B \ominus (b \oplus 0^*) \leq T(v^-) \leq T(v_1)$, thus $B \ominus b_1 < T(v_1)$, as required. We proceed to the harder case of $b_1 = b \oplus 0^*$. Note that this case necessarily implicates that Player 1 has the advantage, i.e., $B \in \mathbb{N}^* \setminus \mathbb{N}$. Indeed, otherwise when he does not have the advantage, i.e., $B \in \mathbb{N}$ then $b \in \mathbb{N}$ too (from Observation 3.13), thus he cannot bid $b \oplus 0^*$, which is in $\mathbb{N}^* \setminus \mathbb{N}$, from his budget B . Recall from Definition 3.6 that when Player 1's budget $B = T(v) \in \mathbb{N}^* \setminus \mathbb{N}$, there are two possibilities: (i) $|T(v^+)| + |T(v^-)|$ is odd and $T(v^-) \in \mathbb{N}$, and (ii) $|T(v^+)| + |T(v^-)|$ is even and $T(v^-) \in \mathbb{N}^* \setminus \mathbb{N}$. In Case (i), $T(v) - b = T(v^-)$, hence when Player 1 bids $b \oplus 0^*$, Player 2's response is 0, and Player 1's budget in the next configuration is strictly lower than the threshold.

We conclude with Case (ii). Recall that in this case $b = \lfloor \frac{|T(v^+)| - |T(v^-)|}{2} \rfloor \ominus 0^*$. This case requires a different approach since Player 1 can bid $b \oplus 0^*$ and even if he wins the bidding,

his budget in the next configuration does not fall below the threshold. We define g to follow \tilde{g} . Consider the move of \tilde{g} from $\langle v_1, c \rangle$. If it lets Player 1 win by proceeding to $\langle v_1, B \ominus b \rangle$, then g responds to f in \mathcal{G} by bidding 0. Recall that Player 2's other action in $G_{T,\mathcal{G}}$ corresponds to a bid of $b \oplus 0^*$, and is represented by proceeding to vertex $\langle v', B \ominus (b \oplus 0^*) \rangle$. Then, in \mathcal{G} , we define g to bid $b \oplus 0^*$, thus both players bid $b \oplus 0^*$ and Player 2 wins the tie since Player 1 has the advantage and does not use it. Player 2 proceeds to v' following \tilde{g} . The key idea is that in both cases, we reach the same configuration in \mathcal{G} and $G_{T,\mathcal{G}}$. That is, even though f disagrees with f_T , we extend the two plays π and $\tilde{\pi}$ and restart the proof. As discussed in the beginning of the proof, the plays cannot be infinite, thus eventually f disagrees with f_T in one of the other manners.

Case 3: f does not choose an allowed vertex; $b_1 = b$ and $v_1 \notin A$. Recall that, by definition, the set A of allowed vertices consists of all vertices v' that satisfy $T(v) \ominus b \geq T(v')$. Therefore, Player 2 responds to f by letting Player 1 win by bidding 0. In the resulting configuration, Player 1's budget is strictly less than T , which coincides with the threshold budget, and, as in the above, Player 2 proceeds with a winning strategy. ■

Finally, we verify that $T \leq \text{Th}_{\mathcal{G}}$. We define a function $T' : V \rightarrow [k] \cup \{k+1\}$ as follows. For $v \in V$, when $T(v) > 0$ we define $T'(v) = (k+1) \ominus T(v)$, and $T'(v) = k+1$ otherwise. Lemma 3.10 shows that T' satisfies the average property. We proceed as in the previous construction only from Player 2's perspective. We construct a partial strategy $f_{T'}$ for Player 2 from T' just as f_T is constructed from T , and construct a turn-based parity game $G_{T',\mathcal{G}}$. Let $\text{Th}_{\mathcal{G}}^2$ denote Player 2's threshold function in \mathcal{G} . That is, at a vertex $v \in V$, Player 2 wins when her budget is at least $\text{Th}_{\mathcal{G}}^2(v)$ and she loses when her budget is at most $\text{Th}_{\mathcal{G}}^2(v) \ominus 0^*$. Applying Lemmas 5.2 and 5.4 to Player 2, we obtain the following.

LEMMA 5.5. *If Player 2 wins from every vertex $\langle v, B \rangle$ in $G_{T',\mathcal{G}}$ with $B \in \mathbb{N}^*$ and $B < k+1$, then $T' \geq \text{Th}_{\mathcal{G}}^2$. If $T' \equiv \text{Th}_{\mathcal{G}}^2$, then Player 2 wins from every vertex $\langle v, B \rangle$ of $G_{T',\mathcal{G}}$ with $B \in \mathbb{N}^*$ and $B < k+1$.*

Given a frugal-parity discrete-bidding game $\mathcal{G} = \langle V, E, k, p, S, \text{fr} \rangle$, a vertex $v \in V$, and $\ell \in [k]$, we guess $T : V \rightarrow [k] \cup \{k+1\}$ and verify that it satisfies the average property. Note that the size of T is polynomial in \mathcal{G} since it consists of $|V|$ numbers each of size $O(\log k)$. We construct $G_{T,\mathcal{G}}$ and $G_{T',\mathcal{G}}$, guess memoryless strategies for Player 1 and Player 2, respectively, and verify in polynomial time that they are indeed winning. Finally, we check whether $T(v) \geq \ell$, and answer accordingly. Correctness follows from Lemmas 5.2, 5.4, and 5.5. We thus obtain our main result.

THEOREM 5.6. *The problem of finding threshold budgets in frugal-parity discrete-bidding games is in NP and coNP.*

6. Discussion

We develop two algorithms to find threshold budgets in discrete-bidding games. Our first algorithm shows, for the first time, that thresholds in parity discrete-bidding games satisfy the average property. Previously, only thresholds in reachability discrete-bidding games were known to have this property. We study, for the first time, the problem of computing threshold budgets in discrete-bidding games in which the budgets are given in binary, and establish membership in NP and coNP for reachability and parity objectives. Previous algorithms for reachability and parity discrete-bidding games have exponential running time in this setting. We develop novel building blocks as part of our algorithms, which can be of independent interest. First, we define and study, for the first *frugal* objectives, which are reachability objectives accompanied by an enforcement on a player's budget when reaching the target. Second, our fixed-point algorithm provides a recipe for extending a proof on the structure of thresholds in reachability bidding games to parity bidding games. Third, we develop, for the first time, strategies that can be implemented with linear memory in reachability and parity discrete-bidding games, whereas previous constructions used exponential memory.

We point to the intriguing state of affairs in parity discrete-bidding games. Deciding the winner in a turn-based parity game is a long-standing open problem, which is known to be in NP and coNP but not known to be in P. A very simple reduction from turn-based parity games to parity discrete-bidding games was shown in [1]. Moreover, the reduction outputs a bidding game with a total budget of 0; that is, a discrete bidding game with constant sum of budgets. Our results show that parity discrete-bidding games are in NP and coNP even when the sum of budgets is given in binary. One might expect that such games would be at least exponentially harder than bidding games with constant sum of budgets. But all of these classes of games actually lie in NP and coNP.

References

- [1] M. Aghajohari, G. Avni, and T. A. Henzinger. Determinacy in discrete-bidding infinite-duration games. *Log. Methods Comput. Sci.* 17(1), 2021. [DOI](#) (3, 4, 7, 11, 45)
- [2] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002. [DOI](#) (6)
- [3] O. Amir, E. Kamar, A. Kolobov, and B. J. Grosz. Interactive teaching strategies for agent training. *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 804–811. IJCAI/AAAI Press, 2016. [DOI](#) (4)
- [4] K.R. Apt and E. Grädel. *Lectures in Game Theory for Computer Scientists*. Cambridge University Press, 2011. [DOI](#) (2)
- [5] N. Atzei, M. Bartoletti, and T. Cimoli. A survey of attacks on ethereum smart contracts. *IACR Cryptology ePrint Archive*, 2016:1007, 2016. [DOI](#) (4)
- [6] G. Avni, E. Goharshady, T. A. Henzinger, and K. Mallik. Bidding games with charging. *Proceedings of the 35th International Conference on Concurrency Theory (CONCUR)*, 2024. [DOI](#) (5)
- [7] G. Avni and T. A. Henzinger. A survey of bidding games on graphs. *Proceedings of the 31st International Conference on Concurrency Theory (CONCUR)*, volume 171 of *LIPICs*, 2:1–2:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. (3)

- [8] G. Avni, T. A. Henzinger, and V. Chonev. Infinite-duration bidding games. *J. ACM*, 66(4):31:1–31:29, 2019. DOI (3, 12)
- [9] G. Avni, T. A. Henzinger, and R. Ibsen-Jensen. Infinite-duration poorman-bidding games. *Proceedings of the 14th Conference on Web and Internet Economics (WINE)*, volume 11316 of LNCS, pages 21–36. Springer, 2018. DOI (3)
- [10] G. Avni, T. A. Henzinger, and D. Zikelić. Bidding mechanisms in graph games. *J. Comput. Syst. Sci.* 119:133–144, 2021. (3)
- [11] G. Avni, I. Jecker, and Đ. Žikelić. Infinite-duration all-pay bidding games. *Proceedings of the 32nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 617–636, 2021. DOI (3)
- [12] G. Avni, K. Mallik, and S. Sadhukhan. Auction-based scheduling. *Proceedings of the 30th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 14572 of Lecture Notes in Computer Science, pages 153–172. Springer, 2024. DOI (4)
- [13] G. Avni, T. Meggendorfer, S. Sadhukhan, J. Tkadlec, and Đ. Zikelić. Reachability poorman discrete-bidding games. *Proceedings of the 26th European Conference on Artificial Intelligence (ECAI)*, volume 372 of *Frontiers in Artificial Intelligence and Applications*, pages 141–148. IOS Press, 2023. DOI (4)
- [14] G. Avni and S. Sadhukhan. Computing threshold budgets in discrete-bidding games. *42nd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2022, December 18–20, 2022, IIT Madras, Chennai, India*, volume 250 of *LIPIcs*, 30:1–30:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. DOI (1)
- [15] M. Babaioff, T. Ezra, and U. Feige. Fair-share allocations for agents with arbitrary entitlements. *Proceedings of the Twenty-Second ACM Conference on Economics and Computation (EC)*, page 127. ACM, 2021. DOI (4)
- [16] J. Bhatt and S. Payne. Bidding chess. *Math. Intelligencer*, 31:37–39, 2009. DOI (4)
- [17] B. Bordais, P. Bouyer, and S. Le Roux. From local to global determinacy in concurrent graph games. *Proceedings of the 41st Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 213 of *LIPIcs*, 41:1–41:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. DOI (11)
- [18] K. Chatterjee, A. K. Goharshady, and Y. Velner. Quantitative analysis of smart contracts. *Proceedings of the 27th European Symposium on Programming (ESOP)*, pages 739–767, 2018. DOI (4)
- [19] A. Condon. The complexity of stochastic games. *Inf. Comput.* 96(2):203–224, 1992. DOI (3, 12)
- [20] M. Develin and S. Payne. Discrete bidding games. *The Electronic Journal of Combinatorics*, 17(1):R85, 2010. DOI (3, 4, 7, 10, 11, 23)
- [21] H. Everett. *Recursive Games*. Princeton University Press, 1957. DOI (24)
- [22] K. Vrieze J. Filar. *Competitive Markov Decision Process*. Springer New York, 2012. DOI (24)
- [23] O. Kupferman and M. Y. Vardi. Weak alternating automata and tree automata emptiness. *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing (STOC)*, pages 224–233. ACM, 1998. DOI (5)
- [24] 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. DOI (2, 11, 12)
- [25] A. J. Lazarus, D. E. Loeb, J. G. Propp, and D. Ullman. Richman games. *Games of No Chance*, 29:439–449, 1996. DOI (2, 11, 12)
- [26] R. Meir, G. Kalai, and M. Tennenholtz. Bidding games and efficient allocations. *Games and Economic Behavior*, 112:166–193, 2018. DOI (4)
- [27] S. Muthukrishnan. Ad exchanges: research issues. *Proceedings of the Fifth Workshop on Internet and Network Economics (WINE)*, pages 1–12, 2009. DOI (4)
- [28] Y. Peres, O. Schramm, S. Sheffield, and D. B. Wilson. Tug-of-war and the infinity laplacian. *J. Amer. Math. Soc.* 22:167–210, 2009. DOI (3)
- [29] A. Pnueli and R. Rosner. On the synthesis of a reactive module. *Proc. 16th POPL*, pages 179–190, 1989. DOI (2)
- [30] R. Kant Rai, U. Larsson, and N. Patel. Discrete richman-bidding scoring games. *Int. J. Game Theory*, 50(3):695–728, 2021. DOI (4)
- [31] W. Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theor. Comput. Sci.* 200(1-2):135–183, 1998. DOI (5)