Hierarchical Cost-Parity Games*

Laura Bozzelli[a], Aniello Murano[a], Giuseppe Perelli[b], Loredana Sorrentino[a]

[a] *University of Naples Federico II*
[b] *University of Oxford*

Abstract

Cost-parity games are a fundamental tool in system design for the analysis of reactive and distributed systems that recently have received a lot of attention from the formal methods research community. They allow to reason about the time delay on the requests granted by systems, with a bounded consumption of resources, in their executions.

In this paper, we contribute to research on cost-parity games by combining them with *hierarchical systems*, a successful method for the succinct representation of models. We show that determining the winner of a *Hierarchical Cost-parity Game* is PSPACE-complete, thus matching the complexity of the proper special case of *Hierarchical Parity Games*. This shows that reasoning about temporal delay can be addressed at a free cost in terms of complexity.

Keywords: Formal methods, Parity games, Cost parity games, Hierarchical systems

1. Introduction

Parity games and Parity games with costs. In formal system design and verification [1, 2, 3, 4], *Parity Games* represent a fundamental machinery for the automatic synthesis and verification of concurrent and reactive systems [5, 6, 7, 8, 9]. The determinacy and the memorylessness of parity games is crucial in various theoretical areas useful in formal verification, among which we mention automata theory, temporal and modal logics, and monadic second-order logics. For instance, the emptiness problem of alternating tree automata [10] as well as model checking and

*This paper is an extended version of the results published in the Proceedings of the 24th International Symposium on Temporal Representation and Reasoning, TIME 2017, LIPIcs 90, pages 1–17, 2017.

Email addresses: lr.bozzelli@gmail.com (Laura Bozzelli), murano@na.infn.it (Aniello Murano), giuseppe.perelli@cs.ox.ac.uk (Giuseppe Perelli), loredana.sorrentino@unina.it (Loredana Sorrentino)

satisfiability in modal μ-calculus [11] can be reduced to deciding the winner of a parity game. In particular, model checking μ-calculus is equivalent via linear time reduction to this problem [12].

As pointed out in [13, 14, 15], the parity winning condition corresponds to a qualitative *request-response condition* [16]: Player 0 wins a play of infinite duration if all but finitely many odd colors (which we think of as requests) are followed by larger even colors (which we think of as responses). In this setting, there is no bound on the wait time, *i.e.*, the number of steps that elapse between a request and its first response in the play. On the other hand, in many applications, it is important to bound the wait time. In the last decade, many papers have focused on quantitative aspects, in particular boundedness requirements, of formal verification [17, 18, 19], including parity games [19, 13, 14, 15]. In [18], the authors introduce Prompt LTL, an extension of standard LTL [20] with the prompt-eventually operator $\mathtt{F}_p$: a finite system satisfies a Prompt LTL formula φ iff there is a bound on the wait time for all the prompt-eventually subformulas of φ in all the computations of the system. The automata-theoretic counterpart of the $\mathtt{F}_p$ operator has been investigated in [17]. Parity games extended with promptness requirements, the so-called finitary parity games, have been studied in [19]. The finitary parity condition [19] extends the parity condition by additionally requiring the existence (along the given play) of a bound k such that almost every odd color is answered within at most k steps. Surprisingly, finitary parity games are solvable in polynomial time, and thus simpler than parity games (according to the state-of-the-art). A meaningful generalization of finitary games is represented by the class of parity games with costs [13] (in the following, referred as cost-parity games). In such games, transitions are labeled by non-negative integers (costs). The cost of traversing a transition can be used to model resource consumption. The goal of Player 0 consists then in ensuring the underlying parity condition by using bounded resources: a play is winning for Player 0 if there is a bound k such that almost every odd color is followed by a larger even color that is reached with cost at most k. On the other hand, Player 1's goal is to exhaust the resources by making the cost unbounded. Note that Player 1's objective is not an ω-regular property, and in general, Player 1 needs infinite memory to win such games. However, cost-parity games enjoy some nice properties: Player 0 has memoryless winning strategies and determining the winner lies in NP $\cap$ coNP. This upper bound has been recently improved to UP $\cap$ coUP in [15], proving thus that the increased expressiveness with respect to parity conditions comes at a free cost in terms of complexity.

In the recent years, many other quantitative extensions of parity games have been introduced. Among them we would like to mention Mean-Payoff Parity Games [21], whose winning condition is a combination of a parity and a mean-payoff objective, and Energy Parity Games [22]. These last ones are played over weighted arenas, and the winning condition extends the parity condition by additionally requiring that the sum of the weights along a play (interpreted as level of energy, or resource usage) remains always positive.

Hierarchical verification. A well-known issue in formal verification is that the translation of a high-level description of a system into a formal model, typically given by a finite-state machine (FSM), often involves an exponential blow-up in the size of the FSM, thus affecting the efficiency of the analysis procedures both in theory and practice. Several sources of this blow-up have been identified in the literature. A well-studied one is the ability of components in the system to work in parallel and communicating with each other, possibly using variables. The impact of the concurrent setting on analysis problems is well-known: it costs an exponential, leading to the so called state-explosion problem. Another source of the blow-up in the translation of systems into FSMs is that in high-level sequential programming, one can specify components only once and then can reuse them in different contexts, leading to modularity and succinct system representation. A smart way to represent such modularity is by means of *hierarchical FSM*, where some of the states of the FSM are boxes (superstates) which correspond to nested FSMs (the reused components). The naive approach to model checking such systems is to 'flatten' them by repeatedly substituting references to sub-structures with copies of them. This results in a flat FSM whose size is exponential in the nesting depth of the hierarchical system. However, differently from the concurrent setting, a wiser approach avoiding flattening, for the case of model checking against temporal logics like LTL, CTL and the more expressive modal μ-calculus, is beneficial in terms of complexity [23, 24, 25, 5, 26]. Parity games have also been investigated under the hierarchical setting. In [5], Aminof et al. prove that deciding the winner in a *Hierarchical Parity Game* (*HPG*) is a PSPACE-complete problem. The technique used in [5] is based on the observation that even though a sub-arena may appear in different contexts, it is possible to extract information about the sub-arena that is independent of the context in which it appears.

Our contribution. In this paper, we further investigate the power of hierarchical representation by introducing and studying *Cost-parity Games over Hierarchical Systems* (*HCPG*). As main result, we establish that the problem of solving *HCPG* is PSPACE-complete, which matches the complexity of the proper special case of hierarchical parity games (*HPG*). The proposed approach for solving the considered problem generalizes in a non-trivial and sophisticated manner the one exploited in [5] for solving *HPG*, and is based on the notion of summary function for a memoryless strategy σ of Player 0 in a given sub-arena. Such a function records in a finite and efficient way the overall behavior of all the finite plays of σ leading to exit states of the sub-arena with respect to requests and responses, by finitely abstracting the set of associated costs and delays. The algorithm for solving *HCPG* then solves a sequence of flat cost-parity games obtained by replacing sub-arenas by simple gadgets (depending only on the set of colors and exit states of the sub-arena) that implement the summary functions.

The sequel of the paper is structured as follows. In Section 2, we first recall the framework of cost-parity games. Then, we introduce hierarchical cost-parity games and outline our solution approach. In Sections 3–7 we illustrate in detail the proposed approach for solving *HCPG* by providing full proofs of our results. Finally, we give few conclusions and future work directions in Section 8.

2. Preliminaries

Let $\mathbb{N}$ be the set of natural numbers. For all $i, j \in \mathbb{N}$, with $i \leq j$, $[i, j]$ denotes the set of natural numbers h such that $i \leq h \leq j$. We fix a non-empty finite set C of natural numbers of the form $[0, j]$ for some $j \in \mathbb{N}$, which represents the set of colors for the given cost-parity winning condition. We denote by C_e and C_o the sets of even and odd colors in C, respectively. We assume that the maximal color j in C, denoted by C_o^{max}, is odd.

For an alphabet Σ, and a non-empty finite or infinite word w over Σ, we denote by $|w|$ the length of w (we set $|w| = \infty$ if w is infinite). Moreover, for all $i, j \geq 1$, with $i \leq j$, $w(i)$ is the i-th letter of w, while $w[i, j]$ denotes the finite subword of w given by $w(i) \cdots w(j)$, and w^i the suffix of w from position i, i.e., the word $w(i)w(i+1) \ldots$.

2.1. Cost-Parity Games

We recall the framework of Cost-parity games [13] which are two-player turn-based games played on finite graphs equipped with a Cost-parity winning condition. In such a setting, Player 0 wins a play of infinite duration if there is a bound $\ell \in \mathbb{N}$ such that almost all odd colors (which we think of as requests) are followed by larger even colors (which we think of as responses) that are reached with cost at most ℓ.

A state-transition graph or FSM is a tuple $\langle \mathrm{S}, \mathrm{R}, in \rangle$ consisting of a finite set S of states, a transition relation $\mathrm{R} \subseteq \mathrm{S} \times \mathrm{S}$, and an initial state $in \in \mathrm{S}$. For a state $s \in \mathrm{S}$, we write $\mathrm{R}(s) = \{s' \in \mathrm{S} \mid (s, s') \in \mathrm{R}\}$ for the set of successors of s. A path in the FSM is a non-empty finite or infinite word π over S such that $\pi(i+1) \in \mathrm{R}(\pi(i))$ for all $i \in [1, |\pi| - 1]$.

An *arena* is a tuple $\mathcal{A} = \langle \mathrm{S}, \mathrm{S}_0, \mathrm{S}_1, \mathrm{R}, in \rangle$ consisting of an FSM $\langle \mathrm{S}, \mathrm{R}, in \rangle$ and a partition $\{\mathrm{S}_0, \mathrm{S}_1\}$ of S into the states of Player 0 (drawn as circles) and the states of Player 1 (drawn as rectangles). The size $|\mathcal{A}|$ of $\mathcal{A}$ is $|\mathrm{S}| + |\mathrm{R}|$. A play of a game over $\mathcal{A}$ proceeds by moving a token on the states of $\mathcal{A}$, starting at some state. If the token is placed on a state $s \in \mathrm{S}_0$ (resp., $s \in \mathrm{S}_1$), then the play ends if s has no successors (we call such a state a *terminal state*); otherwise, Player 0 (resp., Player 1) chooses a successor s' of s and moves the token to s'. Formally, a *play* of $\mathcal{A}$ is a *maximal* path of $\mathcal{A}$, i.e., a path π in the underlying FSM such that either π is infinite, or π is finite and ends at a terminal state.

Let $p \in \{0, 1\}$ and S_p^N be the set of non-terminal states of Player p. A *strategy* for Player p is a mapping $\sigma : \mathrm{S}^* \cdot \mathrm{S}_p^N \mapsto \mathrm{S}$ assigning to each non-empty sequence of states $w \cdot s \in \mathrm{S}^* \cdot \mathrm{S}_p^N$ leading to a non-terminal state s of Player p, a successor of s. A play π is consistent with the strategy σ if for all $k \in [1, |\pi| - 1]$ such that $\pi(k) \in \mathrm{S}_p^N$, it holds that $\pi(k+1) = \sigma(\pi[1, k])$. The strategy σ is *memoryless* if its output does not depend on the whole prefix of the play, but only on the last position, i.e, if for all $w \cdot s \in \mathrm{S}^* \cdot \mathrm{S}_p^N$, $\sigma(w \cdot s) = \sigma(s)$. We can thus represent a memoryless strategy as a mapping $\sigma : \mathrm{S}_p^N \to \mathrm{S}$.

A (zero-sum) game is a pair $\langle \mathcal{A}, \mathrm{Win} \rangle$ consisting of an arena $\mathcal{A} = \langle \mathrm{S}, \mathrm{S}_0, \mathrm{S}_1, \mathrm{R}, in \rangle$ and a subset Win of infinite plays which are winning for Player 0. An infinite play π is *winning for Player* 1 if it is not winning for Player 0. A *finite* play π is winning for Player p if π ends at a state of the opponent Player $1 - p$. A strategy σ for Player p is *winning from a state* s if all the plays π starting from s which are consistent with the strategy σ are winning for Player p. In such a case, we say that state s is winning for Player p. A game is determined if for each state s, s is winning for one of the players. Note that since for all strategies σ^0 and σ^1 of Player 0 and Player 1, respectively, there is a unique play starting from s which is consistent with both σ^0 and σ^1, in (zero-sum) games, a state s cannot be winning for both the players. Solving a game consists in checking whether the initial state is winning for Player 0.

Cost-parity winning conditions. We, now, recall the class of cost-parity winning conditions. A *cost-parity arena* $\mathcal{G} = \langle \mathcal{A}, \mathrm{Cost}, \Omega \rangle$ over the set C of colors consists of an arena $\mathcal{A} = \langle \mathrm{S}, \mathrm{S}_0, \mathrm{S}_1, \mathrm{R}, in \rangle$, a transition-labeling $\mathrm{Cost} : \mathrm{R} \mapsto \{0, 1\}$ (cost function), and a coloring mapping $\Omega : \mathrm{S} \mapsto C$ assigning to each state a color in C. Note that according to [13], the definition of transition-labeling only allows cost 0 or 1 on a transition. Having arbitrary costs in $\mathbb{N}$ would not change our results, as we are interested in boundedness questions only. We extend the cost function Cost to the set of paths π by counting the number of increment transitions (i.e., 1-labeled transitions) traversed along the path, i.e., $\mathrm{Cost}(\pi) = \sum_{i=2}^{i=|\pi|} \mathrm{Cost}(\pi(i-1), \pi(i))$ if $|\pi| > 1$, and $\mathrm{Cost}(\pi) = 0$ otherwise. Note that $\mathrm{Cost}(\pi) \in \mathbb{N} \cup \{\infty\}$.

The pair (Cost, Ω) induces a winning condition for Player 0, where the occurrence of an odd color along a play π is interpreted as a *request*, for which there has to be a *response* later on the play by a higher even color. Formally, let π be a finite or infinite path of $\mathcal{A}$. A *request in* π is a position k along π such that $\pi(k)$ has *odd* color. For an odd color c, a c-request in π is a request k in π such that $\Omega(\pi(k)) = c$. Moreover, we define $\mathrm{Ans}(c) = \{c' \in C_e \mid c' \geq c\}$, i.e., the set of even colors that answer a request of color c. For a request k in π, let r_k be the smallest position $k' \geq k$ that answers to request k, i.e., such that $\Omega(\pi(k')) \in \mathrm{Ans}(\Omega(\pi(k)))$, if such positions k' exist, and let $r_k = |\pi|$ otherwise. In the first (resp., second) case, we say that the request k is *answered* (resp., *unanswered*) in π. The *delay* of the request k in π, denoted by $\mathrm{dl}(\pi, k)$, then is defined as the cost of the infix of π from the request k to position r_k, i.e., $\mathrm{Cost}(\pi[k, r_k])$ if $r_k \neq \infty$, and $\mathrm{Cost}(\pi^k)$ otherwise. The *cost-parity* winning condition induced by (Cost, Ω), written $\mathrm{CostParity}(\mathrm{Cost}, \Omega)$, is then the set of infinite plays π such that there is $n \geq 1$ and a bound $\ell \in \mathbb{N}$ so that for all requests k in π with $k \geq n$, $\mathrm{dl}(\pi, k) \leq \ell$ and the request k is answered in π. Thus, an infinite play $\pi \in \mathrm{CostParity}(\mathrm{Cost}, \Omega)$ iff there is bound ℓ such that all but finitely many requests are answered with cost less than ℓ. Note that $\mathrm{CostParity}(\mathrm{Cost}, \Omega)$ is *prefix-independent*, i.e., for all infinite plays π and $k \geq 1$, $\pi \in \mathrm{CostParity}(\mathrm{Cost}, \Omega)$ iff $\pi^k \in \mathrm{CostParity}(\mathrm{Cost}, \Omega)$. We recall the following known result.

Theorem 1 ([13]). *Cost-parity games are determined and Player* 0 *has memoryless winning strategies from the winning Player* 0 *states. Moreover, solving a cost-parity game* $\mathcal{G} = \langle \mathcal{A}, Cost, \Omega \rangle$ *with* k *colors can be done in time* $|\mathcal{A}|^{O(k \cdot \log k)}$ *and in polynomial space.*

For technical convenience, we also consider a generalization of cost-parity arenas, called *partial cost-parity arenas*, where one considers as additional input a subset *Exit* of the set of terminal states, called *exit states*. Finite plays ending at states in *Exit* are assumed to be *non-winning* for either player and have an undefined value. In this setting, a *non-losing strategy for Player p from state s* is a strategy σ for Player p such that each play starting from s which is consistent with σ and does not lead to an exit state is winning for Player p. A *non-losing strategy* is a non-losing strategy for Player 0 from the initial state in. For a strategy σ for Player 0, an *exit play* of σ is a finite play starting from in and ending at an exit state which is consistent with σ. For $s \in \mathit{Exit}$, an s-exit play of σ is an exit play of σ leading to s. Two partial cost-parity arenas $\mathcal{G} = \langle \mathcal{A}, \mathrm{Cost}, \Omega, \mathit{Exit} \rangle$ and $\mathcal{G}' = \langle \mathcal{A}', \mathrm{Cost}', \Omega', \mathit{Exit}' \rangle$ *have the same interface* if $\mathit{Exit} = \mathit{Exit}'$, $\mathcal{G}$ and $\mathcal{G}'$ have the same initial state in, and for each $s \in \{in\} \cup \mathit{Exit}$, the players of state s in $\mathcal{G}$ and $\mathcal{G}'$ coincide. Note that we do *not* impose any constraint on the colors of the states in $\{in\} \cup \mathit{Exit}$.

2.2. Hierarchical Cost-Parity Games

A *Hierarchical Cost-Parity Game* is a cost-parity game played over a (flat) arena induced by a *hierarchical arena*. The latter is a standard hierarchical FSM [24] in which the set of nodes of each of the underlying FSMs is partitioned into nodes belonging to Player 0 and nodes belonging to Player 1. We refer to the underlying FSMs as modular sub-arenas. Formally, a hierarchical arena is a tuple $\mathcal{V} = \langle \mathcal{V}_1, \ldots, \mathcal{V}_n \rangle$ of modular sub-arenas, where each $\mathcal{V}_i$ is in turn a tuple of the form $\langle \mathrm{N}_i, \mathrm{N}_i^0, \mathrm{N}_i^1, \mathrm{B}_i, in_i, \mathit{Exit}_i, \mathrm{Y}_i, \mathrm{E}_i \rangle$ consisting of the following components:

- A finite set N_i of nodes which is partitioned into a set N_i^0 of nodes of Player 0 and a set N_i^1 of nodes of Player 1, and a finite set B_i of *boxes*. We assume that $\mathrm{N}_1, \ldots, \mathrm{N}_n, \mathrm{B}_1, \ldots, \mathrm{B}_n$ are pairwise disjoint.
- An initial node or entry $in_i \in \mathrm{N}_i$, [1] and a subset Exit_i of N_i called *exit-nodes*. We assume that $\mathit{Exit}_1 = \emptyset$, i.e., the top-level sub-arena $\mathcal{V}_1$ has no exits.
- An indexing function $\mathrm{Y}_i : \mathrm{B}_i \to \{i+1, \ldots, n\}$ that maps each box b of $\mathcal{V}_i$ to an index $\mathrm{Y}_i(b) > i$. The box b represents a reference to the definition of the sub-arena $\mathcal{V}_{\mathrm{Y}_i(b)}$.

[1] We assume a single entry for each sub-arena. Multiple entries can be handled by duplicating sub-arenas.

- An edge relation E_i. Each edge in E_i is a pair (u, v) such that: *(i)* the source u is either a node of $\mathcal{V}_i$ or a pair (b, e), where b is a box of $\mathcal{V}_i$ and e is an exit-node of the sub-arena that b refers to, and *(ii)* the target v is either a node or a box of $\mathcal{V}_i$.

Define $\mathrm{N} = \bigcup_{i=1}^{n} \mathrm{N}_i$ (the set of $\mathcal{V}$-nodes), $\mathrm{E} = \bigcup_{i=1}^{n} \mathrm{E}_i$ (the set of $\mathcal{V}$-edges), and $\mathit{Exit} = \bigcup_{i=1}^{n} \mathit{Exit}_i$ (the set of $\mathcal{V}$-exit-nodes). In a modular sub-arena, the edges connect nodes and boxes with one another. Edges entering a box implicitly lead to the unique entry-node of the sub-arena that the box refers to. On the other hand, an edge exiting a box needs to explicitly specify the identity of the exit-node among the possible exit-nodes of the sub-arena associated with that box. The *size* $|\mathcal{V}_i|$ of a modular sub-arena $\mathcal{V}_i$ is $|\mathrm{N}_i| + |B_i| + |\mathrm{E}_i|$. The size $|\mathcal{V}|$ of $\mathcal{V}$ is $\sum_{i=1}^{i=n} |\mathcal{V}_i|$. The *nesting depth* of $\mathcal{V}$ is the length of the longest chain $i_1, i_2, \ldots, i_j$ of indices in $[1, n]$ such that a box of $\mathcal{V}_{i_l}$ is mapped to i_{l+1} for all $l \in [1, j-1]$. Note that the fact that boxes of a sub-arena can only refer to sub-arenas with a greater index implies that the nesting depth of $\mathcal{V}$ is finite. Such a restriction does not exist in the *recursive* setting [27].

A *Hierarchical Cost-Parity Arena* (*HCPA*, for short) over C is a tuple $\mathcal{H} = \langle \mathcal{V}, \mathrm{Cost}, \Omega \rangle$ consisting of a hierarchical arena $\mathcal{V} = \langle \mathcal{V}_1, \ldots, \mathcal{V}_n \rangle$ equipped with a cost function $\mathrm{Cost} : \mathrm{E} \mapsto \{0, 1\}$ for the set of $\mathcal{V}$-edges, and a coloring mapping $\Omega : \mathrm{N} \mapsto C$ for the set of $\mathcal{V}$-nodes. We can associate to $\mathcal{H}$ an ordinary cost-parity arena (called its *flat expansion*) by recursively substituting each box by a copy of the modular sub-arena it refers to. Since different boxes can refer to the same sub-arena, nodes may appear in different contexts. In general, a state of the flat expansion is a vector whose last component is a node, and the remaining components are boxes that specify the context. Formally, for each modular sub-arena $\mathcal{V}_i$, we inductively define its flat expansion as the *partial* Cost-parity arena $\mathcal{H}_i^F = \langle \mathcal{A}_i, \mathrm{Cost}_i, \Omega_i, \mathit{Exit}_i \rangle$, with $\mathcal{A}_i = \langle \mathrm{S}_i, \mathrm{S}_i^0, \mathrm{S}_i^1, \mathrm{R}_i, \mathit{in}_i \rangle$, defined as follows:

- The set of states S_i is inductively defined as follows: *(i)* if u is a node in $\mathcal{V}_i$, then $u \in \mathrm{S}_i$, and *(ii)* if b is a box of $\mathcal{V}_i$ and $s \in \mathrm{S}_{\mathrm{Y}_i(b)}$, then $(b, s) \in \mathrm{S}_i$.

- S_i^0 (resp., S_i^1) is the set of states in S_i whose node-component belongs to Player 0 (resp., Player 1), and the coloring function Ω_i assigns to each state s of $\mathcal{A}_i$, the color $\Omega(u)$ of the node-component u of s.

- The transition relation R_i and the cost function Cost_i are inductively defined as follows.

– If $(u, v) \in \mathrm{E}_i$ and the target v is a node, then $(u, v) \in \mathrm{R}_i$ and $\mathrm{Cost}_i(u, v) = \mathrm{Cost}(u, v)$. If $(u, b) \in \mathrm{E}_i$ and the target b is a box, then $(u, (b, in_{\mathrm{Y}_i(b)})) \in \mathrm{R}_i$ and $\mathrm{Cost}_i(u, (b, in_{\mathrm{Y}_i(b)})) = \mathrm{Cost}(u, b)$.

– If b is a box of $\mathcal{V}_i$ and $(s, s') \in \mathrm{R}_{\mathrm{Y}_i(b)}$, then $((b, s), (b, s')) \in \mathrm{R}_i$ and $\mathrm{Cost}_i((b, s), (b, s')) = \mathrm{Cost}_{\mathrm{Y}_i(b)}(s, s')$.

Note that since $Exit_1 = \emptyset$, $\mathcal{H}_1^F$ is an ordinary Cost-parity arena (i.e., it is not partial), called the flat expansion of $\mathcal{H}$. Moreover, observe that each state of $\mathcal{H}_1^F$ is a vector of length at most the nesting depth d of $\mathcal{V}$, and that the number of states in $\mathcal{H}_1^F$ can be exponential in d. Solving the game on the *HCPA* $\mathcal{H}$ consists in checking whether the initial state in_1 of the cost-parity arena $\mathcal{H}_1^F$ is winning for Player 0.

To get familiar with the notion of *HCPA*, let us consider the following indexed family of hierarchical arenas. For a fixed $n \in \mathbb{N}$ with $n \geq 2$, let $\mathcal{V} = \langle \mathcal{V}_1, \ldots, \mathcal{V}_n \rangle$ be the hierarchical arena defined as follows:

- $\mathcal{V}_1 = \langle \mathrm{N}_1, \mathrm{N}_1^0, \mathrm{N}_1^1, \mathrm{B}_1, in_1, Exit_1, \mathrm{Y}_1, \mathrm{E}_1 \rangle$ where $\mathrm{N}_1^0 = \{in_1, s_1^1, s_1^2\}$, $\mathrm{N}_1^1 = \emptyset$, $\mathrm{B}_1 = \{b_1^1, b_1^2\}$, $Exit_1 = \emptyset$, $Y_1(b_1^1) = Y_1(b_1^2) = 2$, and $\mathrm{E}_1 = \{(in_1, s_1^1), (in_1, s_1^2), (s_1^1, b_1^1), (s_1^2, b_1^2), ((b_1^1, ex_2), in_1), ((b_1^2, ex_2), in_1)\}$;

- $\mathcal{V}_i = \langle \mathrm{N}_i, \mathrm{N}_i^0, \mathrm{N}_i^1, \mathrm{B}_i, in_i, Exit_i, \mathrm{Y}_i, \mathrm{E}_i \rangle$ where $\mathrm{N}_i^0 = \{in_i, s_i^1, s_i^2, ex_i\}$, $\mathrm{N}_i^1 = \emptyset$, $\mathrm{B}_i = \{b_i^1, b_i^2\}$, $Exit_i = \{ex_i\}$, $Y_i(b_i^1) = Y_i(b_i^2) = i{+}1$, and $\mathrm{E}_i = \{(in_i, s_i^1), (in_i, s_i^2), (s_i^1, b_i), (s_i^i, b_i), ((b_i^1, ex_{i+1}), ex_i), ((b_i^2, ex_{i+1}), ex_i)\}$, for every $1 < i < n$;

- $\mathcal{V}_n = \langle \mathrm{N}_n, \mathrm{N}_n^0, \mathrm{N}_n^1, \mathrm{B}_n, in_n, Exit_n, \mathrm{Y}_n, \mathrm{E}_n \rangle$ where $\mathrm{N}_n^0 = \{ex_n\}$, $\mathrm{N}_n^1 = \{in_n, s_n^1\}$, $\mathrm{B}_n = \emptyset$, $Exit_n = \{ex_n\}$, $Y_n = \emptyset$, and $\mathrm{E}_n = \{(in_n, s_n^1), (s_n^1, s_n^1), (s_n^1, ex_n)\}$.

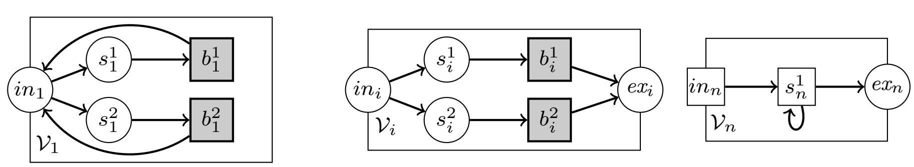

Figure 1: Graphical representation of the modular sub-arenas in $\mathcal{V} = \langle \mathcal{V}_1, \ldots, \mathcal{V}_n \rangle$. Circled nodes are controlled by Player 0, rectangle nodes are controlled by Player 1, and grey-filled nodes are the boxes in every module.

In Figure 1 we depict the graphical representations of the modular sub-arenas introduced above. A game played on this hierarchical arena starts at the node in_1 and enters every module in the hierarchy. At every such level j, Player 0 chooses to move either to node s^1_j or to node s^2_j before entering the lower level. When module $\mathcal{V}_n$ is reached, Player 1 takes over in the game. At this point the game can either loop indefinitely in s^1_n or eventually move to the exit ex_n, which brings the process back to the starting point to start all over again. Note that the arena has $4n-2$ nodes (four nodes for every level, except for levels 1 and n, that have three nodes) and $2(n-1)$ boxes (two boxes for every level except for the bottom). On the other hand, due to the branching feature at every level of the hierarchy, the flat expansion of $\mathcal{V}$ has the shape of a full binary tree of height n. Therefore, the number of its states is exponential in n. In particular, note that there are 2^n instances of the module $\mathcal{V}_n$ occurring in the flat expansion of $\mathcal{V}$.

Now, let us consider a *HCPA* over $\mathcal{V}$. We need to specify both the cost and colour mappings, respectively. Regarding the cost function Cost, we set it up as follows: $\text{Cost}(s^1_n, s^1_n) = 1$ and $\text{Cost}(e) = 0$ for every other $e \in \text{E} \setminus \{(s^1_n, s^1_n)\}$. That is, the only transition with a non-zero cost is the loop in the module $\mathcal{V}_n$. Regarding the colour mapping Ω, we have the following:

- for the initial nodes of the modules: $\Omega(in_1) = 2n$; $\Omega(in_i) = 0$, for every $1 < i < n$; and $\Omega(in_n) = 0$.
- the exit nodes have color 0;
- for the other nodes of the modules: $\Omega(s^1_i) = 2i-1$ for every $1 \leq i < n$; $\Omega(s^2_i) = 2i+1$, for every $1 \leq i < n$; and $\Omega(s^1_n) = 0$.

The game starts from in_1, then it flows through the modules accumulating a number of odd requests whose biggest is lower or equal to $2n-1$, when the game enters the module $\mathcal{V}_n$. At this point, the game can evolve in two different ways: it either stays in s^1_n forever, thus making 0 the only colour occurring infinitely often, or exits the modules one by one and reaches the node in_1 again, thus starting a new iteration. Since the starting node is labelled with the colour $2n$, we have that all the odd requests occurred in the previous walk along the arena are answered and so the parity condition is verified whatsoever. However, Player 1 can force the loop over s^1_n to be longer every time the game enters $\mathcal{V}_n$, thus making the request cost arbitrarily large. Therefore, Player 1 has a winning strategy in the game.

2.3. *Outline of the main results*

In this section we outline the proposed approach for solving hierarchical cost-parity games. Note that the naive method for solving games on *HCPA* $\mathcal{H}$ consisting in applying Theorem 1 on the flat expansion of $\mathcal{H}$ would lead to an exponential space procedure. In this paper, we show that solving hierarchical cost-parity games is PSPACE-complete. Our approach is based on the notion of *summary function* for a strategy σ of Player 0 in a partial cost-parity arena, which records in a finite and efficient way the overall behavior of all the exit plays of σ with respect to requests and responses. The proposed algorithm for solving the game on the given *HCPA* $\mathcal{H}$ then solves a sequence of partial cost-parity games, obtained by replacing each box b referring to a sub-arena $\mathcal{V}_i$ with simple partial-cost parity arenas (*summary-gadget arenas*) having the same interface as the flat expansion $\mathcal{H}_i^F$ of $\mathcal{V}_i$ and depending only on the set of colors and exit states. These gadgets represent the behavior of Player 0 as a choice among the possible summary functions associated with the non-losing memoryless strategies in $\mathcal{H}_i^F$, and also take into account the possibility that the game will stay forever in the sub-arena $\mathcal{V}_i$ for the given context b. The rest of the paper is organized as follows: in Section 3, we introduce the notions of summary and summary-gadget arena, and in Section 4 we show how to check that a summary is associated with non-losing memoryless strategies. Finally, in Section 5, we illustrate the proposed algorithm for solving *HCPA* games, and in Sections 6 and 7, we demonstrate completeness and correctness of the algorithm, respectively.

3. Summaries in partial cost-parity games

In this section, for a given partial cost-parity arena $\mathcal{G}$, we show how to define a finite abstraction of the set of non-losing strategies (of Player 0). Such an abstraction is based on the notion of summary for a strategy σ of Player 0, which is a mapping assigning to each exit state s a value ranging over a finite set (depending only on the set of colors). Such a value summarizes the overall behavior of all the s-exit plays of σ with respect to requests and responses by finitely abstracting the set of associated costs and delays. Then, we associate to each summary $\mathcal{S}$ of $\mathcal{G}$ a simple partial-cost parity arena $Gad(\mathcal{G}, \mathcal{S})$ – exposing the same interface as $\mathcal{G}$ (the initial state and the set of exit states) – which depends only on the set of colors and exit states, and is independent of the set of 'internal' states in $\mathcal{G}$. The set of summary-gadget arenas $Gad(\mathcal{G}, \mathcal{S})$

such that $\mathcal{S}$ is achieved by some non-losing memoryless strategy is 'context-equivalent' to $\mathcal{G}$, i.e., for each memoryless strategy σ achieving some summary $\mathcal{S}$, $\mathcal{G}$ can be equivalently replaced with $Gad(\mathcal{G}, \mathcal{S})$ in any hierarchical context where $\mathcal{G}$ is exploited as a sub-arena and Player 0 chooses strategy σ when entering $\mathcal{G}$. [2]

For ensuring correctness of our approach, we also need to select only summaries which are associated with non-losing memoryless strategies (of Player 0). On the other hand, checking whether a summary is associated with a non-losing memoryless strategy is not an easy task since we have to check the fulfillment of unboundedness conditions. However, we can get around the problem by introducing a binary relation between summaries which gives an indication whether a summary is not worse than another one for Player 0 in order to achieve winning strategies in arbitrary contexts. Then, it suffices to select only the summaries $\mathcal{S}$ which are *relevant*, i.e. such that there is a non-losing memoryless strategy of Player 0 whose summary is not worse than $\mathcal{S}$. As we will see in Section 4, by exploiting monotonicity properties of the cost-parity winning conditions, checking whether a summary is relevant can be done in polynomial space.

The rest of the section is organized as follows. In Subsection 3.1, we introduce the notion of summary for a strategy σ of Player 0 in the given partial cost-parity arena $\mathcal{G}$. Then, in Subsection 3.2, we associate to each summary of $\mathcal{G}$ a gadget $Gad(\mathcal{G}, \mathcal{S})$ having the same interface as $\mathcal{G}$, and show that the *unique* strategy of Player 0 in the gadget has $\mathcal{S}$ as summary. Finally, in Subsection 3.3, we introduce the notion of *relevance* of summaries.

In the following, we fix a partial cost-parity arena $\mathcal{G} = \langle \mathcal{A}, \mathrm{Cost}, \Omega, Exit\rangle$ over the set C of colors, where $\mathcal{A} = \langle \mathrm{S}, \mathrm{S}_0, \mathrm{S}_1, \mathrm{R}, in\rangle$ and $Exit$ is the designated set of exit states.

In order to describe the relative merit of colors, we define an ordering $\succeq_0$ over the given set C of colors by letting $c \succeq_0 c'$ when c is better for Player 0 than c'. Formally, $c \succeq_0 c'$ if: either *(i)* c and c' are even and $c \geq c'$, or *(ii)* c and c' are odd and $c' \geq c$, or *(iii)* c' is odd and c is even. Hence, $\preceq_0$ induces the following total ordering on the set of colors:

$$C_o^{max} \prec_0 C_o^{max} - 2 \prec_0 \ldots \prec_0 1 \prec_0 0 \prec_0 2 \prec_0 \ldots \prec_0 \mathsf{max}(C_e) - 2 \prec_0 \mathsf{max}(C_e)$$

Define $\tilde{C} = C \setminus \{C_o^{max}\}$ and $\tilde{C}_o = (C_o \setminus \{C_o^{max}\}) \cup \{0\}$.

[2]the formal proof of such a context-equivalence is postponed to Section 5.

3.1. *Exit values and summaries of Player 0 strategies*

In this Subsection, for the given partial cost-parity arena $\mathcal{G}$, we introduce the notion of summary for a strategy σ of Player 0. Given an exit node s of $\mathcal{G}$, the summary for σ represents finite information about the delays and costs of requests and responses over the (possibly infinite) set of s-exit plays of σ. Since the cost-parity winning condition is prefix-independent, this information is, in particular, exploited for taking account of scenarios where σ is part, within a context $\mathcal{C}$ (a sequence of boxes), of a winning 'global' strategy σ_g for Player 0 and there are plays π consistent with σ_g which visit infinitely many times the global state $(\mathcal{C}, s)$ (hence, π visits infinitely many times s-exit plays of σ in the context $\mathcal{C}$). Since we can assume that σ_g is memoryless (hence, σ is memoryless as well), all the s-exit plays of σ have a role for ensuring that σ_g is winning. Indeed, each play obtained from π by replacing s-exit plays of σ occurring in π in the context $\mathcal{C}$ with arbitrary s-exit plays of σ is still consistent with the global strategy σ_g.

In order to formalize the notion of summary for a strategy σ of Player 0, we consider various cost measures with respect to the requests and the responses along the exit plays of σ. For this, we extend the cost function Cost to (possibly infinite) sets Π of *finite* paths of $\mathcal{G}$. Formally, $\text{Cost}(\Pi)$ is the least upper bound over the costs of the paths in Π, i.e., $\text{Cost}(\Pi) = \mathsf{sup}\{\text{Cost}(\nu) \mid \nu \in \Pi\}$ where $\mathsf{sup}\,\emptyset = 0$. Note that $\text{Cost}(\Pi) \in \mathbb{N} \cup \{\infty\}$. For a *finite* path ν of $\mathcal{G}$ and an even color $c_e \in C_e$, a *c_e-response in ν* is a position k of ν such that $\nu(k)$ has color c_e. For such a response k, the *cost of response k in ν* is the cost of the prefix of ν leading to position k, i.e., $\text{Cost}(\nu[1, k])$. The *c_e-response cost* of ν, denoted by $\text{ResCost}(\nu, c_e)$, is the cost $\text{Cost}(\nu[1, k])$ of the prefix of ν up to the minimal c'_e-response k in ν for some even color $c'_e \geq c_e$ if such c'_e-responses exist, and it is 0 otherwise. The *maximal even color* of the path ν is the maximal even color visited by ν if ν visits some even color, and it is 0 otherwise (note that a 0-response cannot answer to any request). We exploit the following cost measures for the (possibly infinite) set of exit plays of a given strategy σ of Player 0 leading to a designated exit state.

Definition 1 (Cost measures of Player 0 strategies)**.** *Let $s \in Exit$, σ a strategy of Player 0, Π_s the (possibly empty) set of exit plays of σ leading to s, and $c_e \in C_e$ an even color. We consider the following cost measures:*

- Cost of σ w.r.t. s, *denoted $Cost(\sigma, s)$: it is $Cost(\Pi_s)$.*
- Request-cost of σ w.r.t. s, *denoted $ReqCost(\sigma, s)$: it is the least upper bound over the*

delays associated with the requests along the exit plays in Π_s*, i.e.,* $\mathsf{sup}\{\mathrm{dl}(\nu,k) \mid \nu \in \Pi_s$ *and* k *is a request in* $\nu\}$*.*

- Even c_e-cost of σ w.r.t. s*, denoted* $Cost_e(\sigma, s, c_e)$*: it is* $Cost(\Pi_{c_e})$*, where* Π_{c_e} *is the (possibly empty) set of exit plays in* Π_s *whose maximal even color is at most* c_e*.*
- c_e-response cost of σ w.r.t. s*, denoted* $ResCost(\sigma, s, c_e)$*: it is the least upper bound over the* c_e*-response costs of the exit plays in* Π_s*, i.e.,* $\mathsf{sup}\{ResCost(\nu, c_e) \mid \nu \in \Pi_s\}$*.*

Note that $\mathrm{Cost}_e(\sigma, s, -)$ is monotonic in the third argument, i.e., $\mathrm{Cost}_e(\sigma, s, c'_e) \geq \mathrm{Cost}_e(\sigma, s, c_e)$ for all $c_e, c'_e \in C_e$ such that $c'_e \geq c_e$. We, now, introduce the notion of summary for a (memoryless) strategy σ of Player 0 which records for each exit state s, a value, called *exit value*, ranging over a finite set depending only on the set of colors. This value summarizes the overall behavior of the exit plays of σ leading to s. We distinguish three situations (recall that $C_o^{max} = \mathsf{max}(C)$ and C_o^{max} is odd).

The best scenario for Player 0 is when there is no exit play of σ leading to s. We represent this situation by exploiting the special symbol $\vdash$.

On the opposite side, the worst scenario is when the request-cost of σ w.r.t. s is infinite, or there is an s-exit play of σ having a C_o^{max}-request. Indeed if σ is part, within a context $\mathcal{C}$, of a global memoryless strategy σ_g of Player 0 such that some play consistent with σ_g visits infinitely many times the global state $(\mathcal{C}, s)$ (hence, π visits infinitely many times s-exit plays of σ in the context $\mathcal{C}$), then σ_g cannot be winning for Player 0 . We use the color C_o^{max} to describe the worst scenario.

If none of the two previous conditions is fulfilled, then the exit value is a sextuple of elements:

- the first element (*cost value*) finitely summarizes the cost of σ w.r.t. s. We exploit the symbol bnd_0 when the cost $\mathrm{Cost}(\sigma, s)$ is zero, the symbol bnd_1 when such a cost is finite but not zero, and the symbol unb when such a cost is infinite. We denote by $\succeq_b$ the ordering on $\{bnd_0, bnd_1, unb\}$ defined as: $bnd_0 \succeq_b bnd_1$ and $bnd_1 \succeq_b unb$. Intuitively, $bnd_0 \succeq_b bnd_1$ and $bnd_1 \succeq_b unb$ express that bounded zero-cost is better for Player 0 then non-zero bounded cost, the latter being in turn better than unbounded cost.
- The second element (*parity value*) keeps track of the worst color for Player 0 over the maximal colors along the s-exit plays of σ.
- The third element (*odd value*) represents the maximal odd color, if any, associated with

an unanswered request (note that being the considered scenario distinct from the worst one, such a color is distinct from C_o^{max} and the delays of the unanswered requests along the s-exit plays of σ are bounded).

- The last three elements in the tuple summarize the overall response behavior of the s-exit plays of σ. We keep track of the greatest even color c_e, if any, whose response-cost is bounded and such that each s-exit play visits an even color which is at least c_e (*even-left value*). Moreover, we also record the greatest color c_e, if any, whose response-cost is bounded and such that the cost over the s-exit plays visiting only even colors smaller than c_e is bounded (*even-right value*). Finally, we keep track of the smallest even color c_e, if any, less or equal to the right even-right value such that the cost over the s-exit plays visiting only even colors smaller or equal to c_e is bounded but not null (*even-middle value*).

Formally, the exit values for a strategy of Player 0 are defined as follows.

Definition 2 (Exit values of Player 0 strategies)**.** *Let $s \in Exit$, σ a strategy of Player 0, and Π_s the set of exit plays of σ leading to s. The* exit value $value(\sigma, s)$ of strategy σ w.r.t. s *is defined as follows. If $\Pi_s = \emptyset$, then $value(\sigma, s) = \vdash$. If instead either $ReqCost(\sigma, s) = \infty$ or there is $\nu \in \Pi_s$ having a C_o^{max}-request, then $value(\sigma, s) = C_o^{max}$. Otherwise, $value(\sigma, s) = (value_{Cost}(\sigma, s), value_{pr}(\sigma, s), value_o(\sigma, s), value_e^L(\sigma, s), value_e^M(\sigma, s), value_e^R(\sigma, s)) \in \{bnd_0, bnd_1, unb\} \times \tilde{C} \times \tilde{C}_o \times C_e \times (C_e \cup \{\bot\}) \times C_e$, and the following holds:*

- Cost value $value_{\mathrm{Cost}}(\sigma, s)$*:* (i) $value_{Cost}(\sigma, s) = unb$ *if* $Cost(\Pi_s) = \infty$, (ii) $value_{Cost}(\sigma, s) = bnd_0$ *if* $Cost(\Pi_s) = 0$, *and* (iii) $value_{Cost}(\sigma, s) = bnd_1$ *otherwise.*
- Parity value $value_{pr}(\sigma, s)$*: it is* $\mathsf{min}_{\preceq_0}\{c \in C \mid c$ *is the maximal color of some* $\nu \in \Pi_s\}$.
- Odd value $value_o(\sigma, s)$*: it is the* greatest *odd color $c_o \in C_o$ such that for some $\nu \in \Pi_s$, ν has an unanswered c_o-request if such an odd color c_o exists; otherwise, it is* 0.
- Even-left value $value_e^L(\sigma, s)$*: it is the* greatest *even color $c_e \in C_e$ such that $ResCost(\sigma, s, c_e) \neq \infty$ and for each $\nu \in \Pi_s$, the maximal even color in ν is at least c_e, if such an even color c_e exists; otherwise, it is* 0.
- Even-right value $value_e^R(\sigma, s)$*: it is the* greatest *even color $c_e \in C_e$ such that $ResCost(\sigma, s, c_e) \neq \infty$ and for each $c'_e \in C_e$ with $c'_e < c_e$, $Cost_e(\sigma, s, c'_e) \neq \infty$, if such an even color c_e exists; otherwise, it is* 0.
- Even-middle value $value_e^M(\sigma, s)$*: it is the* smallest *even color $c_e \leq value_e^R(\sigma, s)$ such that*

$Cost_e(\sigma, s, c_e) \in \mathbb{N} \setminus \{0\}$ if such a color c_e exists, and $value_e^M(\sigma, s) = \bot$ otherwise ($\bot$ is for 'undefined').

We now provide further intuitions on the meaning of the exit value $value(\sigma, s)$ when $value(\sigma, s) \notin \{\vdash, C_o^{max}\}$. We also give some examples which illustrate that the information recorded by $value(\sigma, s)$ is necessary for distinguishing situations which behave differently in the same context.

The parity value $value_{pr}(\sigma, s)$ intuitively represents the worst color for Player 0 offered by an arbitrary infinite sequence of s-exit plays of σ in an arbitrary context. For parity winning conditions, the parity value $value_{pr}(\sigma, s)$ suffices for summarizing the s-exit behavior of strategy σ [5]. For cost-parity winning conditions, we also need to keep track of the maximal odd color $value_o(\sigma, s)$ associated with unanswered requests. Note that $value_o(\sigma, s) \preceq_0 value_{pr}(\sigma, s)$, and $value_o(\sigma, s) \prec_0 value_{pr}(\sigma, s)$ entails that the maximal unanswered request is associated with s-exit plays whose maximal color is even.

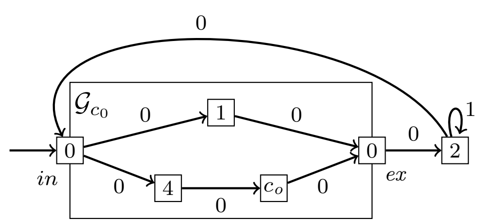

As an example, let us consider the sub-arena $\mathcal{G}_{c_o}$ - parametric in the color c_o - in the figure on the left. Note that all the states are controlled by Player 1 (hence, there is a unique strategy of Player 0). The exit values of the instances $\mathcal{G}_1$ and $\mathcal{G}_3$ of $\mathcal{G}_{c_o}$ with respect to the exit node ex differ only for the odd value. In particular, $\mathcal{G}_1$ and $\mathcal{G}_3$ have parity value 1, and odd value 1 and 3, respectively. While by using $\mathcal{G}_1$, all the plays starting from state in are winning for Player 0, the same does not hold by using $\mathcal{G}_3$ since in this case, there are plays where the request 3 is answered in an unbounded way.

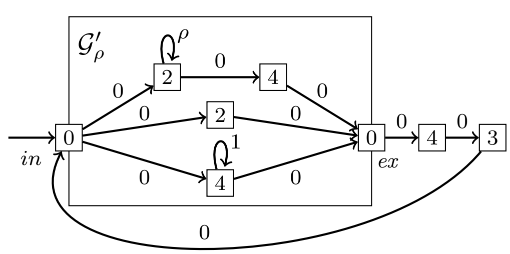

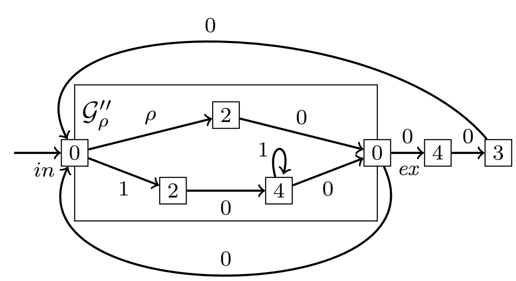

For what concerns the even values, the even-left value $value_e^L(\sigma, s)$ represents, intuitively, the maximal even color that the s-exit plays of σ offer for answering – in a bounded way – to previous requests in an arbitrary context. The even-right value $value_e^R(\sigma, s)$ is at least the even-left value (by Proposition 1(1), $value_e^R(\sigma, s) \geq value_e^L(\sigma, s)$) and is useful only when the cost of the s-exit plays of σ is unbounded (in the other cases, $value_e^R(\sigma, s)$ is always the maximal even color). Intuitively, it represents the greatest even color c_e such that a c_e-answer to a request cannot unboundedly delayed by an arbitrary s-exit play of σ. As an example, let us consider the sub-arena $\mathcal{G}'_\rho$ – parametric in the cost ρ of the self-loop on the state with color 2 – in the left part of the previous figure, where all the states are controlled by Player 1. The instances $\mathcal{G}'_0$ and $\mathcal{G}'_1$ of $\mathcal{G}'_\rho$ have even-left value 2, and even-right value 4 and 2, respectively (note that the exit values of $\mathcal{G}'_0$ and $\mathcal{G}'_1$ differ only for the even-right value). While for $\mathcal{G}'_0$, all the plays starting from state in are winning for Player 0, the same does not hold for $\mathcal{G}'_1$, since in this case, there are plays where the external request 3 is answered in an unbounded way. Finally, in order to illustrate the importance of the even-middle value, let us consider the sub-arena $\mathcal{G}''_\rho$ in the right part of the previous figure, where again all the states are controlled by Player 1. The instances $\mathcal{G}''_0$ and $\mathcal{G}''_1$ of $\mathcal{G}''_\rho$ have even-left value 2, right-even value 4, and even-middle value $\perp$ and 2, respectively. While for $\mathcal{G}''_0$, all the plays starting from in are winning for Player 0, for $\mathcal{G}''_1$, there are plays where the external request 3 is answered in an unbounded way. We make the following observations which easily follow from Definition 2.

Proposition 1. *Let σ be a strategy of Player 0 in $\mathcal{G}$ and $s \in Exit$ such that $value(\sigma, s) = (f, c_{pr}, c_o, c_e^L, c_e^M, c_e^R)$. Then:*

1. $c_o \preceq_o c_{pr}$, $c_e^L \leq c_{pr}$, $c_e^L \leq c_e^R$, *and* $c_e^M \in [c_e^L, c_e^R]$ *if* $c_e^M \neq \perp$.
2. $c_e^M = \perp$ *if* $f = bnd_0$, *and* $c_e^M \neq \perp$ *if* $f = bnd_1$.
3. $c_e^R = \mathsf{max}(C_e)$ *if* $f \neq unb$, *and* $c_e^R < \mathsf{max}(C_e)$ *if* $f = unb$ *and* $c_e^M = c_e^R$.
4. *if* $c_{pr} \in C_e$ *and either* $f \neq unb$, *or* $c_e^L < c_e^R$, *or* $c_e^M = c_e^R$, *then it holds that* $c_e^L = c_{pr}$.

Proof. Let Π_s be the set of exit plays of strategy σ leading to s.

Proof of Property 1. For the inequality $c_o \preceq_o c_{pr}$, we assume on the contrary that $c_o \succ_o c_{pr}$, and derive a contradiction. Since $c_o \in C_o \cup \{0\}$, by definition of the ordering $\succeq_0$, we have that $c_{pr} \in C_o$ and $c_{pr} > c_o$. By definition of parity value, there is an exit play $\nu \in \Pi_s$ whose maximal color is c_{pr}. Since c_{pr} is odd, there is an unanswered c_{pr}-request along ν. By definition of odd value, we obtain that $c_o \geq c_{pr}$, which is a contradiction. Hence, $c_o \preceq_o c_{pr}$.

Now, we show that $c_e^L \leq c_{pr}$. By definition of parity value, there is an exit play $\nu \in \Pi_s$ whose maximal color is c_{pr}. By definition of even-left value, the maximal even color along ν is at least c_e^L. Hence, the result follows.

Next, let us prove that $c_e^L \leq c_e^R$. If $c_e^L = 0$, the result is obvious. Otherwise, by definition of even-left value, $\text{ResCost}(\sigma, s, c_e^L) \neq \infty$ and for each exit play $\nu \in \Pi_s$, the maximal even color in ν is at least c_e^L. Hence, for each even color c_e with $c_e < c_e^L$, $\text{Cost}_e(\sigma, s, c_e) = 0$. Hence, by definition of even-right value, the result follows.

For concluding the proof of Property 1, we need to show that $c_e^M \in [c_e^L, c_e^R]$ if $c_e^M \neq \bot$. Assume that $c_e^M \neq \bot$. By Definition 2, $c_e^M \leq c_e^R$ and $\text{Cost}_e(\sigma, s, c_e^M) \neq 0$. Hence, there is an exit play $\nu \in \Pi_s$ whose maximal color is at most c_e^M. Such a maximal color is at least the even-left value c_e^L. Hence $c_e^L \leq c_e^M$, and the result follows.

Proof of Property 2. If $f = bnd_0$, then $\text{Cost}(\Pi_s) = 0$. Hence, for each even color c_e, $\text{Cost}_e(\sigma, s, c_e) = 0$. By definition of even-middle value, we obtain that $c_e^M = \bot$. Now, assume that $f = bnd_1$. Hence, $\text{Cost}(\Pi_s) \in \mathbb{N} \setminus \{0\}$, $c_e^R = \mathsf{max}(C_e)$, and $\text{Cost}_e(\sigma, s, c_e^R) \in \mathbb{N} \setminus \{0\}$. It follows that $c_e^M \neq \bot$, which concludes the proof of Property 2.

Proof of Property 3. The first part is obvious (i.e., $c_e^R = \mathsf{max}(C_e)$ if $f \neq unb$). For the second part, assume that $f = unb$ and $c_e^M = c_e^R$. We need to show that $c_e^R < \mathsf{max}(C_e)$. We assume that $c_e^R = \mathsf{max}(C_e)$ and derive a contradiction. Since no play in Π_s visits the odd color C_o^{max}, we have that Π_s coincides with the set of exit plays of strategy σ leading to s whose maximal even color is at most c_e^R. Since $f = unb$, we obtain that $\text{Cost}_e(\sigma, s, c_e^R) = \infty$. On the other hand, by hypothesis, $c_e^M = c_e^R$. By Definition 2, it follows that $\text{Cost}_e(\sigma, s, c_e^R) \neq \infty$, which is a contradiction. Hence, $c_e^R < \mathsf{max}(C_e)$.

Proof of Property 4. Assume that $c_{pr} \in C_e$. We need to show that $c_e^L = c_{pr}$ whenever either $f \neq unb$, or $c_e^L < c_e^R$, or $c_e^M = c_e^R$.

- $f \neq unb$: since c_{pr} is even, by definition of parity value, for each play $\nu \in \Pi_s$, the maximal color of ν is even and is at least c_{pr}. Since $f \neq unb$, by Definition 2, we obtain that $c_e^L \geq c_{pr}$. By Property 1, $c_e^L \leq c_{pr}$. Hence, $c_e^L = c_{pr}$, and the result holds.

- $f = unb$ and $c_e^L < c_e^R$: since $c_e^L < c_e^R$, by definition of even-left and even-right value, the set of even colors c_e such that $c_e < c_e^R$ and there is an s-exit play of σ whose maximal even color is c_e is not empty, and c_e^L is the minimum over such colors. Hence, there exists a play $\nu \in \Pi_s$ whose maximal even color is c_e^L. Since c_{pr} is even, by definition of parity value, the maximal color of ν is at least c_{pr}. Hence, $c_e^L \geq c_{pr}$. Thus, by applying Property 1, we obtain that $c_e^L = c_{pr}$.

- $f = unb$ and $c_e^L = c_e^M = c_e^R$. By definition of even-middle value, there is a play $\nu \in \Pi_s$ whose maximal even color is c_e^M. Thus, since $c_e^L = c_e^M$, by proceeding as in the previous case, the result follows.

This concludes the proof of Proposition 1. □

Definition 3 (Summaries of Player 0 strategies)**.** *The set $\mathcal{E}_C$ of* exit values *for the set C of colors is the finite set $\{\vdash, C_o^{max}\} \cup \mathcal{E}'_C$, where $\mathcal{E}'_C$ is the set of tuples $(f, c_{pr}, c_o, c_e^L, c_e^M, c_e^R) \in \{bnd_0, bnd_1, unb\} \times \tilde{C} \times \tilde{C}_o \times C_e \times (C_e \cup \{\bot\}) \times C_e$ satisfying Conditions (1)–(4) in Proposition 1.*

A summary *of $\mathcal{G}$ is a mapping $\mathcal{S} : Exit \mapsto \mathcal{E}_C$ associating to each exit state an exit value. The* summary $\mathcal{S}(\sigma)$ of a strategy σ of Player 0 *in $\mathcal{G}$ is the summary of $\mathcal{G}$ associating to each $s \in Exit$, the exit value $value(\sigma, s)$.*

3.2. Summary gadgets

In this Subsection, for each summary $\mathcal{S}$ of the given partial cost-parity arena $\mathcal{G}$, we define a simple partial-cost parity game $Gad(\mathcal{G}, \mathcal{S})$, exposing the same interface as $\mathcal{G}$ and independent of the set of 'internal' states in $\mathcal{G}$, such that there is a unique strategy σ_S of Player 0 in $Gad(\mathcal{G}, \mathcal{S})$. Moreover, σ_S is non-losing and the exit values of σ_S correspond to the exit values of any strategy of Player 0 in $\mathcal{G}$ having $\mathcal{S}$ as summary (Proposition 2).

Definition 4 (Summary-Gadget Arena)**.** *Let $\mathcal{S}$ be a summary of $\mathcal{G}$. Given $ex \in Exit$, we first define the sub-gadget $Gad(\mathcal{G}, \mathcal{S}, ex)$ of $\mathcal{G}$ for summary $\mathcal{S}$ and ex, which is the partial cost-parity game with set of states $\mathrm{S}_{ex} \cup \{\mathcal{S}, ex\}$ and set of edges R_{ex}, where:*

- *All the states in $\mathrm{S}_{ex} \cup \{\mathcal{S}\}$ are controlled by Player 1, $\mathcal{S}$ has color 0 and is the initial state, ex is the unique exit state and has color 0, and the player of state ex is as in $\mathcal{G}$.*

Moreover, if $\mathcal{S}(ex) = \vdash$, then $\mathrm{S}_{ex} = \emptyset$, and $\mathrm{R}_{ex} = \emptyset$. On the opposite side, if $\mathcal{S}(ex) = C_o^{max}$, then S_{ex} consists of a unique state s having color C_o^{max}, and R_{ex} consists of two edges, one from *state $\mathcal{S}$ to state s with cost 0, and the other one from s to ex with cost 0 as well. Otherwise, let $\mathcal{S}(ex) = (f, c_{pr}, c_o, c_e^L, c_e^M, c_e^R)$. Let d_o and d_e be the colors in C defined as follows:*

$$d_o = \begin{cases} c_{pr} & \textit{if } c_o \in \{0, c_{pr}\} \\ c_o & \textit{otherwise} \end{cases}$$

$$d_e = \begin{cases} c_e^L & \textit{if } c_o \in \{0, c_{pr}\} \\ \mathsf{max}(C_e) & \textit{otherwise} \end{cases}$$

Then, we distinguish six cases, where in the figures illustrating the construction, we assume that ex is controlled by Player 0.

Case $f = \boldsymbol{bnd_0}$.

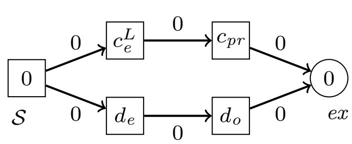

In this case, we have $c_e^M = \bot$ and $c_e^R = \mathsf{max}(C_e)$. The sub-gadget $Gad(\mathcal{G}, \mathcal{S}, ex)$ for this case is a DAG and is illustrated on the left. Note that the cost of any path from state $\mathcal{S}$ to the exit state ex is 0.

Case $f = \boldsymbol{bnd_1}$.

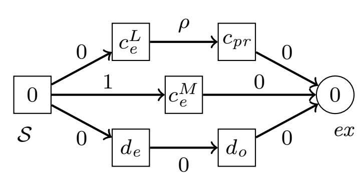

In this case, we have that $c_e^M \in C_e$, $c_e^M \in [c_e^L, c_e^R]$, and $C_e^R = \mathsf{max}(C_e)$. The associated sub-gadget is a DAG and it is illustrated on the right, where $\rho = 0$ if $c_e^L < c_e^M$, and $\rho = 1$ otherwise. Note that the overall cost of all paths from state $\mathcal{S}$ to the exit state ex is 1. Moreover, according to the definition of even-middle value, c_e^M represents the smallest even color c_e such that the overall cost of all exit plays leading to ex and having maximal even color c_e is finite and non-null. Additionally, if $c_e^L < c_e^M$, according to the definition of even-left value, there are exit plays leading to ex whose maximal even color is c_e^L, and the overall cost of such exit plays is 0.

Case $f = \boldsymbol{unb}$, $\boldsymbol{c_e^M = \bot}$, and $\boldsymbol{c_e^L = c_e^R}$.

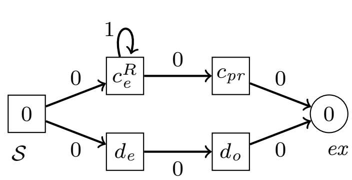

The sub-gadget $Gad(\mathcal{G}, \mathcal{S}, ex)$ for this case is illustrated on the left. When $f = unb$, the overall cost of all exit plays leading to ex is infinite. This is implemented by a self-loop with cost 1 on the state having color c_e^R. Note that for a strategy σ of Player 0 with $value_{\text{Cost}}(\sigma, ex) = unb$ and $value_e^R(\sigma, ex) = c_e^R$, the overall cost of all ex-exit plays having maximal even color at most c_e^R may be finite. However, in this case, $c_e^R < \mathsf{max}(C_e)$ and $\text{ResCost}(\sigma, ex, c_e^R + 2) = \infty$. Thus, the self-loop with cost 1 in the sub-gadget above takes into account also these possible scenarios. Additionally, note that by Proposition 1, this is the unique case where assuming that c_{pr} is even, c_e^L may be strictly smaller than c_{pr} (in all the other case, if c_{pr} is even then $c_e^L = c_{pr}$). Thus, the self-loop with cost 1 in the sub-gadget above also ensures that for the summarized strategies, the even-left value and right-even value coincide with c_e^L.

Case $f = unb$, $c_e^M = \bot$, and $c_e^L < c_e^R$.

This case is similar to the previous one. The unique difference is that now $c_e^L < c_e^R$. Thus, the associated sub-gadget – illustrated on the right – summarizes strategies σ of Player 0 for which, in particular, $value_e^L(\sigma, ex) = c_e^L$ and there are exit plays leading to ex whose maximal even color is c_e^L, and the overall cost of such exit plays is 0.

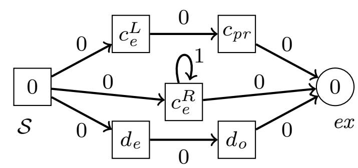

Case $f = unb$, $c_e^M \in C_e$, and $c_e^M < c_e^R$. This case is similar to the previous one, but now $c_e^M \in C_e$, hence, $c_e^M \in [c_e^L, c_e^R]$. The associated sub-gadget is illustrated in the left part of the figure below, where $\rho = 0$ if $c_e^L < c_e^M$, and $\rho = 1$ otherwise.

Case $f = unb$, $c_e^M \in C_e$, and $c_e^M = c_e^R$. We have that $c_e^M \in [c_e^L, c_e^R]$ and $c_e^R < \mathsf{max}(C_e)$. The associated sub-gadget is illustrated on the right of the figure below, where $c_+^R = c_e^R + 2$, $\rho = 0$ if $c_e^L < c_e^M$, and $\rho = 1$ otherwise. In this case there is an even color, namely c_+^R, whose response-cost with respect to ex is infinite. This is consistent with the fact that for all strategies σ of Player 0 such that $value_{\text{Cost}}(\sigma, ex) = unb$, $value_e^R(\sigma, ex) = c_e^R$, and $\text{Cost}_e(\sigma, ex, c_e^R) \neq \infty$, it holds that $\text{ResCost}_e(\sigma, ex, c_e^R + 2) = \infty$.

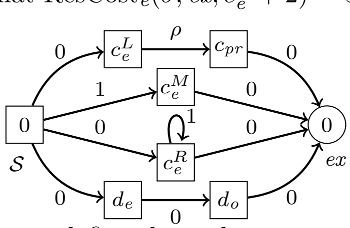

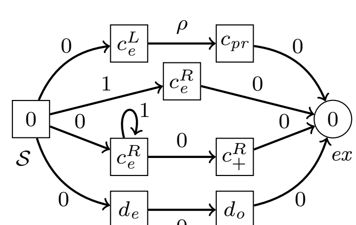

We now define the gadget arena $Gad(\mathcal{G}, \mathcal{S})$ for the given summary $\mathcal{S}$, which is intuitively obtained by merging the sub-gadgets $\mathit{Gad}(\mathcal{G}, \mathcal{S}, ex)$ for the various exit states $ex \in \mathit{Exit}$ and by adding the state in. Formally, assuming that $S_{ex} \cap S_{ex'} = \emptyset$ (i.e., sub-gadgets associated with distinct exit states share only state $\mathcal{S}$), $\mathit{Gad}(\mathcal{G}, \mathcal{S})$ has the same interface as $\mathcal{G}$ and satisfies the following: the set of states of $\mathit{Gad}(\mathcal{G}, \mathcal{S})$ is $\{in, \mathcal{S}\} \cup \mathit{Exit} \cup \bigcup_{ex \in \mathit{Exit}} \mathrm{S}_{ex}$ and the set of transitions is $\{(in, \mathcal{S})\} \cup \bigcup_{ex \in \mathit{Exit}} \mathrm{R}_{ex}$, where transition $(in, \mathcal{S})$ has cost 0.

Remark 1. *Note that in a summary-gadget arena $\mathit{Gad}(\mathcal{G}, \mathcal{S})$, every state which is not in $\{in\} \cup \mathit{Exit}$ is controlled by Player 1. In particular, there is exactly one strategy of Player 0, and such a strategy is non-losing.*

By construction, we deduce that the exit values of the unique strategy σ_S in the gadget $\mathit{Gad}(\mathcal{G}, \mathcal{S})$ correspond to the exit values of any strategy of Player 0 in $\mathcal{G}$ having $\mathcal{S}$ as summary.

Proposition 2. *Let $\mathcal{G} = \langle \mathcal{A}, \mathit{Cost}, \Omega, \mathit{Exit} \rangle$ be a partial-cost parity arena, $\mathcal{S}$ a summary of $\mathcal{G}$, σ_S the unique strategy of Player 0 in $\mathit{Gad}(\mathcal{G}, \mathcal{S})$, and $s \in \mathit{Exit}$. Then, $\mathit{value}(\sigma_S, s) = \mathcal{S}(s)$.*

Proof. If $\mathcal{S}(s) \in \{\vdash, C_o^{max}\}$, the result directly follows from Definition 4. Now assume that $\mathcal{S}(s) = (f, c_{pr}, c_o, c_e^L, c_e^M, c_e^R)$. By construction, $\mathit{value}(\sigma_S, s) \notin \{\vdash, C_o^{max}\}$. Let Π_s be the non-empty set of exit plays of σ_S leading to s. Moreover, let d_o and d_e be the colors exploited in Definition 4. Recall that

$$d_o = \begin{cases} c_{pr} & \text{if } c_o \in \{0, c_{pr}\} \\ c_o & \text{otherwise} \end{cases}$$

$$d_e = \begin{cases} c_e^L & \text{if } c_o \in \{0, c_{pr}\} \\ \mathsf{max}(C_e) & \text{otherwise} \end{cases}$$

We need to prove that $\mathit{value}(\sigma_S, s) = \mathcal{S}(s)$.

Cost Value. By construction of the sub-gadget $\mathit{Gad}(\mathcal{G}, \mathcal{S}, s)$ (Definition 4), we easily deduce that $\mathit{value}_{\mathrm{Cost}}(\sigma_S, s) = f$.

Parity Value. Now we prove that $\mathit{value}_{pr}(\sigma_S, s) = c_{pr}$. By construction of $\mathit{Gad}(\mathcal{G}, \mathcal{S}, s)$, there are three types of exit plays ν in Π_s:

- *c_{pr}-plays*, whose set of visited colors is $\{0, c_e^L, c_{pr}\}$.
- *d_o-plays*, whose set of visited colors is $\{0, d_e, d_o\}$.
- *even-plays*, which visit only even colors and whose maximal even color is at least c_e^L.

We distinguish two cases:

1. $c_o \in \{0, c_{pr}\}$: hence, $d_o = c_{pr}$ and $d_e = c_e^L$, and the d_o-plays are c_{pr}-plays. By Proposition 1(1), $c_e^L \leq c_{pr}$. Thus, if c_{pr} is odd or $c_{pr} = c_e^L$, we obtain that $\mathit{value}_{pr}(\sigma_S, s) = c_{pr}$. Otherwise, c_{pr} is even, $c_e^L < c_{pr}$, and by Proposition 1(4), $f = \mathit{unb}$, $c_e^L = \bot$, and $c_e^L = c_e^R$. By construction of the sub-gadget $\mathit{Gad}(\mathcal{G}, \mathcal{S}, s)$ for the case $f = \mathit{unb}$, $c_e^L = \bot$, and $c_e^L = c_e^R$, it holds that the maximal color in each exit play in Π_s is c_{pr}. Thus, $\mathit{value}_{pr}(\sigma_S, s) = c_{pr}$.

2. $c_o \notin \{0, c_{pr}\}$: hence, $d_o = c_o$ and $d_e = \mathsf{max}(C_e)$. Since $c_o < C_o^{max}$, we have that the maximal color in the d_o-plays is the maximal even color $\mathsf{max}(C_e)$. This means that the d_o-plays have no role in the calculation of the parity value $\mathit{value}_{pr}(\sigma_S, s)$. Thus, by proceeding as in the previous case, the result follows.

Odd Value. Next we show that $\mathit{value}_o(\sigma_S, s) = c_o$. By construction of $\mathit{Gad}(\mathcal{G}, \mathcal{S}, s)$, we have that $\mathit{value}_o(\sigma_S, s) = \mathsf{min}_{\preceq_0}\{c_{pr}, 0, d_o\}$. Thus, we need to prove that $c_o = \mathsf{min}_{\preceq_0}\{c_{pr}, 0, d_o\}$. We distinguish three cases:

- $c_o = 0$. Hence, $d_o = c_{pr}$. By Proposition 1(1), $c_o \preceq_0 c_{pr}$. This means that c_{pr} is even and $\mathsf{min}_{\preceq_0}\{c_{pr}, 0, d_o\} = 0$, and the result follows.
- $c_o \neq 0$ and $c_o = c_{pr}$. Hence, $d_o = c_{pr}$ and c_{pr} is odd. It follows that $\mathsf{min}_{\preceq_0}\{c_{pr}, 0, d_o\} = c_{pr}$, and the result holds in this case too.
- $c_o \notin \{c_{pr}, 0\}$: hence, $d_o = c_o$ and c_o is odd. Since $c_o \preceq_0 c_{pr}$ (Proposition 1(1)), we obtain that $\mathsf{min}_{\preceq_0}\{c_{pr}, 0, d_o\} = c_o$ as desired.

Even-Left Value. We now show that $\mathit{value}_e^L(\sigma_S, s) = c_e^L$. Recall that $d_e \in \{c_e^L, \mathsf{max}(C_e)\}$ and by Proposition 1, $c_e^L \leq c_e^R$, $c_e^L \leq c_{pr}$, and $c_e^M \in [c_e^L, c_e^R]$ if $c_e^M \neq \bot$. Thus, by construction of $\mathit{Gad}(\mathcal{G}, \mathcal{S}, s)$, the following holds:

(i) the maximal even color in each exit play $\nu \in \Pi_s$ is at least c_e^L;
(ii) for each exit play $\nu \in \Pi_s$, the c_e^L-response cost of ν is at most 1.
(iii) there is some exit play $\nu \in \Pi_s$ whose set of visited colors is $\{0, c_e^L, c_{pr}\}$.

Conditions (i) and (ii) imply that $\mathit{value}_e^L(\sigma_S, s) \geq c_e^L$. In order to show that $\mathit{value}_e^L(\sigma_S, s) \leq c_e^L$, we distinguish two cases:

- either $f \neq \mathit{unb}$, or $c_e^L < c_e^R$, or $c_e^M = c_e^R$: we show that there exists an exit play in Π_s

whose maximal even color is c_e^L, hence, the result follows. By Proposition 1(4), either c_{pr} is odd or $c_{pr} = c_e^L$. Thus, by Condition (iii), the result follows.

- $f = unb$, $c_e^L \geq c_e^R$, and $c_e^M \neq c_e^R$. By Proposition 1(1), we deduce that $f = unb$, $c_e^L = c_e^R$ and $c_e^M = \bot$. If either c_{pr} is odd or $c_{pr} \leq c_e^L$, then we proceed as in the previous case. Otherwise, c_{pr} is even and $c_{pr} > c_e^L$. By construction of the sub-gadget $Gad(\mathcal{G}, \mathcal{S}, s)$ for the case where $f = unb$, $c_e^L = c_e^R$ and $c_e^M = \bot$, it holds that $\text{ResCost}(\sigma_S, s, c_{pr}) = \infty$ and the set of even colors visited by the exit plays in Π_s is $\{0, c_e^L, c_{pr}, d_e\}$. Since $d_e \in \{c_e^L, \mathsf{max}(C_e)\}$, by Condition (iii) it follows that for each even color $c_e > c_e^L$, either there is an exit play whose maximal even color is strictly smaller than c_e, or $\text{ResCost}(\sigma_S, s, c_e) = \infty$. This entails that $value_e^L(\sigma_S, s) \leq c_e^L$, and the result follows.

Even-Right Value. We need to prove that $value_e^R(\sigma_S, s) = c_e^R$. If $f \neq unb$, then being $value_{\text{Cost}}(\sigma_S, s) = f$, by Proposition 1(3), we have that $value_e^R(\sigma_S, s) = c_e^R = \mathsf{max}(C_e)$. Now assume that $f = unb$. By construction of $Gad(\mathcal{G}, \mathcal{S}, s)$ and Proposition 1(1), the following holds:

- for each even color $c_e < c_e^R$, $\text{Cost}_e(\sigma, s, c_e) \neq \infty$;
- for each exit play $\nu \in \Pi_s$, the c_e^R-response cost of ν is at most 1.

The previous conditions entail that $value_e^R(\sigma_S, s) \geq c_e^R$. In order to show that $value_e^R(\sigma_S, s) \leq c_e^R$, we distinguish two cases (recall that $f \neq unb$):

- $f = unb$ and either $c_e^L < c_e^R$, or $c_e^M = c_e^R$: by construction either (i) $\text{Cost}_e(\sigma, s, c_e^R) = \infty$, or (ii) $\text{Cost}_e(\sigma, s, c_+^R) = \infty$ and $\text{ResCost}(\sigma_S, s, c_+^R) = \infty$, where $c_+^R = c_e^R + 2$. Hence, the result follows.

- $f = unb$, $c_e^L \geq c_e^R$, and $c_e^M \neq c_e^R$. By Proposition 1(1), we deduce that $f = unb$, $c_e^L = c_e^R$ and $c_e^M = \bot$. First assume that either c_{pr} is odd or $c_{pr} \leq c_e^R$. In this case, by construction of the sub-gadget $Gad(\mathcal{G}, \mathcal{S}, s)$ for the case where $f = unb$, $c_e^L = c_e^R$ and $c_e^M = \bot$, it holds that $\text{Cost}_e(\sigma, s, c_e^R) = \infty$ and the result holds. Otherwise, c_{pr} is even, $c_{pr} > c_e^R$, and by construction, (i) for each even color $c_e \in [c_e^R + 2, c_{pr}]$, $\text{ResCost}(\sigma_S, s, c_e) = \infty$ and (ii) for each even color $c_e \geq c_{pr}$, $\text{Cost}_e(\sigma, s, c_e) = \infty$. Hence, the result holds in this case as well.

Even-Middle Value. Finally, we show that $value_e^M(\sigma_S, s) = c_e^M$. We proceed according to the cases exploited in the construction of $Gad(\mathcal{G}, \mathcal{S}, s)$:

- Case $f = bnd_0$: in this case $value_{\text{Cost}}(\sigma_S, s) = bnd_0$, and by Proposition 1(2), $value_e^M(\sigma_S, s) = c_e^M = \bot$.
- Case $f = bnd_1$: by Proposition 1(4), $c_{pr} = c_e^L$ if c_{pr} is even. Thus, by construction, we have that $\text{Cost}_e(\sigma, s, c_e^M) = 1$ and for each even color $c_e < c_e^M$, $\text{Cost}_e(\sigma, s, c_e) = 0$. This entails that $value_e^M(\sigma_S, s) = c_e^M$.
- Case $f = unb$, $c_e^M = \bot$, and $c_e^L = c_e^R$: by construction for each even color $c_e < c_e^R$, $\text{Cost}_e(\sigma, s, c_e) = 0$, and either (i) $\text{Cost}_e(\sigma, s, c_e^R) = 0$ (in case c_{pr} is even and $c_{pr} > c_e^R$), or (ii) $\text{Cost}_e(\sigma, s, c_e^R) = \infty$. Hence, $value_e^M(\sigma_S, s) = \bot$.
- Case $f = unb$, $c_e^M = \bot$, and $c_e^L < c_e^R$: by Proposition 1(4), $c_{pr} = c_e^L$ if c_{pr} is even. By construction, it follows that for each even color $c_e < c_e^R$, $\text{Cost}_e(\sigma, s, c_e) = 0$. Moreover, $\text{Cost}_e(\sigma, s, c_e^R) = \infty$. Hence, $value_e^M(\sigma_S, s) = \bot$.
- Case $f = unb$, $c_e^M \in C_e$, and $c_e^M < c_e^R$: this case is similar to the case $f = bnd_1$.
- Case $f = unb$, $c_e^M \in C_e$, and $c_e^M = c_e^R$: this case is similar to the case $f = bnd_1$.

This concludes the proof of Proposition 2. □

By Proposition 2 and Definition 4, we deduce the following result (Proposition 3). Proposition 3 will be exploited in Section 5 to show that the existence of a winning memoryless strategy σ_g for Player 0 in a hierarchical arena $\mathcal{V}$ where $\mathcal{G}$ is exploited as a sub-arena within a context $\mathcal{C}$ ensures the existence of a winning memoryless strategy for Player 0 in the arena obtained form $\mathcal{V}$ by replacing $\mathcal{G}$ in the context $\mathcal{C}$ with the gadget $Gad(\mathcal{G}, \mathcal{S})$ where $\mathcal{S}$ is the summary of the non-losing strategy of $\mathcal{G}$ induced by the global strategy σ_g within the context $\mathcal{C}$.

Proposition 3. *Let $\mathcal{G} = \langle \mathcal{A}, Cost, \Omega, Exit \rangle$ be a partial-cost parity arena, σ a strategy of Player 0, σ_S the unique strategy of Player 0 in $Gad(\mathcal{G}, \mathcal{S}(\sigma))$, and $s \in Exit$. Then, $value(\sigma, s) = value(\sigma_S, s)$. Moreover, if $value(\sigma, s) \notin \{\vdash, C_o^{max}\}$, the following holds:*

1. *Let ν be an s-exit play of σ_S with maximal even color c_e. Then, either* (i) *$Cost_e(\sigma, s, c_e) \geq Cost_e(\sigma_S, s, c_e)$, and there is an s-exit play ν' of σ whose maximal even color is at most c_e, or* (ii) *$Cost_e(\sigma_S, s, c_e) = \infty$ and $ResCost(\sigma, s, c'_e) = \infty$ for some even color $c'_e \in C_e$ such that $c'_e \leq c_e + 2$.*
2. *For each $c_e \in C_e$, $ResCost(\sigma_S, s, c_e) = \infty$ entails that $ResCost(\sigma, s, c_e) = \infty$.*

Proof. By Proposition 2, $value(\sigma_S, s) = \mathcal{S}(\sigma)(s) = value(\sigma, s)$. Hence, the first part of the proposition holds. For the second part, assume that $value(\sigma, s) \notin \{\vdash, C_o^{max}\}$. Hence, $value(\sigma, s)$

is of the form $(f, c_{pr}, c_o, c_e^L, c_e^M, c_e^R)$. Let d_o and d_e be the colors exploited in Definition 4. Since $\mathcal{S}(\sigma)(s) = (f, c_{pr}, c_o, c_e^L, c_e^M, c_e^R)$, we have that:

$$d_o = \begin{cases} c_{pr} & \text{if } c_o \in \{0, c_{pr}\} \\ c_o & \text{otherwise} \end{cases}$$

$$d_e = \begin{cases} c_e^L & \text{if } c_o \in \{0, c_{pr}\} \\ \mathsf{max}(C_e) & \text{otherwise} \end{cases}$$

Proof of Property 1. let ν be an s-exit play of σ_S with maximal even color c_e. We need to show that either Condition (i) or Condition (ii) of Property 1 holds. By construction of the gadget $Gad(\mathcal{G}, \mathcal{S}(\sigma))$ in Definition 4 and since $\mathcal{S}(\sigma)(s) = (f, c_{pr}, c_o, c_e^L, c_e^M, c_e^R)$, one of the following cases occur according to the set $Col(\nu)$ of colors visited by the exit play ν.

- **Case** $Col(\nu) = \{0, c_{\boldsymbol{e}}^L, c_{\boldsymbol{pr}}\}$**:** we first assume that either $f \neq unb$, or $c_e^L < c_e^R$ or $c_e^M = c_e^R$. By Proposition 1(4), either c_{pr} is odd, or $c_{pr} = c_e^L$. Hence, the maximal even color c_e in ν is c_e^L. Moreover, since either $f \neq unb$, or $c_e^L < c_e^R$ or $c_e^M = c_e^R$, by definition of even values, it easily follows that there is an s-exit play ν' of σ whose maximal even color is c_e^L. Moreover, by construction of the sub-gadget $Gad(\mathcal{G}, \mathcal{S}(\sigma), s)$ for the cases where either $f \neq unb$, or $c_e^L < c_e^R$ or $c_e^M = c_e^R$, it holds that *either* (i) $\text{Cost}_e(\sigma_S, s, c_e^L) = 0$, *or* (ii) $\text{Cost}_e(\sigma_S, s, c_e^L) = 1$ and $c_e^L = c_e^M$. Since $c_e = c_e^L$, in the first case Condition (i) of Property 1 trivially follows. In the second case, since $c_e^L = c_e^M$, by definition of even-middle value, it holds that $\text{Cost}_e(\sigma, s, c_e^L) \geq 1$. Hence, Condition (i) of Property 1 holds in this case as well.

  Now assume that $f = unb$, $c_e^L \geq c_e^R$ and $c_e^M \neq c_e^R$. By Proposition 1(1), it follows that $c_e^L = c_e^R$ and $c_e^M = \perp$. We distinguish two sub-cases:

  - either c_{pr} is odd or $c_{pr} \leq c_e^R$: hence, the maximal even color c_e in ν is c_e^R. Moreover, by construction of the sub-gadget $Gad(\mathcal{G}, \mathcal{S}(\sigma), s)$ for the case where $f = unb$, $c_e^M = \perp$, and $c_e^L = c_e^R$ (and being either c_{pr} is odd or $c_{pr} \leq c_e^R$), we obtain that $\text{Cost}_e(\sigma_S, s, c_e^R) = \infty$. Furthermore, by definition of even-right value, either $\text{Cost}_e(\sigma, s, c_e^R) = \infty$ (this entails that there are s-exit plays of σ whose maximal color is at most c_e^R) or $\text{ResCost}(\sigma, s, c_e^R + 2) = \infty$. Hence, being $c_e = c_e^R$, Property 1 follows.

– c_{pr} is even and $c_{pr} > c_e^L$: hence, the maximal even color c_e in ν is c_{pr}. By construction of the sub-gadget $Gad(\mathcal{G}, \mathcal{S}(\sigma), s)$ for the case where $f = unb$, $c_e^M = \bot$, and $c_e^L = c_e^R$ (and being c_{pr} even and $c_{pr} > c_e^R$), we obtain that $\mathrm{Cost}_e(\sigma_S, s, c_{pr}) = \infty$. Since $value_{pr}(\sigma, s) = c_{pr}$, by definition of parity value, the maximal color of any s-exit play of σ is even and it is at least c_{pr}. Thus, for each even color $c'_e < c_{pr}$, it holds that $\mathrm{Cost}_e(\sigma, s, c'_e) = 0$. Moreover, being $c_e^R = value_e^R(\sigma, s)$ and $c_{pr} > c_e^R$, by definition of even-right value, we deduce that $\mathrm{ResCost}_e(\sigma, ex, c_{pr}) = \infty$. Hence, being $c_e = c_{pr}$, Condition (ii) of Property 1 directly follows.

- **Case** $Col(\nu) = \{0, c_e^M\}$ **with** $c_e^M \neq \bot$**:** hence $c_e = c_e^M$ and by construction of the sub-gadget $Gad(\mathcal{G}, \mathcal{S}(\sigma), s)$ for the cases where $c_e^M \neq \bot$, it holds that $\mathrm{Cost}_e(\sigma_S, s, c_e^M) = 1$. Since $value_e^M(\sigma, s) = c_e^M$, by definition of even-middle value, $\mathrm{Cost}_e(\sigma_S, s, c_e^M) \geq 1$ (in particular, there is an s-exit play ν' of σ whose maximal even color is at most c_e^M). Since $c_e = c_e^M$, Condition (i) of Property 1 trivially follows.
- **Case** $Col(\nu) = \{0, c_e^R\}$ **with** $f = unb$ **and either** $c_e^L < c_e^R$ **or** $c_e^M \neq \bot$**:** hence, $c_e = c_e^R$. Moreover, by Proposition 1(1), either $c_e^L < c_e^R$ or $c_e^L = c_e^M = c_e^R$. It easily follow that there is some s-exit play ν' of σ whose maximal even color is at most c_e^L (hence, the maximal even color in ν' is at most c_e^R). By construction of the sub-gadget $Gad(\mathcal{G}, \mathcal{S}(\sigma), s)$ for the cases where $f = unb$ and either $c_e^L < c_e^R$ or $c_e^M \neq \bot$, it holds that either (i) $\mathrm{Cost}_e(\sigma_S, s, c_e^R) = \infty$, or (ii) $\mathrm{Cost}_e(\sigma_S, s, c_e^R) = 1$ and $c_e^M = c_e^R$. In the first case, if $\mathrm{Cost}_e(\sigma, s, c_e^R) = \infty$, then Condition (i) of Property 1 holds with $c_e = c_e^R$. Otherwise, by definition of even-right value and since $value_e^R(\sigma, s) = c_e^R$, we have that $\mathrm{ResCost}(\sigma, s, c_e^R + 2) = \infty$, and Condition (ii) of Property 1 with $c_e = c_e^R$ directly follows. Now, assume that $\mathrm{Cost}_e(\sigma_S, s, c_e^R) = 1$ and $c_e^M = c_e^R$. Since $value_e^M(\sigma, s) = c_e^M$, by definition of even-middle value, it holds that $\mathrm{Cost}_e(\sigma_S, s, c_e^M) \geq 1$. Hence, being $c_e = c_e^M$, Condition (i) of Property 1 holds.
- **Case** $Col(\nu) = \{0, c_e^R, c_+^R\}$ **with** $f = unb$**,** $c_+^R = c_e^R + 2$**, and** $c_e^M = c_e^R$**:** hence, $c_e = c_e^R + 2$. By construction of the sub-gadget $Gad(\mathcal{G}, \mathcal{S}(\sigma), s)$ for the case where $f = unb$ and $c_e^M = c_e^R$, we have that $\mathrm{Cost}_e(\sigma_S, s, c_e^R + 2) = \infty$. Since $value_e^M(\sigma, s) = c_e^M = c_e^R$, by definition of even-middle value, it holds that $\mathrm{Cost}_e(\sigma, s, c_e^R) \neq \infty$. Moreover, since $value_e^R(\sigma, s) = c_e^R$, by definition of even-right value, we deduce that $\mathrm{ResCost}(\sigma, s, c_e^R + 2) = \infty$. Hence, being $c_e = c_e^R + 2$, Condition (ii) of Property 1 holds.
- **Case** $Col(\nu) = \{0, d_e, d_o\}$**:** hence, either (i) $d_e = c_e^L$ and $d_o = c_{pr}$, or (ii) $d_e = \mathsf{max}(C_e)$

and $d_o = c_o$. In the first case, we have that $Col(\nu) = \{0, c_e^L, c_{pr}\}$ and this case has already been examined. Otherwise, $Col(\nu) = \{0, \mathsf{max}(C_e), c_o\}$, and since C_o^{max} never occurs in the sub-gadget $Gad(\mathcal{G}, \mathcal{S}(\sigma), s)$, we have that $c_e = \mathsf{max}(C_e)$ and $\text{Cost}_e(\sigma_S, s, c_e)$ coincides with the cost $\text{Cost}(\sigma_S, s)$ of σ_S w.r.t. s. Moreover, by construction, either (i) $f = bnd_0$ and $\text{Cost}(\sigma_S, s) = 0$, or (ii) $f = bnd_1$ and $\text{Cost}(\sigma_S, s) = 1$, or (iii) $f = unb$ and $\text{Cost}(\sigma_S, s) = \infty$. Thus, since $value_{\text{Cost}}(\sigma, s) = f$ and the maximal color of each exit play of σ leading to s is at most $\mathsf{max}(C_e)$, by definition of cost-value, the result easily follows.

Proof of Property 2. Let $c_e \in C_e$ such that $\text{ResCost}(\sigma_S, s, c_e) = \infty$. We need to show that $\text{ResCost}(\sigma, s, c_e) = \infty$. By construction of the gadget $Gad(\mathcal{G}, \mathcal{S}(\sigma))$ (Definition 4), one of the following two conditions holds:

- $f = unb$, $c_e^M = \bot$, $c_e^L = c_e^R$, c_{pr} is even, $c_e^R < c_{pr}$, and $c_e \in [c_e^R{+}2, c_{pr}]$: since $value_{pr}(\sigma, s) = c_{pr}$, by definition of parity value, the maximal color of any s-exit play of σ is even and it is at least c_{pr}. Thus, for each even color $c'_e < c_{pr}$, it holds that $\text{Cost}_e(\sigma, s, c'_e) = 0$. Since $c_e \in [c_e^R + 2, c_{pr}]$, then $\text{Cost}_e(\sigma, s, c'_e) = 0$ for each even color $c'_e < c_e$. Moreover, being $c_e^R = value_e^R(\sigma, s)$ and $c_e > c_e^R$, by definition of even-right value, we deduce that $\text{ResCost}_e(\sigma, ex, c_e) = \infty$, and Property 2 holds.
- $f = unb$, $c_e^M = c_e^R$, and $c_e = c_e^R + 2$: hence, $value_e^R(\sigma, s) = value_e^M(\sigma, s) = c_e^R$. By definition of even-middle value, we deduce that $c_e^R < \mathsf{max}(C_e)$ and $\text{Cost}_e(\sigma, ex, c_e^R) \neq \infty$. Thus, by definition of even-right value, it follows that $\text{ResCost}_e(\sigma, ex, c_e^R + 2) = \infty$, and Property 2 holds.

This concludes the proof of Proposition 3. □

3.3. Relevant summaries

Not all the summaries of the given partial cost-parity arena $\mathcal{G}$ are associated with non-losing (memoryless) strategies (of Player 0). On the other hand, checking whether a summary is associated with a non-losing memoryless strategy is not an easy task since we have to check the fulfillment of unboundedness conditions. However, we can get around the problem by exploiting monotonicity properties of the cost-parity winning conditions. In this section, we define a reflexive and transitive relation $\sqsupseteq$ over the set of summaries. Intuitively, $\mathcal{S} \sqsupseteq \mathcal{S}'$ when $\mathcal{S}$ is not worse than $\mathcal{S}'$ for Player 0. A summary $\mathcal{S}$ is then *relevant* if $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$ for some non-losing *memoryless* strategy σ. As we will see in Section 4, checking whether a summary is relevant can be done in polynomial space.

Definition 5 (Relevant summaries)**.** *Let $\sqsupseteq$ be a binary relation over $\mathcal{E}_C$ defined as follows:*

- $\vdash \sqsupseteq ev$ *for all* $ev \in \mathcal{E}_C$*;*
- $ev \sqsupseteq C_o^{max}$ *for all* $ev \in \mathcal{E}_C$*;*
- $(f, c_{pr}, c_o, c_e^L, c_e^M, c_e^R) \sqsupseteq (\tilde{f}, \tilde{c_{pr}}, \tilde{c_o}, \tilde{c_e^L}, \tilde{c_e^M}, \tilde{c_e^R})$ *if* $f \succeq_b \tilde{f}$*,* $c_{pr} \succeq_0 \tilde{c_{pr}}$*,* $c_o \succeq_0 \tilde{c_o}$*,* $c_e^L \geq \tilde{c_e^L}$*,* $c_e^R \geq \tilde{c_e^R}$*, and the following holds:*
  - *if* $c_e^M \neq \bot$*, then either* $\tilde{c_e^M} \neq \bot$ *and* $c_e^M \geq \tilde{c_e^M}$*, or* $\tilde{c_e^M} = \bot$ *and* $c_e^M \geq \tilde{c_e^R}$*.*

Given two summaries $\mathcal{S}$ and $\mathcal{S}'$ of $\mathcal{G}$, we say that $\mathcal{S}$ is not worse than $\mathcal{S}'$ for Player 0*, written $\mathcal{S} \sqsupseteq \mathcal{S}'$, if $\mathcal{S}(s) \sqsupseteq \mathcal{S}'(s)$ for all $s \in Exit$. A summary $\mathcal{S}$ of $\mathcal{G}$ is* relevant *iff there is a memoryless non-losing strategy σ in $\mathcal{G}$ such that $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$.*

Remark 2. *The binary relation $\sqsupseteq$ over the set of summaries is reflexive and transitive.*

Note that if $\mathcal{G}$ has no exits, then the unique summary is the empty set, and such a summary is relevant iff there is a memoryless winning strategy of Player 0 from *in*. By construction, we obtain the following result, which represents the converse of Proposition 3. Proposition 4 will be exploited in Section 5 to show that in a hierarchical arena $\mathcal{V}$ where $\mathcal{G}$ is exploited as a sub-arena in a context $\mathcal{C}$, the existence of winning memoryless strategies for Player 0 in the arena obtained from $\mathcal{V}$ by replacing $\mathcal{G}$ in the context $\mathcal{C}$ with the gadget $Gad(\mathcal{G}, \mathcal{S})$ for some relevant summary $\mathcal{S}$ of $\mathcal{G}$ implies the existence of winning strategies for Player 0 in the original arena.

Proposition 4. *Let $\mathcal{G} = \langle \mathcal{A}, Cost, \Omega, Exit \rangle$ be a partial-cost parity arena, $\mathcal{S}$ a summary of $\mathcal{G}$, σ a strategy of Player 0 such that $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$, σ_S the unique strategy of Player 0 in $Gad(\mathcal{G}, \mathcal{S})$, and $s \in Exit$. Then $value(\sigma, s) \sqsupseteq value(\sigma_S, s)$. Moreover, if $\mathcal{S}(s) \notin \{\vdash, C_o^{max}\}$, the following holds:*

1. *Let ν be an s-exit play of σ, c_e the maximal even color of ν, and $Cost_e(\sigma, s, c_e) = m \in \mathbb{N} \cup \{\infty\}$. Then, either* (i) *$Cost_e(\sigma_S, s, c_e) = m'$ where $m' > 0$ if $m > 0$, and $m' = \infty$ if $m = \infty$, and there is a s-exit play ν' of σ_S whose maximal even color is at most c_e, or* (ii) *$Cost_e(\sigma, s, c_e) = \infty$ and $ResCost(\sigma_S, s, c_e') = \infty$ for some even color $c_e' \in C_e$ such that $c_e' \leq c_e + 2$.*
2. *For each $c_e \in C_e \setminus \{0\}$, if $ResCost(\sigma, s, c_e) = \infty$, one of the following holds:*
   - *either $ResCost(\sigma_S, s, c_e') = \infty$ for some even color $c_e' \leq c_e$,*

- *or there is an even color $c'_e \leq value^R_e(\sigma, s) < c_e$ such that $Cost_e(\sigma_S, s, c'_e) = \infty$.*

Proof. By Proposition 2, $value(\sigma_S, s) = \mathcal{S}(s)$. Thus, since $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$, we obtain that $value(\sigma, s) \sqsupseteq value(\sigma_S, s)$, and the first part of the proposition holds. For the second part, assume that $\mathcal{S}(s) \notin \{\vdash, C^{max}_o\}$. Hence, $\mathcal{S}(s)$ is of the form $(f, c_{pr}, c_o, c^L_e, c^M_e, c^R_e)$.

Proof of Property 1. let ν be an s-exit play of σ with maximal even color c_e and such that $\text{Cost}_e(\sigma, s, c_e) = m \in \mathbb{N} \cup \{\infty\}$. Note that by definition of even-left value, the maximal even color of ν is at least $value^L_e(\sigma, s)$. Hence, $c_e \geq value^L_e(\sigma, s)$. We need to show that either Condition (i) or Condition (ii) of Property 1 holds. We distinguish the following cases according to whether the even c_e-cost $\text{Cost}_e(\sigma, s, c_e)$ is zero, finite and non-null, or infinite.

- $\text{Cost}_e(\sigma, s, c_e) = 0$: since $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$, we have that $c_e \geq value^L_e(\sigma, s) \geq c^L_e$. By construction of the gadget $Gad(\mathcal{G}, \mathcal{S})$ (Definition 4), there exists an exit play ν' of σ_S whose set of visited colors is $\{0, c^L_e, c_{pr}\}$. If the maximal even color of ν' is c^L_e, then being $c_e \geq c^L_e$, Condition (i) of Property 1 directly follows. Otherwise, c_{pr} is even and the maximal even color of ν' is c_{pr}. In this case, being $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$ and c_{pr} even, we have that $value_{pr}(\sigma, s) \succeq_o value_{pr}(\sigma_S, s) = c_{pr}$ and $value_{pr}(\sigma, s)$ is even. Being $value_{pr}(\sigma, s)$ even and c_e be the maximal even color of the s-exit play ν of σ, by definition of parity value, we deduce that $c_e \geq value_{pr}(\sigma, s)$. Hence, $c_e \geq c_{pr}$, and Condition (i) of Property 1 holds in this case as well.
- $\text{Cost}_e(\sigma, s, c_e) \in \mathbb{N} \setminus \{0\}$: hence, by definition of even-middle value, *either* $value^M_e(\sigma, s) \neq \bot$ and $c_e \geq value^M_e(\sigma, s)$, *or* $value^M_e(\sigma, s) = \bot$ and $c_e > value^R_e(\sigma, s)$. Since $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$, it follows that either (*) $c^M_e \neq \bot$ and $c_e \geq c^M_e$, or (**) $c^M_e = \bot$ and $c_e \geq c^R_e$. If Condition (*) holds, by construction of the gadget $Gad(\mathcal{G}, \mathcal{S})$, there is an a s-exit play ν' of σ_S whose maximal even color is c^M_e and $\text{Cost}_e(\sigma_S, s, c^M_e) \geq 1$. Hence, being $c^M_e \leq c_e$, Condition (i) of Property 1 directly follows. Now, assume that Condition (**) holds. Since $c^M_e = \bot$, $\text{Cost}_e(\sigma, s, c_e) \neq 0$, and $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$, we deduce that $f = unb$. Let c'_e be the even color given by c^R_e if c_{pr} is either odd or $c_{pr} \leq c^R_e$, and by c_{pr} otherwise. By construction of the sub-gadgets $Gad(\mathcal{G}, \mathcal{S}, s)$ for the cases where $f = unb$ and $c^M_e = \bot$, it follows that $\text{Cost}_e(\sigma_S, s, c'_e) = \infty$ (in particular, there exists an s-exit play ν' of σ_S whose maximal even color is at most c'_e). If $c'_e = c^R_e$, being $c_e \geq c^R_e$, then Condition (i) of Property 1 holds. Otherwise, $c'_e = c_{pr}$ and c_{pr} is even. Being c_{pr} even and $value_{pr}(\sigma, s) \succeq_o$

$value_{pr}(\sigma_S, s) = c_{pr}$, by definition of parity value, we deduce that $value_{pr}(\sigma, s)$ is even and $c_e \geq value_{pr}(\sigma, s) \geq c_{pr}$ (recall that c_e is the maximal even color in the s-exit play ν of σ). Hence, being $c'_e = c_{pr}$, Condition (i) of Property 1 holds in this case too.

- $\text{Cost}_e(\sigma, s, c_e) = \infty$: hence, $c_e \geq value^R_e(\sigma, s)$. Since $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$, it holds that $value^R_e(\sigma, s) \geq c^R_e$ and $f = unb$. By construction of the sub-gadget $Gad(\mathcal{G}, \mathcal{S}, s)$ for the cases where $f = unb$, we deduce that *either* $\text{Cost}_e(\sigma_S, s, c^R_e) = \infty$, *or* $c^R_e + 2 \in C_e$ and $\text{ResCost}(\sigma_S, s, c^R_e + 2) = \infty$. In the first case, being $c_e \geq c^R_e$, Condition (i) of Property 1 directly follows. In the second case, we set $c'_e = c^R_e + 2$, and being $c'_e \leq c_e + 2$, Condition (ii) of Property 1 follows.

Proof of Property 2. Let $c_e \in C_e$ such that $c_e \neq 0$ and $\text{ResCost}(\sigma, s, c_e) = \infty$. By definition of right-even value, we deduce that $value^R_e(\sigma, s) < c_e$. Moreover, since $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$, we have that $f = unb$ and $c^R_e \leq value^R_e(\sigma, s)$. By construction of the sub-gadget $Gad(\mathcal{G}, \mathcal{S}, s)$ for the cases where $f = unb$, it follows that *either* $\text{Cost}_e(\sigma_S, s, c^R_e) = \infty$ *or* $c^R_e + 2 \in C_e$ and $\text{ResCost}(\sigma_S, s, c^R_e + 2) = \infty$. In the first case, we set $c'_e = c^R_e$ and being $c^R_e \leq value^R_e(\sigma, s) < c_e$, we get $c'_e \leq value^R_e(\sigma, s) < c_e$ and $\text{Cost}_e(\sigma_S, s, c'_e) = \infty$. Hence, Property 2 follows. In the second case, we set $c'_e = c^R_e + 2$ and being $c^R_e < c_e$, we get $c'_e \leq c_e$ and $\text{ResCost}(\sigma_S, s, c'_e) = \infty$, and the result follows. □

Note that the set of relevant summaries in $\mathcal{G}$ is empty iff there does not exist a memoryless non-losing strategy in $\mathcal{G}$. By Theorem 1, checking this condition can be done in polynomial space. In this case, we associate with $\mathcal{G}$ a simple partial cost-parity arena (bad gadget), where Player 0 always loses.

Definition 6 (Bad-Gadget Arena)**.** *The* bad-gadget arena $BadGad(\mathcal{G})$ *of a partial cost-parity arena $\mathcal{G}$ is the partial cost-parity game having the same interface as $\mathcal{G}$ and defined as follows: $BadGad(\mathcal{G})$ has a unique 'internal' state $s \notin \{in\} \cup Exit$, which is controlled by Player 0, and a unique transition, namely (in, s), which has cost 0. Moreover, each state has color 0.*

4. Checking relevance of summaries

We reduce the problem of checking summary relevance in partial cost-parity arenas to verifying the existence of *memoryless* strategies in cost-parity arenas under a simple imperfect-information setting. Formally, an *observation-based cost-parity arena* (*OCPA*) is a cost-parity arena $\mathcal{G} = \langle \mathcal{A}, \text{Cost}, \Omega, \mathit{Obs} \rangle$ equipped with an observability equivalence relation $\mathit{Obs} \subseteq \text{S} \times \text{S}$ over the set of states. An *observation-based memoryless* strategy of Player 0 is a memoryless strategy σ of Player 0 such that, for all non-terminal states s and s' controlled by Player 0, $(s, s') \in \mathit{Obs} \Rightarrow (\sigma(s), \sigma(s')) \in \mathit{Obs}$. First, we observe the following.

Theorem 2. *Let $\mathcal{G} = \langle \mathcal{A}, \mathit{Cost}, \Omega, \mathit{Obs} \rangle$ be an OCPA. Checking the existence of a winning observation-based memoryless strategy of Player 0 from the initial state can be done in polynomial space.*

Proof. Let *ObsWin* be the set of *OCPA* such that there is a winning observation-based memoryless strategy of Player 0 from the initial state. We show that $\mathit{ObsWin} \in \text{NP}$. Hence, being $\text{NP} \subseteq \text{PSPACE}$, the result follows. The following is a nondeterministic polynomial-time algorithm for deciding *ObsWin*: given an *OCPA* $\mathcal{G}$, (i) guess an observation-based memoryless strategy σ of $\mathcal{G}$, and (ii) check whether σ is winning for Player 0 from the initial state. We need to show that the second step can be done in polynomial time. The strategy σ can be represented as a subgraph $\mathcal{G}_\sigma$ of $\mathcal{G}$. This subgraph coincides with $\mathcal{G}$ except that all the transitions (s, s') where s is a state controlled by Player 0 and $s' \neq \sigma(s)$ are removed (i.e., for a non-terminal state controlled by Player 0, we only keep the outgoing transitions referred to by σ). Given $\mathcal{G}_\sigma$, we need to check that the following conditions do *not* hold:

1. there exists a finite path of $\mathcal{G}_\sigma$ starting from the initial state and leading to a terminal state controlled by Player 0.

2. there exists an infinite path of $\mathcal{G}_\sigma$ starting from the initial state which does not satisfy the parity-cost condition associated with $\mathcal{G}$. Since $\mathcal{G}_\sigma$ is finite, one can easily check that this condition is satisfied if and only if there exists a cycle in $\mathcal{G}_\sigma$ reachable from the initial state which falls into one of the following two types:
   - *unanswered cycle:* i.e., a cycle having an unanswered request;
   - *unbounded cycle:* i.e., a cycle ρ starting and leading to a state with an odd color c_o such that ρ has a prefix of the form $\rho_1 \cdot \rho_2$, where ρ_2 is a cycle of non-null cost and the odd color c_o is unanswered in the prefix $\rho_1 \cdot \rho_2$ (by pumping the subcycle ρ_2 an arbitrary number of times and by concatenating the resulting cycles, one obtains an

infinite path which does not satisfy the parity-cost condition).

It is not difficult to show that checking the existence of unanswered cycles or unbounded cycles can be done in polynomial time. Hence, the result follows. □

Theorem 3 (Checking summary relevance)**.** *Let $\mathcal{G} = \langle \mathcal{A}, Cost, \Omega, Exit \rangle$ be a partial cost-parity arena over C with $\mathcal{A} = \langle \mathrm{S}, \mathrm{S}_0, \mathrm{S}_1, \mathrm{R}, in \rangle$ and $\mathcal{S}$ a summary of $\mathcal{G}$. Then, one can check in polynomial space whether $\mathcal{S}$ is relevant.*

Proof. We build in polynomial time an *OCPA* $\mathcal{G}_{\mathcal{S}}$ such that there is a winning observation-based memoryless strategy of Player 0 in $\mathcal{G}_{\mathcal{S}}$ from the initial state iff $\mathcal{S}$ is relevant in $\mathcal{G}$. We first construct a partial *OCPA* $\mathcal{G}'$ obtained from $\mathcal{G}$ by extending every state of $\mathcal{G}$ with additional information which keeps tracks of the maximal even color and the maximal *unanswered* odd color visited in the current play-prefix from in and a flag indicating whether such a prefix has cost zero. Formally, $\mathcal{G}' = \langle \mathcal{A}', \mathrm{Cost}', \Omega', Exit', Obs \rangle$ where $\mathcal{A} = \langle \mathrm{S}', \mathrm{S}'_0, \mathrm{S}'_1, \mathrm{R}', in' \rangle$ and:

- $\mathrm{S}' = \mathrm{S} \times C_e \times \tilde{C}_o \times \{0,1\}$, $Exit' = Exit \times C_e \times \tilde{C}_o \times \{0,1\}$, $in' = (in, 0, 0, 0)$, $\Omega'((s, c_e, c_o, d)) = \Omega(s)$, and $((s, c_e, c_o, d), (s', c'_e, c'_o, d')) \in Obs$ *iff* $s = s'$. Moreover, the player of each state (s, c_e, c_o, d) is the player of s in $\mathcal{G}$ if $s \notin Exit$, Player 0 if $s \in Exit$ and $\mathcal{S}(s) = \vdash$, and Player 1 otherwise.
- $((s, c_e, c_o, d), (s', c'_e, c'_o, d') \in \mathrm{E}'$ *iff* (i) $(s, s') \in \mathrm{E}$, (ii) $c'_e = \mathsf{max}_{\succeq_0}(\{c_e, \Omega(s)\})$, (iii) $c'_o = 0$ if $\Omega(s') \in C_e$ and $\Omega(s') \geq c_o$, and $c'_o = \mathsf{min}_{\succeq_0}(\{c_o, \Omega(s)\})$ otherwise, and (iv) $d' = 0$ if $d = 0$ and $\mathrm{Cost}(s, s') = 0$, and $d' = 1$ otherwise;
- $\mathrm{Cost}'((s, c_e, c_o, d), (s', c'_e, c'_o, d')) = \mathrm{Cost}(s, s')$.

Note that by construction, there is a bijection, denoted by *Obs*, between the memoryless strategies σ of Player 0 in $\mathcal{G}$, and the observation-based memoryless strategies of Player 0 in $\mathcal{G}'$. Formally, for each non-terminal state (s, c_e, c_o, d) of $\mathcal{G}'$ controlled by Player 0, $Obs(\sigma)((s, c_e, c_o, d))$ is the unique successor of (s, c_e, c_o, d) having as S-component $\sigma(s)$.

For each $ex \in Exit$, let $Exit'_{ex}$ be the set of exit states of $\mathcal{G}'$ having ex as S-component. $\mathcal{G}_{\mathcal{S}}$ is obtained from $\mathcal{G}'$ by adding for each exit state $ex \in Exit$ such that $\mathcal{S}(ex) \notin \{\vdash, C_o^{max}\}$, a gadget (subgraph) consisting of states controlled by Player 1 that connect the exit states of $\mathcal{G}'$ in $Exit'_{ex}$ with the initial state $in' = (in, 0, 0, 0)$, and an additional terminal state □ which is controlled by Player 0. If $\mathcal{S}(ex) = \{\vdash\}$, then for every strategy σ of Player 0 in $\mathcal{G}$, $\mathcal{S}(\sigma)(ex) = \vdash$

if $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$, and our choice (states in $Exit'_{ex}$ are controlled by Player 0) allows to capture only the non-losing strategies σ of $\mathcal{G}$ for which there is no exit play leading to ex. On the other hand, if $\mathcal{S}(ex) = C_o^{max}$, then for each strategy σ of Player 0 in $\mathcal{G}$, it holds that $\mathcal{S}(\sigma)(ex) \sqsupseteq C_o^{max}$, and accordingly, states in $Exit'_{ex}$ are controlled by Player 1.

Now, we describe the construction of the gadget for ex when $\mathcal{S}(ex) \notin \{\vdash, C_o^{max}\}$, i.e. $\mathcal{S}(ex)$ is of the form $(f, c_{pr}, c_o, c_e^L, c_e^M, c_e^R) \in \{bnd_0, bnd_1, unb\} \times \tilde{C} \times \tilde{C}_o \times C_e \times (C_e \cup \{\bot\}) \times C_e$. We distinguish the following cases:

Case $f = unb$ and $c_o > c_e^R$. Note that for each strategy σ of Player 0, it holds that $value_{\text{Cost}}(\sigma, ex) \succeq_b unb$. The gadget for this case is obtained by adding 2 new states *controlled by Player* 1, namely ex_e and ex_o^R, an additional terminal state $\Box$ controlled by Player 0, and new transitions. State $\Box$ has color 0, state ex_e has the even color $c_o + 1$, and state ex_o^R has color 0 if $c_e^R = 0$, and the odd color $c_e^R - 1$ otherwise. The new transitions have cost 0 and are as follows:

- for each $s = (ex, c'_e, c'_o, d) \in Exit'_{ex}$ such that one of the following bad conditions is satisfied, we add the transition $(s, \Box)$.
  - *Bad conditions:* either (i) $\mathsf{max}(\{c'_o, c'_e\}) \prec_o c_{pr}$, or (ii) $c'_e < c_e^L$, or (iii) $d = 1$, $c_e^M = \bot$, and $c'_e < c_e^R$, or (iv) $d = 1$, $c_e^M \neq \bot$, and $c'_e < c_e^M$.
- the transitions (ex_e, ex_o^R) and (ex_o^R, in'), and for each $s \in Exit'_{ex}$, the transition (s, ex_e).

The transitions having as target state $\Box$ are exploited to capture the strategies σ of Player 0 in $\mathcal{G}$ satisfying the following: (i) $value_{pr}(\sigma, ex) \succeq_o c_{pr}$, (ii) in each exit play ν of σ leading to ex, the maximal even color of ν is at least c_e^L and (iii) if $value_e^M(\sigma, ex) \neq \bot$, then either $c_e^M = \bot$ and $value_e^M(\sigma, ex) \geq c_e^R$, or $c_e^M \neq \bot$ and $value_e^M(\sigma, ex) \geq c_e^M$.

Moreover, given a memoryless strategy σ of Player 0 in $\mathcal{G}$, the chains of transitions (s, ex_e), (ex_e, ex_o^R) and (ex_o^R, in') entering the initial state in', where $s \in Exit'_{ex}$, are responsible of cycles consistent with $Obs(\sigma)$ of the form $\nu \cdot ex_e \cdot ex_o^R \cdot in'$, where ν is an arbitrary exit play of $Obs(\sigma)$ leading to some exit state $s \in Exit'_{ex}$. By concatenating these cycles, one obtains infinite plays consistent with $Obs(\sigma)$ which are winning for Player 0 iff $value(\sigma, ex) \neq C_o^{max}$ (the request cost of σ w.r.t. 0 is finite), $value_o(\sigma, ex) \succeq_0 c_o$, and $value_e^R(\sigma, ex) \geq c_e^R$.

Case $f = unb$ and $c_o \leq c_e^R$. The construction is similar to the previous case, but we replace the even color $c_o + 1$ assigned to the new state ex_e in the previous case with the even color c_e^R. Moreover, for each state $s = (ex, c'_e, c'_o, c, d) \in Exit'_{ex}$ satisfying the bad condition $c'_o > c_o$, we add the transition $(s, \Box)$. In this way, we capture the strategies σ of Player 0 in $\mathcal{G}$ such that $value_o(\sigma, ex) \succeq_0 c_o$.

Case $f = \boldsymbol{bnd}_0$. Hence, $c_e^M = \bot$ and $c_e^R = \mathsf{max}(C_e)$. In this case, the gadget consists of a unique new terminal state $\Box$ having color 0 and controlled by Player 0. Moreover, for each $s = (ex, c'_e, c'_o, d) \in Exit'_{ex}$ such that one of the following bad conditions is satisfied, we add the transition $(s, \Box)$ with cost 0.

- *Bad conditions:* either (i) $\mathsf{max}(\{c'_o, c'_e\}) \prec_o c_{pr}$, or (ii) $c'_o > c_o$, or (iii) $c'_e < c_e^L$, or (iv) $d = 1$.

These transitions are exploited to capture the strategies σ of Player 0 in $\mathcal{G}$ satisfying the following: (i) $value_{pr}(\sigma, ex) \succeq_o c_{pr}$, (ii) $value_o(\sigma, ex) \succeq_0 c_o$, (iii) $value_e^L(\sigma, ex) \geq c_e^L$, and (iv) $value_{\text{Cost}}(\sigma, ex) = bnd_0$ (i.e., each exit play of σ leading to ex has cost 0). Note that for each strategy σ of Player 0 such that $value_{\text{Cost}}(\sigma, ex) = bnd_0$, it holds that $value_e^M(\sigma, ex) = \bot$ and $value_e^R(\sigma, ex) = \mathsf{max}(C_e)$. Hence, correctness of the construction easily follows.

Case $f = \boldsymbol{bnd}_1$. Hence, $c_e^M \neq \bot$ and $c_e^R = \mathsf{max}(C_e)$. The gadget for this case is obtained by adding 2 new states *controlled by Player* 1, namely ex_e^{max} and ex_o^{max}, the terminal state $\Box$ controlled by Player 0, and new transitions. State $\Box$ has color 0, state ex_o^{max} has the maximal color C_o^{max}, and state ex_e^{max} has the even color $C_o^{max} + 1$ (i.e., the smaller even color which answers to a C_o^{max}-request). The new transitions have cost 0 and are as follows:

- the transitions (ex_e^{max}, ex_o^{max}) and (ex_o^{max}, in'), and for each $s \in Exit'_{ex}$, the transition (s, ex_e^{max}).
- for each $s = (ex, c'_e, c'_o, d) \in Exit'_{ex}$ such that one of the following bad conditions is satisfied, we add the transition $(s, \Box)$.
  - *Bad conditions:* either (i) $\mathsf{max}(\{c'_o, c'_e\}) \prec_o c_{pr}$, or (ii) $c'_o > c_o$, or (iii) $c'_e < c_e^L$, or (iv) $d = 1$ and $c'_e < c_e^M$.

Given a memoryless strategy σ of Player 0 in $\mathcal{G}$, the chains of transitions (s, ex_e^{max}), (ex_e, ex_o^{max}) and (ex_o^{max}, in') entering the initial state in', where $s \in Exit'_{ex}$, are responsible of cycles consistent with $Obs(\sigma)$ of the form $\nu \cdot ex_e^{max} \cdot ex_o^{max} \cdot in'$, where ν is an arbitrary exit play of $Obs(\sigma)$ leading to some exit state $s \in Exit'_{ex}$. By concatenating these cycles, one obtains infinite plays consistent with $Obs(\sigma)$. Being $C_o^{max} + 1$ (resp., C_o^{max}) the color of ex_e^{max} (resp., ex_o^{max}), it follows that these additional infinite plays are winning for Player 0 iff the cost value of σ w.r.t. ex is in

Algorithm 1	CheckRel

`Input:` $\mathcal{H} = \langle\langle \mathcal{V}_1, \ldots, \mathcal{V}_n\rangle, \mathrm{Cost}, \Omega\rangle$
`return` CheckRel($\mathcal{H}$,$\mathcal{V}_1$,$\emptyset$)
	`Input:` $\mathcal{H}$, $\mathcal{V}_i$, $\mathcal{S}$ ($\mathcal{S}$ is summary of $\mathcal{H}_i^F$)
	`for each` substitution g of $\mathcal{V}_i$
	    $f \leftarrow$ `true`
	    `for each` box b of $\mathcal{V}_i$ with $g(b) \neq bad$
	        $f \leftarrow f \wedge$ CheckRel($\mathcal{H}$, $\mathcal{V}_{\Upsilon_i(b)}$, $g(b)$)
	    `if` $f = true$ and $\mathcal{S}$ is relevant in *Simplify*($\mathcal{V}_i, g$)
	        `return true`
	`return false`

Figure 2: PSPACE procedure

$\{bnd_0, bnd_1\}$.

Moreover, the transitions to $\Box$ are exploited to capture the strategies σ of Player 0 in $\mathcal{G}$ satisfying the following: (i) $value_{pr}(\sigma, ex) \succeq_o c_{pr}$, (ii) $value_o(\sigma, ex) \succeq_0 c_o$, (iii) $value_e^L(\sigma, ex) \geq c_e^L$, and (iv) either $value_e^M = \bot$, or $value_e^M \geq c_e^M$. Note that for each strategy σ of Player 0 such that $value_{\mathrm{Cost}}(\sigma, ex) \neq unb$, it holds that $value_e^R(\sigma, ex) = \mathsf{max}(C_e)$, and $value_e^M(\sigma, ex) = \bot$ iff $value_{\mathrm{Cost}}(\sigma, ex) = bnd_0$.

By construction, it easily follows that for each memoryless strategy σ of Player 0 in $\mathcal{G}$, σ is non-losing and $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$ *iff* $Obs(\sigma)$ is winning for Player 0 from state in'. Thus, since Obs is a bijection between the memoryless strategies of Player 0 in $\mathcal{G}$ and the observation-based memoryless strategies of Player 0 in $\mathcal{G}_{\mathcal{S}}$, by Theorem 2, Theorem 3 follows. $\Box$

5. Algorithm for solving games on *HCPA*

In this section, by exploiting the summary-gadget arena construction of Section 3, we derive a polynomial space procedure for solving hierarchical cost-parity games. The outline of the procedure, called Algorithm 1, is given in Fig. 2.

Given an *HCPA* $\mathcal{H} = \langle \mathcal{V}, \mathrm{Cost}, \Omega\rangle$ with $\mathcal{V} = \langle \mathcal{V}_1, \ldots, \mathcal{V}_n\rangle$, Algorithm 1 checks that the unique summary (the empty one) in the flat expansion $\mathcal{H}_1^F$ of the highest level sub-arena $\mathcal{V}_1$ (recall that $\mathcal{V}_1$ has no exits) is relevant for $\mathcal{H}_1^F$. For this, it exploits the auxiliary procedure *CheckRel* that takes as input a modular sub-arena $\mathcal{V}_i$ of $\mathcal{H}$ and a summary $\mathcal{S}$ in the flat expansion $\mathcal{H}_i^F$ of $\mathcal{V}_i$

(note that the summaries of $\mathcal{H}_i^F$ can be constructed directly from $\mathcal{V}_i$ without constructing the flat expansion $\mathcal{H}_i^F$). The procedure *CheckRel*$(\mathcal{H}, \mathcal{V}_i, \mathcal{S})$ examines the *substitution mappings* of $\mathcal{V}_i$, i.e., mappings g assigning to each box b of $\mathcal{V}_i$ either a summary in the flat expansion of the lower level sub-arena $\mathcal{V}_{\Upsilon_i(b)}$, or the special symbol *bad*. For the currently precessed substitution g, the procedure recursively checks by nested calls on the lower level modular sub-arenas associated with the boxes b of $\mathcal{V}_i$ (recall that $\Upsilon_i(b) > i$) that the mapping g is *relevant*, i.e., for each box b of $\mathcal{V}_i$ such that $g(b) \neq$ *bad*, $g(b)$ is a relevant summary in the flat expansion of $\mathcal{V}_{\Upsilon_i(b)}$. If the check is positive, then the procedure verifies that $\mathcal{S}$ is a relevant summary of the partial cost-parity arena obtained by applying the operation *Simplify* (simplification) to the modular sub-arena $\mathcal{V}_i$ and the relevant substitution g. If this second check is positive, the procedure returns `true`. The essence of the operation *Simplify* is to replace each box b of $\mathcal{V}_i$ with a copy of the summary-gadget arena associated with the flat expansion $\mathcal{H}^F_{\Upsilon_i(b)}$ and the summary $g(b)$ if $g(b) \neq$ *bad*, and a copy of the bad-gadget arena *BadGad*$(\mathcal{H}^F_{\Upsilon_i(b)})$ if $g(b) =$ *bad*. Note that if $i = n$, g is empty ($\mathcal{V}_n$ has no box) and *Simplify*$(\mathcal{V}_n, g)$ coincides with $\mathcal{H}_n^F$. As we will show (Corollary 1), the relevant summaries in the flat expansion $\mathcal{H}_i^F$ can be obtained by applying the simplification operation to $\mathcal{V}_i$ and the *relevant* substitutions g for $\mathcal{V}_i$. Hence, *CheckRel*$(\mathcal{H}, \mathcal{V}_i, \mathcal{S})$ returns `true` if and only if $\mathcal{S}$ is a relevant summary of $\mathcal{H}_i^F$. Note that the symbol *bad* is used for taking into account situations where the set of relevant summaries associated to the flat expansion of a modular sub-arena is empty. We now formally define the simplification operation.

Definition 7 (Simplification)**.** *Let $\mathcal{H} = \langle \mathcal{V}, Cost, \Omega \rangle$ be an HCPA with $\mathcal{V} = \langle \mathcal{V}_1, \ldots, \mathcal{V}_n \rangle$, $i \in [1, n]$, and b a box of $\mathcal{V}_i$ with $\Upsilon_i(b) = k$. For a summary $\mathcal{S}$ of the flat expansion $\mathcal{H}_k^F$ of $\mathcal{V}_k$, we denote by $Gad_b(\mathcal{H}_k^F, \mathcal{S})$ the copy of the summary-gadget arena $Gad(\mathcal{H}_k^F, \mathcal{S})$ associated with $\mathcal{H}_k^F$ and $\mathcal{S}$ obtained by replacing each state s in $Gad(\mathcal{H}_k^F, \mathcal{S})$ with the copy (b, s). The b-copy $BadGad_b(\mathcal{H}_k^F)$ of the bad-gadget arena $BadGad(\mathcal{H}_k^F)$ for $\mathcal{H}_k^F$ is defined in a similar way. Note that the copies of the states in $\{in_k\} \cup Exit_k$ are states in the flat expansion $\mathcal{H}_i^F$ of $\mathcal{V}_i$.*

Let g be a substitution mapping for $\mathcal{V}_i$. The simplification *Simplify*$(\mathcal{V}_i, g, b)$ of $\mathcal{V}_i$ w.r.t. the substitution g and the box b *is the* partial cost-parity arena *obtained from the flat expansion $\mathcal{H}_i^F$ of $\mathcal{V}_i$ as follows, where $\mathcal{G}$ is the gadget $BadGad_b(\mathcal{H}_k^F)$ if $g(b) = bad$, and $\mathcal{G}$ is the gadget $Gad_b(\mathcal{H}_k^F, g(b))$ otherwise:*

- *all the states in $\mathcal{H}_i^F$ of the form (b, s) such that $s \notin \{in_k\} \cup Exit_k$ are removed together with* *the associated transitions, and all the states in $\mathcal{G}$ are added together with the associated transitions. Moreover, the states (b, s) such that $s \in \{in_k\} \cup Exit_k$ have color 0, i.e., the same color of such states in $\mathcal{G}$.*

The simplification *Simplify*$(\mathcal{V}_i, g)$ of $\mathcal{V}_i$ w.r.t. the substitution g *is the partial cost-parity arena obtained from $\mathcal{H}_i^F$ by applying for each box b of $\mathcal{V}_i$, the simplification operation w.r.t. g and b.*

Note that the arena *Simplify*$(\mathcal{V}_i, g)$ has the same interface as $\mathcal{H}_i^F$ (hence, initial state in_i and set of exit states $Exit_i$) and can be constructed directly from $\mathcal{V}_i$ without constructing the flat expansion $\mathcal{H}_i^F$. By Propositions 3 and 4, we deduce that the *Simplify* operation preserves the set of relevant summaries. In particular, we establish the following result, where for a partial cost-parity arena $\mathcal{G}$, $RS(\mathcal{G})$ is the set of relevant summaries in $\mathcal{G}$.

Theorem 4 (Correctness and completeness of simplification)**.** *Let $\mathcal{H} = \langle \mathcal{V}, Cost, \Omega \rangle$ with $\mathcal{V} = \langle \mathcal{V}_1, \ldots, \mathcal{V}_n \rangle$ be an HCPA, $i \in [1, n]$, and b a box of $\mathcal{V}_i$ with $\Upsilon_i(b) = k$. Then:*

1. *for each $\mathcal{S} \in RS(\mathcal{H}_i^F)$, $\mathcal{S} \in RS(Simplify(\mathcal{V}_i, g, b))$ for some* relevant *substitution g for $\mathcal{V}_i$ (completeness);*
2. *for each* relevant *substitution g for $\mathcal{V}_i$, $RS(Simplify(\mathcal{V}_i, g, b)) \subseteq RS(\mathcal{H}_i^F)$ (correctness).*

The proof of Property 1 (resp., Property 2) in Theorem 4 is postponed to Section 6 (resp., Section 7). By Theorem 4, we easily obtain the following corollary.

Corollary 1. *Let $\mathcal{H} = \langle \mathcal{V}, Cost, \Omega \rangle$ with $\mathcal{V} = \langle \mathcal{V}_1, \ldots, \mathcal{V}_n \rangle$ be an HCPA and $i \in [1, n]$. Then, for each summary $\mathcal{S}$ of $\mathcal{H}_i^F$, $\mathcal{S} \in RS(\mathcal{H}_i^F)$ if and only if $\mathcal{S} \in RS(Simplify(\mathcal{V}_i, g))$ for some* relevant *substitution g for $\mathcal{V}_i$.*

Proof. By a straightforward induction on the number of boxes b of $\mathcal{V}_i$, the result easily follows from Theorem 4 and the fact that the simplification *Simplify*$(\mathcal{V}_i, g, b)$ corresponds to the flat expansion of a modular sub-arena $\mathcal{V}_i'$ in an *HCPA* of the form $\mathcal{H} = \langle \mathcal{V}', \mathrm{Cost}', \Omega' \rangle$, where $\mathcal{V} = \langle \mathcal{V}_1, \ldots, \mathcal{V}_{i-1}, \mathcal{V}_i', \mathcal{V}_{i+1}, \ldots, \mathcal{V}_n \rangle$ and the boxes of $\mathcal{V}_i'$ consist of the boxes of $\mathcal{V}_i$ distinct from b. □

By Corollary 1, we deduce the main result of this paper.

Theorem 5. *Solving hierarchical cost-parity games is* PSPACE*-complete.*

Proof. PSPACE-hardness directly follows from PSPACE-hardness of solving hierarchical parity games [5]. For the matching upper bound, we show that Algorithm 1 solves the considered problem in polynomial space. Let $\mathcal{H} = \langle \mathcal{V}, \mathrm{Cost}, \Omega \rangle$ be an *HCPA* with $\mathcal{V} = \langle \mathcal{V}_1, \ldots, \mathcal{V}_n \rangle$. First, we show that Algorithm 1 is correct, i.e. it accepts the input $\mathcal{H}$ iff Player 0 wins in $\mathcal{H}_1^F$. For this, we prove that for all $i \in [1, n]$ and summaries $\mathcal{S}$ in the flat expansion $\mathcal{H}_i^F$ of $\mathcal{V}_i$, the procedure *CheckRel* with input $\mathcal{V}_i$ and $\mathcal{S}$ returns `true` if and only if $\mathcal{S}$ is a *relevant* summary of $\mathcal{H}_i^F$. Hence, the result follows. The proof is by induction on $n - i$. For the base case where $i = n$, $\mathcal{V}_n$ has no box and the unique substitution mapping for $\mathcal{V}_n$ is the empty one. Hence, *Simplify*$(\mathcal{V}_n, \emptyset)$ coincides with $\mathcal{H}_n^F$ and by construction, *CheckRel*$(\mathcal{H}, \mathcal{V}_n, \mathcal{S})$ returns `true` iff $\mathcal{S}$ is a relevant summary of *Simplify*$(\mathcal{V}_n, \emptyset)$. Now, assume that $i < n$. By construction and the induction hypothesis, *CheckRel*$(\mathcal{H}, \mathcal{V}_i, \mathcal{S})$ return `true` iff there is a relevant substitution g for $\mathcal{V}_i$ such that $\mathcal{S}$ is a relevant summary of *Simplify*$(\mathcal{V}_i, g)$ iff (by Corollary 1) $\mathcal{S}$ is a relevant summary of $\mathcal{H}_i^F$.

Finally, we show that Algorithm 1 can be implemented in polynomial space. Recall that for the given input $\mathcal{H}$, an index $1 \leq i \leq n$, and a summary $\mathcal{S}$ of $\mathcal{H}_i^F$, the procedure *CheckRel*$(\mathcal{H}, \mathcal{V}_i, \mathcal{S})$ goes over all the substitutions g for $\mathcal{V}_i$. We can assume that the set of substitutions for $\mathcal{V}_i$ is lexicographically ordered and there is a polynomial space procedure that given an order index i (encoded in binary) returns the i^{th} substitution. For each such substitution g, the procedure checks that g is relevant. This check is done by recursive calls to *CheckRel*$(\mathcal{H}, \mathcal{V}_{\mathrm{Y}_i(b)}, g(b))$ for each box b of $\mathcal{V}_i$ such that $g(b) \neq$ *bad*. If the check is positive, the procedure additionally verifies that $\mathcal{S}$ is a relevant summary of the simplification *Simplify*$(\mathcal{V}_i, g)$. Thus, the procedure *CheckRel*$(\mathcal{H}, \mathcal{V}_i, \mathcal{S})$ needs to remember the current substitution g which requires space $O(|\mathcal{H}|)$. Moreover, since the simplification *Simplify*$(\mathcal{V}_i, g)$ has size $O(|\mathcal{H}|)$ (recall that the gadget arena for a specific summary has size linear in the number of exits), by Theorem 3, the memory required for checking that $\mathcal{S}$ is relevant in *Simplify*$(\mathcal{V}_i, g)$ is polynomial in $|\mathcal{H}|$. Hence, since the depth of the recursive calls to *CheckRel* is at the most the nesting depth of the hierarchical system, membership in PSPACE of solving hierarchical cost-parity games follows. This concludes the proof of Theorem 5. □

6. Completeness of simplification

In this section, we provide a proof of the following result which corresponds to Property 1 of Theorem 4.

Theorem 6. *[Completeness of simplification] Let $\mathcal{H} = \langle \mathcal{V}, Cost, \Omega \rangle$ be an HCPA with $\mathcal{V} = \langle \mathcal{V}_1, \dots, \mathcal{V}_N \rangle$, $i \in [1, N]$, b a box of $\mathcal{V}_i$, and $\mathcal{S}$ a relevant summary of $\mathcal{H}_i^F$. Then, there exists a relevant substitution mapping g for $\mathcal{V}_i$ such that $\mathcal{S}$ is a relevant summary of the simplification Simplify$(\mathcal{V}_i, g, b)$.*

In the rest of this section, by exploiting Proposition 3, we provide a proof of Theorem 6. We fix an *HCPA* $\mathcal{H} = \langle \mathcal{V}, \mathrm{Cost}, \Omega \rangle$ with $\mathcal{V} = \langle \mathcal{V}_1, \dots, \mathcal{V}_N \rangle$, $i \in [1, N]$, a box b of $\mathcal{V}_i$ with $\mathrm{Y}_i(b) = k$, and a relevant summary $\mathcal{S}$ of $\mathcal{H}_i^F$.

Since $\mathcal{S}$ is a relevant summary of $\mathcal{H}_i^F$, there exists a non-losing memoryless strategy σ in $\mathcal{H}_i^F$ such that $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$. The strategy σ induces a memoryless strategy σ_b of Player 0 in $\mathcal{H}_k^F$. Formally, for each non-terminal state s of Player 0 in $\mathcal{H}_k^F$, $\sigma_b(s) = \sigma((b, s))$.

We first consider the case where σ_b is *not* a non-losing strategy for Player 0 in $\mathcal{H}_k^F$ (from state in_k).

Lemma 1. *If the strategy σ_b is* not *a non-losing strategy in $\mathcal{H}_k^F$, then for each relevant substitution mapping g for $\mathcal{V}_i$ such that $g(b) = bad$, $\mathcal{S}$ is a relevant summary of the simplification Simplify$(\mathcal{V}_i, g, b)$.*

Proof. Assume that σ_b is not a non-losing strategy for Player 0 in $\mathcal{H}_k^F$. Being σ non-losing and memoryless, it easily follows that every play of $\mathcal{H}_i^F$ starting from in_i which is consistent with σ cannot visit states of the form (b, s). In particular, every exit play of σ cannot visit states of the form (b, s). By Definition 6, there is a unique strategy of Player 0 in $BadGad_b(\mathcal{H}_k^F)$. Hence, the strategy σ induces a unique memoryless strategy σ_{simp} in *Simplify*$(\mathcal{V}_i, g, b)$ which coincides with σ on the states which are not of the form (b, s). By the above observations, σ_{simp} is a non-losing strategy (of Player 0) in *Simplify*$(\mathcal{V}_i, g, b)$ and $\mathcal{S}(\sigma_{simp}) = \mathcal{S}(\sigma)$. Hence, being $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$, we obtain that $\mathcal{S}$ is a relevant summary of *Simplify*$(\mathcal{V}_i, g, b)$, and the result follows. □

By Lemma 1, we can assume that the memoryless strategy σ_b of Player 0 in $\mathcal{H}_k^F$ is non-losing. In the following, we denote by $\mathcal{S}_b$ the relevant summary associated with σ_b, i.e. $\mathcal{S}_b = \mathcal{S}(\sigma_b)$, and by g any relevant substitution mapping for $\mathcal{V}_i$ such that $g(b) = \mathcal{S}_b$. By construction, in the gadget arena $Gad(\mathcal{H}_k^F, \mathcal{S}_b)$, there is exactly one strategy of Player 0. Hence, the strategy σ and the summary $\mathcal{S}_b$ uniquely induce a memoryless strategy σ_{simp} of Player 0 in $Simplify(\mathcal{V}_i, g, b)$, defined as follows for each non-terminal state s of $Simplify(\mathcal{V}_i, g, b)$ controlled by Player 0:

- if s is a state of the form (b, s') and s' is not an exit state of $\mathcal{H}_k^F$, then $\sigma_{simp}((b, s')) = (b, s'')$, where s'' is the unique successor of s' in $Gad(\mathcal{H}_k^F, \mathcal{S}_b)$;
- otherwise, $\sigma_{simp}(s) = \sigma(s)$.

We demonstrate in Lemmata 2 and 3 below that σ_{simp} is a non-losing strategy of $Simplify(\mathcal{V}_i, g, b)$ such that $\mathcal{S}(\sigma_{simp}) \sqsupseteq \mathcal{S}(\sigma)$. Hence, since $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}$, we obtain that $\mathcal{S}$ is a relevant summary of the simplification $Simplify(\mathcal{V}_i, g, b)$. This result together with Lemma 1 completes the proof of Theorem 6.

In the following, in order to simplify the notation, we use some shorthands: $\mathcal{H}_b$ is for $\mathcal{H}_k^F$, $\mathcal{H}^F$ is for $\mathcal{H}_i^F$, $\mathcal{H}_{simp}^F$ is for $Simplify(\mathcal{V}_i, g, b)$, $Gad(\mathcal{S}_b)$ is for $Gad(\mathcal{H}_k^F, \mathcal{S}_b)$, $Gad_b(\mathcal{S}_b)$ is for $Gad_b(\mathcal{H}_k^F, \mathcal{S}_b)$, in is for in_i, and in_b is for in_k. Moreover, we identify $Gad(\mathcal{S}_b)$ (resp., $Gad_b(\mathcal{S}_b)$) with the unique strategy of Player 0 in $Gad(\mathcal{S}_b)$ (resp., in $Gad_b(\mathcal{S}_b)$). Moreover, we exploit the following definitions:

- A state s of $\mathcal{H}_{simp}^F$ is called *b-state* if it is of the form (b, s') (i.e., it is a state in $Gad_b(\mathcal{S}_b)$), and is called *context-state* otherwise.
- A *b-segment* is an exit play of $Gad_b(\mathcal{S}_b)$.
- A *b-play of* σ is a finite play ν of $\mathcal{H}^F$ consistent with σ visiting only b-states, starting from (b, in_b), and leading to a state of the form (b, s) for some exit state s of $\mathcal{H}_b$. Note that if we remove the b-component to every state occurring in ν, we obtain an exit play of σ_b.

Lemma 2. *σ_{simp} is a non-losing strategy of $\mathcal{H}_{simp}^F$.*

Proof. Let π be a play of $\mathcal{H}_{simp}^F$ from in which is consistent with the strategy σ_{simp} such that π is not a finite play leading to an exit state (i.e., an exit node in $Exit_i$). We need to show that π is winning for Player 0. First, assume that π is finite. Then, by construction of $Gad(\mathcal{S}_b)$, either π leads to context-state, or π leads to a state in $Gad_b(\mathcal{S}_b)$ which is controlled by Player 1, or π leads to an exit state of $Gad_b(\mathcal{S}_b)$. Since $\mathcal{S}_b = \mathcal{S}(\sigma_b)$, for each b-segment ρ in π, there is a b-play of σ leading to the last state of ρ. Thus, by definition of σ_{simp} and since σ is non-losing, we obtain that π is winning for Player 0. Now, assume that π is infinite. If there is a suffix of

π which visits only context-states, then by construction, such a suffix is consistent with the non-losing strategy σ. Hence, π is evidently winning for Player 0. On the other hand, if π gets trapped in the gadget arena $Gad_b(\mathcal{S}_b)$, there is a suffix of π which is consistent with the unique non-losing strategy of Player 0 in $Gad_b(\mathcal{S}_b)$. Hence, also in this case, π is winning for Player 0. It remains to consider the most difficult case, when π visits infinitely often b-segments. Hence, π can be factorized in the form:

$$\pi = \pi'_0 \cdot \pi_0 \cdot \pi'_1 \cdot \pi_1 \ldots$$

such that for all $n \geq 0$, π'_n visits only context-states, while π_n is a b-segment. Moreover, since the cost-parity winning condition is prefix-independent, without loss of generality, we can assume that for each exit state (b, s) of $Gad_b(\mathcal{S}_b)$ occurring in π, there are infinitely many occurrences of (b, s) in π (in other terms, if for some $\ell \geq 0$, the b-segment π_ℓ leads to (b, s), then there are infinitely many n such that π_n lead to (b, s)). First, we observe that for each of such states (b, s), $\mathcal{S}_b(s) \neq C_o^{max}$. Otherwise, since $\mathcal{S}_b = \mathcal{S}(\sigma_b)$, we can construct, starting from π, an infinite play of $\mathcal{H}^F$ from in which is consistent with σ and is losing for Player 0. We need to show that π is winning for Player 0. We assume the contrary and derive a contradiction. Hence, one of the following two conditions is fulfilled:

(1) There are infinitely many unanswered requests along π.

(2) There are infinitely many answered requests along π whose set of associated delays is unbounded.

First, let us examine case (1). Hence, there is a state s_o with odd color c_o such that for infinitely many positions n, n is an unanswered c_o-request along π associated with state s_o. We show that there exists a play consistent with σ which is losing for Player 0. Fix a b-segment π_ℓ of π leading to an exit state (b, s), and let c_ℓ be the maximal color in π_ℓ. Since $\mathcal{S}_b = \mathcal{S}(\sigma_b)$ and $\mathcal{S}_b(s) \neq C_o^{max}$, by Proposition 3, $value_{pr}(\sigma_b, s) = value_{pr}(\sigma_S, s)$, where σ_S is the unique strategy of Player 0 in $Gad(\mathcal{S}_b)$. Hence, there exists a b-play ν_ℓ of σ leading to (b, s) whose maximal color c'_ℓ satisfies: $c'_\ell \preceq_0 c_\ell$. Note that if π_ℓ visits state s_0 (s_0 is a b-state), then c'_ℓ is an odd color such that $c'_\ell \geq c_o$. Otherwise, if c'_ℓ is even, then c'_ℓ cannot answer to the odd color c_o (i.e., $c'_\ell < c_o$). Hence, by replacing each b-segment π_ℓ in π with ν_ℓ, we obtain a play consistent with σ which contains infinitely many unanswered requests, and the result follows.

Now, let us examine case (2). By hypothesis, there must exist an infinite set $\{[q_n, r_n]\}_{n\geq 0}$ of non-overlapping position-intervals such that for all $n \geq 0$:

- the cost of $\pi[q_n, r_n]$ is at least n;
- $q_n < q_{n+1}$;
- q_n (resp., r_n) is a request (resp., response) along π, and r_n is the smallest response answering the request q_n along π.

We show that it is possible to replace each b-segment π_ℓ in π by a b-play of σ leading to the last state of π_ℓ in such a way that the resulting infinite sequence is a play consistent with σ which is losing for Player 0. Let us consider a b-segment π_ℓ along π, and let (b, s) be the last state of π_ℓ. We distinguish the following cases:

- π_ℓ does not overlap any segment $\pi[q_n, r_n]$: we replace π_ℓ with an arbitrary b-play of σ leading to (b, s) (note that such a b-play exists).
- there is a segment $\pi[q_n, r_n]$ for some $n \geq 0$ such that q_n corresponds to a request in π_ℓ with odd color c_o: by construction of $Gad(\mathcal{S}_b)$, $\pi[q_n + 1] = (b, s)$ and each edge in $Gad_b(\mathcal{S}_b)$ leading to (b, s) has zero-cost. Since $\mathcal{S}_b = \mathcal{S}(\sigma_b)$ and $\mathcal{S}(\sigma_b)(s) \neq C_o^{max}$, by Proposition 3, the odd value of σ_b w.r.t. s coincides with the odd value of $Gad(\mathcal{S}_b)$ w.r.t. s. Hence, there is a b-play ν of σ leading to (b, s) having a c'_o-unanswered request for some $c'_o \geq c_o$. We replace the b-segment π_ℓ with ν.
- there is a segment $\pi[q_n, r_n]$ for some $n \geq 0$ which strictly contains π_ℓ. Let c_e be the maximal even color in π_ℓ. Note that c_e is strictly smaller than the even color associated with the response r_n. By Proposition 3(1), there exists a b-play ν of σ leading to (b, s) such that either (i) the maximal even color in ν is at most c_e and the cost of ν is at least the cost of π_ℓ, or (ii) there is an even color $c'_e \in C_e$ such that $c'_e \leq c_e + 2$ and the c'_e-response cost of ν is at least the cost of $\pi[q_n, r_n]$. We replace the b-segment π_ℓ with ν.
- there is a segment $\pi[q_n, r_n]$ for some $n \geq 0$ such that r_n corresponds to a response in π_ℓ and $\pi[q_n, r_n]$ is not a sub-path of π_ℓ. Let c_e (resp., c_o) be the color associated with the response r_n (resp., the request q_n) and c_e^{min} be the smallest even color answering to the request q_n. Moreover, let Π_s be the set of b-segments $\pi_{\ell'}$ leading to state (b, s) satisfying the same condition as π_ℓ, i.e. such that there is a segment $\pi[q_{n'}, r_{n'}]$ so that $r_{n'}$ corresponds to a response in $\pi_{\ell'}$, $\pi[q_{n'}, r_{n'}]$ is not a sub-path of $\pi_{\ell'}$, and c_e (resp., c_o) is the color associated with the response $r_{n'}$ (resp., the request $r_{n'}$). Let $\mathrm{Cost}_{c_e} := \{\mathrm{ResCost}(\pi_{\ell'}, c_e) \mid \pi_{\ell'} \in \Pi_s\}$,

i.e., the set of c_e-response costs of the b-segments in Π_s. If Cost_{c_e} is finite, we can safely replace the b-segment π_ℓ with an arbitrary b-play of σ leading to (b, s) (note that such a b-play exists). Now, assume that $\text{Cost}_{c_e} = \infty$. By construction, for each $\pi_{\ell'} \in \Pi_s$, the smaller response k in $\pi_{\ell'}$ answering the odd color c_o has color c_e. Thus, since $\text{Cost}_{c_e} = \infty$, it follows that $\text{ResCost}_e(\sigma_S, s, c_e^{\mathsf{min}}) = \infty$, where σ_S is the unique strategy of Player 0 in $Gad(\mathcal{S}_b)$. Since $\mathcal{S}_b = \mathcal{S}(\sigma_b)$ and $\mathcal{S}(\sigma_b)(s) \notin \{\vdash, C_o^{max}\}$, by applying Proposition 3(2), it holds that $\text{ResCost}(\sigma_b, s, c_e^{\mathsf{min}}) = \infty$. Hence, there exists a b-play ν of σ leading to (b, s) such that the c_e^{min}-response cost in ν is at least the cost of $\pi[q_n, r_n]$. We replace the b-segment π_ℓ with ν.

Thus, we obtain a play consistent with σ which is losing for Player 0, which is a contradiction. This concludes the proof of Lemma 2. □

It remains to show that $\mathcal{S}(\sigma_{simp}) \sqsupseteq \mathcal{S}(\sigma)$.

Lemma 3. *Let ex be an exit state of $\mathcal{H}_F$. Then $\mathcal{S}(\sigma_{simp})(ex) \sqsupseteq \mathcal{S}(\sigma)(ex)$.*

Proof. If either $\mathcal{S}(\sigma)(ex) = C_o^{max}$ or $\mathcal{S}(\sigma_{simp})(ex) = \vdash$, the result trivially follows. Moreover, note that since $\mathcal{S}_b = \mathcal{S}(\sigma_b)$, for each b-segment leading to (b, s), there is a b-play of σ leading to (b, s). Hence, if $\mathcal{S}(\sigma)(ex) = \vdash$, then $\mathcal{S}(\sigma_{simp})(ex) = \vdash$ as well. Hence, we can assume that $\mathcal{S}(\sigma)(ex) \notin \{\vdash, C_o^{max}\}$ and $\mathcal{S}(\sigma_{simp})(ex) \neq \vdash$. First, we show that $\mathcal{S}(\sigma_{simp})(ex) \neq C_o^{max}$ as well. We assume on the contrary that $\mathcal{S}(\sigma_{simp})(ex) = C_o^{max}$, and derive a contradiction. We exploit the fact that since $\mathcal{S}(\sigma)(ex) \neq C_o^{max}$ and $\mathcal{S}_b = \mathcal{S}(\sigma_b)$, for each exit play ν of σ_{simp} leading to ex, there is no exit state (b, s) of $Gad_b(\mathcal{S}_b)$ occurring in ν such that $\mathcal{S}_b(s) = C_o^{max}$. Since $\mathcal{S}(\sigma_{simp})(ex) = C_o^{max}$, one of the following two conditions is fulfilled.

(1) *either* there is an exit play ν of σ_{simp} leading to ex having a C_o^{max}-request,

(2) *or* there must exist an infinite set of exit plays $\{\nu_n\}_{n\geq 0}$ of σ_{simp} leading to ex and an infinite set $\{[q_n, r_n]\}_{n\geq 0}$ of position-intervals such that for all $n \geq 0$:
   - either (i) q_n is an unanswered request in ν_n and r_n is the last position of ν_n, or (ii) q_n (resp., r_n) is a request (resp., response) along ν_n, and r_n is the smallest response answering the request q_n along ν_n.
   - the cost of $\nu_n[q_n, r_n]$ is at least n.

First, let us examine case (1). Since there is no exit state (b, s) of $Gad_b(\mathcal{S}_b)$ occurring in ν such that $\mathcal{S}_b(s) = C_o^{max}$, we deduce that the C_o^{max}-request in ν is associated with a context-state. It follows that there exists an exit play of σ leading to ex having a C_o^{max}-request, which contradicts the hypothesis $\mathcal{S}(\sigma)(ex) \neq C_o^{max}$. Now, let us examine case (2). Since for each exit state (b, s) of $Gad_b(\mathcal{S}_b)$ along ν_n, $\mathcal{S}_b(s) \neq C_o^{max}$, by proceeding as in the part of the proof of Lemma 2 concerning the treatment of Condition (2), we deduce the following.

- For each $n \geq 0$, it is possible to replace each b-segment ρ in ν_n by a b-play of σ leading to the last state of ρ in such a way that the resulting sequence ν'_n is an exit play of σ leading to ex, and the following holds: there are two positions q'_n and r'_n in ν_n with $q'_n < r'_n$ such that
  - the cost of $\nu'_n[q'_n, r'_n]$ is at least n;
  - either (i) q'_n is an unanswered request in ν'_n and r'_n is the last position of ν'_n, or (ii) q'_n (resp., r'_n) is a request (resp., response) along ν'_n, and r'_n is the smallest response answering the request q'_n along ν'_n.

Hence, we obtain that $\mathcal{S}(\sigma)(ex) = C_o^{max}$, which contradicts the hypothesis. Therefore, in the rest of the proof, we can assume that $\mathcal{S}(\sigma)(ex) \notin \{\vdash, C_o^{max}\}$ and $\mathcal{S}(\sigma_{simp})(ex) \notin \{\vdash, C_o^{max}\}$. Then, the result directly follows from Claims 1–6 below.

Claim 1. $value_{\mathrm{Cost}}(\sigma_{simp}, ex) \succeq_b value_{\mathrm{Cost}}(\sigma, ex)$.

Proof of Claim 1. Since $\mathcal{S}(\sigma_{simp})(ex) \neq C_o^{max}$, for each exit play ν of σ_{simp} leading to ex, there is no exit state (b, s) of $Gad_b(\mathcal{S}_b)$ occurring in ν such that $\mathcal{S}_b(s) = C_o^{max}$. Being $\mathcal{S}_b = \mathcal{S}(\sigma_b)$, by Proposition 3(1), it holds that for each b-segment ρ leading to (b, s) such that $\mathcal{S}_b(s) \neq C_o^{max}$, there is a b-play of σ leading to (b, s) whose cost is at least the cost of ρ. Hence, the result trivially follows. □

Claim 2. $value_{pr}(\sigma_{simp}, ex) \succeq_0 value_{pr}(\sigma, ex)$.

Proof of Claim 2. Since $\mathcal{S}_b = \mathcal{S}(\sigma_b)$, by Proposition 3, for each b-segment ρ leading to (b, s) such that $\mathcal{S}_b(s) \neq C_o^{max}$, the parity value of σ_b w.r.t. s coincides with the parity value of $Gad(\mathcal{S}_b)$ w.r.t. s. It follows that for each b-segment ρ leading to (b, s) such that $\mathcal{S}_b(s) \neq C_o^{max}$, there is a b-play of σ leading to (b, s) whose maximal color c' satisfies $c' \preceq_0 c$. Hence, the result easily follows. □

Claim 3. $value_o(\sigma_{simp}, ex) \leq value_o(\sigma, ex)$.

Proof of Claim 3. If $value_o(\sigma_{simp}, ex) = 0$, the result is obvious. Otherwise, $value_o(\sigma_{simp}, ex) = c_o$

for some odd color c_o. We assume that $c_o > value_o(\sigma, ex)$, and derive a contradiction. Thus, since $value_o(\sigma_{simp}, ex) = c_o$, there must exist an exit play ν of σ_{simp} leading to ex having an unanswered c_o-request. Assume that the maximal c_o-request of ν is inside a b-segment ρ_{c_o} of ν (the other case, where the maximal c_o-request of ν is associated with a context-state, being simpler). Since $\mathcal{S}(\sigma_{simp})(ex) \neq C_o^{max}$, there is no exit state (b, s) of $Gad_b(\mathcal{S}_b)$ along ν such that $\mathcal{S}_b(s) = C_o^{max}$. We show how to construct, starting from ν, an exit play ν' of σ leading to ex having an unanswered c'_o-request with $c'_o \geq c_o$. Let us consider a b-segment ρ along ν leading to a state (b, s). We distinguish the following cases:

- ρ strictly precedes the b-segment ρ_{c_o}: we replace ρ with any b-play of σ leading to (b, s) (note that such a b-play exists).
- ρ corresponds to ρ_{c_o}: we replace ρ with any b-play of σ leading to (b, s) having an unanswered c'_o-request for some $c'_o \geq c_o$. Since $\mathcal{S}_b(s) \neq C_o^{max}$ and $\mathcal{S}_b = \mathcal{S}(\sigma_b)$, the existence of such a b-play in ensured by Proposition 3.
- ρ strictly follows the b-segment ρ_{c_o}: let c_e be the maximal even color occurring in ρ. Since $\mathcal{S}_b(s) \neq C_o^{max}$ and $\mathcal{S}_b = \mathcal{S}(\sigma_b)$, by Proposition 3(1), the following holds: either (i) there exists a b-play of σ leading to s whose maximal even color is c_e, or (ii) for some even color $c'_e \leq c_e + 2$, $\text{ResCost}(\sigma_b, s, c'_e) = \infty$. Condition (ii) cannot hold. Otherwise, we deduce that the request-cost of σ with respect to ex is ∞, which would contradict our assumption that $\mathcal{S}(\sigma)(ex) \neq C_o^{max}$. Hence, we replace ρ with any b-play of σ satisfying Condition (i).

Therefore, we obtain an exit play ν' of σ leading to ex having an unanswered c'_o-request for some $c'_o \geq c_o$. This is a contradiction since we have assumed that $c_o > value_o(\sigma, ex)$. □

Claim 4. $value_e^L(\sigma_{simp}, ex) \geq value_e^L(\sigma, ex)$.

Proof of Claim 4. We assume that $value_e^L(\sigma_{simp}, ex) < value_e^L(\sigma, ex)$, and derive a contradiction. Let $c_e^L = value_e^L(\sigma, ex)$. By hypothesis, one of the following two conditions holds.

1. either there is an exit play ν of σ_{simp} leading to ex such that the maximal even color in ν is strictly smaller than c_e^L,
2. or there is a family $\{\nu_n\}_{n \geq 0}$ of exit plays of σ_{simp} leading to ex and an infinite sequence of natural numbers $\{r_n\}_{n \geq 0}$ such that: (i) r_n is the smaller response in ν_n associated with an even color $c_e \geq c_e^L$, and (ii) the response-cost of r_n in ν_n is at least n.

First, let us examine case (1). By hypothesis, for each b-segment ν_b along ν, the maximal even color in ν is strictly smaller than c_e^L. Since $\mathcal{S}(\sigma_{simp})(ex) \neq C_o^{max}$, for each exit state (b, s) of $Gad_b(\mathcal{S}_b)$ occurring in ν, $\mathcal{S}_b(s) \neq C_o^{max}$. Thus, since $\mathcal{S}(\sigma_b) = \mathcal{S}_b$, by applying Proposition 3, it holds that for each exit state (b, s) of $Gad_b(\mathcal{S}_b)$ occurring in ν, $value_e^L(\sigma_b, s) < c_e^L$. Hence, we easily obtain that $value_e^L(\sigma, ex) < c_e^L$, which is a contradiction.

Now, let us examine case (2). Fix $n \geq 0$. It suffices to show that it is possible to replace each b-segment ρ in ν_n by a b-play of σ leading to the last state of ρ in such a way that the resulting sequence ν'_n is an exit play of σ leading to ex satisfying the following: either (i) the maximal even color in ν'_n is strictly smaller than c_e^L, or (ii) the response-cost of any c_e-response in ν'_n with $c_e \geq c_e^L$ is at least n. Since for each exit state (b, s) of $Gad_b(\mathcal{S}_b)$ along ν_n, $\mathcal{S}_b(s) \neq C_o^{max}$, we proceed as in the part of the proof of Lemma 2 concerning the third and fourth case in the treatment of Condition (2). □

Claim 5. If $value_e^M(\sigma_{simp}, ex) \neq \bot$, then the following holds:

- *either* $value_e^M(\sigma, ex) \neq \bot$ and $value_e^M(\sigma_{simp}, ex) \geq value_e^M(\sigma, ex)$,
- *or* $value_e^M(\sigma, ex) = \bot$ and $value_e^M(\sigma_{simp}, ex) \geq value_e^R(\sigma, ex)$.

Proof of Claim 8. Let $value_e^M(\sigma_{simp}, ex) = c_e \in C_e$. Hence, $\text{Cost}_e(\sigma_{simp}, ex, c_e) \in \mathbb{N} \setminus \{0\}$ and there is an exit play of σ_{simp} leading to ex having a non-null cost and whose maximal even color is c_e. Let ν be any exit play of σ_{simp} leading to ex whose maximal even color is c_e. Since $\mathcal{S}(\sigma_{simp})(ex) \neq C_o^{max}$, for each exit state (b, s) of $Gad_b(\mathcal{S}_b)$ occurring in ν, $\mathcal{S}_b(s) \neq C_o^{max}$. Moreover, for each b-segment ρ occurring in ν and leading to a state (b, s), it holds that $\text{Cost}_e(\sigma_S, s, c'_e) \neq \infty$, where σ_S is the unique strategy of Player 0 in $Gad(\mathcal{S}_b)$ and c'_e is the maximal even color in ρ. Since $\mathcal{S}_b = \mathcal{S}(\sigma_b)$, by applying Proposition 3(1), for such a b-segment ρ, there exists a b-play ρ' of σ leading to (b, s) whose maximal color is at most c'_e and whose cost is at least the cost of ρ. Hence, there exists an exit play ν' of σ leading to ex whose maximal even color is at most c_e and such that $\text{Cost}(\nu') \geq \text{Cost}(\nu)$. If follows that $\text{Cost}_e(\sigma, ex, c_e) \in (\mathbb{N} \setminus \{0\}) \cup \{\infty\}$. If $value_e^M(\sigma, ex) \neq \bot$ then $value_e^M(\sigma, ex) \leq c_e$. Otherwise, we have that $c_e \geq value_e^R(\sigma, ex)$, and the result follows. □

Claim 6. $value_e^R(\sigma_{simp}, ex) \geq value_e^R(\sigma, ex)$.

Proof of Claim 6. Let $c_e^R = value_e^R(\sigma, ex)$ and $\mathcal{K}_{simp}^F$ (resp., $\mathcal{K}^F$) be the partial cost-parity arena obtained from $\mathcal{H}_{simp}^F$ (resp., $\mathcal{H}^F$) by replacing the color of the exit state ex with c_e^R. Evidently, the even-*left* value of σ w.r.t. ex' in $\mathcal{K}^F$ is c_e^R. Thus, it suffices to show that the even-left value of strategy σ_{simp} w.r.t. ex in $\mathcal{K}^F_{simp}$ is at least the even-left value of σ w.r.t. ex in $\mathcal{K}^F$. The proof of this is similar to the proof of Claim 4.

This concludes the proof of Lemma 3. □

7. Correctness of simplification

In this section, we provide a proof of the following result which corresponds to Property 2 of Theorem 4.

Theorem 7. *[Correctness of simplification] Let $\mathcal{H} = \langle \mathcal{V}, Cost, \Omega \rangle$ be an HCPA with $\mathcal{V} = \langle \mathcal{V}_1, \ldots, \mathcal{V}_N \rangle$, $i \in [1, N]$, b a box of $\mathcal{V}_i$, g a relevant substitution mapping for $\mathcal{V}_i$, and $\mathcal{S}$ a relevant summary of the simplification Simplify$(\mathcal{V}_i, g, b)$. Then, $\mathcal{S}$ is a relevant summary of $\mathcal{H}^F_i$.*

In the rest of this section, by exploiting Proposition 4, we provide a proof of Theorem 7. We fix an *HCPA* $\mathcal{H} = \langle \mathcal{V}, \mathrm{Cost}, \Omega \rangle$ with $\mathcal{V} = \langle \mathcal{V}_1, \ldots, \mathcal{V}_N \rangle$, $i \in [1, N]$, a box b of $\mathcal{V}_i$ with $\mathrm{Y}_i(b) = k$, a relevant substitution mapping g for $\mathcal{V}_i$, and a relevant summary $\mathcal{S}$ of the simplification *Simplify*$(\mathcal{V}_i, g, b)$. We first consider the case where $g(b) = bad$.

Lemma 4. *If $g(b) = bad$, then $\mathcal{S}$ is a relevant summary of $\mathcal{H}^F_i$.*

Proof. Let $g(b) = bad$ and $\mathcal{S}$ be a relevant summary of *Simplify*$(\mathcal{V}_i, g, b)$. Therefore, there exists a non-losing memoryless strategy σ_{simp} of *Simplify*$(\mathcal{V}_i, g, b)$ such that $\mathcal{S}(\sigma_{simp}) \sqsupseteq \mathcal{S}$. By Definition 6, the bad-gadget arena $BadGad_b(\mathcal{H}^F_k)$ has a unique internal state which is controlled by Player 0, and a unique transition: this transition has as source state in_k and as target state the internal state. Hence, being σ_{simp} non-losing, each play starting from in_i and consistent with σ_{simp} cannot visit states of $BadGad_b(\mathcal{H}^F_k)$. In particular, every exit play of σ_{simp} cannot visit states of $BadGad_b(\mathcal{H}^F_k)$. It follows that any memoryless strategy σ of Player 0 in $\mathcal{H}^F_i$ which coincides with σ_{simp} on the states which are not of the form (b, s) is non-losing and satisfies $\mathcal{S}(\sigma) = \mathcal{S}(\sigma_{simp})$. Hence, $\mathcal{S}$ is a relevant summary of $\mathcal{H}^F_i$, and the result follows. □

By Lemma 4, we can assume that $g(b) = \mathcal{S}_b$ for some relevant summary $\mathcal{S}_b$ of $\mathcal{H}^F_k$. Since $\mathcal{S}$ is a relevant summary of *Simplify*$(\mathcal{V}_i, g, b)$, there exist a non-losing memoryless strategy σ_{simp} of *Simplify*$(\mathcal{V}_i, g, b)$ such that $\mathcal{S}(\sigma_{simp}) \sqsupseteq \mathcal{S}$, and a non-losing memoryless strategy σ_b of $\mathcal{H}^F_k$ such that $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$. The strategies σ_{simp} and σ_b induce a memoryless strategy σ in $\mathcal{H}^F_i$, defined as follows for each non-terminal state s of $\mathcal{H}^F_i$ controlled by Player 0:

- if s is state of the form (b, s') such that s' is not an exit state of $\mathcal{H}_k^F$, then $\sigma((b, s')) = (b, \sigma_b(s'))$;
- otherwise, $\sigma(s) = \sigma_{simp}(s)$.

We demonstrate in Lemmata 5 and 6 below that σ is a non-losing strategy of $\mathcal{H}_i^F$ such that $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}(\sigma_{simp})$. Since $\mathcal{S}(\sigma_{simp}) \sqsupseteq \mathcal{S}$, this result together with Lemma 4 completes the proof of Theorem 7.

In the following, in order to simplify the notation, we use some shorthands: $\mathcal{H}_b$ is for $\mathcal{H}_k^F$, $\mathcal{H}^F$ is for $\mathcal{H}_i^F$, $\mathcal{H}_{simp}^F$ is for $\mathit{Simplify}(\mathcal{V}_i, g, b)$, $\mathit{Gad}(\mathcal{S}_b)$ is for $\mathit{Gad}(\mathcal{H}_k^F, \mathcal{S}_b)$, $\mathit{Gad}_b(\mathcal{S}_b)$ is for $\mathit{Gad}_b(\mathcal{H}_k^F, \mathcal{S}_b)$, in is for in_i, and in_b is for in_k. Recall that in the gadget arena $\mathit{Gad}_b(\mathcal{S}_b)$, there is exactly one strategy of Player 0. Thus, in the following, we identify $\mathit{Gad}_b(\mathcal{S}_b)$ with such a strategy.

Moreover, we exploit the following definitions:

- A state s of $\mathcal{H}_i^F$ is called *b-state* if it is of the form (b, s'), and is called *context-state* otherwise.
- A *b-segment* is a finite play ν of $\mathcal{H}^F$ consistent with σ visiting only b-states, starting from (b, in_b), and leading to a state of the form (b, s) for some exit state s of $\mathcal{H}_b$. Note that if we remove the b-component to every state occurring in ν, we obtain an exit play of σ_b.[3]

Lemma 5. *σ is a non-losing strategy of $\mathcal{H}^F$.*

Proof. Let π be a play of $\mathcal{H}^F$ from in which is consistent with the strategy σ such that π is not a finite play leading to an exit state (i.e., an exit node in Exit_i). We need to show that π is winning for Player 0. First, assume that π is finite. Then, since σ_b and σ_{simp} are non-losing, by construction of σ, it easily follows that π leads to a state controlled by Player 1. Hence, π is winning for Player 0. Now, assume that π is infinite. If there is a suffix of π which visits only context-states, then such a suffix π_s is consistent with the non-losing strategy σ_{simp}. Since $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$, for each b-segment leading to a state (b, s), there is an exit play of $\mathit{Gad}_b(\mathcal{S}_n)$ leading to (b, s). Hence, there exists an infinite play of $\mathcal{H}_{simp}^F$ starting from the initial state in which is consistent with σ_{simp} and has π_s as suffix. Hence, π is winning for Player 0. On the other hand, if π gets trapped in $\mathcal{H}_b$, there is a suffix of π corresponding to an infinite play of $\mathcal{H}_b$ starting

[3] Note that we have inverted the terminology with respect to the proof of Theorem 6.

from the initial state in_b which is consistent with the non-losing strategy σ_b. Thus, also in this case, π is winning for Player 0. It remains to consider the most difficult case, when π visits infinitely often b-segments. Hence, π can be factorized in the form:

$$\pi = \pi'_0 \cdot \pi_0 \cdot \pi'_1 \cdot \pi_1 \ldots$$

such that for all $n \geq 0$, π'_n visits only context-states, while π_n is a b-segment. Moreover, without loss of generality, we can assume that for each exit state (b, s) of $Gad_b(\mathcal{S}_b)$ occurring in π, there are infinitely many occurrences of (b, s) in π (in other terms, if for some $\ell \geq 0$, the b-segment π_ℓ leads to (b, s), then there are infinitely many n such that π_n lead to (b, s)). We assume that π is not winning for Player 0, and derive a contradiction. Hence, one of the following two conditions is fulfilled:

1. There are infinitely many unanswered requests along π.

2. There are infinitely many answered requests along π whose set of associated delays is unbounded.

First, we observe that for each exit state (b, s) of $Gad_b(\mathcal{S}_b)$ occurring in π, $\mathcal{S}_b(s) \neq C_o^{max}$; otherwise, by construction of $Gad_b(\mathcal{S}_b)$, we can construct from π an infinite play of $\mathcal{H}^F_{simp}$ from in which is consistent with σ_{simp} and is not winning for Player 0, a contradiction.

Let us examine case (1). Hence, there is a state s_o with odd color c_o such that for infinitely many positions n, n is an unanswered c_o-request along π associated with state s_o. We show that there exists a play consistent with σ_{simp} which is losing for Player 0. Fix a b-segment π_ℓ of π leading to an exit state (b, s), and let c_ℓ be the maximal color in π_ℓ. Since $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$ and for each exit state (b, s) of $Gad_b(\mathcal{S}_b)$ occurring in π, $\mathcal{S}_b(s) \neq C_o^{max}$, by Proposition 4, $value_{pr}(\sigma_b, s) \succeq_o value_{pr}(\sigma_S, s)$, where σ_S is the unique strategy of Player 0 in $Gad(\mathcal{S}_b)$. Hence, there exists an exit play ν_ℓ of $Gad_b(\mathcal{S}_b)$ leading to (b, s) whose maximal color c'_ℓ satisfies: $c'_\ell \preceq_0 c_\ell$. Note that if π_ℓ visits state s_0 (s_0 is a b-state), then c'_ℓ is an odd color such that $c'_\ell \geq c_o$. Otherwise, if c'_ℓ is even, then c'_ℓ cannot answer to the odd color c_o (i.e., $c'_\ell < c_o$). Hence, by replacing each b-segment π_ℓ in π with ν_ℓ, we obtain a play consistent with σ_{simp} which contains infinitely many unanswered requests, and the result follows.

Now, let us examine case (2). Recall that for each exit state (b, s) of $Gad_b(\mathcal{S}_b)$ occurring in π, $\mathcal{S}_b(s) \neq C_o^{max}$, hence $\mathcal{S}(\sigma_b)(s) \neq C_o^{max}$ as well (being $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$). Thus, by hypothesis there must exist an infinite set $\{[q_n, r_n]\}_{n\geq 0}$ of non-overlapping position-intervals such that for all $n \geq 0$;

- the cost of $\pi[q_n, r_n]$ is at least n;
- $q_n < q_{n+1}$;
- q_n (resp., r_n) is a request (resp., response) along π, and r_n is the smallest response answering the request q_n along π.
- there is no b-segment π_ℓ such that $\pi[q_n, r_n]$ is a sub-path of π_ℓ.

We show that it is possible to replace each b-segment π_ℓ in π, by an exit play of $Gad_b(\mathcal{S}_b)$ leading to the last state of π_ℓ in such a way that the resulting infinite sequence is a play consistent with σ_{simp} which is losing for Player 0. Let us consider a b-segment π_ℓ along π leading to a state (b, s). We distinguish the following cases:

1. π_ℓ does not overlap any segment $\pi[q_n, r_n]$: we replace π_ℓ with an arbitrary exit play of $Gad_b(\mathcal{S}_b)$ leading to (b, s) (note that since $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$, such an exit play exists).
2. There is a segment $\pi[q_n, r_n]$ for some $n \geq 0$ such that q_n corresponds to a request in π_ℓ with odd color c_o: since $\pi[q_n, r_n]$ is not a sub-path of π_ℓ, q_n is an unanswered request in π_ℓ. Since $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$ and $\mathcal{S}_b(s) \neq C_o^{max}$, by Proposition 4, the request cost of σ_b w.r.t. s is bounded and there is an exit play ν of $Gad_b(\mathcal{S}_b)$ leading to (b, s) having a c'_o-unanswered request, with $c'_o \geq c_o$. We replace the b-segment π_ℓ with ν.
3. there is a segment $\pi[q_n, r_n]$ for some $n \geq 1$ which strictly contains π_ℓ. Let c_e be the maximal even color in π_ℓ. Note that c_e is strictly smaller than the color associated with response r_n, Since $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$ and $\mathcal{S}_b(s) \neq C_o^{max}$, by Proposition 4(1), one of the following conditions is satisfied:
   (a) there exists an exit play ν of $Gad_b(\mathcal{S}_b)$ leading to (b, s) whose cost is at least the cost of π_ℓ and whose maximal even color in ν is at most c_e. We replace the b-segment π_ℓ with ν.
   (b) $\text{Cost}_e(\sigma_b, s, c_e) \in \mathbb{N} \setminus \{0\}$ and $\text{Cost}_e(\sigma_S, s, c_e) > 0$, where σ_S is the unique strategy of Player 0 in $Gad_b(\mathcal{S}_b)$: in this case, since $\text{Cost}_e(\sigma_b, s, c_e)$ is finite, it suffices to show that one can replace π_ℓ with an exit play ν of $Gad_b(\mathcal{S}_b)$ leading to (b, s) having a non-zero cost and such that the maximal even color in ν is at most c_e. Since, $\text{Cost}_e(\sigma_S, s, c_e) > 0$, the result trivially follows.
   (c) there exists an exit play ν of $Gad_b(\mathcal{S}_b)$ leading to (b, s) and an even color c'_e with

$c'_e \leq c_e + 2$ such that the c'_e-response cost in ν is at least the cost of $\pi[q_n, r_n]$. We replace the b-segment π_ℓ with ν.

4. There is a segment $\pi[q_n, r_n]$ for some $n \geq 1$ such that r_n corresponds to a response in π_ℓ. Let c_e (resp., c_o) be the color associated with the response r_n (resp., the request q_n) and c_e^{min} be the smallest even color answering to the request q_n. Note that $c_e \geq c_e^{\mathsf{min}} > 0$. Moreover, let Π_s be the set of b-segments $\pi_{\ell'}$ leading to state (b, s) satisfying the same condition as π_ℓ, i.e. such that there is a segment $\pi[q_{n'}, r_{n'}]$ so that $r_{n'}$ corresponds to a response in $\pi_{\ell'}$ and c_e (resp., c_o) is the color associated with the response $r_{n'}$ (resp., the request $r_{n'}$). Let $\text{Cost}_{c_e} := \{\text{ResCost}(\pi_{\ell'}, c_e) \mid \pi_{\ell'} \in \Pi_s\}$, i.e., the set of c_e-response costs of the b-segments in Π_s. If Cost_{c_e} is finite, we can safely replace the b-segment π_ℓ with an arbitrary exit play of $Gad_b(\mathcal{S}_b)$ leading to (b, s). Now, assume that $\text{Cost}_{c_e} = \infty$. Since the segments $\pi[q_{n'}, r_{n'}]$ cannot be sub-paths of the b-segments along π, for each $\pi_{\ell'} \in \Pi_s$, the smaller response k in $\pi_{\ell'}$ answering the odd color c_o has color c_e. Thus, since $\text{Cost}_{c_e} = \infty$, it follows that $\text{ResCost}_e(\sigma_b, s, c_e^{\mathsf{min}}) = \infty$. Being $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$, $\mathcal{S}_b(s) \neq C_o^{max}$, and $c_e^{\mathsf{min}} > 0$, by Proposition 4(2), one of the following two conditions is satisfied:
   (a) There exists an exit play ν of $Gad_b(\mathcal{S}_b)$ leading to (b, s) and an even color c'_e with $c'_e \leq c_e^{\mathsf{min}}$ such that the c'_e-response cost in ν is at least the cost of $\pi[q_n, r_n]$. We replace the b-segment π_ℓ with ν.
   (b) there exists an exit play ν of $Gad_b(\mathcal{S}_b)$ leading to (b, s) whose cost is at least the cost of $\pi[q_n, r_n]$ and whose maximal even color c_e^S satisfies $c_e^S < c_e^{\mathsf{min}}$. Since c_e^S does not answer to request q_n, we can safely replace π_ℓ with ν.

Thus, we obtain a play consistent with σ_{simp} which is losing for Player 0. This concludes the proof of Lemma 5. □

It remains to show that $\mathcal{S}(\sigma) \sqsupseteq \mathcal{S}(\sigma_{simp})$.

Lemma 6. *Let ex be an exit state of $\mathcal{H}_F$. Then $\mathcal{S}(\sigma)(ex) \sqsupseteq \mathcal{S}(\sigma_{simp})(ex)$.*

Proof. If either $\mathcal{S}(\sigma_{simp}) = C_o^{max}$ or $\mathcal{S}(\sigma)(ex) = \vdash$, the result trivially follows. Moreover, note that since $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$, for each b-segment ρ, there is an exit play of $Gad_b(\mathcal{S}_b)$ leading to the last state of ρ. Hence, if $\mathcal{S}(\sigma_{simp})(ex) = \vdash$, then $\mathcal{S}(\sigma)(ex) = \vdash$ as well. Hence, we can assume that $\mathcal{S}(\sigma_{simp})(ex) \notin \{\vdash, C_o^{max}\}$ and $\mathcal{S}(\sigma)(ex) \neq \vdash$. First, we show that $\mathcal{S}(\sigma)(ex) \neq C_o^{max}$ as well.

We assume on the contrary that $\mathcal{S}(\sigma)(ex) = C_o^{max}$, and derive a contradiction. We exploit the fact that since $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$ and $\mathcal{S}(\sigma_{simp})(ex) \neq C_o^{max}$, for each exit play ν of σ leading to ex and for each exit state (b, s) of $Gad(\mathcal{S}_b)$ occurring in ν, it holds that $\mathcal{S}_b(s) \neq C_o^{max}$ (hence, $\mathcal{S}(\sigma_b)(s) \neq C_o^{max}$ as well). Since $\mathcal{S}(\sigma)(ex) = C_o^{max}$, one of the following two conditions is fulfilled.

1. *either* there is an exit play ν of σ leading to ex having a C_o^{max}-request,
2. *or* there exists an infinite set of exit plays $\{\nu_n\}_{n\geq 0}$ of σ leading to ex and a infinite set $\{[q_n, r_n]\}_{n\geq 0}$ of position-intervals such that for all $n \geq 0$;
   - either (i) q_n is an unanswered request in ν_n and r_n is the last position of ν_n, or (ii) q_n (resp., r_n) is a request (resp., response) along ν_n, and r_n is the smallest response answering the request q_n along ν_n.
   - the cost of $\nu_n[q_n, r_n]$ is at least n;
   - there is no b-segment ν_b in ν_n such that $\nu_n[q_n, r_n]$ is a sub-path of ν_b.

First, let us examine case (1). Since for each exit state (b, s) of $Gad_b(\mathcal{S}_b)$ along ν it holds that $\mathcal{S}(\sigma_b)(s) \neq C_o^{max}$, we deduce that the C_o^{max}-request in ν is associated with a context-state. It follows that there exists an exit play of σ_{simp} leading to ex having a C_o^{max}-request, which contradicts our assumption that $\mathcal{S}(\sigma_{simp})(ex) \neq C_o^{max}$. Now, let us examine case (2). Since for each exit state (b, s) of $Gad_b(\mathcal{S}_b)$ along ν_n, $\mathcal{S}_b(s) \neq C_o^{max}$, by proceeding as in the part of the proof of Lemma 5 concerning the treatment of Condition (2), we deduce the following.

- For each $n \geq 0$, it is possible to replace each b-segment ρ in ν_n by an exit play of $Gad_b(\mathcal{S}_b)$ leading to the last state of ρ in such a way that the resulting sequence ν'_n is an exit play of σ_{simp} leading to ex, and the following holds: there are two positions q'_n and r'_n in ν'_n, and a bound $m_n \geq 0$ such that
  - the cost of $\nu'_n[q'_n, r'_n]$ is at least m_n;
  - either (i) q'_n is an unanswered request in ν'_n and r'_n is the last position of ν'_n, or (ii) q'_n (resp., r'_n) is a request (resp., response) along ν'_n, and r'_n is the smallest response answering the request q'_n along ν'_n.

  Moreover, and, more importantly, the set of the bounds m_n with $n \geq 0$ is infinite.

Hence, we obtain that $\mathcal{S}(\sigma_{simp})(ex) = C_o^{max}$, which contradicts the hypothesis. Therefore, in the rest of the proof, we can assume that $\mathcal{S}(\sigma)(ex) \notin \{\vdash, C_o^{max}\}$ and $\mathcal{S}(\sigma_{simp})(ex) \notin \{\vdash, C_o^{max}\}$. Then, the result directly follows from Claims 1–6 below.

Claim 1. $value_{\mathrm{Cost}}(\sigma, ex) \succeq_b value_{\mathrm{Cost}}(\sigma_{simp}, ex)$.

Proof of Claim 1. Since $\mathcal{S}(\sigma_{simp})(ex) \neq C_o^{max}$ and $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$, for each exit play ν of σ leading to ex, and for each exit state (b, s) of $Gad_b(\mathcal{S})$ along ν, condition $\mathcal{S}_b(s) \neq C_o^{max}$ holds. By Proposition 4, for such exit state (b, s) of $Gad_b(\mathcal{S})$, it holds that $value_{\mathrm{Cost}}(\sigma_b, s) \succeq_b value_{\mathrm{Cost}}(\sigma_S, s)$, where σ_S is the unique strategy of Player 0 in $Gad(\mathcal{S}_b)$. Hence, the result easily follows. □

Claim 2. $value_{pr}(\sigma, ex) \succeq_0 value_{pr}(\sigma_{simp}, ex)$.

Proof of Claim 2. Recall that since $\mathcal{S}(\sigma_{simp})(ex) \neq C_o^{max}$ and $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$, for each exit play ν of σ leading to ex, there is no exit state (b, s) of $Gad_b(\mathcal{S})$ along ν such that $\mathcal{S}_b(s) = C_o^{max}$. By Proposition 4, for such exit state (b, s) of $Gad_b(\mathcal{S})$, it holds that $value_{pr}(\sigma_b, s) \succeq_0 value_{pr}(\sigma_S, s)$, where σ_S is the unique strategy of Player 0 in $Gad(\mathcal{S}_b)$. Hence, for each b-segment ρ leading to (b, s) such that $\mathcal{S}_b(s) \neq C_o^{max}$ and whose maximal color is c, there is an exit play of $Gad_b(\mathcal{S}_b)$ leading to (b, s) whose maximal color c' satisfies $c' \preceq_0 c$. Therefore, the result trivially follows. □

Claim 3. $value_o(\sigma, ex) \leq value_o(\sigma_{simp}, ex)$.

Proof of Claim 3. If $value_o(\sigma, ex) = 0$, the result is obvious. Otherwise, $value_o(\sigma, ex) = c_o$ for some odd color c_o. We assume that $c_o > value_o(\sigma_{simp}, ex)$, and derive a contradiction. Since $value_o(\sigma, ex) = c_o$, there must exist an exit play ν of σ leading to ex having an unanswered c_o-request. Assume that the maximal c_o-request of ν is inside a b-segment ρ_{c_o} of ν (the other case, where the maximal c_o-request of ν is associated with a context-state being simpler). Since $\mathcal{S}(\sigma_{simp})(ex) \neq C_o^{max}$ and $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$, there is no exit state (b, s) of $Gad_b(\mathcal{S})$ along ν such that $\mathcal{S}_b(s) = C_o^{max}$. We show how to construct, starting from ν, an exit play ν' of σ_{simp} leading to ex having an unanswered c'_o-request with $c'_o \geq c_o$. Let us consider a b-segment ρ along ν leading to a state (b, s). We distinguish the following cases:

- ρ strictly precedes the b-segment ρ_{c_o}: we replace ρ with any exit play of $Gad_b(\mathcal{S}_b)$ leading to (b, s) (note that such an exit play exists).
- ρ corresponds to ρ_{c_o}: we replace ρ with any exit play of $Gad_b(\mathcal{S}_b)$ leading to (b, s) having an unanswered c'_o-request for some $c'_o \geq c_o$. Since $\mathcal{S}_b(s) \neq C_o^{max}$ and $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$, the existence of such an exit play in ensured by Proposition 4.
- ρ strictly follows the b-segment ρ_{c_o}: let c_e be the maximal even color of ρ. Since $\mathcal{S}_b(s) \neq$

C_o^{max} and $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$, by Proposition 4(1): either (i) there exists an exit play ν of $Gad_b(\mathcal{S}_b)$ leading to s whose maximal even color is c_e, or (ii) there is an even color $c'_e \in C_e$ such that $c'_e \leq c_e + 2$ and $\text{ResCost}(\sigma_S, s, c'_e) = \infty$, where σ_S is the unique strategy of Player 0 in $Gad_b(\mathcal{S}_b)$. Condition (ii) cannot hold. Otherwise, we deduce that the request-cost of σ_{simp} with respect to ex is ∞, which would contradict our assumption that $\mathcal{S}(\sigma_{simp})(ex) \neq C_o^{max}$. Hence, we replace ρ with any exit play of $Gad_b(\mathcal{S}_b)$ satisfying Condition (i).

Thus, we obtain an exit play ν' of σ_{simp} leading to ex having an unanswered c'_o-request for some $c'_o \geq c_o$. This is a contradiction since we have assumed that $c_o > value_o(\sigma_{simp}, ex)$. □

Claim 4. $value_e^L(\sigma, ex) \geq value_e^L(\sigma_{simp}, ex)$.

Proof of Claim 4. We assume that $value_e^L(\sigma, ex) < value_e^L(\sigma_{simp}, ex)$, and derive a contradiction. Let $c_e^L = value_e^L(\sigma_{simp}, ex)$. By hypothesis, one of the following two conditions holds.

1. either there is an exit play ν of σ leading to ex such that the maximal even color in ν is strictly smaller than c_e^L,
2. or there is a family $\{\nu_n\}_{n \geq 0}$ of exit plays of σ leading to ex and an infinite sequence of natural numbers $\{r_n\}_{n \geq 0}$ such that: (i) r_n is the smaller response in ν_n associated with an even color $c_e \geq c_e^L$, and (ii) the response-cost of r_n in ν_n is at least n.

First, let us examine case (1). By hypothesis, for each b-segment ν_b along ν, the maximal even color in ν is strictly smaller than c_e^L. Moreover, since $\mathcal{S}(\sigma_{simp})(ex) \neq C_o^{max}$ and $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$, we deduce that for the last state (b, s) of such a b-segment ν_b, $\mathcal{S}_b(s) \neq C_o^{max}$. Thus, by applying Proposition 4, we deduce that for each exit state (b, s) of $Gad_b(\mathcal{S}_b)$ occurring in ν, $value_e^L(\sigma_S, s) < c_e^L$, where σ_S is the unique strategy of Player 0 in $Gad_b(\mathcal{S}_b)$. Hence, we easily obtain that $value_e^L(\sigma_{simp}, ex) < c_e^L$, which is a contradiction.

Now, let us examine case (2). Fix $n \geq 0$. We show that it is possible to replace each b-segment ν_b in ν_n by an exit play of $Gad_b(\mathcal{S}_b)$ leading to the last state of ν_b in such a way that the resulting sequence ν'_n is an exit play of σ_{simp} leading to ex , and the following holds: either (i) for some $n \geq 0$, the maximal even color in ν'_n is strictly smaller than c_e^L, or (ii) for infinitely many n, the response-cost of any c_e-response in ν'_n with $c_e \geq c_e^L$ is m_n, and the set of such costs m_n is infinite. Since for each exit state (b, s) of $Gad_b(\mathcal{S}_b)$ occurring in ν_n, $\mathcal{S}_b(s) \neq C_o^{max}$, we proceed as in the part of the proof of Lemma 5 concerning the cases 3 and 4 in the treatment of Condition (2). □

Claim 5. If $value_e^M(\sigma, ex) \neq \bot$, then the following holds:

- *either* $value_e^M(\sigma_{simp}, ex) \neq \bot$ and $value_e^M(\sigma, ex) \geq value_e^M(\sigma_{simp}, ex)$,
- *or* $value_e^M(\sigma_{simp}, ex) = \bot$ and $value_e^M(\sigma, ex) \geq value_e^R(\sigma_{simp}, ex)$.

Proof of Claim 5. Let $value_e^M(\sigma, ex) = c_e \in C_e$. Hence, $\mathrm{Cost}_e(\sigma, ex, c_e) \in \mathbb{N} \setminus \{0\}$ and there is an exit play of σ leading to ex having a non-null cost and whose maximal even color is c_e. Let ν be any exit play of σ leading to ex whose maximal even color is c_e. Since $\mathcal{S}(\sigma_{simp})(ex) \neq C_o^{max}$ and $\mathcal{S}(\sigma_b) \sqsupseteq \mathcal{S}_b$, for each b-segment ρ occurring in ν leading to a state (b, s), $\mathcal{S}_b(s) \neq C_o^{max}$. Moreover, since $\mathrm{Cost}_e(\sigma, ex, c_e)$ is finite, for each b-segment ρ occurring in ν leading to a state (b, s), it holds that $\mathrm{Cost}_e(\sigma_b, ex, c'_e) \neq \infty$, where $c'_e \leq c_e$ is the maximal even color in ρ. By applying Proposition 4, for such a b-segment ρ, there exists an exit play ρ' of $Gad_b(\mathcal{S}_b)$ leading to the last state of ρ such that $\mathrm{Cost}(\rho') > 0$ if $\mathrm{Cost}(\rho) > 0$, and the maximal even color in ρ' is at most the maximal even color in ρ. Hence, there exists an exit play ν' of σ_{simp} leading to ex whose maximal even color is at most c_e and such that $\mathrm{Cost}(\nu') > 0$ if $\mathrm{Cost}(\nu) > 0$. It follows that $\mathrm{Cost}_e(\sigma_{simp}, ex, c_e) \in (\mathbb{N} \setminus \{0\}) \cup \{\infty\}$. If $value_e^M(\sigma_{simp}, ex) \neq \bot$ then $value_e^M(\sigma_{simp}, ex) \leq c_e$. Otherwise, we have that $c_e \geq value_e^R(\sigma_{simp}, ex)$, and the result follows. □

Claim 6. $value_e^R(\sigma, ex) \geq value_e^R(\sigma_{simp}, ex)$.

Proof of Claim 6. Let $c_e^R = value_e^R(\sigma_{simp}, ex)$ and $\mathcal{K}_{simp}^F$ (resp., $\mathcal{K}^F$) be the partial cost-parity arena obtained from $\mathcal{H}_{simp}^F$ (resp., $\mathcal{H}^F$) by replacing the color of the exit state ex with c_e^R. Evidently, the even-*left* value of σ_{simp} w.r.t. ex in $\mathcal{K}_{simp}^F$ is c_e^R. Thus, it suffices to show that the even-left value of strategy σ w.r.t. ex in $\mathcal{K}^F$ is at least the even-left value of σ_{simp} w.r.t. ex in $\mathcal{K}_{simp}^F$. The proof of this is similar to the proof of Claim 4. □

This concludes the proof of Lemma 6. □

8. Conclusion

Cost-parity games represent a powerful machinery for the verification of temporal requirements that are bounded in time. As in many settings, the representation of systems by means of cost-parity games is affected by an exponential blow-up in the size of the resulting game. To overcome this, many techniques exploiting system regularities have been successfully applied. Among them, hierarchical systems deserve a special mention. In this paper, we have introduced and investigated the problem of solving cost-parity games over hierarchical FSMs, showing that the problem is PSPACE-complete, thus not harder than solving parity games over hierarchical models. As future work, we aim to adapt the proposed approach to all the other winning bounded conditions introduced in [15]. Moreover, it would be interesting to investigate cost-parity conditions over concurrent game structures, the last one being a suitable formalism for modelling strategic environments where there is simultaneous interaction between multiple players [9, 28]. Other relevant research directions include the study of cost-parity games in the imperfect information setting as well as for infinite-state systems.

References

[1] E. Clarke, E. Emerson, Design and Synthesis of Synchronization Skeletons Using Branching-Time Temporal Logic., in: Proc. of LP'81, LNCS 131, Springer, 1981, pp. 52–71.

[2] E. Clarke, O. Grumberg, D. Peled, Model Checking., MIT Press, 2002.

[3] O. Kupferman, M. Vardi, P. Wolper, An Automata Theoretic Approach to Branching-Time Model Checking., J. ACM 47 (2) (2000) 312–360.

[4] J. Queille, J. Sifakis, Specification and Verification of Concurrent Programs in Cesar., in: Proc. of SP'81, LNCS 137, Springer, 1981, pp. 337–351.

[5] B. Aminof, O. Kupferman, A. Murano, Improved Model Checking of Hierarchical Systems, Inf. Comput. 210 (2012) 68–86.

[6] B. Aminof, F. Mogavero, A. Murano, Synthesis of Hierarchical Systems., in: Proc. of FACS'11, LNCS 7253, Springer, 2011, pp. 42–60.

[7] B. Aminof, F. Mogavero, A. Murano, Synthesis of Hierarchical Systems., Sci. Comput. Program. 83 (2014) 56–79.

[8] O. Kupferman, M. Vardi, P. Wolper, Module Checking., Inf. Comput. 164 (2) (2001) 322–344.

[9] V. Malvone, A. Murano, L. Sorrentino, Concurrent Multi-Player Parity Games, in: Proc. of AAMAS'16, 2016, pp. 689–697.

[10] E. Emerson, C. Jutla, The Complexity of Tree Automata and Logics of Programs., SIAM J. Comput. 29 (1) (1999) 132–158.

[11] D. Kozen, Results on the Propositional μ-Calculus., Theor. Comput. Sci. 27 (3) (1983) 333–354.

[12] E. Emerson, C. Jutla, Tree Automata, μ-Calculus and Determinacy, in: Proc. of FOCS'91, 1991, pp. 368–377.

[13] N. Fijalkow, M. Zimmermann, Parity and Streett Games with Costs, Logical Methods in Computer Science 10 (2).

[14] F. Mogavero, A. Murano, L. Sorrentino, On Promptness in Parity Games., in: Proc. of LPAR'13, LNCS 8312, Springer, 2013, pp. 601–618.

[15] F. Mogavero, A. Murano, L. Sorrentino, On Promptness in Parity Games, Fundam. Inform. 139 (3) (2015) 277–305.

[16] M. Jurdzinski, M. Paterson, U. Zwick, A Deterministic Subexponential Algorithm for Solving Parity Games, SIAM J. Comput. 38 (4) (2008) 1519–1532.

[17] S. Almagor, Y. Hirshfeld, O. Kupferman, Promptness in omega-Regular Automata, in: Proc. of ATVA'10, LNCS 7388, 2010, pp. 22–36.

[18] O. Kupferman, N. Piterman, M. Y. Vardi, From Liveness to Promptness, Formal Methods in System Design 34 (2) (2009) 83–103.

[19] K. Chatterjee, T. A. Henzinger, F. Horn, Finitary winning in ω-regular games, ACM Trans. Comput. Logic 11 (1) (2009) 1:1–1:27.

[20] A. Pnueli, The Temporal Logic of Programs., in: Proc. of FOCS'77, IEEE Computer Society, 1977, pp. 46–57.

[21] K. Chatterjee, T. A. Henzinger, M. Jurdzinski, Mean-Payoff Parity Games, in: Proc. of LICS'05, 2005, pp. 178–187.

[22] K. Chatterjee, L. Doyen, Energy Parity Games, Theor. Comput. Sci. 458 (2012) 49–60.

[23] R. Alur, S. Kannan, M. Yannakakis, Communicating hierarchical state machines, in: Proc. of ICALP'99, LNCS 1644, Springer, 1999, pp. 169–178.

[24] R. Alur, M. Yannakakis, Model checking of hierarchical state machines, ACM Trans. Program. Lang. Syst. 23 (3) (2001) 273–303.

[25] A. Murano, M. Napoli, M. Parente, Program complexity in hierarchical module checking, in: Proc. of LPAR'08, LNCS 5330, Springer, 2008, pp. 318–332.

[26] S. Göller, M. Lohrey, Fixpoint logics on hierarchical structures, in: Proc. of FSTTCS'05, LNCS 3821, Springer, 2005, pp. 483–494.

[27] R. Alur, M. Benedikt, K. Etessami, P. Godefroid, T. W. Reps, M. Yannakakis, Analysis of recursive state machines, TOPLAS 27 (4) (2005) 786–818.

[28] B. Aminof, A. Murano, S. Rubin, F. Zuleger, Prompt alternating-time epistemic logics, in: Proc. of KR'16, AAAI Press, 2016, pp. 258–267.