Computing Threshold Budgets in Discrete-Bidding Games

Guy Avni ✉
University of Haifa, Israel

Suman Sadhukhan ✉
University of Haifa, Israel

Abstract

In a two-player zero-sum graph game, the players move a token throughout the 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. We distinguish between *continuous*- and *discrete*-bidding games. In the latter, the granularity of the players' bids is restricted, e.g., bids must be given in cents. Continuous-bidding games are well understood, however, from a practical standpoint, discrete-bidding games are more appealing.

In this paper we focus on discrete-bidding games. We study the problem of finding *threshold budgets*; namely, a necessary and sufficient initial budget for winning the game. Previously, the properties of threshold budgets were only studied for reachability games. For parity discrete-bidding games, thresholds were known to exist, but their structure was not understood. We describe two algorithms for finding threshold budgets in parity discrete-bidding games. The first algorithm is a fixed-point algorithm, and it reveals the structure of the threshold budgets in these games. Second, we show that the problem of finding threshold budgets is in NP and coNP for parity discrete-bidding games. Previously, only exponential-time algorithms where known for reachability and parity objectives. A corollary of this proof is a construction of strategies that use polynomial-size memory.

2012 ACM Subject Classification Theory of computation → Formal languages and automata theory; Theory of computation → Solution concepts in game theory

Keywords and phrases Discrete bidding games, Richman games, parity games, reachability games

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2022.30

Related Version *Full Verrsion*: `https://arxiv.org/abs/2210.02773`

Acknowledgements This research was supported in part by ISF grant no. 1679/21.

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* [21], in which the goal is to synthesize a reactive system that satisfies a given specification no matter how the environment behaves.

Two orthogonal classifications of graphs 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.

42nd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2022).
Editors: Anuj Dawar and Venkatesan Guruswami; Article No. 30; pp. 30:1–30:18
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

We study *bidding graph games* [17, 16], which use 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*, which is a necessary and sufficient initial budget for winning the game.

Continuous-bidding games. Bidding games were largely studied under continuous bidding. We briefly survey the relevant literature. Bidding games were introduced in [17, 16]. 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). It was shown that in continuous-bidding games, a target is necessarily reached, thus double-reachability games essentially coincide with reachability games under continuous-bidding.

Threshold budgets were shown to exist; namely, each vertex v has a value $\mathtt{Th}(v)$ such that if Player 1's budget is strictly greater than $\mathtt{Th}(v)$, he wins the game from v, and if his budget is strictly less than $\mathtt{Th}(v)$, Player 2 wins the game. Moreover, it was shown that the threshold function $\mathtt{Th}$ is a *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, $\mathtt{Th}$ assigns a value in $[0, 1]$ to each vertex such that at the "end points", we have $\mathtt{Th}(t_1) = 0$ and $\mathtt{Th}(t_2) = 1$, and the threshold at every other vertex is the average of two of its neighbors. Uniqueness implies that the problem of finding threshold budgets[1] is in NP and coNP. Moreover, an intriguing equivalence was observed between reachability continuous-bidding games and a class of stochastic game [13] called *random-turn games* [20]. Intricate equivalences between *mean-payoff* continuous-bidding games and random-turn games have been shown in [6, 7, 8, 9] (see also [5]).

Parity continuous-bidding games were studied in [6]. The following key property was identified. Consider a strongly-connected parity continuous-bidding game $\mathcal{G}$. 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. *Discrete-bidding games* are similar to continuous-bidding games only that we fix the sum of the budgets to be $k \in \mathbb{N}$ and bids are restricted to be integers. Ties in biddings need to be handled explicitly. We focus on the tie-breaking mechanism that was defined in [14]: 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 considered in [1].

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

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 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 would be useful in practice.

Bidding games model ongoing and stateful auctions. Such auctions arise in various domains. An immediate example is auctions for online advertisements [19]. As another example, in *blockchain* technology, *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 [12, 4]. Bidding games have been applied as a mechanism for fair allocation of resources [18]. 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 [11]. In all of these applications, the granularity of the bids is restricted.

Previous results. Threshold budgets and their properties were previously only studied for reachability discrete-bidding games [14]. In these games, discrete versions of the properties of continuous-bidding games were observed; namely, discrete versions of threshold budgets exist, they satisfy a discrete version of the average property, and, similar to continuous bidding, winning strategies are constructed in which a player's bid is derived from the threshold budgets. A value-iteration algorithm was developed to compute the thresholds with a worst-case exponential running time, when k is given in binary.

Parity discrete-bidding games were shown to be determined in [1], but the structure of the threshold budgets was not understood. In particular, threshold budgets were not known to have the average property. The previously known algorithm to solve parity discrete-bidding games is naive. Construct an arena based on the exponentially-sized "configuration graph" that corresponds to a bidding game. Determinacy implies that it suffices to solve a turn-based game rather than a concurrent game on this arena. The essence of bidding games is completely lost in this construction; namely, the structure given by the threshold budgets is not used and the bids made by winning strategies have no connection with the thresholds. To make things worse, unlike reachability games, the properties of threshold budgets in parity discrete- and continuous-bidding games are known to differ significantly; namely, there are strongly-connected games in which the maximal parity index is odd and Player 1 loses with any initial budget, and no reduction is known to reachability discrete-bidding games.

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 is based on repeated calls to an algorithm to solve reachability discrete-bidding games in combination with a recursion over the parity indices, similar in spirit to Zielonka [22] and Kupferman and Vardi's [15] algorithms to solve turn-based parity games. An important corollary from the algorithm is that threshold budgets in parity discrete-bidding games satisfy the average property. The worst-case running time of the algorithm is exponential.

Second, we show that the problem of finding threshold budgets in parity discrete-bidding gamesis in NP and coNP. The bound follows to reachability discrete-bidding games for which only an exponential-time algorithm was known. We briefly describe the idea of our proof. We first show that, interestingly, unlike continuous-bidding games, functions that satisfy the discrete average property are not unique, but by definition the threshold budgets are. Thus, one cannot simply guess a function and verify that it satisfies the average property. We overcome this challenge as follows. Given a guess of function T, we first verify that it satisfies the average property. Then, we construct a partial bidding strategy f_T for Player 1 based on T, and construct a turn-based parity game in which a Player 1 strategy corresponds to a bidding strategy f' that agrees with f_T and a Player 2 strategy corresponds to a response to f'. We show that Player 1 wins the turn-based game iff T coincides with the threshold budgets. As a corollary, we show the existence of winning strategies that use memory that is polynomial in the size of the arena. Previously, only strategies that use exponential-sized memory were known.

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. We focus on reachability and parity objectives, which we define later in this section.

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

A (finite) *history* is a sequence $\langle v_0, a_0^1, a_0^2 \rangle, \ldots, \langle v_{n-1}, a_{n-1}^1, a_{n-1}^2 \rangle, v_n \in (V \times A \times A)^* \cdot V$ such that, for each $0 \leq i < n$, we have $v_{i+1} = \delta(v_i, a_i^1, a_i^2)$. A *strategy* is a function from histories to actions, thus it is of the form $\sigma : (V \times A \times A)^* \cdot V \to A$. We restrict attention to *legal* strategies; namely, strategies that for each history $\pi \in (V \times A \times A)^* \cdot V$ that ends in $v \in V$, choose an action in $\lambda(v, i)$, for $i \in \{1, 2\}$. A *memoryless* strategy is a strategy that, for every vertex v, assigns the same action to every history that ends in v.

Two strategies σ_1 and σ_2 for the two players and an initial vertex v_0, give rise to a unique *play*, denoted $\text{play}(v_0, \sigma_1, \sigma_2)$, which is a sequence in $(V \times A \times A)^\omega$ and is defined inductively as follows. The first element of $\text{play}(v_0, \sigma_1, \sigma_2)$ is v_0. Suppose that the prefix of length $j \geq 1$ of $\text{play}(v_0, \sigma_1, \sigma_2)$ is defined to be $\pi^j \cdot v_j$, where $\pi^j \in (V \times A \times A)^*$. Then, at turn j, for $i \in \{1, 2\}$, Player i takes action $a_i^j = \sigma_i(\pi^j \cdot v_j)$, the next vertex is $v^{j+1} = \delta(v_j, a_1^j, a_2^j)$, and we define $\pi^{j+1} = \pi^j \cdot \langle v_j, a_1^j, a_2^j \rangle \cdot v_{j+1}$. The *path* that corresponds to $\text{play}(v_0, \sigma_1, \sigma_2)$ is $v_0, v_1, \ldots$.

Turn-based games are a special case of concurrent games in which the vertices are partitioned between the two players. When the token is placed on a vertex v, the player who *controls* v decides how to move the token. Formally, a vertex v is controlled by Player 1 if for every action $a_1 \in A$, there is a vertex v' such that no matter Player 2 takes which action $a_2 \in A$, we have $v' = \delta(v, a_1, a_2)$. The definition is dual for Player 2. 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 $\mathcal{G} = \langle V, E, k \rangle$, 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 use $N(v)$ to denote its *neighbors*, namely $N(v) = \{u : E(v, u)\}$. The size of the arena 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 stays 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 the tie-breaking mechanism that was defined in [14]: exactly one of the players holds the *advantage* at every 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.

Following [14], we denote the advantage with $*$. Let $\mathbb{N}$ denote the non-negative integers and $\mathbb{N}^*$ denote the set $\{0, 0^*, 1, 1^*, 2, 2^*, \ldots\}$. Throughout the paper, we use k to denote the sum of budgets, and use $[k]$ to denote the set $\{0, 0^*, \ldots, k, k^*\}$. Intuitively, when saying that Player 1 has a budget of $m^* \in [k]$, we mean that Player 1 can choose a bid in $\{0, \ldots, m\}$ and that he has the advantage. Implicitly, we mean that Player 2's budget is $k - m$ and she does not have the advantage.

We define an order $<$ on $\mathbb{N}^*$ by $0 < 0^* < 1 < 1^* < \ldots$. Let $m \in \mathbb{N}^*$. We denote by $|m|$ the integer part of m, i.e., $|m^*| = m$. We define operators $\oplus$ and $\ominus$ over $\mathbb{N}^*$: $x^* \oplus y = (x+y)^*$, $x \oplus y = x + y$ for $x, y \in \mathbb{N}$. Intuitively, we use $\oplus$ to keep track of the budget of a player when they loses a biding to the other player, and as a result their budget gets increased by the other player's bid: therefore, in general $x^* \oplus y^*$ is not well-defined. The definition of $\ominus$ is *almost* dual, we use the notation to keep track of the budget of a player when they wins a bidding: $x \ominus y = x - y$, $x^* \ominus y = (x - y)^*$, and in particular $x^* \ominus y^* = x - y$. For $B \in \mathbb{N}^*$, of exceptional use is $B \oplus 0^*$ and $B \ominus 0^*$, which respectively denote the *successor* and *predecessor* of B in $\mathbb{N}^*$ according to $<$.

Consider an arena $\mathcal{G} = \langle V, E, k \rangle$ of a bidding game. We describe a concurrent game $\mathcal{C}$ that corresponds to it. The vertices in $\mathcal{C}$ consist of *configuration* vertices $C = V \times [k]$ and *intermediate* vertices $\{i_{c,b} : c \in C, b \leq k^*\}$. A vertex $c = \langle v, B \rangle \in (V \times [k])$ represents the configuration of the bidding game in which the token is placed on vertex $v \in V$, Player 1's budget is B, and Player 2's budget is $k^* \ominus B$. Consider a configuration vertex $c = \langle v, B \rangle$. The available actions for a player in c represent the legal bids and the vertex to move to upon winning. Thus, the available actions for Player 1 are $\{0, \ldots, |B|\} \times N(v)$ and the available actions for Player 2 are $\{0, \ldots, k - |B|\} \times N(v)$. Each intermediate vertex is owned by a single player, so we only specify their outgoing transitions below. We describe the transition function. Suppose that the token is placed on a configuration vertex $c = \langle v, B \rangle$ 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. We address the case of a bidding tie, namely the case that $b_1 = b_2 = b$. Assume that Player 1 has the advantage, i.e., $c = \langle v, B_1^* \rangle$, and the other case is dual. The game proceeds to an intermediate vertex $i_{c,b}$ that is controlled by Player 1 and has two outgoing edges. The first edge models Player 1 using the advantage to win the bidding, and directs to the configuration $\langle u_1, B_1 - b_1 \rangle$. The second edge models Player 1 allowing Player 2 to win the bidding and keeping the advantage, and it directs to $\langle u_2, (B_1 + b_2)^* \rangle$. We often say that the player who holds the advantage bids b^*, and we mean that he bids b and uses the advantage if a tie occurs. Note that the size of the arena is $O(|V| \times k)$, which is exponential in the size of the bidding game.

Consider two strategies f and g in $\mathcal{C}$ and an initial configuration $c = \langle v, B \rangle$ We often abuse notation and refer to $\tau = \text{play}(v, f, g)$ as the infinite path in $\mathcal{G}$ that is obtained by removing intermediate vertices from π. We use $inf(\tau)$ to denote the set of vertices that τ visits infinitely often.

2.3 Objectives

An objectives in a bidding game specifies the infinite paths that are winning for Player 1. We consider the following two canonical objectives:

- **Reachability** A game is equipped with a target set $T \subseteq V$. Player 1, the reachability player, wins an infinite play iff it visits T.
- **Parity** Each vertex is labeled by a *parity index*, given by a function $p : V \to \{1, \ldots, d\}$, for $d \in \mathbb{N}$. A play τ is winning for Player 1 iff $\max_{v \in inf(\tau)} p(v)$ is odd.

In addition, we introduce the following extension of the two objectives above.

▶ **Definition 1** (Frugal objectives). *Consider an arena $\langle V, E, k \rangle$. A* frugal objective *is a set $S \subseteq V$ of sink states and a function $\mathit{fr} : S \to [k]$ which assigns a* frugal-target budget *to each sink. Player 1's frugal objective is satisfied if the game reaches some $s \in S$ with Player 1's budget at least $\mathit{fr}(s)$. We consider both* frugal-reachability *and* frugal-parity *games. Player 1 wins a frugal-reachability game if the frugal objective is satisfied. Note that a reachability game is a special case of a frugal-reachability game in which $\mathit{fr} = 0$. A frugal-parity game is $\langle V, E, k, p, S, \mathit{fr} \rangle$, where $p : (V \setminus S) \to \{0, \ldots, d\}$. Player 1 wins a play π if (1) π does not reach S and satisfies the parity objective, or (2) π reaches S and satisfies the frugal objective.*

A *winning strategy* from a configuration $c = \langle v, B \rangle$ for Player 1 is a strategy f such that no matter which strategy g Player 2 chooses, $\text{play}(c, f, g)$ is winning for Player 1. We say that Player 1 *wins* from c if he has a winning strategy. The definition is dual for Player 2.

3 Threshold Budgets and the Average Property

A key quantity in bidding games is the *threshold budget* at a vertex, which intuitively represents 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** (Threshold budgets). *Consider a bidding game $\mathcal{G}$ in which the sum of budgets is $k \in \mathbb{N}$. The* threshold budget *at a vertex v in $\mathcal{G}$, denoted $\mathit{Th}_{\mathcal{G}}(v)$, is such that if Player 1's budget at v is at least $\mathit{Th}_{\mathcal{G}}(v) \in [k]$, then Player 1 wins the game from v, and if his budget is at most $\mathit{Th}_{\mathcal{G}}(v) \ominus 0^*$, then Player 2 wins the game. We refer to the function $\mathit{Th}_{\mathcal{G}}$ as the* threshold budgets.

3.1 Reachability continuous-bidding games

The properties of threshold budgets in discrete-bidding games have only been studied for reachability games [14]. The properties of the threshold budgets and the techniques to prove them are similar to those used in reachability continuous-bidding games [17, 16]. We thus first survey the latter.

▶ **Definition 3** (Continuous threshold budgets). *Normalize the sum of budgets to 1. The* continuous threshold budget *at a vertex v is a budget $\mathit{Th}(v) \in [0, 1]$ such that if Player 1's budget exceeds $\mathit{Th}(v)$, he wins the game from v, and if Player 2's budget exceeds $1 - \mathit{Th}(v)$, she wins the game from v.*

We call the objective that was considered in [17, 16] *double-reachability*. Such a 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. The game ends once one of the targets is reached, and the player whose target is reached is the winner. Every other vertex has a path to both targets. Even though a priori, it is not implicit that one of the player always wins in a double reachability games (because the game might not reach either of the target vertices), here it is explicitly shown that this is not the case [17, 16].

▶ **Definition 4** (Continuous average property). *Consider a double-reachability continuous-bidding game $\mathcal{G} = \langle V, E, t_1, t_2 \rangle$ and a function $T : V \to [0, 1]$. We say that T has the continuous average property if $T(t_1) = 0$ and $T(t_2) = 1$, and for every other $v \in V \setminus \{t_1, t_2\}$, we have $T(v) = 0.5 \cdot \big(T(v^-) + T(v^+)\big)$, where $v^+ := \arg\max_{u \in N(v)} \mathtt{Th}(u)$ and $v^- := \arg\min_{u \in N(v)} \mathtt{Th}(v)$.*

We show the main properties of double-reachability continuous-bidding games. The ideas used in the proof are adapted to the discrete setting in later sections.

▶ **Theorem 5** ([17, 16]). *Consider a double-reachability continuous-bidding game $\langle V, E, t_1, t_2 \rangle$. Continuous threshold budgets exist, and the threshold budgets $\mathtt{Th} : V \to [0, 1]$ is the unique function that has the continuous average property.*

Proof Sketch. Let $\mathtt{Th}$ be a function that satisfies the continuous average property. We prove that $\mathtt{Th}(v)$, for every vertex v, is the continuous threshold budget at v. Uniqueness follows immediately. We omit the proof of existence of $\mathtt{Th}$.

Suppose that Player 1's budget at v is $\mathtt{Th}(v) + \varepsilon$, for $\varepsilon > 0$. We describe a Player 1 winning strategy. Player 1 maintains the following invariant:

Invariant. When the token is on $u \in V$, Player 1's budget is strictly greater than $\mathtt{Th}(u)$.

The invariant implies that Player 1 does not lose; indeed, it implies that if t_2 is reached, Player 1's budget is strictly greater than 1, which violates the assumption that the sum of budgets is 1.

The invariant holds initially, and we show how to maintain it. Suppose that the token is placed on $v \in V$ and Player 1's budget is $B = \mathtt{Th}(v) + \varepsilon$. Recall that $v^+, v^- \in N(v)$ are respectively the neighbors of v with the maximal and minimal threshold budgets. Let $b = 0.5 \cdot \big(\mathtt{Th}(v^+) - \mathtt{Th}(v^-)\big)$. The key observation is that $B + b = \mathtt{Th}(v^+) + \varepsilon$ and $B - b = \mathtt{Th}(v^-) + \varepsilon$. We claim that by bidding b, Player 1 guarantees that the invariant is maintained. Indeed, if he wins the bidding, he moves the token to v^-, and if he loses the bidding, the worst that Player 2 can do is move the token to v^+.

We omit the details of how Player 1 guarantees winning, i.e., forcing the token to t_1.

Finally, we show that Player 2 wins when Player 1's budget is $\mathtt{Th}(v) - \varepsilon$. We intuitively "flip" the game and associate Player 1 with Player 2. Let $\mathcal{G}'$ be the same as $\mathcal{G}$ only 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 $\mathtt{Th}'$ as $\mathtt{Th}'(u) = 1 - \mathtt{Th}(u)$. A key observation is that $\mathtt{Th}'$ satisfies the average property in $\mathcal{G}'$. Now, in order to win from v in $\mathcal{G}$ when Player 1's budget is $\mathtt{Th}(v) - \varepsilon$, Player 2 follows a winning Player 1 strategy in $\mathcal{G}'$ with an initial budget of $1 - \mathtt{Th}(v) + \varepsilon$. ◀

3.2 Constructing strategies based on the discrete average property

We first adapt the definition of the average property (Def. 4) to the discrete setting.

▶ **Definition 6** (Average property). *Consider a discrete-bidding game* $\mathcal{G} = \langle V, E, k, S, \mathit{fr}\rangle$ *with frugal objective. We say that a function* $T : V \to [k] \cup \{k+1\}$ *has the* average property *if for every* $s \in S$*, we have* $T(s) = \mathit{fr}(s)$*, and for every* $v \in V \setminus S$*,*

$$T(v) = \lfloor \frac{|T(v^+)| + |T(v^-)|}{2} \rfloor + \varepsilon,$$

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

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

Note that, the range of T includes $k+1$, which captures the threshold budget of a player at a vertex from where the other player wins the game with budget 0. Consider a function $T : V \to [k] \cup \{k+1\}$ that satisfies the average property. We develop a *partial strategy* f_T, which is a function from histories to $[k] \times 2^V$. An output $\langle b, A\rangle$ of f_T means that the bid is b and upon winning, Player 1 must choose an *allowed vertex* in A to move the token to. A strategy f' that *agrees* with f_T bids in the same manner and upon winning a bidding, it chooses a vertex in A.

We define f_T so that when Player 1 plays according to a strategy that agrees with f_T, an invariant is maintained on Player 1's budget, similar to the invariant in the continuous setting (see the proof of Thm. 5). Suppose that a history ends in a vertex v with Player 1's budget $B \geq T(v)$. Intuitively, as in Thm. 5, when Player 1 wins the bidding in v, he proceeds to some neighbor v^- that attains the minimal value according to T, and when Player 2 wins the bidding, the worst she can do is move to a vertex v^+ that attains the maximal value according to T. Formally, f_T restricts the choice of neighbor of v to a set of *allowed vertices*, denoted $A(v)$, and depend only v. Let $v^+ = \arg\max_{u \in N(v)} T(u)$ and $v^- = \arg\min_{u \in N(v)} T(u)$. We define $A(v) = \{u \in N(v) : T(u) = T(v^-)\}$ when $T(v^-) \in \mathbb{N}$, and $A(v) = \{u \in N(v) : T(u) \leq T(v^-) \oplus 0^*\}$ when $T(v^-) \in \mathbb{N}^* \setminus \mathbb{N}$.

Next, for each $v \in V$, we define a bid that f_T proposes. Suppose that Player 1's budget is $B \in [k]$, for $B \geq T(v)$. We define f_T so that it proposes one of two possible bids b_v^T or $b_v^T \oplus 0^*$, depending on whether Player 1 holds the advantage. Intuitively, Player 1 "attempts" to bid b_v^T at v. This is not possible if $b_v^T \in \mathbb{N}^* \setminus \mathbb{N}$ and $B \in \mathbb{N}$, i.e., Player 1 wants to use the advantage but does not have it. In such a case, Player 1 bids $b_v^T \oplus 0^* \in \mathbb{N}$. Formally, we define

$$b_v^T = \begin{cases} \frac{|T(v^+)| - |T(v^-)|}{2} & \text{When } |T(v^+)| + |T(v^-)| \text{ is even and } T(v^-) \in \mathbb{N} \\ \lfloor \frac{|T(v^+)| - |T(v^-)|}{2} \rfloor & \text{When } |T(v^+)| + |T(v^-)| \text{ is odd and } T(v^-) \in \mathbb{N}^* \setminus \mathbb{N} \\ \frac{|T(v^+)| - |T(v^-)|}{2} \ominus 0^* & \text{When } |T(v^+)| + |T(v^-)| \text{ is even and } T(v^-) \in \mathbb{N}^* \setminus \mathbb{N} \\ \lfloor \frac{|T(v^+)| - |T(v^-)|}{2} \rfloor \oplus 0^* & \text{When } |T(v^+)| + |T(v^-)| \text{ is odd and } T(v^-) \in \mathbb{N} \end{cases} \tag{1}$$

Assuming that the game reaches v with a budget of $B \geq T(v)$, we define f_T to bid $b^T(v, B)$ where $b^T(v, B) = b_v^T$ when both b_v^T and B belong to either $\mathbb{N}$ of $\mathbb{N}^* \setminus \mathbb{N}$, and $b_v^T \oplus 0^*$ otherwise. We formalize the guarantees of f_T in the lemma below, whose proof can be found in the full version[10].

▶ **Lemma 7.** *For every* $v \in V$ *and* $B \geq T(v)$*, we have* $B \ominus b^T(v, B) \geq T(v^-)$ *and* $B \oplus b^T(v, B) \oplus 0^* \geq T(v^+)$*.*

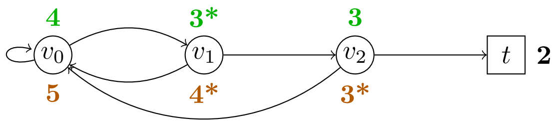

Figure 1 A discrete-bidding reachability game with two functions that satisfy the average property.

A corollary of Lem. 7 is that any strategy that agrees with f_T maintains an invariant on Player 1's budget and is thus a legal strategy, i.e., it never prescribes bids that exceed the available budget.

▶ **Corollary 8.** *Suppose that Player* 1 *plays according to a strategy that agrees with* f_T *starting from configuration* $\langle u, B\rangle$ *having* $B \geq T(u)$*, we have:*
- *When the game reaches* $v \in V$*, Player* 1*'s budget is at least* $T(v)$*.*
- *The bid* b *prescribed by* f_T *does not exceed the available budget, i.e.,* $b \leq B$*.*

3.3 Properties of functions with the average property

We show that, somewhat surprisingly, unlike in continuous-bidding, functions that satisfy the discrete average property are not unique. That is, there are functions that satisfy the average property but do not coincide with the threshold budgets.

▶ **Theorem 9.** *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 $T : V \to [k]$ 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, whose proof can be found in the full version [10], is key in developing a winning Player 2 strategy. We intuitively show that the "complement" of T satisfies the average property. The idea is similar to the continuous case (see the last point in the proof of Thm. 5).

▶ **Lemma 10.** *Let* $\mathcal{G} = \langle V, E, k, S, \mathit{fr}\rangle$ *be a discrete-bidding game with a frugal objective. Let* $T : V \to [k] \cup \{k + 1\}$ *be a function that satisfies the average property. We define* $T' : V \to [k] \cup \{k + 1\}$ *as follows. For* $v \in V$*, let* $T'(v) = k^* \ominus \big(T(v) \ominus 0^*\big)$ *when* $T(v) > 0$*, and* $T'(v) = k + 1$ *otherwise.*

Then, T' *satisfies the average property.*

3.4 Frugal-reachability discrete-bidding games

We close this section by extending the results of [14] from reachability to frugal-reachability discrete-bidding games.

▶ **Lemma 11.** *Consider a frugal-reachability discrete-bidding game* $\mathcal{G} = \langle V, E, k, S, \mathit{fr}\rangle$*. If* $T : V \to [k] \cup \{k + 1\}$ *is a function that satisfies the average property, then* $T(v) \leq \mathit{Th}_{\mathcal{G}}(v)$ *for every* $v \in V$*.*

Proof. We show that if for some vertex v, Player 1 has a budget less than $T(v)$, then Player 2 has a winning strategy, which proves that the threshold budgets for Player 1 cannot be less than $T(v)$, when T is a average property satisfying function.

Given such T that satisfies the average property, we construct T' as in Lem. 10. 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 $T'(s) = k^* \ominus (\texttt{fr}(s) \ominus 0^*)$. We "flip" the game; namely, we associate Player 2 with Player 1, and construct a partial strategy $f_{T'}$ for Player 2 as in Sec. 3.2. 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 Corollary 8, 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. ◀

▶ **Lemma 12.** *Consider a frugal-reachability discrete-bidding game $\mathcal{G} = \langle V, E, k, S, \texttt{fr} \rangle$. There is a function T that satisfies the average property with $T(v) \geq \texttt{Th}_\mathcal{G}(v)$, for every $v \in V$.*

Proof. The proof is similar to the one in [14]. We illustrate the main ideas. For $n \in \mathbb{N}$, we consider the *truncated game* $\mathcal{G}[n]$, which is the same as $\mathcal{G}$ only 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 $u \in V$, since Player 1 cannot win from u in 0 steps, we have $T_0(u) = k + 1$. For $s \in S$, we have $T_0(s) = \texttt{fr}(s)$. Clearly, $T_0 = \texttt{Th}_{\mathcal{G}[0]}$. For the inductive step, suppose that T_{n-1} is computed. For each vertex v, we define $T_n(v) = \lfloor \frac{|T_{n-1}(v^+)| + |T_{n-1}(v^-,k)|}{2} \rfloor + \varepsilon$ as in Def. 6. Following a similar argument to Thm. 5, it can be shown that if Player 1's budget is $T_n(v)$, he can bid 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 $\texttt{Th}_{\mathcal{G}[n]}(v) = T_n(v)$, for every $v \in V$. For every vertex v, let $T(v) = \lim_{n \to \infty} T_n(v)$. It is not hard to show that T satisfies the average property and that $T(v) \geq \texttt{Th}_\mathcal{G}(v)$, for every $v \in V$. ◀

Let T be a function that results from the fixed-point computation from the proof of Lem. 12. Since it satisfied the average property, we apply Lem. 11 to show that Player 2 wins from v when Player 1's budget is $T(v) \ominus 0^*$. We thus conclude the following.

▶ **Theorem 13.** *Consider a frugal-reachability discrete-bidding game $\mathcal{G} = \langle V, E, k, S, \texttt{fr} \rangle$. Threshold budgets exist and satisfy the average property. Namely, there exists a function $T : V \to [k] \cup \{k+1\}$ such that for every vertex $v \in V$*

- *if Player 1's budget is $B \geq T(v)$, then Player 1 wins the game, and*
- *if Player 1's budget is $B < T(v)$, then Player 2 wins the game*

Moreover, there is an exponential-time algorithm for finding such a T.

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. As a corollary, we show, for the first time, that threshold budgets in parity discrete-bidding games satisfy the average property.

For the remainder of this section, fix a frugal-parity game $\mathcal{G} = \langle V, E, k, p, S, \texttt{fr} \rangle$. Denote the maximal parity index by $d \in \mathbb{N}$ and let $F_d = \{v : p(v) = d\}$. For a bidding game $\mathcal{G}$, instead of $\texttt{Th}_\mathcal{G}$, we sometimes use $\texttt{FrRe-Th}_\mathcal{G}$ and $\texttt{FrPa-Th}_\mathcal{G}$ to highlight that $\mathcal{G}$ is respectively a frugal-reachability and frugal-parity game.

A description of the algorithm

The algorithm recurses over the parity indices. The base case is when only one parity index is used. Then, $\mathcal{G}$ is a frugal-reachability game and we use the algorithm in Thm. 13.

For the induction step, suppose that $d > 1$. We describe the key idea. For ease of presentation, we assume that the maximal parity index 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. Since d is even, in order for Player 1 to win, it is necessary (but not sufficient) to visit F_d only finitely often.

We iteratively define and solve a sequence of frugal-parity games $\mathcal{G}_0, \mathcal{G}_1, \ldots$. For $i \geq 0$, the arena of $\mathcal{G}_i$ is obtained from $\mathcal{G}$ by setting F_d to be sinks. The games differ in the frugal target in the "new" sinks F_d. We set the frugal target budgets in $\mathcal{G}_i$ so that, for every vertex v that is not a sink, $\texttt{Th}_{\mathcal{G}_i}(v)$ is a necessary and sufficient initial budget for winning in $\mathcal{G}$ by visiting F_d at most i times. Since $\mathcal{G}_i$ has only $d-1$ parity indices, we solve it recursively.

The definition of the frugal target budgets in $\mathcal{G}_0$ is immediate: since no visits to F_d are allowed, simply set the frugal target budget to be $k+1$ in these vertices. Since the sum of budgets is k, a play that ends in F_d in $\mathcal{G}_0$ is necessarily losing for Player 1, thus in order to win, he must satisfy the parity or frugal objective without visiting F_d.

In order to define the frugal target budgets in $\mathcal{G}_1$, we first construct a frugal-reachability game from $\mathcal{G}$, which we denote $\mathcal{R}_0$. The sinks in $\mathcal{R}_0$ are $V \setminus F_d$. The frugal target budget at $u \in (V \setminus F_d)$ is defined to be $\texttt{fr}(u) = \texttt{Th}_{\mathcal{G}_0}(u)$. Thus, when the game starts at $v \in F_d$ with a Player 1 budget of $\texttt{Th}_{\mathcal{R}_0}(v)$, Player 1 can guarantee that the game eventually reaches $V \setminus F_d$ with a budget that suffices for winning in $\mathcal{G}$ without visiting F_d again.

We define the frugal target budgets in $\mathcal{G}_1$. For each $v \in F_d$ we define $\texttt{fr}(v) = \texttt{Th}_{\mathcal{R}_0}(v)$. Then, when $\mathcal{G}$ starts from $u \in (V \setminus F_d)$ with a Player 1 budget of $\texttt{Th}_{\mathcal{G}_1}(u)$, Player 1 plays as follows. He first follows a winning strategy in $\mathcal{G}_1$ to ensure either (1) the parity condition is satisfied, (2) an "old" sink $s \in S$ is reached with budget at least $\texttt{fr}(s)$, or (3) the game reaches $v \in F_d$ with a budget of at least $\texttt{Th}_{\mathcal{R}_0}(v)$. Cases (1)-(2) are winning in $\mathcal{G}$. In Case (3), Player 1 plays as described above to guarantee winning in $\mathcal{G}$ without visiting F_d again. The construction of $\mathcal{R}_1, \mathcal{R}_2, \ldots$ and $\mathcal{G}_2, \mathcal{G}_3, \ldots$ follows the same idea.

Formally, for $i \geq 0$, we define $\mathcal{G}_i = \langle V \setminus F_d, E', p', S \cup F_d, \texttt{fr}_i \rangle$, where E' is obtained from E by removing outgoing edges from vertices in F_d, i.e., $E'(v,u)$ iff $E(v,u)$ and $v \notin F_d$, the parity function p' coincides with p but is not defined over F_d, and $\texttt{fr}_i$ is defined below. Note that p' assigns at most $d-1$ parity indices. The function $\texttt{fr}_i$ coincides with $\texttt{fr}$ on the "old" sinks; namely, $\texttt{fr}_i(s) = \texttt{fr}(s)$, for every $s \in S$.

For the "new" sinks F_d, for $i \geq 0$, the definition $\texttt{fr}_i$ is inductive. Let $v \in F_d$. We define $\texttt{fr}_0(v) = k+1$, meaning that Player 1 is not allowed to visit F_d at all in $\mathcal{G}_0$. For $i \geq 1$, assume that $\texttt{fr}_i$ has been defined and we define $\texttt{fr}_{i+1}$ as follows. Since $\mathcal{G}_i$ has less parity indices than $\mathcal{G}$, we can recursively run the algorithm to obtain $\texttt{FrPa-Th}_{\mathcal{G}_i}(u)$, for each $u \in V \setminus F_d$. As we prove formally below, a budget of $\texttt{FrPa-Th}_{\mathcal{G}_i}(v)$ suffices for Player 1 to win $\mathcal{G}$ while visiting F_d only i times. For $i \geq 0$, we construct a frugal-reachability game $\mathcal{R}_i = \langle V, E'', V \setminus F_d, \texttt{FrPa-Th}_{\mathcal{G}_i} \rangle$, where E'' is obtained from E by removing outgoing edges from every vertex in $(V \setminus F_d)$. That is, in $\mathcal{R}_i$, the only vertices that are *not* sinks are the vertices in F_d, and in order to win in a sink $u \in (V \setminus F_d)$, Player 1's budget should be at least $\texttt{FrPa-Th}_{\mathcal{G}_i}(u)$. We can now define the the frugal target $\texttt{fr}_{i+1}$ of $\mathcal{G}_{i+1}$: for each $v \in F_d$, define $\texttt{fr}_{i+1}(v) = \texttt{FrRe-Th}_{\mathcal{R}_i}(v)$.

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

Algorithm 1 Frugal-Parity-Threshold($\mathcal{G}$).

---

if $\mathcal{G}$ uses one parity index **then**
  Return Frugal-Reachability-Threshold($\mathcal{G}$)
$\texttt{fr}_0(v) = k+1$, for $v \in F_d$
for $i = 0, 1, \ldots$ **do**
  $\texttt{FrPa-Th}_{\mathcal{G}_i} \leftarrow$ Frugal-Parity-Threshold($\mathcal{G}_i$)
  $\texttt{FrRe-Th}_{\mathcal{R}_i} \leftarrow$ Frugal-Reachability-Threshold($\mathcal{R}_i$)
  For each $v \in F_d$, define $\texttt{fr}_{i+1}(v) = \texttt{FrRe-Th}_{\mathcal{R}_i}(v)$
  if $\texttt{fr}_i(v) = \texttt{fr}_{i+1}(v)$, for all $v \in F_d$ **then**
    Define $\texttt{FrPa-Th}_{\mathcal{G}}(v) = \texttt{fr}_i(v)$ for $v \in F_d$
    Define $\texttt{FrPa-Th}_{\mathcal{G}}(u) = \texttt{FrRe-Th}_{\mathcal{G}_i}(u)$ for $u \in V \setminus F_d$.
    Return $\texttt{FrPa-Th}_{\mathcal{G}}$

Correctness

The intuition for the proof of the following lemma is described above, and we formally prove this below. The formal proof can be found in the full version [10].

▶ **Lemma 14.** *Let d be the maximal parity index in $\mathcal{G}$. Define $j = 1$ when d is even and $j = 2$ when d is odd. For $i \geq 0$, let $\texttt{FrPa-Th}_{\mathcal{G}_i}$ be the threshold budget of Player j in the game $\mathcal{G}_i$. Then, for $v \in V \setminus F_d$, a budget of $\texttt{FrPa-Th}_{\mathcal{G}_i}(v)$ suffices for Player j to win $\mathcal{G}$ while visiting F_d at most i times.*

It follows from the following lemma, whose proof can be found in the full version [10], that Alg. 1 reaches a fixed point and terminates.

▶ **Lemma 15.** *For every $v \in F_d$ and $i \geq 0$, we have $\texttt{fr}_i(v) \geq \texttt{fr}_{i+1}(v)$.*

Next, we show that the budgets returned by Alg. 1 are necessary for Player 1 to win.

▶ **Lemma 16.** *Consider an output T of Algorithm 1 and let $u \in V \setminus F_d$. If Player 1's budget is $T(u) \ominus 0^*$, then Player 2 wins $\mathcal{G}$ from v.*

Proof. The proof is by induction on the number of parity indices in $\mathcal{G}$. When there is one parity index, $\mathcal{G}$ is a frugal-reachability game and the proof follows from Thm. 13. For the induction step, let $i \geq 0$ be the index at which the algorithm reaches a fixed point. That is, for every $u \in V \setminus F_d$, we have $T(u) = \texttt{FrPa-Th}_{\mathcal{G}_i}(u)$. Since $\mathcal{G}_i$ is a game with less parity indices than $\mathcal{G}$, by the induction hypothesis, a budget of $\texttt{FrPa-Th}_{\mathcal{G}_i}(u)$ is necessary for Player 1 to win from u in $\mathcal{G}_i$. That is, if Player 1's budget is less than $\texttt{FrPa-Th}_{\mathcal{G}_i}(u)$, then Player 2 has a strategy that guarantees that an infinite play satisfies Player 2's parity objective, and a finite play that ends in $v \in S \cup F_d$ violates Player 1's frugal objective $\texttt{fr}_i(v)$.

We construct a winning Player 2 strategy in $\mathcal{G}$. Player 2 initially follows a winning strategy in $\mathcal{G}_i$. Assume that Player 1 plays according to some strategy and let π be the resulting play. If π is infinite or ends in S, then π is a play in $\mathcal{G}$ and is thus winning for Player 2 in both games. Suppose that π ends in $v \in F_d$. Player 2's strategy guarantees that Player 1's budget at v is at most $\texttt{fr}_i(v) \ominus 0^* = \texttt{FrRe-Th}_{\mathcal{R}_i}(v) \ominus 0^*$. Suppose that Player 2 follows a winning strategy from v in $\mathcal{R}_i$, Player 1 follows some strategy, and let π' be the resulting play. Since Player 2's strategy is winning, there are two cases. First, π' remains in F_d, thus π' is a play in $\mathcal{G}$ that is winning for Player 2 since d is even. Second, π' ends in a vertex $u' \in V \setminus F_d$ is reached with Player 1's budget at most $\texttt{FrPa-Th}_{\mathcal{G}_{i-1}}(u') \ominus 0^*$. Note that since the algorithm terminates at a fixed point, we have $\texttt{FrPa-Th}_{\mathcal{G}_{i-1}}(u') = \texttt{FrPa-Th}_{\mathcal{G}_i}(u')$.

Player 2 switches to a winning strategy in $\mathcal{G}_i$ and repeats. Note that an infinite play in which Player 2 alternates infinitely often between a winning strategy in $\mathcal{G}_i$ and a winning strategy in $\mathcal{R}_i$ necessarily visits F_d infinitely often and is thus winning for Player 2 in $\mathcal{G}$. ◀

We conclude with the following theorem. Proving that threshold budgets satisfy the average property is done by induction on the parity indices.

▶ **Theorem 17.** *Given a frugal-parity discrete-bidding game $\mathcal{G}$, Alg. 1 terminates and returns $\mathtt{Th}_{\mathcal{G}}$. Moreover, $\mathtt{Th}_{\mathcal{G}}$ satisfies the average property.*

5 Finding threshold budgets is in NP and coNP

We consider the following decision problem.

▶ **Definition 18** (Finding threshold budgets). *Given a bidding game $\mathcal{G} = \langle V, E, k\rangle$, a vertex $v \in V$, and $\ell \in [k]$, decide whether $\mathtt{Th}_{\mathcal{G}}(v) \geq \ell$.*

Consider a frugal-parity discrete-bidding game $\mathcal{G}$ with vertices V and consider a function $T : V \to [k] \cup \{k + 1\}$. We describe a polynomial-time algorithm to check whether $T = \mathtt{Th}_{\mathcal{G}}$. Given such an algorithm, it is not hard to show that the problem of finding threshold budgets is in NP and coNP. Indeed, given $\mathcal{G}$, a vertex v, and $\ell \in [k]$, guess T and verify, using the algorithm, that $T = \mathtt{Th}_{\mathcal{G}}$. Then, check whether $T(v) \geq \ell$, and answer accordingly.

Our algorithm is based on a reduction to turn-based games. We first verify that T satisfies the average property, and if it does not, we reject, following Thm. 17. Next, we construct the partial strategy f_T based on T as in Sec. 3.2. Recall that given a history that ends in v, the function f_T prescribes a bid and a set $A(v) \subseteq V$ of allowed vertices for Player 1 to choose from upon winning the bidding. A strategy f' agrees with f_T if it bids in the same manner and always chooses only allowed vertices. In order to verify whether T is a correct guess, i.e., $T = \mathtt{Th}_{\mathcal{G}}$, we construct a parity turn-based game $G_{T,\mathcal{G}}$ in which a Player 1 strategy corresponds to a strategy f' that agrees with f in $\mathcal{G}$ and a Player 2 strategy corresponds to a response to f' in $\mathcal{G}$. We show that the procedure is sound and complete; namely, if Player 1 wins in $G_{T,\mathcal{G}}$, he wins in $\mathcal{G}$ (hence $T \geq \mathtt{Th}_{\mathcal{G}}$), and if $T \geq \mathtt{Th}_{\mathcal{G}}$, then Player 1 wins in $G_{T,\mathcal{G}}$. Finally, in order to verify that $T \leq \mathtt{Th}_{\mathcal{G}}$, we apply the same procedure from Player 2's perspective.

5.1 From bidding games to turn-based games

Consider a function T that satisfies the average property, and let f_T be the partial strategy as constructed in Sec. 3.2. We construct a parity turn-based game $G_{T,\mathcal{G}}$ such that if Player 1 wins in every vertex in $G_{T,\mathcal{G}}$, then $T \geq \mathtt{Th}_{\mathcal{G}}$.

Intuitively, $G_{T,\mathcal{G}}$ simulates $\mathcal{G}$. In each turn, Player 1's bid is determined according to f_T. Suppose that Player 1 bids b. Player 2's actions in $G_{T,\mathcal{G}}$ represent responses to b. Namely, Player 2 can choose between winning the bidding by bidding $b \oplus 0^*$ and in addition choosing the successor vertex in $\mathcal{G}$, or losing the bidding by bidding 0 and letting Player 1 decide how the game proceeds.

A key challenge is keeping the size of $G_{T,\mathcal{G}}$ polynomial in the size of $\mathcal{G}$. Consider a history that ends in a configuration $\langle v, B\rangle$, for $B \geq T(v)$. Recall that f_T can prescribe one of two possible bids: the bid that f_T prescribes at $\langle v, B\rangle$ coincides either with the bid that it prescribes in $\langle v, T(v)\rangle$ or $\langle v, T(v) \oplus 0^*\rangle$, depending on which of the two agrees with B on which player has the advantage. We obtain a polynomial-sized arena by representing every configuration $\langle v, B\rangle$, for $B > T(v) \oplus 0^*$, with a vertex $\langle v, \top\rangle$.

See an example of the construction in the App. A.

Formally, we define $G_{T,\mathcal{G}} = \langle V_1, V_2, E, \gamma\rangle$, where for $i \in \{1,2\}$, the vertices V_i are controlled by Player i, $E \subseteq (V_1 \cup V_2) \times (V_1 \cup V_2)$ is a collection of edges, and $\gamma : V_1 \cup V_2 \to \{1, 2, , \ldots d\}$ assigns parity indices to the vertices. The vertices of $G_{T,\mathcal{G}}$ are $V_2 = \{\langle v, T(v)\rangle, \langle v, T(v) \oplus 0^*\rangle, \langle v, \top\rangle : v \in (V \cup S)\}$ and $V_1 = \{w^1 : w = \langle v, B\rangle \in V_2, B \neq \top\}$. We define the edges in $G_{T,\mathcal{G}}$. The sinks in $G_{T,\mathcal{G}}$ are of the form $\langle s, B\rangle$, for $s \in S$, or $\langle v, \top\rangle$, for $v \in V$. Sinks have self loops. We describe the other edges. Let $w = \langle v, B\rangle \in V_2$, where $B \neq \top$ and $v \notin S$. Intuitively, reaching w in $G_{T,\mathcal{G}}$ means that $\mathcal{G}$ is in configuration $\langle v, B_1\rangle$, where $B_1 \geq B$ and they agree on which player has the advantage. Thus, f_T bids $b^T(v, B)$. Outgoing edges from w model Player 2's two options. First, Player 2 can choose to win the bidding by bidding $b^T(v, B) \oplus 0^*$ (any higher bid is wasteful on her part) and moving the token to $u \in N(v)$. The next configuration is intuitively $\langle u, B'\rangle$, where $B' = B \oplus b^T(v, B) \oplus 0^*$. If $T(u) \leq B' \leq T(u) \oplus 0^*$, then $\langle u, B'\rangle \in V_2$ and we define $E(\langle v, B\rangle, \langle u, B'\rangle)$. Otherwise, we truncate Player 1 budget by defining $E(\langle v, B\rangle, \langle u, \top\rangle)$. We disallow transitions in which Player 2's bid exceeds her available budget, i.e., when $b^T(v, B) \oplus 0^* > k^* \ominus B$. Second, Player 2 can choose to move to w^1 modeling Player 2 allowing Player 1 to win the bidding, e.g., by bidding 0. In this case, Player 1's budget is updated to $B' = B \ominus b^T(v, B)$ and he moves the token to some vertex in $A(u)$. Thus, the neighbors of w^1 are $\langle u, B'\rangle$, for $u \in A(v)$. By the proof of Corollary 8, $B' \in \{T(u), T(u) \oplus 0^*\}$, thus $\langle u, B'\rangle \in V_2$. Note that all paths in $G_{T,\mathcal{G}}$ are infinite. Finally, we define the parity indices. A non-sink vertex in $G_{T,\mathcal{G}}$ "inherits" its parity index from the vertex in $\mathcal{G}$; namely, for $w = \langle v, B\rangle \in V_2$, we define $\gamma(w) = \gamma(w^1) = p(v)$. We define γ so that Player 1 wins in sinks, thus we set the parity index of a sink to be odd.

5.2 Correctness

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

▶ **Lemma 19.** *If Player 1 wins from every vertex in $G_{T,\mathcal{G}}$, then $T \geq \mathtt{Th}_{\mathcal{G}}$.*

Proof. We describe the main ideas of the proof and the details can be found in the full version [10]. Assume that Player 1 wins from every vertex in $G_{T,\mathcal{G}}$ and fix some memoryless winning strategy f'. We describe a winning strategy f^* of Player 1 in $\mathcal{G}$, which he can play when he has a budget of at least $T(v)$ at vertex v for all vertices v of $\mathcal{G}$. This proves that $T \geq \mathtt{Th}_{\mathcal{G}}$.

Suppose that $\mathcal{G}$ starts from $c_0 = \langle v, T(v)\rangle$. We initiate $G_{T,\mathcal{G}}$ from $v_0 = \langle v, T(v)\rangle$. Suppose that Player 2 plays according to a strategy g^* in $\mathcal{G}$. Player 1 simulates a Player 2 strategy g in $G_{T,\mathcal{G}}$ so that when $\mathcal{G}$ reaches a configuration $c = \langle u, B_1\rangle$, the vertex in $G_{T,\mathcal{G}}$ is $w = \langle u, B\rangle$, where $B \in \{T(u), T(u) \oplus 0^*\}$ agrees with B_1 on the advantage. We describe how we simulate g^* with g, and how f^* simulates f'. Suppose that $\mathcal{G}$ is in configuration $c = \langle u, B_1\rangle$ and $G_{T,\mathcal{G}}$ is in $w = \langle u, B\rangle$. We define f^* to agree with f_T and bid $b_T(u, B)$. If g^* loses the bidding, then we define g to proceed to w^1 and we define f^* to match the move of f' from w^1. If g^* wins the bidding in $\mathcal{G}$, we define g to win the bidding and move the token as g^* does. In the full version [10], we show that the correspondence between the games is maintained.

Let π be the play in $G_{T,\mathcal{G}}$ that results from f' and g and π^* the play in $\mathcal{G}$ that results from f^* and g^*. Since f' is winning, π satisfies Player 1's objective. We distinguish between three cases. In the first case, π does not reach a sink in $G_{T,\mathcal{G}}$. Then, the play π^* matches π up to repetitions. Indeed, by removing occurrences of Player 1 vertices in $G_{T,\mathcal{G}}$ from π and projecting both plays on V, we obtain the same path. Recall that a Player 1 vertex

w^1 has the same parity index of its predecessor $w \in V_2$. It follows that since π satisfies Player 1's parity objective, so does π^*. In the second case, π reaches a sink $\langle s, T(s)\rangle$, where $s \in S$. Then π^* reaches a configuration $\langle s, B\rangle$. Since T satisfies the average property, we have $B \geq \mathtt{fr}(s)$, and Player 1 wins $\mathcal{G}$. In the final case, π reaches a sink $\langle v, \top\rangle$. Let $\langle u, B_1\rangle$ be the corresponding configuration in $\mathcal{G}$. Note that B_1 is strictly greater than $T(u) \oplus 0^*$. Let $B' < B_1$ that agrees with B_1 on the advantage and $\langle u, B'\rangle \in V_2$. Player 1 intuitively adds $B_1 - B'$ to his "spare change" account and plays as if his budget is B' by restarting $G_{T,\mathcal{G}}$ from $\langle u, B'\rangle$. Since the sum of budgets is fixed and restarting f^* in this manner strictly increases his spare change, it follows that this last case can only occur finitely often. Thus, eventually either of the first two cases must occur implying that f^* is winning in $\mathcal{G}$. ◀

▶ **Corollary 20.** *In the proof of Lem. 19, we construct a winning Player 1 strategy f^*. Note that f^* only keeps 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 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 21.** *If $T = \mathtt{Th}_{\mathcal{G}}$, then Player 1 wins from every vertex in $G_{T,\mathcal{G}}$.*

Proof. We describe the idea of the proof and the details can be found in the full version [10]. Assume towards contradiction that $T \equiv \mathtt{FrPa\text{-}Th}$ and there is $\langle v, B\rangle \in V_2$ that is losing for Player 1. Since $T(v) \leq B$, Player 1 has a winning strategy f^* in $\mathcal{G}$ starting from $\langle v, B\rangle$. Let g be a memoryless winning strategy for Player 2 in $G_{T,\mathcal{G}}$ starting from $\langle v, B\rangle$. Based on g, we construct a Player 2 strategy g^* in $\mathcal{G}$ and show that it is winning against f^*, which contradicts our assumption that f^* is a Player 1's winning strategy. Note that since f^* is fixed, when g^* selects a bid, it is in response to the bid chosen by f^*. Let us consider an arbitrary Player 1 strategy f in $G_{T,\mathcal{G}}$, which by our earlier assumption is losing for Player 1 from vertex $\langle v, B\rangle$ of V_2 (because all Player 1 strategies are losing from there). Intuitively, we construct g^* such that it follows g as long as f^* follows f. By doing so, Player 2 maintains a similar correspondence between the play in $\mathcal{G}$ and $G_{T,\mathcal{G}}$ as in the above: when $\mathcal{G}$ is in configuration $\langle v, B_1\rangle$, then $G_{T,\mathcal{G}}$ is in vertex $\langle v, B_1\rangle$. Since the two plays traverse the same vertices in V and g is winning, if f^* always agrees with f, the resulting play will be winning for Player 2. Thus, f^* must not agree with f at some point. Either he bids differently from $b^T(v, B)$, or, upon winning he chooses a vertex u which is not in the allowed set of vertices. First assume that he bids b at $\langle u, B\rangle$. If $b > b(u, B)$, then g^* bids 0, intuitively causing Player 1 to pay too much for a bidding win. If $b < b^T(u, B)$, g^* bids $b^T(u, B)$ intuitively buying a win cheaply. In both cases, we show that $\mathcal{G}$ reaches a configuration $c' = \langle u', B'\rangle$ for $B' < T(u)$. Since we assume $T \equiv \mathtt{Th}_{\mathcal{G}}$, Player 2 has a winning strategy from c', which she uses to win $\mathcal{G}$. Second, we show that when he chooses some vertex $u' \notin A(v)$ following a winning bid of $b^T(v, B)$, the game reaches a configuration $\langle u', B'\rangle$, such that $B' < T(u')$, and again Player 2 uses a winning strategy to win $\mathcal{G}$. ◀

Finally, we verify that $T \leq \mathtt{Th}_{\mathcal{G}}$. We define a function $T' : V \to [k] \cup \{k+1\}$ as follows. For $v \in V$, when $T(v) > 0$ we define $T(v) = k^* \ominus (T(v) \ominus 0^*)$, and $T'(v) = k+1$ otherwise. Lem. 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 $\mathtt{Th}^2_{\mathcal{G}}$

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 $\mathtt{Th}^2_{\mathcal{G}}(v)$ and she loses when her budget is at most $\mathtt{Th}^2_{\mathcal{G}}(v) \ominus 0^*$. Existence of $\mathtt{Th}^2_{\mathcal{G}}$ follows from Thm. 17. Applying Lemmas 19 and 21 to Player 2, we obtain the following.

▶ **Lemma 22.** *If Player 2 wins from every vertex in $G_{T',\mathcal{G}}$, then $T' \geq \mathtt{Th}^2_{\mathcal{G}}$. If $T' \equiv \mathtt{Th}^2_{\mathcal{G}}$, then Player 2 wins from every vertex of $G_{T',\mathcal{G}}$.*

Given a frugal-parity discrete-bidding game $\mathcal{G} = \langle V, E, k, p, S, \mathtt{fr} \rangle$, a vertex $v \in V$, and $\ell \in [k]$, we guess $T : V \to [k] \cup \{k+1\}$ 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 winning strategies for Player 1 and Player 2, respectively. We check whether $T(v) \geq \ell$, and answer accordingly. Correctness follows from Lemmas 19, 21, and 22. We thus obtain our main result.

▶ **Theorem 23.** *The problem of finding threshold budgets in frugal-parity discrete-bidding games is in* NP *and* coNP*.*

6 Discussion

We study, for the first time, the problem of computing threshold budgets in discrete-bidding games in which the budgets are given in binary. Previous algorithms for reachability and parity discrete-bidding games have exponential running time in this setting. We developed two algorithms for finding threshold budgets. The algorithms are complementary, and mirror the situation in the continuous setting; there too, there are two proof techniques to show results for threshold budgets, a fixed-point technique and an NP algorithm that relies on knowledge of the structure of threshold budgets. Prior to this work, a fixed-point algorithm was only known for reachability discrete-bidding games [14]. While our fixed-point algorithm for parity discrete-bidding games has exponential worst case running time, it sheds light on the structure of threshold budgets in these games. A structure that was crucial for our NP and coNP membership proof. This latter proof adds to the previously observed good news on discrete-bidding games: parity discrete-bidding games are not only a sub-class of concurrent games that is determined [1], we show that it is a sub-class of concurrent games that are represented in an exponentially-succinct manner and can still be solved in NP and coNP.

We leave open the exact complexity of finding threshold budgets. For the lower bound, it was shown in [1] that turn-based parity games reduce to parity discrete-bidding games with constant sum of budgets. Since solving turn-based parity games is a long-standing open problem, we expect that it will be challenging to find a polynomial-time algorithm for solving parity discrete-bidding games. Still, improved upper bounds might be possible to obtain. For example, a quasi-polynomial time algorithm for parity discrete-bidding games or a polynomial-time algorithm for reachability or Büchi discrete-bidding games.

References

1. M. Aghajohari, G. Avni, and T. A. Henzinger. Determinacy in discrete-bidding infinite-duration games. In *Proc. 30th CONCUR*, volume 140 of *LIPIcs*, pages 20:1–20:17, 2019.
2. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002.
3. O. Amir, E. Kamar, A. Kolobov, and B. J. Grosz. Interactive teaching strategies for agent training. In *Proc. 25th IJCAI*, pages 804–811. IJCAI/AAAI Press, 2016.
4. N. Atzei, M. Bartoletti, and T. Cimoli. A survey of attacks on ethereum smart contracts. *IACR Cryptology ePrint Archive*, 2016:1007, 2016.

5 G. Avni and T. A. Henzinger. A survey of bidding games on graphs. In *Proc. 31st CONCUR*, volume 171 of *LIPIcs*, pages 2:1–2:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
6 G. Avni, T. A. Henzinger, and V. Chonev. Infinite-duration bidding games. *J. ACM*, 66(4):31:1–31:29, 2019.
7 G. Avni, T. A. Henzinger, and R. Ibsen-Jensen. Infinite-duration poorman-bidding games. In *Proc. 14th WINE*, volume 11316 of *LNCS*, pages 21–36. Springer, 2018.
8 G. Avni, T. A. Henzinger, and Đ. Žikelić. Bidding mechanisms in graph games. In *In Proc. 44th MFCS*, volume 138 of *LIPIcs*, pages 11:1–11:13, 2019.
9 G. Avni, I. Jecker, and Đ. Žikelić. Infinite-duration all-pay bidding games. In *Proc. 32nd SODA*, pages 617–636, 2021.
10 G. Avni and S. Sadhukhan. Computing threshold budgets in discrete-bidding games. *CoRR*, abs/2210.02773, 2022. arXiv:2210.02773.
11 J. Bhatt and S. Payne. Bidding chess. *Math. Intelligencer*, 31:37–39, 2009.
12 K. Chatterjee, A. K. Goharshady, and Y. Velner. Quantitative analysis of smart contracts. In *Proc. 27th ESOP*, pages 739–767, 2018.
13 A. Condon. The complexity of stochastic games. *Inf. Comput.*, 96(2):203–224, 1992.
14 M. Develin and S. Payne. Discrete bidding games. *The Electronic Journal of Combinatorics*, 17(1):R85, 2010.
15 O. Kupferman and M. Y. Vardi. Weak alternating automata and tree automata emptiness. In *Proc. 30th STOC*, pages 224–233. ACM, 1998.
16 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.
17 A. J. Lazarus, D. E. Loeb, J. G. Propp, and D. Ullman. Richman games. *Games of No Chance*, 29:439–449, 1996.
18 R. Meir, G. Kalai, and M. Tennenholtz. Bidding games and efficient allocations. *Games and Economic Behavior*, 112:166–193, 2018. doi:10.1016/j.geb.2018.08.005.
19 S. Muthukrishnan. Ad exchanges: Research issues. In *Proc. 5th WINE*, pages 1–12, 2009.
20 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.
21 A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. 16th POPL*, pages 179–190, 1989.
22 Wieslaw Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theor. Comput. Sci.*, 200(1-2):135–183, 1998. doi:10.1016/S0304-3975(98)00009-7.

A Example of the construction of $G_{T,\mathcal{G}}$

▶ **Example 1.** We apply the construction to the frugal-reachability discrete-bidding game $\mathcal{G}_1$ depicted in Fig. 1. We use the two functions $T_1, T_2 : V \to V \to [k] \cup \{k+1\}$ that are specified in the figure, both of which satisfy the average property and correspond to the vectors $T_1 = \langle 4, 3^*, 3, 2\rangle$ and $T_2 = \langle 5, 4^*, 3^*, 2\rangle$.

The turn-based game $G_{T_1,\mathcal{G}_1}$ is depicted in Fig. 2a and $G_{T_2,\mathcal{G}_1}$ in Fig. 2b. In these two games, the allowed actions given by f_{T_1} and f_{T_2} from each vertex are singletons. Thus, Player 1 has no choice of successor vertex when winning a bidding, and so we omit Player 1 vertices from the figure. That is, all vertices are controlled by Player 2. Since the games are reachability games, we omit the parity indices from the vertices. Player 1's goal in both games is to reach a sink. We label each edge with the bids of the two players that it represents. 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. There are exceptions like $\langle v_1, 5\rangle$ in $G_{T_2,\mathcal{G}_1}$ where $b_1 = 1$ and Player 2 cannot bid $b_1 \oplus 0^* = 1^*$ since it exceeds her available budget when $k = 5$.

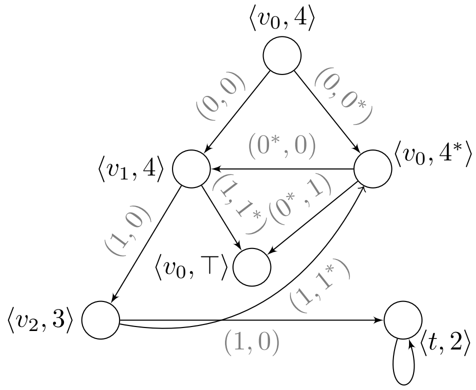

(a) Corresponding to T_1.

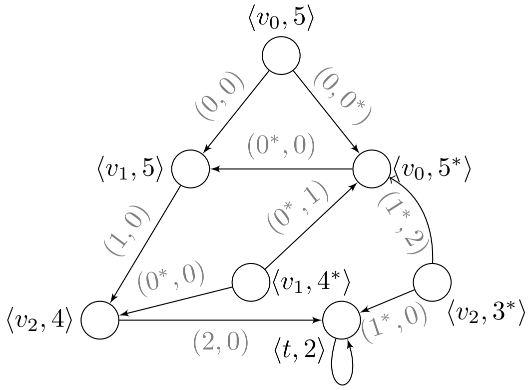

(b) Corresponding to T_2.

Figure 2 Turn-based games corresponding to G_1 and two different guesses.

Note that in $G_{T_1,\mathcal{G}_1}$ 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, which means that $T_2 = \mathtt{Th}_{\mathcal{G}_1}$. ◀